Socrates, Trust and the Internet

Kieron O'Hara

Dept of Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
kmo@ecs.soton.ac.uk

Abstract: Socrates, one of the world's greatest philosophers, never wrote anything, and confined all his philosophy to spoken debate. The important issues for Socrates were trust and control: he felt the radical decontextualisation that resulted from the portability and stasis of written forms would obscure the author's intentions, and allow the misuse of the written outside of the local context. Trust has once more become a central problem, both politically and epistemologically, but since Socrates' day, various technologies have undermined his distinction, making the relationship between trustworthiness and linguistic mode more complex. In this paper, I review the state of the art in Internet technologies, showing (a) how developers and authors attempt to establish trust in their websites or e-commerce processes, and (b) how new work in dynamic content creation further blurs the spoken/written and global/local distinctions.

Introduction

The interrelations between technology, society and psychology are marked and deep. Artefacts play a central role in the development of both social constructs and individual psychology. Accordingly, technological change can be highly disruptive of settled forms of life; even very mundane technologies can be highly transformative. Indeed, it may perhaps be more accurate to say that especially mundane technologies are transformative. Though the invention of movable type seems, with historians' hindsight, to be a straightforward and rather dull extension of existing technology, it goes without saying that efficient and cheap printing, mundane as it is, has proved about as transformative as it gets.

The central importance of the development of technologies for writing – indeed writing as a technology itself – has long been acknowledged, if not fully understood (McLuhan 1962, Ong 1982, Havelock 1963). Plato must be seen as the pioneer, the explorer who began to develop the virgin ground, in a series of brilliant, often contradictory, works, the *Republic*, the *Phaedrus*, the *Ion*, the (possibly fake) Seventh Letter. His instincts, as ever, were correct; Western philosophy is a series of footnotes to Plato precisely because of his unparalleled ability to locate the permanent fault lines of a controversy. Much of his work involves deep meditation on the likely effects of the shift in Ancient Greece from a primarily oral culture to one which might reasonably be called literate (Havelock 1963).

As linguistic storage and dissemination technologies have proliferated, the distinctions between speech and writing have become blurred – if indeed they were ever very clear (Hughes 1996). And each new linguistic technology has created a dilemma, exactly the sort of dilemma anticipated by Plato. On the one hand, the technology provides an opportunity for people to externalise some linguistic or psychological faculty (in the

case of writing, for example, the effect is – in the modern argot – to *outsource* memory), thereby conserving psychological resources. On the other hand, questions are inevitably raised about whether such developments *dehumanise* us in some crucial respect, and whether the technological implementation of the faculty is *trustworthy*, or suitably sensitive to its no doubt highly nuanced social context (O'Hara 2004).

The latest in this stream of linguistic technologies is the Internet (O'Hara 2002a). The Internet may or may not have a big effect on society; clearly in a world where half the population has never made a telephone call, its effects will be limited (though equally there are an estimated 654 million people online). Nevertheless, in the decade or so since the development of the World Wide Web lifted it out of the academic ghetto of computer nerds, the Internet has begun to influence many lives, not only in the developed world, quite dramatically (for good or ill). More information is now stored and disseminated over the Internet than we can really imagine.

With this phenomenal resource, trust is naturally an issue. How much of this information, how many of the computing services that the Internet provides, are trustworthy? How can we possible investigate over such a giant domain? In this paper, I would like to argue that, happily, many of the latest Internet technologies coming on stream now, many still in prototype form, will help provide some answers to questions of trust. In particular, the capabilities afforded by the *Semantic Web*, a new extension of the World Wide Web, will allow some rigorous interrogation of documents and various other automated interlocutors. To see how this helps, we will need to review the technologies underlying the Semantic Web, and show how they can help meet some of the problems that linguistic technologies perennially face.

Accordingly, the paper has the following structure. We begin with a quick review of the relevant technologies that make up the Internet and the World Wide Web, and then move on to describe how they have been extended and improved in the development of the concept of the Semantic Web. The next section shows how trust fits into the original conception of the Semantic Web; the issue of trust was always anticipated as being of central concern here. Having introduced trust in the context of modern technology, we will then rewind several centuries, and review Socrates' arguments about trust, orality and literacy in Plato's dialogue the *Phaedrus*. What is interesting about the Semantic Web in this context is that the technology begins to blur the contrast between the spoken and the written, and therefore in some ways provides an avenue to online trust by addressing Socrates' arguments, and we outline this with some examples from the latest computer science research. In the final section, we will discuss how far we can take these conclusions, and also the prospects for Semantic Web technology.

The Internet and the World Wide Web

We mentioned above that mundane technologies could be transformative – well, they don't come more mundane than the Internet. Originally a military project to set up communications networks that would withstand nuclear attack, all the Internet is is a protocol for transferring data down telephone lines from one computer to another. A file of data is sliced up into *packets*. The packets get sent from the originating computer, routed to other computers across the world. The path to the target computer is calculated

dynamically, individually for each packet, rather than there being any fixed route; by these means it is extremely hard to disrupt Internet traffic, because if some key line gets cut, dynamic routing will simply work out another route to send the packets. Once all the packets arrive – in the blink of an eye – the target computer reassembles them in the right order, and *voila*, it has an exact copy of the original data file.

That is all we are going to say about the Internet as a whole, except to remark that it *is* a very mundane technology, and it has had undeniably giant effects. For the rest of this section, we will focus on the *World Wide Web* (WWW).

The WWW is part of the Internet, and so functions in the same way. It is the multimedia part of the Internet, and it is arguably the great success of the WWW in the early 1990s that lifted the Internet out of the computer science department, and into public and commercial life. Information is presented on the WWW in webpages, which can show text, pictures, sounds, movies etc together. The ability of the Internet to move large amounts of information around quickly and seamlessly has helped commercial development; firms can put pictures and text about their goods and services online, and people can send orders through immediately. Customers can find out about a whole range of goods, comparing prices and quality, without leaving their desks, and so save themselves large amounts of shoe leather; this is what economists call cutting *transaction costs*.

The WWW is based on a particular computer language, HTML (the Hypertext Markup Language). Webpages are created by commands in HTML – you can see raw HTML if you look at a webpage, and click on the menu command 'view source' (this command may be slightly different in different webpage browsers). Because HTML is the common WWW language, webpages are always more or less alike, and can all be seen by the main browsers, such as Microsoft's Internet Explorer, or Netscape's Navigator. HTML tells your computer how to arrange all the information in a webpage on a screen, where the text should go, how it should be formatted, where the pictures fit, how the different panels of text should fit together, and so on; it controls the look of the page.

The key to the usability of the WWW is that users should be able to find the page they want easily; this is not a trivial task, given that there are about 550 *billion* web documents to search through, containing enough information to fill 7.5 billion books of the size of this one. The mechanism used is that of web addresses, doled out by an organisation called Internet Corporation for Assigned Names and Numbers (ICANN); these addresses are the bits of text like http://www.kansaigaidai.ac.jp which you can see at the top of your browser. Each webpage is stored on a computer (called a *server*), and is given a web address. If you type in a particular address into your web browser, you can see that webpage. This gives the illusion that all 550 billion webpages are held on your machine, whereas actually when you type the address in, your computer merely sends a message to the server via the Internet, and the server responds by sending all the information held on the webpage back across the Internet to you. Because the page is written in HTML, the HTML instructions tell your computer how to reconstruct the webpage for itself.

