Knowledge Creation and Collective Memory in the Era of Web 2.0 and Network Non-Neutrality

McIver Jr., William J.

Publisher's version / Version de l'éditeur:

Community Informatics Research Network (CIRN), 2006 [Proceedings], 2007

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

NRC·CNRC

*Knowledge Creation and Collective Memory in the Era of Web 2.0 and Network Non-Neutrality **

W. McIver, Jr.
2007

* published through the Community Informatics Research Network (CIRN) 2007. NRC 48815

Canada

KNOWLEDGE CREATION AND COLLECTIVE MEMORY IN THE ERA OF WEB 2.0 AND NETWORK NON-NEUTRALITY

William J. McIver, Jr.

National Research Council Canada, Institute for Information Technology
Fredericton, New Brunswick
Bill.McIver@nrc.gc.ca

Abstract

This paper examines knowledge creation and ownership issues raised by the nexus of Web 2.0 technologies and business models, changes to network neutrality policy and operation, and the advent of next generation Internet architectures.

Context

One constant in the production of human knowledge is the evolution of modalities of knowledge representation and knowledge sharing. At certain transitions, existing societal norms regarding these modalities have been challenged at fundamental levels. Notable examples can be seen in the area of communication. For example, d'Arcy (1969) derived the influential concept of a *right to communicate* in the 1960s out of a recognition that newer technologies – satellite-based communication at that time – had raised issues that were not easily addressable by existing norms.[1]

The Internet has itself been a long, steady fount of changing perspectives not only on communication issues, but also the relatively less well-examined process of managing externalized knowledge. A nexus of recent engineering developments and changing business practices, known commonly as *Web 2.0*, is producing socio-technical shifts in the loci and control of both individual and collective representations of knowledge. New web application architectures are encouraging the caching of ever increasing amounts of personal information in external spaces that are only nominally controlled by the users that create it. Corresponding developments – again, both in engineering and business – are leading users to existences in post-Internet *milieux,* in or through which attempts are being made to construct a tiered model in the regulation of access to Web 2.0-constructed knowledge.

[1] A right to communicate has not been accepted in international law; however, d'Arcy's introduction of the concept resulted in considerable efforts in the 1970s and 1980s being made in Unesco and other international bodies to examine the issue.

New Memory Augmentation Frameworks

In examining a context defined by Web 2.0 and post-Internet architectures, it is instructive to examine memory and knowledge from a computational point of view, in particular the representation and processing of memory and the communication of its contents.

From a computational perspective, the fundamental operations on memory can be classified into several categories, including:

(*Type I*) the storage of information into memory,

(*Type II*) the search for and retrieval of information in memory relevant to the task at hand, which often depends on the creation of indices.

A distinct category of operations related to *Type I* is that which is necessary if information retrieved from memory is to be made useful:

(*Aux Type I*) operations that facilitate the communication of information to appropriate agents for processing.

Here agents are understood to include people and computational processes. Haseman, Nazareth, and Souren (2005, p. 592) define memory in the context of decision support technology as “the facility to retain, recall, and manipulate past or present”.

The relationships between these memory and communications operations and individual and community knowledge can be seen in how they support the cyclic process of social communication, due to Ó Siochrú (2005, p. 33), which is depicted in figure 1. This model is the most recent effort in the direction set by d'Arcy with respect to technology and communication rights and is meant to show the roles and risk points for communication in the process of knowledge creation. Nodes (1) through (5) depict different phases of knowledge creation, management, and use. At a fundamental level, the transitions from each node are facilitated in our computational perspective of memory by the storage and indexing of data within each node and the communication of that data to subsequent nodes.

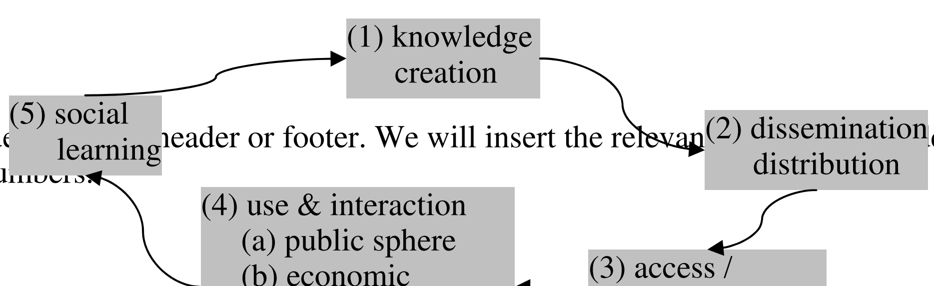

Figure 1. The cyclic process of social communication.

