International Journal of Communication 5 (2011), 651-664 1932-8036/20110651

IJoC
special section

The Ever Evolving Web: The Power of Networks

WENDY HALL
University of Southampton

In this article, we consider the Web as a network of networks and reflect on its evolution. We do this by analyzing the reasons why it became the first truly ubiquitous hypertext system against all competitors, and then by looking both at the way it has evolved from a network of linked documents to a system that facilitates social networking on a scale previously unimaginable, and at how it will evolve in the future as a network of linked data and beyond. The study of the Web—its evolution and its impact on society, on business, and on government—is referred to as Web Science. We consider some of the major challenges of Web Science and discuss possible Web worlds of the future.

Introduction

There have been many profound developments in communication technology over the centuries, but the World Wide Web is possibly the most significant. The speed with which the Web has penetrated and influenced society is unprecedented. Its operation is also a testament to the power and reach of networks. The Web sits on top of the network of computers that is enabled by Internet technology, is itself a network of interconnected or linked documents, and has also given rise to a social networking phenomenon that has taken us all by surprise. The use and popularity of these networks and applications now make possible the emergence and evolution of a new network—the Web of Linked Data. What impact will this emerging network have on society, on business, on government, and what will future Web worlds be like?

Looking back in time, scholars envisioned worlds of inter-connected information centuries before the Web existed. The ideas of mapping and indexing associations between ideas, facts, and documents long pre-date the existence of computers, and in many ways, they reflect the sophisticated way in which information is indexed by the human brain. When our only means of communicating information was via the written or printed word on paper, inter-connecting pieces of information on different pages was very difficult other than via manually created indexes. With the advent of machines, people were able to imagine a time when the machine could support vast quantities of similar inter-connections.

Wendy Hall: wh@ecs.soton.ac.uk
Date submitted: 2011–02–20

The invention of the terms hypertext and hypermedia (using computers to link related items of information) is credited to Ted Nelson in 1965. The terms are often used interchangeably, although, strictly speaking, hypertext refers to linking text documents, and hypermedia refers more broadly to linking media of any type. Nelson's vision of a universal hypermedia system, Xanadu, is most fully explored in his book *Literary Machines* (Nelson, 1981). He defines hypertext as *non-sequential writing* and views hypertext as a literary medium, but the term encapsulates a wider set of meanings that includes both cross-referencing—linking between documents—and the more general associations between ideas.

Examples of paper-based "hypertext" systems are readily seen in ancient and modern scholarly texts, as well as in reference works such as dictionaries and encyclopedias. We are all familiar with using libraries: following references from one book to another, looking up definitions of terms that are unfamiliar to us, and seeking alternative explanations for difficult concepts. The idea behind hypertext is that this process can be automated, and that we can harness the power of computers to make the search and query process easier. Nelson argued that it would become possible to store digitally anything that anyone would write, record, photograph, or film, and to produce a system that could connect any piece of information to any other piece of information. This is the vision he called Xanadu. To quote from *Literary Machines*:

> There are no intellectual subjects. For someone used to learning, to grabbing vocabulary and ideas, the elements of a new subject can come quickly. The more diagrams you have seen, the more words you know, the more theories you have heard, the more easily you can grasp the next one and assimilate it to the snowball of ideas already rolling around in your head. (1981)

Here, Nelson is presenting his vision of the type of knowledge management environment that hypertext can help to support.

Nelson readily admits that inspiration for his ideas comes from the writings of Vannevar Bush, as well as the pioneering work of Douglas Engelbart. Bush, who was scientific advisor to President Roosevelt during the Second World War, proposed a theoretical design for a system that we would now call a hypertext system. In his seminal paper, "As We May Think" (Bush, 1945), he foresaw the explosive increase in the production of scientific information and predicted the need for a machine to help scientists follow developments in their discipline. Bush called his system Memex (memory extender), which he described as "a sort of mechanised private file and library." His design was based on mechanical technology, but the principles behind the design are, in many ways, even more valid today because of developments in information technology.

> It (a Memex) affords an immediate step to associative indexing, the basic idea of which is a provision whereby an item may be caused at will to select immediately and automatically another. This is the essential feature of the Memex. (1945, section 7)

Bush goes on to describe *trails*, which users build as they move through the information, so that their paths of discovery can be saved and recalled later or passed on to other researchers. Over 60 years later, much of what Bush foresaw has been achieved, although possibly on a scale that even he didn't contemplate. It can be argued that we have yet to achieve his grander vision of systems that enable associative indexing. We will return to this later.