In ye olden days, when the WWW started, you had to type in each web address, and remember them from one time to the next; as web addresses became very long and complex, that was of course very tedious. The web would get nowhere that way. Two essential improvements made it possible to navigate through the giant quantity of information on the web. The first were *hyperlinks*. HTML allows the writer of the webpage to insert links to other webpages; they generally appear in blue on a webpage. If the user clicks the mouse on a hyperlink, then he will move to the new web address. This allows a set of pages to be linked *associatively*, so reading becomes non-linear and there are many different routes through material. A writer can also link to other people's pages, so, for example, an online news report about the Liberal Democratic Party could link to other articles about the party, Mr Koizumi's personal webpage, pages explaining Japan's political system, and so on. This allows very interesting methods of commentary; http://eliotswasteland.tripod.com/, for example, contains the text of Eliot's *Waste Land* with links between not only the various images deployed through the poem, but also out to other relevant pages (e.g. to an online exegetical Bible for the biblical references, to an online version of Dante's *Commedia*, etc).

The second innovation was the development, in the mid-90s, of *search engines*, which could search, more or less accurately, through the information on the web. These engines, which have their own web addresses, and which include Google and Yahoo!, allow you to suggest a key word or phrase, and then will find all the webpages that contain that phrase. Artfully chosen key words cut down the search time for users massively.

The Semantic Web

Anything wrong with that? Currently, not quite. But the WWW is actually reaching the limits of capacity. Google, the market leader, is a brilliant search engine, that actually looks around intelligently for the best pages to refer you to (see O'Hara 2004, Chapter 5 for a discussion of Google's PageRank system). But the more information that appears on the web, the less use a simple key word system is; it will give you too many irrelevant hits. The trouble with key words is that they are *uninterpreted*; if your key word is, say, Bush, a key word system is unable to tell the difference between George Bush, Kate Bush, the African Bush or the metal lining of an axle-hole (known as a bush, believe it or not).

Hence, as the WWW reaches its limits, work is already commencing on its successor, the *Semantic Web* (SW). The aim of the SW is to give the computer the ability to interpret the key words for you, to search much more intelligently (Berners-Lee et al 2001, O'Hara 2002a, 49-52). In this section, we'll discuss briefly how this will work, and how it has extended the concepts underlying the WWW. Warning: technophobes should take this section slowly; it will come to resemble alphabet soup very quickly!

Recall that the underlying language of the WWW was HTML. Playing that role in the SW is a language called *XML*, the eXtensible Markup Language. XML differs from HTML in that it allows users to define little sub-languages for describing objects. Hence, whereas HTML tells your computer how to arrange the content on the page, XML allows you to tell your computer what the things named by the content *are*.

So, for example, if we take a webpage such as the home page of my department, shown in Figure 1, HTML tells the computer to take the text ‘University of Southampton’ and to display it in a certain position in a certain font. XML will allow you to tell the computer not only that, but also that ‘University of Southampton’ refers to a university.

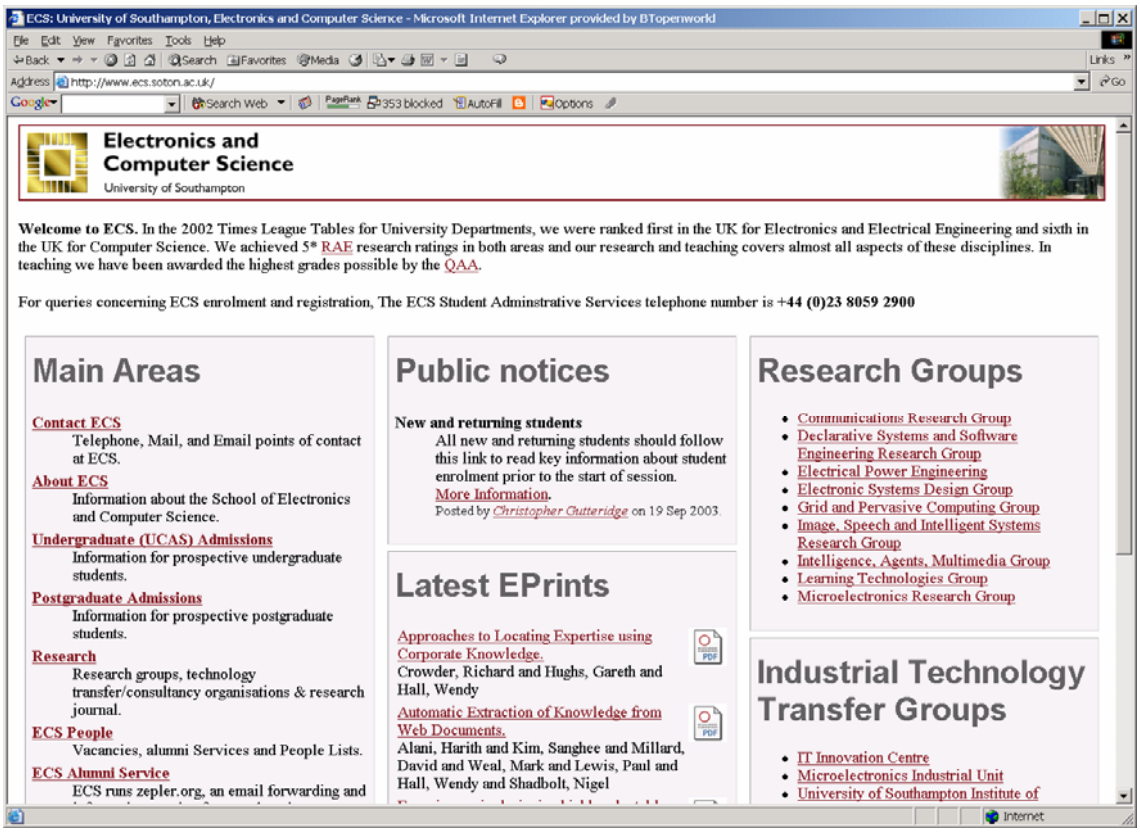

Figure 1: http://www.ecs.soton.ac.uk

However, that is not enough. The computer can do very little more with that information than it could do before. Therefore the SW needs another bit of kit, the *Resource Description Framework* (RDF). RDF is a framework that brings together three things, two objects, and a relation between them. So, for example, the two objects might be ‘University of Southampton’, and the picture in the top right of Figure 1 (which will be a file, perhaps called something like ‘soton.jpg’). RDF lets you tell the computer that there is a relationship between the two: the relation might be called ‘picture-of’, and so RDF lets us assert that ‘soton.jpg picture-of “University of Southampton”’ – i.e. that the picture at the top right is a picture of the University of Southampton. This is obvious to the human reader of the page, but it will be news to the computer.

So XML lets you tell your computer what things are, and RDF lets you tell it how these things are related. A third element is still needed, because the computer still does not know (in a metaphorical sense of ‘know’) what these terms mean. This is achieved with the use of *ontologies*.