Thus, all three categories of operations are necessary in community contexts – at biological, social, computational, and other levels – for deriving new knowledge. The information systems concept of *group memory* (or *collective memory*), wherein multiple agents are able to manipulate the same information or contribute individual items to a collection, adds complexity to these operations. In a broad study of the psychology of groups, Wegner (1986, p. 186) defined a specific class of group memory -- called *transactive memory* -- that has been influential in decision support systems research as "a set of individual memory systems in combination with the communication that takes place between individuals".

Again from a computational perspective, coordination mechanisms are required to prevent multiple agents from overwriting information in a group memory system. To prevent overwriting, it may be necessary to have a *versioning* mechanism that can maintain the multiple versions of the same information that different agents may wish to contribute. Transitively, if multiple versions exist, then depending on the use of the information, mechanisms may be necessary to make certain that only the appropriate versions of information are seen by certain agents.

This computational perspective is necessary for contrasting this new era of hyper-disembodied knowledge from earlier forms of memory augmentation that were devised to represent and share human knowledge. Knowledge represented in these earlier approaches was also fundamentally disembodied; however, not in as many dimensions nor to as great a degree as this new era.

Early memory augmentation technologies largely involved the use of physical objects for recording discrete units of information.[1] Writing took this further by creating an external storage system for natural language, thereby reducing dependence on fault-prone biological memory. The development of alphabetization and book and citation indexing facilitated more efficient searches of and for written information.[2] Printing took this further by allowing for relatively flawless replication of writing. Postal communications, telegraphy, telephony, FAX, and many other subsequent technologies introduced new possibilities for the communication of information to distant agents. The evolution described here was somewhat analog in nature, with each innovation seen, more or less, as a continuous extension of the previous ones.

In contrast, the advent of electronic digital computing and, relatively soon thereafter, data communication networks were evolutionary steps of a transilient nature when seen from the perspective of human memory. As is well known, digital computers and networks -- unlike earlier memory augmention schemes -- enable the more rapid, voluminous, and flawless recording, replication, and transmission of representations of human and community knowledge. Additionally, computation has made possible the previously intractable tasks of searching and linking vast amounts of information.

The extent to which knowledge has been externalized in the preceding technological epoch was mediated by the speed and cost of printing, the limits of human productivity in the use of physical media, and the limits of transportation. This socio-technical dynamic has now reached a far more advanced state where the distributed sharing of personal knowledge is becoming ever more pervasive, with new technologies making possible the recording of ever greater levels of detail and the storage, replication, and distribution of ever greater volumes of data. Most importantly, this has enabled an increasing loss of control and dependency creation in the management of personal and community knowledge. The means of production here are the technologies and business models of the so-called "Web 2.0" era. This era is leading in major part to the creation of services that enable the externalization of individual and community knowledge through digital and distributed implementations of memory operations *Type I* and *Type II*.

[1]The Ishango Bone from the geographic area of what is now the Democratic Republic of the Congo has been dated at least 9,000 B.C.E. Is an example of an early object used to record mathematical quantities for faster calculation. See J. de Heinzelin, Ishango, Scientific American, 206:6 (June 1962) 105--116.

[2]Witty (1965) cited research that found the use of alphabetization from Hellenistic Egypt for the listing of taxpayers. This period extended from 323 BCE with the death of Alexander the Great to 146 BCE when Rome annexed Greece. According to Weinberg (1997) book citation indexes were created in Europe at least as early as 1522.

A New “Mind-Body” Problem

Web 2.0 refers simultaneously to a highly useful collection of new technical approaches to managing and delivering web content to users and an evolving set of business models that depend on these techniques. The term was coined out of efforts to identify the characteristics of those companies that survived the collapse in late 2001 of the dot-com industry (O'Reilly, 2005). O'Reilly (section 2, para. 24), perhaps the most influential proponent of the Web 2.0 concept, stated:

> If an essential part of Web 2.0 is harnessing collective intelligence, turning the web into a kind of global brain, the blogosphere is the equivalent of constant mental chatter in the forebrain, the voice we hear in all of our heads.