Pre-Web Hypertext Systems

Douglas Engelbart, one of the early pioneers in the computing industry, is credited with the invention of word processing, screen windows, and the mouse, and thus, with inspiring the developments in graphical user interfaces that took place in the 1970s and 1980s. In 1962, Engelbart, working at Stanford University, started work on his Augment project (Engelbart, 1963). He foresaw a world of instant text access on screens, interconnections that could be made and shared, a new style of shared work amongst colleagues, and the use of computers to augment the human intellect. One part of the Augment project was NLS (for oN-Line System), which had several hypertext features and was first publicly demonstrated in 1968. This demonstration, which was filmed for posterity and is now readily available on the Web—search for "mother of all demos"—has become iconic for its prescience, illustrating the sort of integrated, desktop environment that we take for granted today.

The development of hypertext systems stayed in the realm of the research lab for most of the 1970s and 1980s. The systems grew in sophistication as the underlying technologies became more powerful. First-generation systems were implemented on mainframes and were largely text-based. Work on the so-called second-generation systems took advantage of the more advanced user interfaces of 1980s workstation technology. These systems supported graphics and animations, as well as fully formatted text documents. They were also able to provide graphical overviews of the hypertext structure and multi-user support. For more information about such systems, the reader is referred to Conklin's classic survey (Conklin, 1987).

The Intermedia project deserves a special mention here. Intermedia was a hypermedia system developed by a team of researchers at Brown University between 1985 and 1990. As members of the team state in the retrospective paper published after the end of the project (Haan, Kahn, Riley, Coombs, & Meyrowitz, 1992), Intermedia was distinct from other hypermedia systems of the time. The team's intention was to create a model for hypermedia functionality handling at the system level, where linking would be available for all participating applications. In Intermedia, information about anchors in documents, and links between the anchors, were maintained in a database management system. This separation of link data and document data was a distinctive feature of the Intermedia design. The system demonstrated how it was possible to integrate hypermedia technology across many different types of desktop applications, and thus, to create an integrated hypermedia environment accessible to both authors and readers. The separation of links and documents allows the application of multiple sets of links across common sets of documents.

Unfortunately, just as Intermedia was made available commercially, an up-grade in the operating system on which the implementation was based created incompatibilities that made the system inoperable, and the project died—a fate that is unlikely to befall the Web because of the nature of its design. Nonetheless, the development of Intermedia pointed the way toward a more open and efficient system for creating and storing hypermedia relationships. As will be discussed below, these ideas are only now re-emerging as the key to the next generation of innovation on the Web.

With the advent of the PC in the mid-1980s came the emergence of a new generation of hypertext systems that were much easier for individuals to use—notably Hyperties, OWL's Guide, and Apple's HyperCard. HyperCard was made available for free on every Macintosh computer in 1987, and it did more to popularize hypertext in the late 1980s than any other event. The first ACM Hypertext conference was held in the same year. But such systems could only be used to build relatively small applications. The hypertext research community continued to explore the development of hypertext systems that handled information on a large scale and in distributed environments. Notable examples were the Hyper-G system from the University of Graz (Andrews, Kappe, & Maurer, 1995), the Microcosm system from the University of Southampton (Fountain, Hall, Heath, & Davis, 1990), and of course, the Web itself (Berners-Lee, et al., 1994).

Like the Web, the Hyper-G system was both designed to run on the Internet and based on a client-server architecture, but unlike the Web, it stored links in databases, rather than embedding them in documents. As with Intermedia, this decision was made to ensure the consistency and integrity of links, as well as to allow for the application of different sets of links to the same set of documents. However, its hypertext functionality could only be accessed through dedicated Hyper-G viewers. In this sense, Hyper-G can be likened to Intermedia on the Internet.