An ontology is a specification of the language and concepts of a restricted domain of discourse, and gives as it were the conceptual background to the words and phrases used in XML and RDF. These words and phrases are therefore defined in terms of each other, which specifies a little web of interrelated concepts. Naturally the technology is too weak to say that the computer *understands* anything (in other than metaphorical senses), but an ontology will provide the beginning of a little web of terms in the sense of Quine and Ullian (1970). So the computer could be told, for instance, that a *university* is a type of *educational establishment*, that it has *students* and *lecturers*, an *address*, a *website*, a *telephone number*, that the lecturers will include *professors*, *senior lecturers* and *readers*, that the students will include *undergraduates* and *postgraduates*, that the postgraduates will have *degrees*, that these degrees will be in *subjects*, and will have been awarded by *educational establishments*, and so on. The ontology links together all the concepts and terms that will help the computer to make holistic sense of the terms like 'university' that were being used in the XML specification of the webpage.

In short, where in the WWW, HTML told the computer how to arrange the content on the screen, in the SW, XML + RDF + ontologies tell it not only how to arrange the content, but also *what it is all about*. The effect is suddenly to provide the computer with a richer characterisation of the domain. Figure 2 shows how the new expressivity can allow the computer to see much more. On the left, we see the basic WWW view of the computing world; the resources are whatever is held at web addresses, and they are connected by hypertext links written in HTML. But on the right, the same domain is seen in much more detail. With an XML characterisation of the domain, the computer can see that some of the resources are pieces of software, others are documents, others are persons and so on. The links are made more meaningful too by RDF. For instance, we can see that one document is a version of another, that the creator of one of the documents is a particular person, and so on. The formalisms of the SW allow you to tell the computer so much more about the domains you are describing.

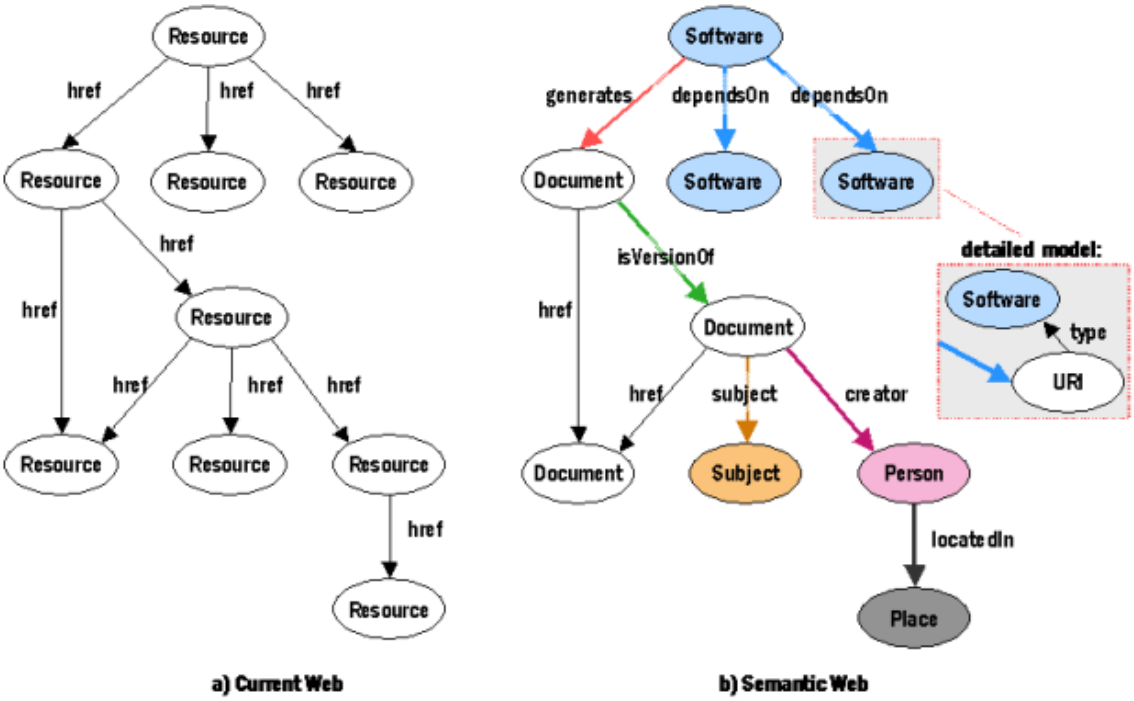

Figure 2: From the WWW to the SW (taken from Koivunen & Miller 2001)

Trust and the Semantic Web

What, then, does this have to do with trust? Actually, the role of trust was built into the standard conception of the SW right from its beginnings. Figure 3 was developed by Tim Berners-Lee in order to articulate the layers of expression, comprehension and inferential power that would be required by his conception of an extension of the WWW that supported intelligence.

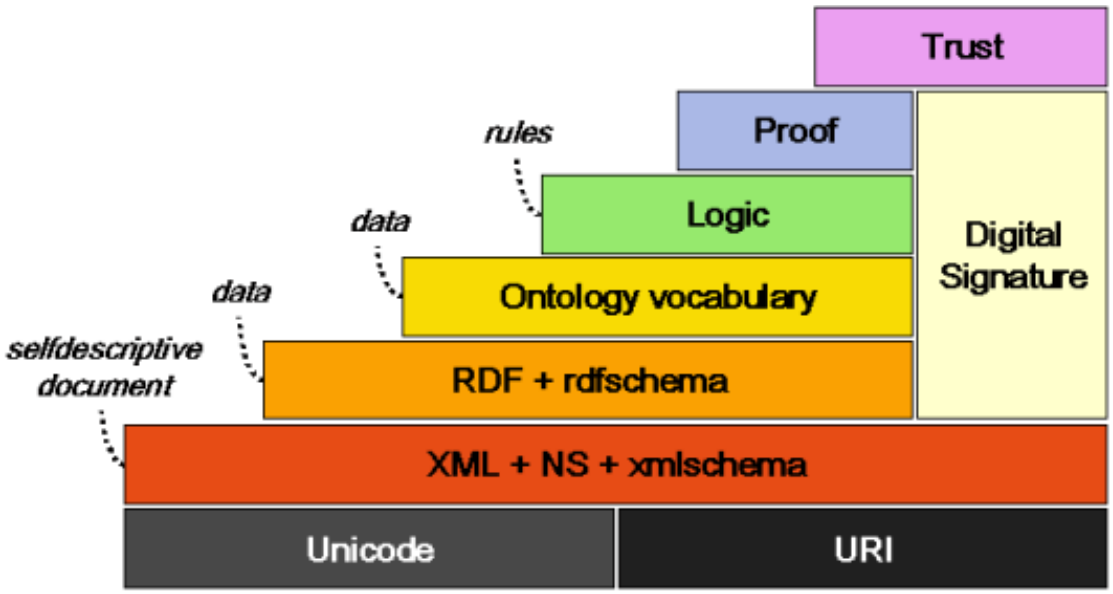

Figure 3: The layered view of the Semantic Web (taken from Koivunen & Miller 2001)

How do we interpret these layers? At the bottom are Unicode (a standardised system for encoding data) and URIs (Uniform Resource Identifiers – web addresses). These are the nuts and bolts of the SW. But there is no point having these unless you have XML to tell you what they refer to. And there is no point having XML unless you have RDF to tell you how those things relate to each other. Similarly, there is no point having XML and RDF unless you have ontologies to explain the significance of the XML classes and RDF relations. There is no point having ontologies without logic to provide methods of inferring one thing from another. There is no point inferring things without a theory of proof to tell you that the inferences are sound.