The blogosphere and other segments of the Web 2.0 *milieu*, such as web-based e-mail services, Internet chat facilities, and free web content hosting services, all have a critical dependence on the externalization of personal or community knowledge. That is, they exist only through individuals or groups of people contributing and managing personal data via their web services. O'Reilly (section 1, para. 21) stresses this in elaborating on a central tenet of Web 2.0: “the service automatically gets better the more people use it.”

Arguably the term “Web 2.0” is backed by more hype than it deserves, which would not be a new phenomenon in discussing ICT. In its original sense, however, Web 2.0 does represent a paradigmatic shift in two dimensions: business practices and approaches to implementing web application software. The latter dimension has been the development of the AJAX (Asynchronous JavaScript and XML) approach to achieving enhanced interactivity in web client interfaces. Two factors taken together argue for AJAX as a paradigmatic shift. First, it has enabled the implementation of user interface interactivity at a level necessary for web applications to emulate the behaviour of traditional client-side software applications, such as word processors and spreadsheet editors. Second, the life cycle costs of implementing, managing, distributing, and evolving AJAX-based applications can be far lower than those of traditional application development projects. Indeed, many so-called Web 2.0 companies now refer to the “eternal Beta” in the sense that this new paradigm makes it practical and relatively safe for them, from a business model perspective, to maintain their web applications in a beta or test version throughout its life cycle. That is, as bugs are fixed and functionality is changed or extended, new versions of applications of this type can be deployed almost instantaneously – via the Web. Software revision cycles are reportedly much faster as well. Bug fixes, functionality extensions, and distribution are all more costly and time-consuming in the context of traditional application development. This differential in the cost of software development has itself enabled or reinforced the business dimension of this paradigmatic shift: it has allowed more flexibility for many companies to operate on revenue models that allow “free” access to their applications, which has had the transitive effect of attracting more users to their products and services.

Even before this era of pervasive connectivity and massive sharing of personal knowledge, it was not farfetched to say that legions of users of personal computers, personal digital assistants, and smart phones came to see their devices as extensions of their memory: places to keep and process knowledge in the form of notes, appointments, and addresses. A modest amount of knowledge-sharing could be performed, to synchronize with other devices -- usually owned by the same person – or to selectively send the information to others. In contrast, massive numbers of users now regularly externalize (i.e. post) detailed personal knowledge in the form of text, pictures, video, and audio on web services owned and controlled by someone else. Technorati.com, a highly recognized monitor of the “blogosphere,” stated as of May 26, 2006 that they were monitoring 41.2 million blogging sites having 2.4 billion links.[1]

This trend is not entirely new, however. Even before *Web 1.0*, users of the network-based communication and knowledge sharing systems began the process of ceding control of information they created to external organizations. Usenet news, a system of bulletin boards started in 1979, is a primary example. Archives of postings to Usenet, which is still in operation, extend back to 1981 and are maintained by such organizations as Google – a Web 2.0 era company.

If these new technologies can be seen as extending individual and community knowledge, examining the loci of control and ownership of Web 2.0-managed data leads to new analogs of the mind-body problem. O'Reilly (section 3, para. 2) eventually asks: “Who owns the data?”[2] This was also the subject of a high-level discussion at the 2006 W3C Technical Plenary and WG Week, the main forum for working groups and interest groups within the World-Wide Web Consortium (W3C, 2006).

---

[1] See < http://www.technorati.com/about/ >.
[2] O'Reilly's discourse up to and following this question in his article can reasonably be interpreted as emanating less from a focus on individual and community rights and more from an effort to recognize business strategies based on ownership of specialized data. He does address user privacy and copyright concerns, and predicts the rise of a “Free Data” movement.