A good review of the design and features of Hyper-G can be found in BYTE (Flohr, 1995). Here Hyper-G is described as representing

> [A]n advance over the Web as we've known it because it provides real hypermedia. It supports tools for structuring, maintaining, and serving, heterogeneous multimedia data. Hyper-G guarantees automatic hyperlink consistency, and it supports hyperlinks among multimedia documents, full-text retrieval, a UNIX-like security system, and client gateways to Gopher and Web browsers such as Netscape, Mosaic and MacWeb. (1995, p. 59)

Interestingly, Hyper-G was seen as a more sophisticated hypertext system than the Web, yet it was the Web that became the universal system. We discuss this more below.

Unlike Hyper-G and the Web, Microcosm was based on a peer-to-peer architecture. Like Intermedia, it was one of the class of so-called "open hypermedia systems" which separate the link data from the document data, and which embody some form of link service to enable hypermedia functionality to be integrated into the general computing environment and allow linking from all applications on the desktop. The focus in the Microcosm project was on links and linking—how to create them, store them, apply them across different file formats and different applications, manage them, and filter them. A full account of the design, functionality, and features of Microcosm can be found in Hall, Davis, and Hutchings (1996). Links in Microcosm described relationships between objects in the system, and they were stored in databases called *linkbases*. As will be discussed below, links were, in fact, triples consisting of a source, a destination, and a description of the relationship between the objects, rather than just entities providing point-to-point information.

Microcosm introduced the concept of generic links. A generic link may be followed from any occurrence of a particular object, such as a particular text string, wherever it occurs in any document. For example, the linkbase might contain generic links on my name, “Wendy Hall,” to information about my background or qualifications. When an author includes my name in a document, Microcosm automatically provides the reader with access to these links. This is a very powerful feature, as it enables the easy generation of links based on metadata associated with documents. This, in turn, enables a large corpus of information to be given a simple hypertext structure very quickly. Authoring links was also extremely simple in Microcosm—you simply pointed to where you wanted the link to begin, then to where you wanted it to end, and lastly, you chose the type of link you wanted to create. Furthermore, anybody could create links on any information, whether they owned it or not. This last feature sometimes created problems, as confused owners of information objected to third-party links being applied to their information without permission—a foretaste of some of the debates that rage on the Internet today about who can do what with whose information.

Microcosm had filters for creating and analyzing links, and for determining the links which users were presented with in response to a particular query. It offered both proprietary viewers for displaying links and the potential for adding Microcosm linking functionality to third-partly applications. The team experimented with adding links to and from any multimedia document format, adding hypertext functionality to a variety of desktop applications, such as spreadsheets and databases, or GIS and 3D modeling systems, and also experimented with the integration of the Microcosm link service and the Web (Carr, De Roure, Hall, & Hill, 1995). The peer-to-peer architecture that was the basis of the original Microcosm design was fully implemented to support a distributed version of Microcosm in 1994 (Hill & Hall, 1994), but it was too late. By that time, the Web was well on the way to becoming the ubiquitous Internet-based hypertext system, and in the next section we examine the reasons why.

The Growth of the Web

It is important to understand that, as systems like Hyper-G, Microcosm, and the Web were emerging in the early 1990s, the age of the Internet was upon us. Many research laboratories were already using email as a standard means of communication, and users were anxiously awaiting the ability to share files over the network. Systems such as Gopher and WAIS were being developed to enable files to be downloaded easily, without the need for expert technical knowledge. The world was hungry to share information over the Internet, and the time was ripe for the emergence of an easy-to-use, Internet-based hypertext system to facilitate these activities.

In 1989, Tim Berners-Lee started work at CERN, the high-energy physics laboratory in Switzerland, on the development of the system he and Robert Cailliau would eventually call the World Wide Web. Berners-Lee already had experience building other hypertext systems (Berners-Lee, 1999). The aim of the project was to provide a distributed hypertext environment to enable physicists to easily share and distribute information. The main features of the design included ease of use, accessibility from anywhere—decentralization—and the provision of open protocols and universal standards.

The open protocols on which the client-server model of the Web is based—the Hypertext Transfer Protocol (HTTP) and Hypertext Mark-up Language (HTML)—were the cornerstones of its success. The original Web viewer at CERN worked over line-oriented telnet connections, meaning it could be used essentially from any computer in the world. Early viewers implemented by Berners-Lee at CERN were also editors, which enabled easy creation of html documents by users. However, it was the development of the Mosaic browser from NCSA (the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign) that really made the Web so universally popular. Suddenly the Web, via Mosaic (and subsequent browsers, such as Netscape and, later, Microsoft's Internet Explorer), became the user-friendly interface to the Internet, and the rest, as they say, is history.