And, finally, there is no point having a system of proof unless those who will use it (in the case of the SW, this means the 654m Internet users) have confidence in it. Without trust of the users in the system, the SW will never get off the ground.

Trust and orality in the *Phaedrus*

So trust is essential to the SW. This brings us back round to the question of what trust is, how it operates, and how it can be fostered (Misztal 1996, Fukuyama 1995, O'Hara 2004). These questions have long been pondered over; in this paper, I would like to focus on some interesting and still-relevant work by Plato. In his early work the *Phaedrus*, Plato represents Socrates as arguing as follows.

SOCRATES: You know, Phaedrus, writing shares a strange feature with painting. The offspring of painting stand there as if they are alive, but if anyone asks them anything, they remain most solemnly silent. The same is true of written words. You'd think they were speaking as if they had some understanding, but if you question anything that has been said because you want to learn more, it continues to signify just that very same thing forever. When it has once been written down, every discourse roams about everywhere, reaching indiscriminately those with understanding no less than those who have no business with it, and it doesn't know to whom it should speak and to whom it should not. And when it is faulted and attacked unfairly, it always needs its father's support; alone, it can neither defend itself nor come to its own support. (Plato 1997, 275de, p.552)

Socrates, of course, wrote nothing; what we know of him is filtered through the testimonies of contemporaries, notably Plato and Xenophon. However, my feeling is that this passage, and the whole of the discussion of orality, literacy and trust in the *Phaedrus* may well be a representation of Socrates' authentic voice.

This argument that writing is untrustworthy, compared to speech, has five components. First, a listener is able to control the associations that follow from a speech. And by entering into dialogue with the speaker, the listener can turn a conversation in what for him is an interesting direction; a reader, however, is stuck with the associations that the author deems significant, and sometimes – notoriously with poetry in translation – even these associations can easily be lost. Second, a speaker can assess his interlocutor and personalise the content, make the message relevant for that person, explain when the interlocutor has misunderstood, etc. A writer cannot do this; the message is 'one size fits all'. Further, a speaker can control the composition of his audience; if he wishes to prevent someone from hearing the message, all he has to do is shut up, or whisper. A writer cannot; once the written word is disseminated, it is very hard to prevent it falling into the wrong hands.

Third, once the written word goes out, it is more or less fixed (this is even more the case since the introduction of printing – Ong 1982), whereas a spoken message, since it has to be repeated anew each time it is reproduced, can alter with time, responding to changes in context. Fourth, with the spoken, the audience can engage in interrogation of the speaker, determining whether the speaker 'really' knows what he is talking about, clarifying points and raising objections. With the written, the audience cannot do this. Fifth, the spoken is fixed in time, it is anchored to a context, and can be properly evaluated only in reference to its context. The written is similarly context-bound, but it is trivial to take the written out of context.

This model of trust is very interesting. Of course, in a small society, such as a Greek city-state such as Socrates' Athens, trust can work easily on this model. But even now, this distinction between spoken and written underlies many models of trust. In a court, for example, evidence has to be given orally, and there must be an opportunity for witnesses to enter into an interrogative conversation with a properly-appointed and trained representative of the defendant in order that the witness' evidence will be trusted. Written evidence, even in the form of a signed affidavit, is not trusted, and is usually only admitted in the case of trivial or uncontested evidence. Politics is another oral 'society'; recall the election for the Master of a Cambridge College in C.P. Snow's novel *The Masters*, where all the key electioneering is conducted through the medium of the spoken, and the written plays a minor and often obstructive role. And in education, the oral remains a key mode of evaluation; the oral examination of a PhD thesis is a central part of the assessment of a candidate's ability to perform *written* research!

However, such a model of trust can be quite inhibiting for more complex or distributed societies, as argued for example in (Fukuyama 1995), which claims that the lack of social structures able to support the trust of non-kin is a key explanatory variable for the relatively poor economic performance of countries such as Italy and China, compared to Japan or the USA. The analogy between the spoken and the written, and different types of trust – restricting one's trust to a local context, or trusting more distributed agents with whom one is personally unacquainted – is suggestive; in fact, *pace* Socrates, the written word is often trusted, and quite properly, with the help of various institutions, such as peer review, or the publicity and certification of an author's provenance (O'Hara 2004). Institutions such as this carry many of the information-processing costs of examining and regulating potential trustees' *bona fides*. Trust-bearing institutions have often been seen as essential for the creation of truly modern societies (Fukuyama 1995, Misztal 1996).

This makes the case of the Internet, the postmodern space *par excellence*, extremely surprising. The net has a very sparse institutional structure, and in particular, trust regulation is negligible; this despite many worries about security of property on the net (particularly with respect to e-commerce), and about quality of information (both for data sharing between programs and for ordinary sources of knowledge such as academic papers). Hence, despite the global nature of the space – in theory geography is irrelevant to it – trust on the net is often *local* (O'Hara 2004, Chapter 5). And despite its being a medium for the written, trust is on the model of the spoken.

The development of the Internet as an anarchic and unregulated space is very central to the value consensus at its heart (Lessig 1999). So the creation of institutions to regulate it will meet resistance (even if it is inevitable). However, what I want to argue here is that the increasing sensitivity to the audience that the SW makes possible will affect the nature of trust online. The distinction between the spoken and the written has blurred since Socrates' time, in particular with the inventions of recording media (Hughes 1996); technology has always distorted and affected linguistic modes (McLuhan 1962, Ong 1982). The SW is another technology in this tradition. In the next section, therefore, I will look once more at Socrates' five marks of the oral, and see whether it will be possible to develop technologies that allow the written to be subject to these conditions.

Socrates' marks of the oral: trust, expressivity and intelligence

The first mark is that of the control of associations. Hypertext tries to mimic associative memory, the recall phenomenon where someone remembers some very different experience as a result of present stimuli. The classic literary example of associative memory, of course, is in Proust's *Swann's Way*, where the narrator, upon consumption of a coconut madeleine dipped in tea, suddenly remembers his childhood stays in Combray with his Aunt Leonie, who also used to give him madeleines served that way. Hypertext allows that sort of associative link to be made in a webpage; the reader clicks on any word, highlighted in blue, that interests him, and is taken to a relevant page.

However, hyperlinks in an HTML document don't exactly model associative memory. They are written by the author of the page; hence continuing the analogy with Proust, the experience would be more like the narrator biting into the madeleine and remembering some incident from the *baker's* past. However interesting this might be, it is hardly an exact model of associative memory.