Access to one's own information stored in many of these services may not be proprietary in the sense that open technologies like XHTML and XML may be used to represent information and HTTP or RSS may be used to retrieve the information. Users are not, however, guaranteed that the knowledge they store in these services is easily transportable to other services or to their own private computing environments. Other major issues related to security and privacy are raised as well.[1]

The business value in Web 2.0 for companies that support knowledge sharing services comes in a major part from creating a dependency on their particular implementations of transactive memory. Unlike many common software applications, one cannot easily migrate their information to other services. For example, Blogger.com, one of the most popular blogging sites in the world, states:

> Blogger does not have an export or download function. However, you can use the following instructions to create a single file with all your posts which you may publish and then copy to your own computer for use as desired. [2]

This is followed by a 10 step process that requires users to author XML fragments and change a number of configuration settings on their blog account. Arguably such knowledge management arrangements are not geared toward portability of data for the average user. It should be noted that well-known software systems, such as MovableType, used to support blogging services provide system administrators with tools to conveniently migrate data from one server to another.[3] These are not, however, user-level facilities and more importantly they are not necessarily geared toward interoperability between different blogging systems in the same way that HTML files are readable using any Web browser.

As Web 2.0 services are enabling dependency creation in the externalization of individual and collective knowledge, engineering and business dynamics may force major changes in access to that knowledge through aspects of the concept of *network neutrality*. A potential move to non-neutral (or content-sensitive) networks would relate to *Aux Type I* operations (data communications) in the context of both the implementation and running of major parts of the Internet.

---

[1] Security and privacy issues in this context are beyond the scope of this particular paper.
[2] See <http://help.blogger.com/bin/answer.py?answer=130&query=backup&topic=0&type=f >.
[3] See <http://www.sixapart.com/movabletype/ >.

Network Neutrality

While Web 2.0 technologies and services have enabled the loss of control and dependency creation in the management of personal and community knowledge, current developments in networking have the potential to limit access to knowledge via commodity networks, through discrimination in how data are communicated (i.e. *Aux Type I* memory operations).

Network neutrality issues exist at multiple levels. On one level, data communications networks can be viewed from both technical and economic perspectives in terms of how they treat the content of the messages they transmit. On another level, the interconnection of multiple data communications networks is at issue. These interconnections can also be viewed from both technical and economic perspectives, but in this case the concern is how the networks agree or disagree to cooperate in transferring data between users or computers when it must cross network boundaries -- making use of the interconnections.

Networks that make up the Internet have generally operated in a content-neutral manner. This practice derives in a major part from an adherence to the *end-to-end principle* in the design of computer systems and, in this case, the placement of functionality relative to data communications subsystems (Saltzer, Reed, & Clark, 1984).

The end-to-end argument addresses the software engineering dilemma of where to place data communications services, such as message ordering, filtering or error detection, when they can either be implemented in a network or be made the responsibility of the systems that use the network. The principle holds that as long as sufficient basic data communications functionality exists at a lower layer, advanced functionality can and should be added later at higher layers. Further, such functionality can be implemented in an optimal manner in this way since it can be tailored to the particular application for which it is being created. Saltzer, Reed, and Clark showed that this approach is necessary to achieve optimal modularity and economy in developing software, where complex and expensive functionality need not be designed ahead of time and then locked into the lower layers of a network.[1]

[1]The authors described this as “a kind of 'Occam's razor' when it comes to choosing the functions to be provided in a communication subsystem” (p. 9).

The use of the end-to-end principle has been cited as a major factor in enabling the revolution in on-line services on the Internet. The relatively ease with which the transmission control protocol / Internet protocol (TCP/IP) has been able to host higher level protocols, such as the Hypertext Transfer Protocol (HTTP) that is the basis for the Web, is an example of what has been made possible by end-to-end design.

At a higher level, Geist (2005) points out that network neutrality has made access to Web content and service provider sites equal since their content was not discriminated against in terms of the data retrieved or submitted by users. Content neutrality has also been necessary in establishing cooperation in the interconnection of networks and permitting the transit of data across peering points. At a legal level, Geist (2005) also points out that content-neutrality provided ISPs with "deniability" regarding the content to which they provided access.[1]

In contrast to content-neutral networks, a content-sensitive network operator may choose to discriminate in how it treats the information that it is transmitting. It may variously block, impede, or priviledge information based on its content or who is sending or receiving it. The Internet is composed of multiple autonomous networks. Tier one Internet service providers (ISPs) – operators of member networks or individual nodes that make up the Internet – are generally those that cooperate in the delivery of information without regard to whether their sources or destinations belong to their domain or not and do so without paying or charging transit fees. The locations where they connect to each other are called peering points. In such an arrangement, a major Internet service provider will route packets to and from its own subscribers as well as help to transmit packets to and from other service providers, possibly on to other peer networks, without prejudice to some packets over others.