The growth of the Web since then has been phenomenal. It impacts every aspect of the way we live and work, and it is significantly changing our culture and society in ways we don't yet understand. It is a massive network of networks, and it has shown us quite dramatically, and in so many different ways, both how networks shape our lives and the amazing potential of a global hypertext system. But why did the Web take off and beat all the competition to become the universally accepted standard?

With hindsight, we can see what Berners-Lee got right and the rest of us got wrong. First, he argued that the network was everything—so many other hypertext systems at the time ran on stand-alone workstations. Second, he showed us that "scruffy works." It didn't have to be perfect, because we aren't perfect, but it had to be good enough. And third, he argued that the model had to be decentralized, that the protocols had to be non-proprietary, that the standards had to be universal, and that it had to be free to use—either everyone would use the Web, or no one would.

Lessons to Learn

The success of the Web points to two important and related lessons about what makes a network successful. The first is the power of network effects. Network effects are positive feedback effects connected with *Metcalfe's law*, which states that the value of a network is proportional to the square of the number of its users/members. As value increases, more agents join the network to get the benefits, including information that they own into the network, and thus further increasing its value.

A second lesson is the important role that users play in identifying the weaknesses in a particular platform, creating the opportunity for new innovators to address these weaknesses and further the Web's evolution. As the Web grew and took hold, some of its disadvantages versus the more complex, less "scruffy" systems became apparent. For example, the data encoded in Internet hyperlinks does not provide the information that helps a user reason about relationships using context, content, link description, etc. that had been contained in the more sophisticated hypertext systems. As the Web increased in size and its number of links exploded, the ability to find information by simply following hypertext links was lost (serendipity is a wonderful thing, but it's not the search method of choice when you have a deadline to meet). Yet this loss of functionality led to new adaptations. In particular, the search engines grew to fill these gaps. It's hard to remember life before Google these days, let alone life before the Web, but the two systems are completely synergistic—one can't exist without the other (although the policies which have fuelled their parallel growth are diametrically opposed). Google can function because of the way in which the Web encodes hypermedia data, and the Web has overcome some of its inherent disadvantages through the development of Google.

Indeed, the history of the Web's growth and success is also a history of the technological innovations which enabled it to expand its application, reach new users, and overcome obstacles to broader and deeper use. See Figure 1.