New technology, though, allows the reader to regain control of these associative links. In a piece of software called the Conceptual Open Hypermedia Services Environment (COHSE), the reader is presented with a set of links that are created specifically for him, at the moment that he reads the page (and hence these links might change dynamically over time). Using ontologies to understand the subject matter of the webpage and the interests of the reader, and a collection of interesting links that the software keeps up and maintains, a webpage can be given a set of associative hyperlinks that are tailored to the *reader's* interests, not the *author's*. Furthermore, the links created for one reader may well be very different from the links created for another (Carr et al 2001). In this way, the written webpage becomes a little more like a spoken item, as the reader becomes able to control a type of hypertext conversation with the page – the reader, not the author, is able to suggest new topics, as he clicks on the links that interest him (Figure 4).

Such an application can have interesting effects. For example, legal theorist Cass Sunstein has argued that the effects of hyperlinking mean that those for whom the Internet is their chief source of information are now so in control of their information environment that they are never exposed to alternative views, and so extremism may flourish (Sunstein 2001). I have suggested elsewhere (O'Hara 2002b) that open

hypermedia systems, perhaps integrated into a browser, might well be a way to achieve the catholicity of links that Sunstein hopes for, without the weight of regulation that Sunstein expects would be necessary.

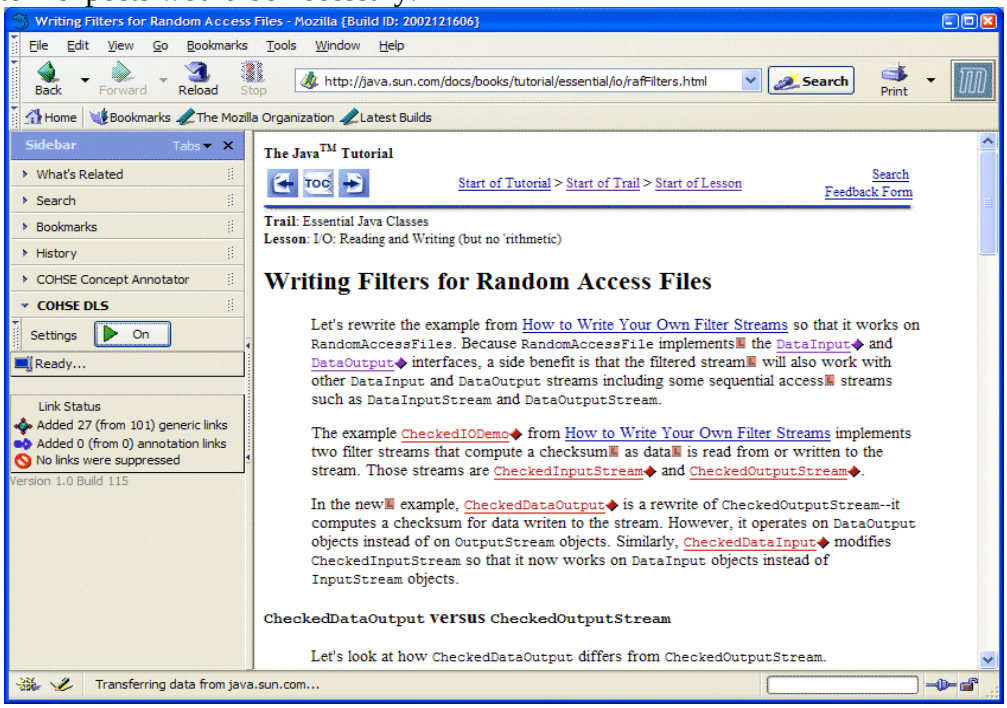

Figure 4: The COHSE interface: coloured diamonds denote links generated by the system for the reader

The second mark of the oral is that of the personalisation of content. Again, ontologies and the nimbleness that the SW provides can ameliorate this issue for the written word. Artequakt is a system that creates biographies of artists (Alani et al 2003). The biographies are created at the moment they are requested, by downloading information from the Internet, matching it against an ontology of art, and reassembling it according to various narrative templates.

Figure 5 shows how the system works. Once the request is received for a biography of a particular artist, a search of the WWW is conducted. Pages are assembled, and a natural language understanding program extracts the key information from these texts (this in itself is not a trivial task) and maps it onto an ontology. This information supplements Artequakt's knowledge base. Narrative templates are then downloaded – the particular template chosen will depend on what the reader requested (e.g. a short narrative, a narrative for a child, etc). This template will then demand information of particular types to fill it in, which the system retrieves from the knowledge base. Figure 6 shows some of Artequakt's output.