Regulatory and business sector changes in a number of countries have resulted in the existence of ISPs which sell not only access to the Internet, but also offer access to their own content. Examples of such content include e-mail, news, audio, and movies. This has created potential conflicts of interest, where providers -- who are expected by consumers and third party content-providers to provide content-neutral Internet access -- want to ensure that access to their own content is not hampered by network traffic to third party sites. Examples of such third party content providers include Google.com and Yahoo.com.

This growing challenge to network neutrality can be seen in proposals by telecommunications companies in various countries, including the U.S. and Canada, to levy tolls for the use of their networks. Vint Cerf (2005), co-developer of TCP/IP, warned the U.S. Congress in a letter that network neutrality was being threatened and cited U.S. Federal Communications Commission (FCC) documents that showed plans by Verizon, a major U.S. telecommunications company, to allocate over 80% of its network capacity for its own services (Yang, 2006).

Several schemes have been suggested by members of the industry, including: the creation of premium services for consumers, where network performance superior to the basic service is guaranteed; or offering corporate content providers premium delivery of their information to consumers. Classes of content that have been discussed here include e-mail and Web pages; however, the focus seems to be on managing the greater bandwidth demands posed by Voice over Internet Protocol (VoIP) and Internet Protocol Television (IPTV) applications. A summary of the position of a majority of the largest telecommunications companies on network neutrality was perhaps best given by Edward Whitacre (Business Week Online, 2006, December 2, para. 25), CEO of SBC Telecommunications, a peer of Verizon, who said of third party content providers:

> How do you think they're going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes?
>
> The Internet can't be free in that sense, because we and the cable companies have made an investment and for a Google or Yahoo! (YHOO) or Vonage or anybody to expect to use these pipes [for] free is nuts!

Opponents of network non-neutrality such as Cerf (Mohammed, 2006, para. 12) claim that non-neutral networks will hinder access and that ISPs are already sufficiently compensated:

> In the Internet world, both ends essentially pay for access to the Internet system, and so the providers of access get compensated by the users at each end, ... My big concern is that suddenly access providers want to step in the middle and create a toll road to limit customers' ability to get access to services of their choice even though they have paid for access to the network in the first place.

Technology critic and essayist Bill Thompson (2006) differentiates between access-based classes of services versus those that are based on content, noting that most people already accept the former. For example, broadband is accepted as being more expensive than dial-up. On the other hand, Thompson

---

[1]This is not the case in all jurisdictions. The author was informed, by a reviewer, of case law in England and Wales that

claims, content-based classes of service are like “building two roads into every town and up to every house, one smooth and well-maintained tarmac and the other a dirt track, and then letting Tesco and Waitrose bid for the right to use the good road” (para. 21).

Major content providers that feel threatened by such proposals, such as Google.com, have responded, along with people like Cerf, in calling on the U.S. Congress to pass legislation to block them (Bray, 2005). Libertarian positions exist as well, such as those of the U.S.-based Cato Institute, which generally hold that the market should determine how access is to be provided (Thompson, 2006).

The lines between content providers and ISPs are not so clear, however. America Online, an ISP, and Yahoo.com, a content provider, announced systems for premium service delivery (Goldenberg, 2006).

The essential point of these developments with respect to individual and community knowledge management is that as people come to depend increasingly on Web 2.0 services, access to “their” knowledge is becoming more complicated and even prohibitive.

New Networks and Post Internets

While these service dependency and network access dynamics play out, the Internet itself – from which these dynamics were born -- may change in fundamental ways. Google is reportedly building its own global network, which some experts believe could eventually have peering points across all of the Internet's tier one networks. This could be interpreted simultaneously as a creative technical and business strategy and as a defensive business strategy to nullify network neutrality threats. Google's efforts could represent a major enhancement to the provisioning of Web 2.0 services. In a separate, long-standing effort, major actors in the Internet engineering community are contemplating a so-called “clean slate” approach, which would attempt to resolve critical technical issues through the creation of an entirely new network. The clean slate community's efforts could result in fundamental changes to Internet access and Internet governance.

---

decided that Internet Service Providers are liable for defamatory content if they are aware of it (see Akdeniz, 1999).