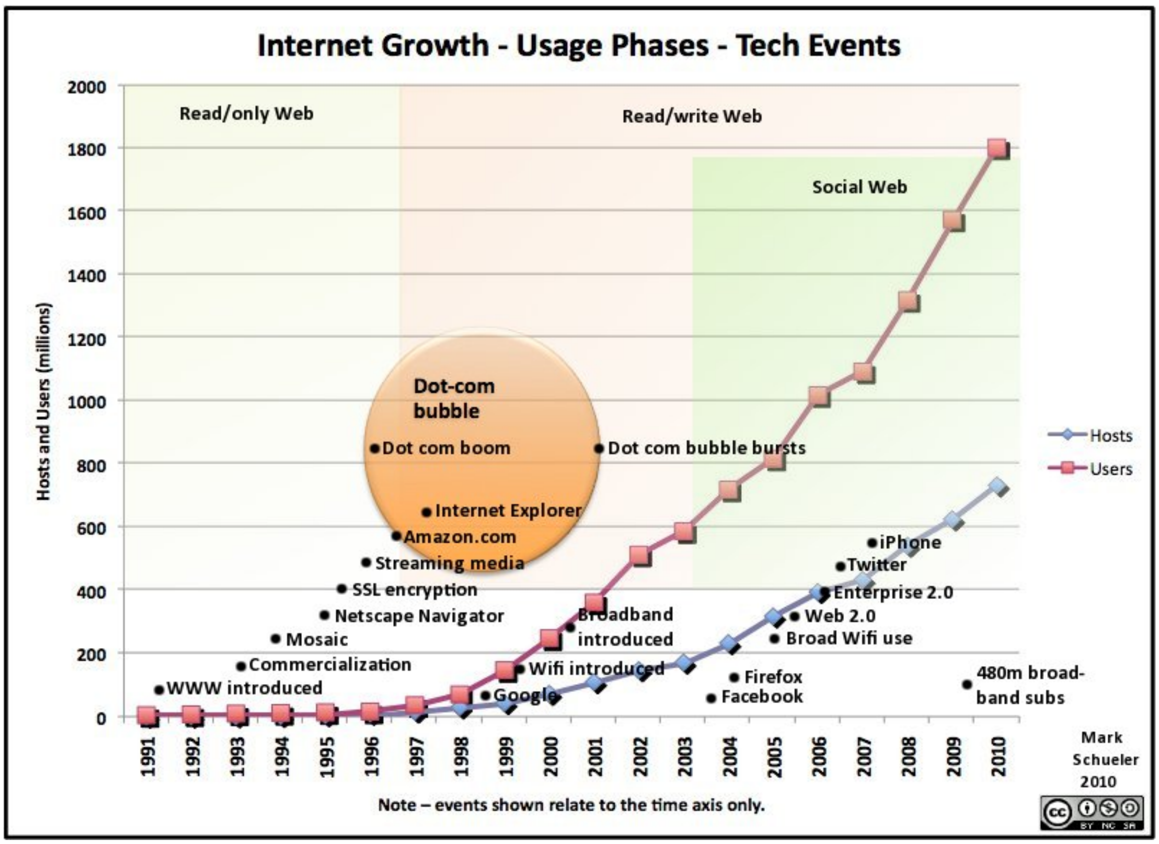

Figure 1. The Growth of the Web.

In the Web's early years, the browsers only enabled users to read Web documents. You could only write to the Web if you could produce html. Also, it wasn't until the late 1990s that we saw broadband became readily available, thus enabling the growth of the home computing industry. With large numbers of users possessing high-speed access to the Internet in their homes, companies trying to sell things on the Web had a market for their products.

As the browser technology developed, users were able to interact more easily with the Web. This lead to the beginnings of the social networking phenomenon, the so-called Web 2.0, that is the center of so much attention today. We saw the emergence of wikis and blogs, and applications such as Flickr and YouTube that enable us to share information so easily on the Web. The emergence of social networking sites, such as Myspace and Facebook, allow us to interact and communicate with each other in ways that until recently were the stuff of science fiction. More recently, the micro-blogging application Twitter has surprised us all in terms of the creative ways people are using it to communicate. The fact that such technology has become the dominant method of communication for many young people is a phenomenon we still seek to understand.

An amazing thing about the Web is the way that the smallest changes in the underlying technology can give rise to massive developments at the social networking level. And of course, there are no boundaries either in terms of personal privacy or geographical borders—so there are many lessons for us to learn.

There are so many interesting things going on here. The first Web—call it Web 1.0 or the Web of linked documents—amazed us all as it emerged because we had never experienced anything like it before. It grew slowly to begin with, and few people understood its potential or why they needed to even bother with it. As the underlying technology developed to enable more people to have access to the Internet, and as more information became available on the Web, the rate of growth began to accelerate until we reached the tipping point—beyond which people began to expect to find the information they were looking for on the Web, rather than either being surprised or hoping it might be there. It had become ubiquitous. And as the technology continued to evolve to enable people without great technical skills to write to the Web, and as more and more people were able to access the Internet from home or mobile devices, we were all amazed again at how quickly the social networking phenomenon took hold. Why did so many people want to share their photos and home movies with the world, and how did they have the time to write all those blogs and wiki entries? Suddenly, the Web wasn't about linking documents, it was about connecting people. And as soon as new trends emerged in specific communities on the Web, the technology developed to enable everyone to follow their lead. And so the technology was shaping society, but at the same time, society was shaping the technology.

Another interesting feature of the evolution of the Web arises from the fact that it was designed to become ubiquitous. The Web is open and free, but the successful companies that emerge at each stage of its evolution become monopolies. It's a winner-take-all world, and the usual rules of market economics don't apply.