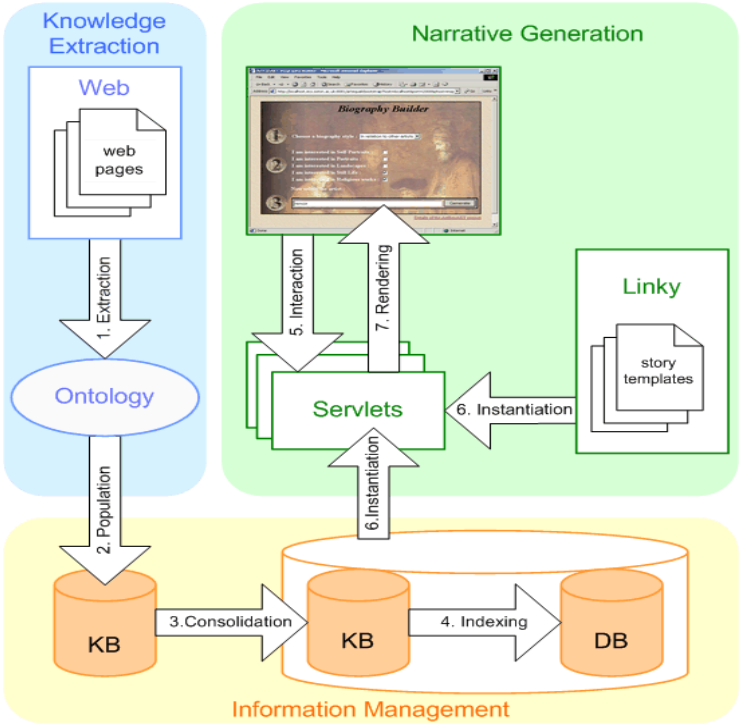

Figure 5: The architecture of Artequakt

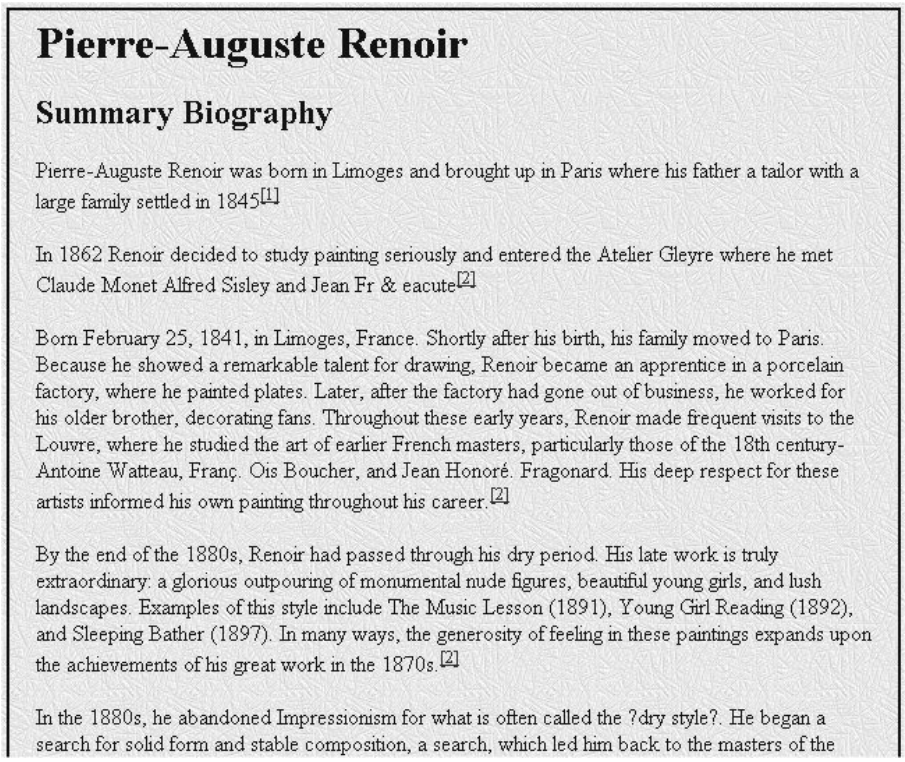

Pierre-Auguste Renoir

Summary Biography

Pierre-Auguste Renoir was born in Limoges and brought up in Paris where his father a tailor with a large family settled in 1845[1]

In 1862 Renoir decided to study painting seriously and entered the Atelier Gleyre where he met Claude Monet Alfred Sisley and Jean Fr & eacute[2]

Born February 25, 1841, in Limoges, France. Shortly after his birth, his family moved to Paris. Because he showed a remarkable talent for drawing, Renoir became an apprentice in a porcelain factory, where he painted plates. Later, after the factory had gone out of business, he worked for his older brother, decorating fans. Throughout these early years, Renoir made frequent visits to the Louvre, where he studied the art of earlier French masters, particularly those of the 18th century-Antoine Watteau, Franç. Ois Boucher, and Jean Honoré. Fragonard. His deep respect for these artists informed his own painting throughout his career.[2]

By the end of the 1880s, Renoir had passed through his dry period. His late work is truly extraordinary: a glorious outpouring of monumental nude figures, beautiful young girls, and lush landscapes. Examples of this style include The Music Lesson (1891), Young Girl Reading (1892), and Sleeping Bather (1897). In many ways, the generosity of feeling in these paintings expands upon the achievements of his great work in the 1870s.[2]

In the 1880s, he abandoned Impressionism for what is often called the ?dry style?. He began a search for solid form and stable composition, a search, which led him back to the masters of the

Figure 6: A dynamically-configured biography of Renoir

The Digital Document Discussion Environment (D3E – Buckingham Shum & Sumner 2001) is a system which allows debate about a document to be added to the document itself. In a way, this takes us back to the days of copying before printing, when marginal

comments might well be incorporated into the main text (Ong 1982), but in this case the authorship of the annotations, and the comments or text to which the annotations refer, are made explicit. The author of the document, of course, can also take part in the debate. In this way, the third mark of the oral, the lack of closure of a discourse, can be mirrored in the written as well. Figure 7 shows a D3E discussion of a website; the panel on the right of the window shows the comments laid out for viewing. The comments are as easily readable as the original document, though the original always retains a privileged presentation in its own window, and furthermore the structure of the original is used to structure the presentation of the comments.

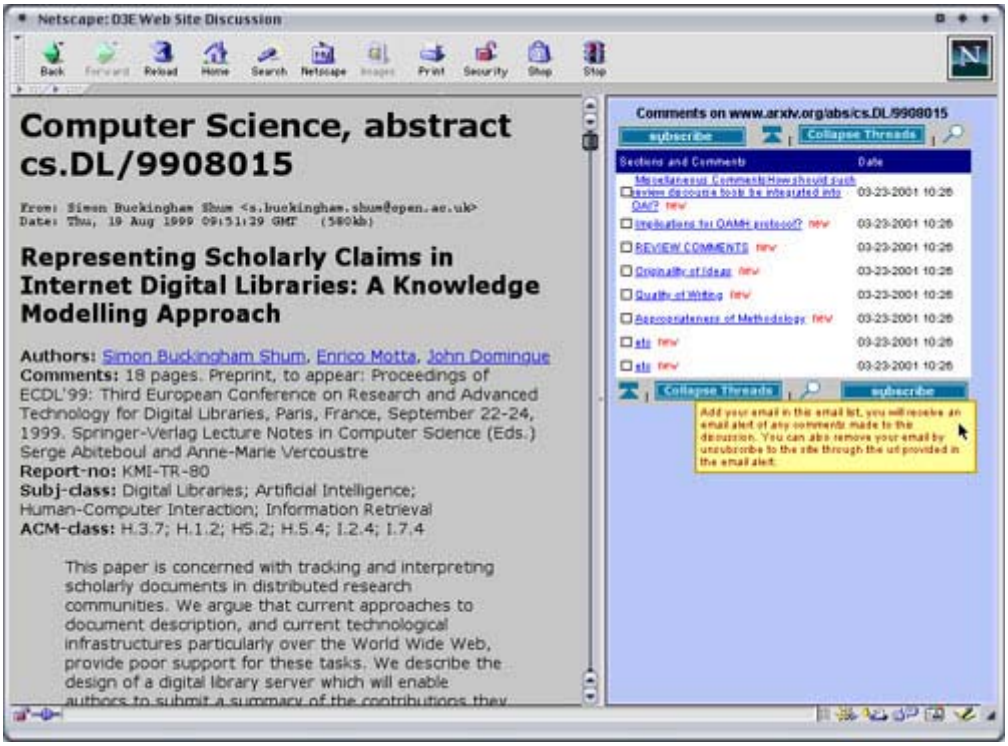

Figure 7: D3E

Systems that allow interrogation of written works are becoming increasingly common; indeed D3E can be seen as an example of this. Use of ontologies and natural language understanding programs underpin a piece of software called AQUA (Vargas Vera et al 2003), which allows the user to interrogate systems using natural language questions (Figure 8).