The interconnection of private networks with the Internet is not a new practice.[1] Large organizations have long joined their multiple, private, physical networks together using various commodity interconnection technologies. Prior to the wide-spread use of TCP/IP-based networking to do this, telecommunications providers offered technologies such as *frame relay*, which allowed the interconnection of local area networks. In the TCP/IP realm, companies such as Akamai began to provide so-called "edge services" in the mid 1990s, which are systems hired by major content providers to cache content in various physical locations on the Internet, ideally at peering points. Internet caching reduces the number of network "hops" needed to transfer content, thereby reducing download times for users. This approach can also build some measure of fault tolerance for content providers because they effectively create multiple access points for users.[2] Yahoo and AOL are known users of such services. The putative efforts by Google would extend their own internal edge architecture, placing them at the same tier one peer status as the providers that threaten them through network non-neutrality.

It is known that Google has been seeking to acquire unused fiber optic network infrastructure "as part of development of a global backbone network" (Google, n.d., para. 5).[3] Cringely (2005) reports that the company is buying up large amounts of "dark fiber" to allow the company to tranfer data between its data centres around the world at very high speeds so as to service customers more efficiently. Cringely also reports that to leverage such a network, Google is creating portable data centres within the form factor of standard shipping containers. As the company's network grows, the theory goes, these portable data centres could be easily positioned at Internet peering points coinciding with access points for its own network. This approach has the potential to provide efficient bandwidth-rich Web 2.0 services, such as video-on-demand, while side-stepping at least some of the barriers being raised by advocates of a non-neutral Internet. This approach would reduce the number of network hops for which Google data on the way to consumers would be exposed to transit fees. Theoretically, if Google is able to implement freely-accessible wireless networks globally, it could also with its own private backbone network circumvent other ISPs entirely. Google is already developing a free wireless network in San Francisco.

In addition to the largely policy and operational issues discussed to this point, influential actors have been attempting to address critical security, scalability, and functionality problems in the Internet which

[1]This is to be differentiated from the creation of private networks such as local area networks (LANs) entirely within organizations. This practice is extremely common. See RFC 1918 at < http://ietf.org/rfc.html >.
[2]This is a crude analog to our earlier discussion on the fault tolerance that printing brought to the written word.
[3]The job posting has no posting data, but new reports of this began as early as January, 2005. See < http://news.com.com/Google+wants+dark+fiber/2100-1034_3-5537392.html >.

may become more acute as networking becomes even more pervasive. Researchers remain concerned about: security vulnerabilities in the domain name system (DNS) on which the Internet is dependent; the lack of dynamic configuration mechanisms, especially in the face of major disasters; and the inability of the Internet to scale up and provide ideal functionality for the ever-expanding numbers and types of users, embedded processors, sensor networks, and mobile platforms (Clark, 2005). While technical change has traditionally occurred in incremental steps through processes managed by the Internet Engineering Task Force (IETF), many believe that it is not possible to address the most difficult and long-standing issues in this way. Thus, some have called for the creation of a new network from a clean slate.

Toward this end, the U.S. National Science Foundation (NSF) has funded the Global Environment for Networking Innovations (2006) (GENI) and Future Internet Design (FIND) projects (Clark, 2005; The Economist, 2006, March 9).[1] GENI states that "One of our goals is to create a robust and trustworthy foundation for future applications, so that society can start to rely on the Future Internet" (p. 36). Both projects have proposed new designs and efforts have been underway to implement them, with GENI intended in its inception as a research testbed.

As a funder of CSNET in the 1980s and, subsequently of NSFNet, NSF's support for this endeavor signals a major development. It is not yet clear, however, how these developments will impact consumer-level use and in what time frame. The potentials are clear, however. The clean slate philosphy suggests the abandonment of many existing IETF standards. GENI has called for an IETF liaison function, apparently not as a bridge to the past, but so that "research results produced by GENI scientists and/or developers will impact the adoption of standards for the future Internet" (p. 87). A future network, as contemplated by GENI, is likely to allow ISPs to move away from a network neutral approach:

> Routing protocols must be redesigned to deal with the range of business policies that ISPs want to express. Issues to be considered include signaling the direction of value flow, provisioning and accounting for higher-level services, dynamic pricing, explicit distance sensitive pricing, and alternatives to the simple interconnection models of peering and transit. (p. 22)