The Return of the Link

We are now looking forward to the next stage in the evolution of the Web—the so-called Semantic Web. This has always been part of Berners-Lee's vision (Berners-Lee, 1999). In fact, he talked about it in his keynote at the first WWW conference in 1994. The aim of the Semantic Web is to shift the emphasis of associative linking from documents to data, facilitating a more comprehensive form of automated reasoning. This shift is desirable for three reasons. First, it will facilitate data reuse, often in new and unexpected contexts. Second, it will help to reduce the amount of relatively expensive human information processing. Third, it will release the large quantity of information, not currently accessible, that is stored in relational databases and Excel spreadsheets by making it directly machine-processable.

The Semantic Web encapsulates a vision of a Web of Linked Data, enabling the automated or semi-automated querying, sharing, and interpretation of data from distributed sources in heterogeneous formats (Shadbolt, Hall, & Berners-Lee, 2006). The basic building blocks of this brave new world are the Universal Resource Identifiers (URIs), which denote pieces of data, be they text, documents, media, or concepts; the Resource Description Framework (RDF), used to describe these data in links that record their association; and ontologies which provide the rules by which these associations may be given explicit, automated semantic interpretation.

It is hoped that the Semantic Web will exhibit the same *network effects* that promoted the growth of the WWW. Just as using the Web became more valuable to more people as Web usage increased, the more that people share data which can be mapped onto URIs and linked to other data, the more valuable that data is. This, like the WWW model, is radically different from other models of the value of information, wherein value is dictated by *scarcity* (copyright, intellectual property restrictions, etc.). In decentralized networks like the Web, the value of information is dictated by *abundance*. Abundantly available information can be placed in new contexts and reused in unanticipated ways. This is the dynamic that enabled the WWW to spread, as the value of Web documents was seen to be greater in information-rich contexts (O'Hara & Hall, 2009).

So what will the impact be of the Semantic Web, or of the Web of Linked Data? It is our hypothesis that it will become the dominant data-sharing and -integration platform, and that its effect on both business and society will be as profound and disruptive as the impact of the first Web. But at this stage, it is only possibly to hypothesize; it is not possible to predict. Just as we did for the first Web—the Web of Linked Documents—we look for the tipping points. Just as we were hungry to share documents in the early 1990s, we are now hungry to share data. We can see that there is data all around us, and that, if we could share that data more easily and integrate large, disparate datasets, there would be huge benefits.

But there is competition in terms of the methodology to do this. Some ask, why do we need a new standard, such as RDF? Many people have invested time and effort in creating datasets in their preferred, or their company's standard, formats. Why should they make the effort to convert these datasets into yet another standard? The benefits are hard to see, because there is not enough data out there, already encoded in the standard, to link to. In this sense, we are in the same position today, with linked data, as we were with linked documents in the mid-1990s. Enough people have to take the calculated leap of faith into the unknown, and be prepared to put their data on the Web, for other people to link from and to, so that the benefits can then begin to emerge.

There is more and more linked data emerging on the Web, and it is interesting that the driving force for development at the moment is the drive for transparent government on both sides of the Atlantic—look at the data.gov and data.gov.uk Web sites to see how much data about how our governments spend our money is being made available in linked data format. Other governments are following suit, and companies are becoming interested—first, as users of the data, and second, as providers of data themselves.

As more organizations put their data on the Web, the technology will evolve to enable us to make sense of this new world of data—to smooth out the scruffiness in the early manifestations of the Semantic Web. With the participation of more companies and individuals, the demands of these early adopters will sow the seeds for the next breakthrough—the company that becomes the next big thing—and we will begin to see the next phase of the Web evolution cycle emerge. As people start to do new things or change the way they do old things using the new technology, we will, at that point, be able to say what Web 3.0 is.

As we seek to understand the origins of the Web, appreciate its current state, and anticipate possible futures, there is an increasing need to address the critical questions that will determine how the Web evolves as both a social and a technical network. When do standards become core, and what are the tipping points that determine when they will be universally adopted? What is the impact that changes in data storage and sharing capability will have on the data creating and sharing habits of individuals and organizations? How do these changes then lead to change in the web of relationships that constitutes the Web and the knowledge it, itself, comprises?

The emerging field of Web Science seeks to understand these issues (Berners-Lee et al., 2006). The Web, and the ways in which people interact with it, change faster than we can observe. This rapid evolution calls for new approaches and new methodologies to enable us to define and implement studies in Web Science. We cannot use the "specify, design, implement, test, refine, release" cycle of traditional software engineering. And as we can't predict human behavior, we can only anticipate what people might do with the systems we build. Web Science must also, by definition, be very multi-disciplinary, bringing together researchers from many different disciplines to help us shed light on what is actually going on in this very complex sociotechnical system.