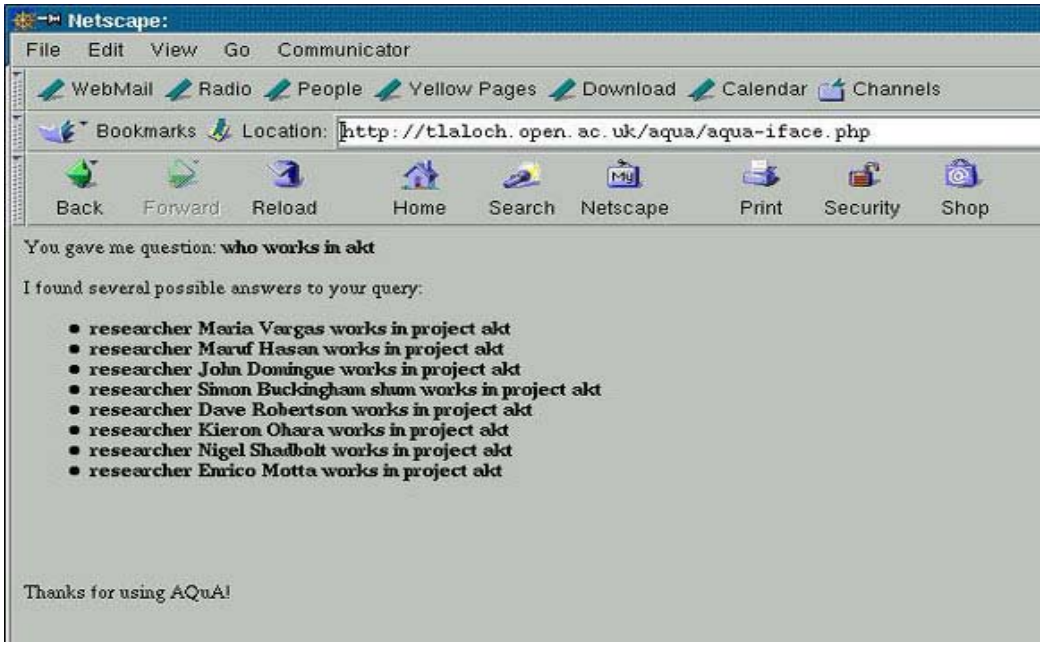

Figure 8: A query to AQUA

Finally, there is the temporal fixedness of speech. This has already been challenged by technologies such as recording and playback methods, but even so contextual cues are neglected. CoAKTinG (Buckingham Shum et al 2002) is a project that is designed to enable the vast amount of knowledge that is generated and held in meetings to be used by corporations. Figure 9 shows part of the CoAKTinG system; the panel at the top right is a map of the people in an organisation, showing who is available for discussion, and who is actually in discussion at any one moment. The other panels show a map of a discussion, based on an ontology of conversational moves – the various symbols on the map show, for example, whether a comment agreed with the previous one, disagreed, provided support, or changed the subject. Figure 10 shows a meeting being replayed. The video screens in the top left are videos of the actual meeting, which took place across three sites; these videos were played as part of the original videoconference. At the top right, the Powerpoint slide that was being shown at the time is displayed. Below, a coloured slider, showing who was speaking, allows the user to find any particular point in the discussion. At the bottom is an abstract specification, again using the ontology of conversational moves, of what was being said at the time and how it related to other themes.

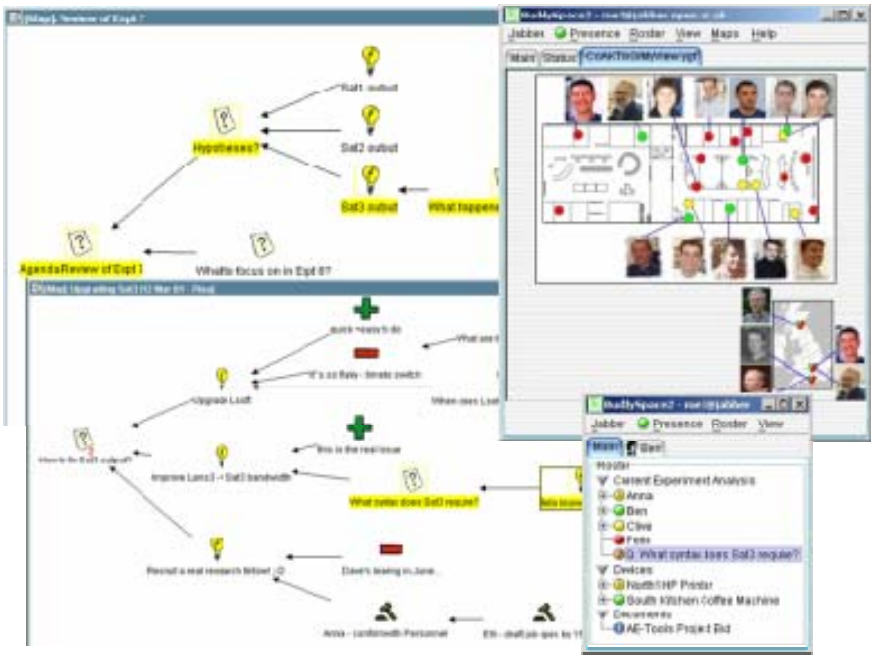

Figure 9: Online presence and conversational mapping

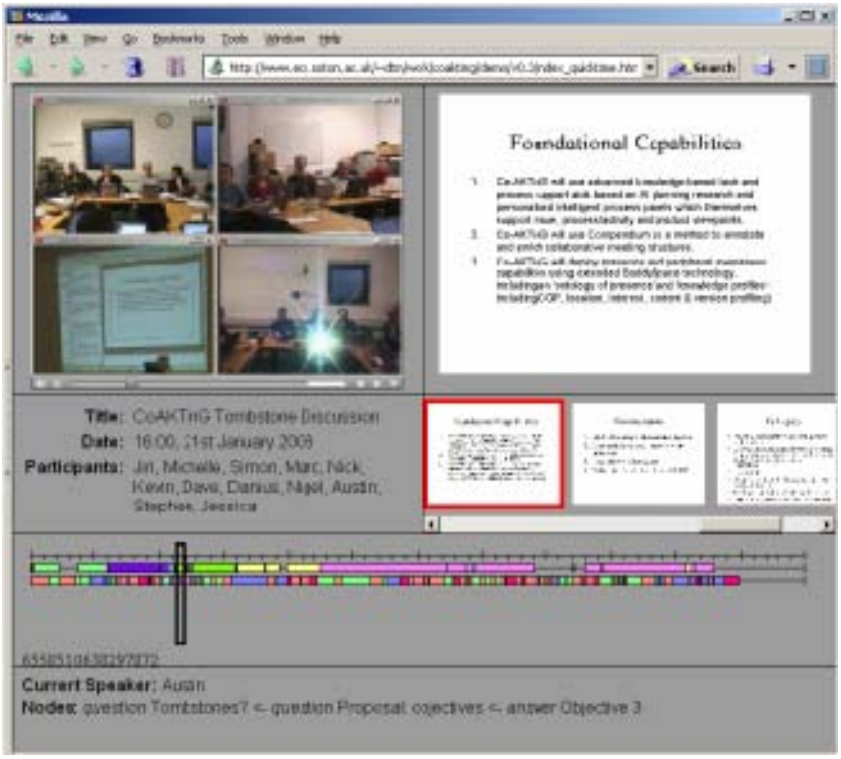

Figure 10: Replay and navigation through a meeting

Conclusion

So, intriguingly, we discover that the marks of the oral, that for Socrates ensured the trustworthiness of the oral, can, using SW technology, be applied to certain written contexts. The possibilities that the SW provides for dynamic, intelligent creation of

content in real time, allows this primarily written medium to mimic the instantaneity and context-sensitivity of the spoken. The Socratic oral criteria of trust is widely accepted in various sectors of society – as we saw, for example, in education or the law – and so the new expressivity enabled by SW technology is a very important step forward in the creation and fostering of online trust.