Arguably, a high level of inertia exists in the Internet with respect to the dispersion of major changes to technical standards. The prime example is the legacy of Internet Protocol version 6 (Ipv6), which was

[1]Clark and others were also part of the earlier "Next-Generation Internet Architecture" project funded by the Defense Advanced Research Projects Agency (DARPA) of the US Department of Defense. See < http://www.isi.edu/newarch/ >.

adopted by the IETF in 1994 as an attempt to resolve major scalability issues in the IPv4 Internet addressing scheme, but which has not yet been widely used. This example might suggest that a GENI-like Internet may be a long way off; however, there are significant counter arguments. IPv6 dispersion was reduced by the introduction of a less costly solution, namely Network Address Translation (NAT). Second, the GENI project has come about in a period where business models are now extremely dependent on the Internet and, thus, more sensitive to its functionality gaps and risks. The large economic impacts of Internet attacks provide a major motivator for industry to adopt a new architecture.

Conclusions

Web 2.0 technologies have been virtualizing the notion of boundary objects (Star & Griesemer, 1989), extending their original *repository* and *standard form* types to enable new individual and community knowledge sharing possibilites, possibly across many intersecting domains of interest, using standardized representations in the form of Web markup technologies and Internet protocols. It might be said that the network is the *meme* following Dawkin's (1989) notion of a genetic-like replicator which acts as a "unit of cultural transmission" or "imitation" (p.192). These very capabilities and the business models used to operate them have produced: an increasingly strong dependency relationship, where access to information on a standalone processor is no longer sufficient to retrieve personal knowledge; the formation of knowledge via the network is becoming normative within the domain of life-critical knowledge creation processes, such as those necessary for education, health care, and emergency response. Moves toward network non-neutrality and new Web 2.0 technologies are posing major barriers to access, thereby exacerbating this dependency relationship. A d'Arcyian moment has presented itself yet again.

Acknowlegements

The author acknowledges support from the Strategic Joint Initiative of the Social Sciences Humanities Research Council and the Law Commission of Canada for supporting this research. An early discussion with Bill St. Arnaud of CANARIE helped the author to further the development of an analysis of the potential impacts of next-generation Internet architectures.