Discussion and Conclusions

We can spend many hours discussing why the Web became the ubiquitous hypermedia system and analyzing its resulting impact on the world. It is indisputable that, from the perspective of the research taking place in the implementation of other hypermedia systems, the Web as it emerged in the mid-1990s was very primitive. Of course, the Web is an open and completely programmable environment that makes it possible to build any hypertext system you might like on top of it. Yet, the next generation of hypermedia management tools has not emerged. In part, this is because we have all been so busy enjoying—indeed, reveling—in a global information system that has robustly and reliably always been there, that less time has been spent on re-implementing our former projects. Yet, there are aspects of the way the Web has evolved that have also curbed the interest in these projects. Indeed, as anyone will know who has tried to develop Web sites of any size, it is difficult to create and maintain links in the Web. As a result, the Web is actually a strangely linkless place.

We have still not achieved Bush's vision of associative indexing (linking), but the development of the Web of Linked Data could be about to change all that. RDF represents information as a subject-predicate-object triple, each of whose component parts is a URI. Each triple is therefore, in fact, a link, and triple stores can be likened to the linkbases of Microcosm. If you add to that the concept of multiple linked ontologies (replacing the simplistic descriptions in the Microcosm links and enabling the inter-linking of the linkbases), you could very quickly implement a richly linked hypertext environment that would combine the universality of the Web with the sophisticated hypertext linking features of Microcosm and support associative indexing in the way envisioned by Bush.

Writing about the Web some 17 years since it first hit our desktops, it is hard to imagine that any of the contenders for the crown of universal hypertext system could have won. But that is no reason to believe things will always stay as they are. The success of the Web is an experiment we cannot rerun, and we cannot experiment in a world without the Web. In fact, its very existence has changed society irrevocably, itself determining the path by which hypertext creation and usage has developed. By this influence, the world is very different today than it was when the Web was born. The original advantages of the Web's scruffiness have been supplanted by many of the innovations that the Web, itself, has inspired.

It is also important to remember that the way we interact with the Web is fundamentally changing. As mobile devices become the standard device for accessing the Internet, the Web will move "behind the screens." We won't use Web browsers to find information; we will increasingly access the information we want directly from the many applications that will be available to us on our mobile phones. Thus, the ease with which users can both directly interact with and create and upload content to the Web may become less of a priority. An article in *Wired* recently proclaimed the death of the Web (Anderson & Wolf, 2010). This is somewhat premature, as countered by Tim Berners-Lee in his article in Scientific American (Berners-Lee, 2010). It is Web technology, including the new linked data technology that is being used to develop the social networking applications that are becoming an alternative user interface to the Web and many of the mobile phone applications. People in developing countries who first encounter the Internet through a mobile device will have a completely different mental model of the Web than those of us who have grown up with it by using a browser. They may not even be aware of that older Web's existence. Their conceptions of hypertext and hypermedia, its best uses and implementations, are thus likely to be substantially different from ours.

Different conceptualizations have different advantages. When we study the work of the early hypertext visionaries, like Bush and Nelson, who were so keen to mimic the associative powers of the brain, to move us away from linear reading and writing into the world of hyperlinking, it is ironic to note today the increasing concern from the neuroscientists and psychologists about the Facebook generation losing the ability to study anything in-depth in a linear fashion (Daily Telegraph, 2010).

In fact, to complete the irony, we read very recent reports of the work of two neuroscientists from the University of Southern California (Thompson & Swanson, 2010), which seems to indicate that the brain works more like the Internet than it does like the "top down" view of brain structure that has held sway since the 19th century. Could it be that building a system that supports associative linking will help us to explain how the brain manages memory?

At the time of writing this paper, we are seeing social revolutions driving political—and indeed, regime—change in the Middle East and Northern Africa. It is arguable that the technical revolutions brought about by the Internet, the Web, and the ubiquity of mobile devices are a major driving force in such developments. Are we changing the nature of democracy? Is it now less about the ballot box, and more about social networking enabling the silent majority to finally find their voice? We are witnessing profound changes in society, and we need to be constantly vigilant to ensure that the technology that is such an important driver in these changes is being developed as a force for social good, rather than as one that can be used to harm or control or manipulate.

How is the Web going to be transformed when billions of people in the developing world gain access to the Internet for the first time and experience the power of the network? Economic, social, political, cultural, and economic agendas all have a role to play in shaping the new networks, and we need to understand the inter-play between those different perspectives to be able to anticipate the nature of the Web worlds of the future. This is why Web Science is such an important new discipline. In order to monitor and analyze such developments, we need to be collecting evidence on a global scale of the changes brought about by the application of Web technology, and we need to be sharing the data we collect to be re-used in increasingly sophisticated analysis. This is a major role for the Web Scientists of the future.

References

Anderson, C., & Wolf, M. (2010, September). The Web is dead. Long live the Internet. *Wired*. Available at: http://www.wired.com/magazine/2010/08/ff_webrip/all/1

Andrews, K., Kappe, F., & Maurer, H. (1995). Hyper-G: Towards the next generation of network information technology. *Journal of Universal Computer Science, 1*(4), 206–220.

Berners–Lee, T. (1999). *Weaving the Web.* London: Texere.

Berners-Lee, T. (2010, December) Long Live the Web: A call for continued open standards and neutrality. *Scientific American*. Available at: http://www.scientificamerican.com/article.cfm?id=long-live-the-web

Berners–Lee, T., Cailliau, R., Nielsen, H. F., & Secret, A. (1994). The World Wide Web. *Communications of the ACM, 37*(8), 76–82.

Berners–Lee, T., Hall, W., Hendler, J. A., O'Hara, K., Shadbolt, N., & Weitzner, D. J. (2006). A framework for Web science. *Foundations and Trends in Web Science*, *1*(1), 1–130.

Bush, V. (1945, July). As we may think. *Atlantic Monthly*, 101–108.

Carr, L. A., De Roure, D., Hall, W., & Hill, G. J. (1995). The distributed link service: A tool for publishers, authors and readers. In the *Fourth International World Wide Web Conference Proceedings* (pp. 647–656). Sebastapol, CA: O'Reilly & Associates.

Conklin, J. (1987). Hypertext: An introduction and survey. *IEEE Computer, 20*(9), 17–41.

Daily Telegraph. (2010, February 11). Students' brains 'rewired' by the Internet. Available at: http://www.telegraph.co.uk/technology/news/7205852/Students-brains-rewired-by-the-internet.html

Engelbart, D. (1963). A conceptual framework for the augmentation of man's intellect. In R. Howerton & D. Weeks (Eds.), *Vistas of information handling* (Vol. 1, pp. 1–29). London: Spartan Books.

Flohr, U. (1995). Hyper-G organizes the Web. *BYTE, 20*(11), 59–64.

Fountain, A., Hall, W., Heath, I., & Davis, H. (1990). Microcosm: An open model for hypermedia with dynamic linking. In *Hypertext: Concepts, Systems and Applications, Proceedings of ECHT'90, Paris, November 1990* (pp. 298–311).

Haan, B. J., Kahn, P., Riley, V. A., Coombs, J. H., & Meyrowitz, N. K. (1992). IRIS hypermedia services. *Communications of the ACM, 35*(1), 36–51.

Hall, W., Davis, H., & Hutchings, G. (1996). *Rethinking hypermedia: The Microcosm approach.* New York: Kluwer Academic Publishers.

Hill, G. J., & Hall, W. (1994). Extending the Microcosm model to a distributed environment. In *ACM Conference on Hypermedia Technology, ECHT'94 Proceedings.*

Nelson, T. (1981). *Literary machines.* Published by the author.

O'Hara, K., & Hall, W. (2009). Semantic Web. In M. J. Bates, M. N. Maack,, & M. Drake (Eds*),* *Encyclopedia of Library and Information Science Second Edition*, London: Taylor & Francis.

Shadbolt, N., Hall, W., & Berners–Lee, T. (2006). The Semantic Web revisited. *IEEE Intelligent Systems 2006, 21*(3), 96–101

Thompson, R. H., & Swanson, L. W. (2010). Hypothesis-driven structural connectivity analysis supports network over hierarchical model of brain architecture. *Proceedings of the National Academy of Sciences*, *107*(34), 15,235–15,239.