References

Akdeniz, Y. (1999). Case Analysis: Laurence Godfrey v. Demon Internet Limited, (1999) *Journal of Civil Liberties*, 4(2), 260-267 (July). Retrieved August 2006 from <http://www.cyber-rights.org/reports/demon.htm >.
At SBC, It's All About "Scale and Scope" (2006, February 2). *Business Week Online* Retrieved December 27, 2006 from <http://www.businessweek.com/print/technology/content/feb2006/tc20060202_061809.htm>
Bray, Hiawatha. (2005, December 13). Telecoms want their products to travel on a faster Internet Major site owners oppose 2-tier system. The Boston Globe. Retrieved January 29, 2006 from <http://www.boston.com/business/technology/articles/2005/12/13/telecoms_want_their_products_to_travel_on_a_faster_internet>.
Cerf, V. (2005, November, 8). [Electronic Version] < http://googleblog.blogspot.com/2005/11/vint-cerf-speaks-out-on-net-neutrality.html >.
Clark, D. D. (2005, September). Contemplating a new Internet: An upcoming NSF initiative. *MIT CSAIL*. Retrieved January 9, 2006 from < http://www.geni.net/NGN_FIND_ddc.pdf >.
Cringely, R. X. (2005). Google-Mart. *I,cringely*. November 17. Retrieved November 18, 2005 from < http://www.pbs.org/cringely/pulpit/pulpit20051117.html >
d'Arcy, J. (1969). Direct broadcast satellites and the right to communicate. In L. S. Harms, J. Richstad, & K. A. Kie (Eds.), *Right to communicate: Collected papers* (p. 1-9). (Originally published in EBU Review, 118, 14-18). Honolulu: University of Hawaii Press.
Dawkins, R. (1989). The Selfish Gene, 2nd edition. Oxford University Press.
Geist, M. (2005, December 18). The Search for Net Neutrality. *MichaelGeist.ca* Retrieved December 28, 2005 from <http://michaelgeist.ca/component/option,com_content/task,view/id,1040/Itemid,0/ >.
Google. (n.d.) Google Jobs: Strategic Negotiator, Global Infrastructure - Mountain View. Retrieved June 6, 2006 from < http://www.google.com/support/jobs/bin/answer.py?answer=23607 >.
Global Environment for Network Investigations (GENI). (2006, January, 10). Design Challenges. *GENI Project Web Site*. Retrieved January 9, 2006 from < http://www.geni.net/research.php >.
Goldenberg, S. (2006, February 6). Internet giants announce plans to charge for speedier emails. *The Guardian.* Retrieved February 6, 2006 from <http://www.guardian.co.uk/ >.
Hansen, E. (2005, January 17). Google wants 'dark fiber' CNET News.com. Retrieved June 6, 2006 from < http://news.com.com/Google+wants+dark+fiber/2100-1034_3-5537392.html?tag=nefd.top >.
Haseman, W. D., Nazareth, D.L.,Souren, P. (2005). Implementation of a group decision support system utilizing collective memory. Information & Management. 42 (2005) 591-605.
Mohammed, A. (2006, February 7). Verizon Executive Calls for End to Google's 'Free Lunch'. *Washington Post.* Retrieved February 8, 2006 from <http://washingtonpost.com >.
Ó Siochrú, S. (2005). Assessing Communication Rights: A Handbook. *CRIS Campaign.*
O'Reilly, T. (2005, September 30). What Is Web 2.0? Design Patterns and Business Models for the Next Generation of Software. *O'Reilly.* Retrieved May 18, 2006, from <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>.
Peterson, L., Anderson, T., Culler, D. and Roscoe, T. (October 2002). A Blueprint for Introducing Disruptive Technology into the Internet. *First Workshop on Hot Topics in Networking (HotNets-I).* < http://www.cs.princeton.edu/nsg/papers/planetlab_hotnets_02/hotnets.pdf >.
Peterson, L., Bavier, A., Fiuczynski, M., Muir, S., and Roscoe, T.. (2005, June). Towards a Comprehensive PlanetLab Architecture. *PlanetLab Design Note 05-030*. Retrieved October 2005 from <http://www.planet-lab.org/PDN/PDN-05-030/ >.
Reinventing the internet. (2006, March 9). *The Economist.*
Saltzer, J.H., Reed, D.P. & Clark, D.D. (1984). END-TO-END ARGUMENTS IN SYSTEM DESIGN. ACM Transactions in Computer Systems 2, 4, November, 1984, pages 277-288. [Electronic version] Retrieved April 6, 2006 from <http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf >

St. Arnaud, B. (2005). CA*net 4 UCLP Roadmap. ONT Workshop - NASA Ames. September 12, 2005. Retrieved January 11, 2006 from < http://www.canarie.ca/canet4/library/recent/ONT_slides_%20NASA_Ames-CANARIE.ppt >
Star, S. L. & Griesemer, J. R. (1989). Institutional Ecology, 'Translations,' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907 - 1939. *Social Studies of Science.* 19: 387-420.
Thompson, B. (2006, February 2). Why the net should stay neutral. *BBC News*. Retrieved February 12, 2005, from < http://news.bbc.co.uk/go/pr/fr/-/2/hi/technology/4700430.stm >.
W3C. (2006). Session 2: My data, your data, the Web's Data: Challenges of Data Ownership. W3C Technical Plenary Day, Mandelieu, France, 1 March 2006. Retreived August 2006 from <http://www.w3.org/2006/03/TP-minutes.html#itemData >.
Wegner, D.M. (1986). Transactive memory: a contemporary analysis of the group mind. IN B. Mullen, G.R. Goethals (eds.). *Theories of Group Behaviour*. New York: Springer-Verlag. 185-208. [Electronic version] Retrieved May 25, 2006 from <http://www.wjh.harvard.edu/~wegner/pdfs/Wegner%20Transactive%20Memory.pdf >.
Weinberg, B. H. (1997). The earliest Hebrew Citation Indexes. Journal of the American Society for Information Science. 48(4):318-330.
Witty, F. J. (1965). *The Library Quarterly*. Vol. XXXV, No. 3, July.
Yang, C. (2006, February 2). Is Verizon a Network Hog? *Business Week Online*. Retrieved February 3, 2006 from <http://www.businessweek.com/print/technology/content/feb2006/tc20060202_061809.htm>.