Nevertheless, it is still important to ask in what circumstances this adoption of the Socratic marks of the oral will facilitate trust. The answer is that, so far at least, the circumstances will be relatively limited ones. First, we must look for contexts where there is little or no institutional support. Institutions exploit the long-distance characteristics of the written: memory, regulation, sanction. Institutions spread trust far more efficiently than mechanisms based on speech and personal acquaintance (Fukuyama 1995). Hence where institutions exist online, with sufficient power over their subjects to enforce trustworthy behaviour, SW technologies are likely to have only a marginal effect on trust.

Second, since the enhancement of interaction is the whole point of the SW innovations, then we must look for contexts where interaction is possible and important. And third, there must be a relative lack of history; where two communities share a history of enmity and mistrust, as for example in Northern Ireland, or in Israel and Palestine, then that mistrust will of course not be removed simply by providing more channels of communication. Substantially more work than that is required for trust to emerge.

There are contexts that meet these characteristics, and in which therefore SW technology might help provide trust. They are fairly limited, and yet widespread enough for the technology to be of a good deal of social value. They include, for instance, the world of science, the world of academic discussion generally, and business. Political problems will take more time and effort to solve; but in science, academe and business there is a lot of goodwill waiting to be propagated, with tangible and clear benefits for all from greater interaction, knowledge sharing and transfers of resources.

Two final points deserve to be made, to act as a proper dampener on expectations. First, to begin with, much of what I have said in this essay seems to chime in with the theory of technology developed in the 1960s by Marshall McLuhan (McLuhan 1962), who predicted that the development of the new electronic media would, by the instantaneity that they support, herald a return to orality. McLuhan's coining of the famous term 'global village' is more significant here than is often remembered; the village is an acoustic space, where speech will always be more important (because more context-sensitive, more nimble) than writing (McLuhan & Fiore 1967). In a similar vein, Paul Levinson has described how on the WWW the written has been transformed, given speech-like qualities (Levinson 1999).

As far as they go, these views are interesting and prescient. However, the development of the SW does *not* – as yet – presage McLuhan's return to orality. Actually what we see on the SW is an opportunistic combination of literacy and orality, as technologies and work and social contexts permit. And it needn't be emphasised that the SW is not yet a sufficiently pervasive medium for its innovations to be socially transformative. Eric Havelock writes of the centuries after Ancient Greece became a literate society for

habits of literacy, the psychology of writing, to become ingrained (Havelock 1963). When the only examples of writing readily available were inscriptions on stone monuments, then that would not be enough to support the internalisation of discourse and the externalisation of memory that, psychologically, were the ultimate effects of writing (Ong 1982). Recall the limited set of contexts in which the SW's linguistic innovations will be effective; the SW is similarly unpervasive, and its social effects will be correspondingly limited.

And second, we must of course note that potential is *not* actuality; the SW is a very exciting series of developments, but it could equally go the way of other great technologies that just didn't catch on. Its catching on, like the WWW, will depend on the timely creation of brilliant tools that bring in new users – tools such as search engines like Google, or the ability to transfer money around securely in commercial transactions. These tools have yet to be developed, and until they are, the SW will remain a ghetto in the computer science department.

The way we develop, the way we use technology to externalise psychological processes – as first described, memorably, by Socrates in the *Phaedrus* – is a process that may well be influenced by unpredictable forces, such as market forces, but equally is at least partly under our control. *Contra* Lyotard, the human/inhuman distinction remains always to be negotiated anew (Lyotard 1991). And, *contra* the hopes of many cybertheorists – most prominently Donna Haraway, who seeks to exploit technology as a way of burying identity and breaking down conceptual barriers (Haraway 1991), and Kevin Warwick (Warwick 2002) – the new flexibility that the SW provides will not, for the foreseeable future, bring the Cyborg concept any closer.

The Semantic Web can be radical, and could be transformative. But predicting the course of such transformation is much more difficult than exaggerating it. As computer scientist Charles Jonscher once wrote:

> I have gleaned two lessons from the history, short as it has been, of electronic technology; they are, inevitably, over-simplifications, but they have been my guide. The first is to regard almost any prediction of the future power of the technology itself as understated. The second is to regard almost any prediction of what it will do to our everyday lives as overstated. (Jonscher 1999, p.248)

References

Alani H., Kim S., Millard D.E., Weal M.J., Hall W., Lewis P.H. & Shadbolt N.R. 2003. 'Automatic ontology-based knowledge extraction from web documents' *IEEE Intelligent Systems* 18(1): 14-21.

Berners-Lee T., Hendler J. & Lassila O. 2001. 'The Semantic Web' *Scientific American* May 2001.

Buckingham Shum S., De Roure D., Eisenstadt M., Shadbolt N. & Tate A. 2002. 'CoAKTinG: collaborative advanced knowledge technologies in the grid' *11th IEEE Symposium on High Performance Distributed Computing*, Edinburgh, July 2002.

Buckingham Shum S. & Sumner T. 2001. 'JIME: an interactive journal for interactive media' *Learned Publishing* 14: 273-285.

Carr L., Bechhofer S., Goble C. & Hall W. 2001. 'Conceptual linking: ontology-based open hypermedia' *10th World Wide Web Conference*, Hong Kong, 2001.

Fukuyama F. 1995. *Trust: The Social Virtues and the Creation of Prosperity.* New York: Free Press.

Haraway D. 1991. *Simians, Cyborgs and Women*. London: Free Association Press.
Havelock E.A. 1963. *Preface to Plato*. Oxford: Basil Blackwell.
Hughes R. 1996. *English in Speech and Writing: Investigating Language and Literature*. London: Routledge.
Jonscher C. 1999. *Wiredlife: Who Are We in the Digital Age?* London: Anchor.
Koivunen, M.-R. & Miller E. 2001. 'W3C Semantic Web activity' *Semantic Web Kick-Off Seminar*, Finland 2001.
Lessig L. 1999. *Code and Other Laws of Cyberspace.* New York: Basic Books.
Levinson P. 1999. *Digital McLuhan: A Guide to the Information Millennium*. London: Routledge.
Lyotard J.-F. 1991. *The Inhuman: Reflections on Time*, Bennington G. & Bowlby R. (trans.). Oxford: Blackwell.
McLuhan M. 1962. *The Gutenberg Galaxy*. Toronto: University of Toronto Press.
McLuhan M. & Fiore Q. 1967. *The Medium is the Massage: An Inventory of Effects*. Harmondsworth: Penguin.
Misztal B. 1996. *Trust in Modern Societies.* Cambridge: Polity Press.
O'Hara K. 2002a. *Plato and the Internet*. Cambridge: Icon Books.
O'Hara K. 2002b. 'The Internet: a tool for democratic pluralism?' *Science as Culture* 11(2): 287-298.
O'Hara K. 2004. *Trust: From Socrates to Spin.* Cambridge: Icon Books.
Ong W.J. 1982. *Orality and Literacy: The Technologizing of the Word*. London: Methuen & Co.
Plato. 1997. 'Phaedrus' (A. Nehamas & P. Woodruff, trans.) in J.M. Cooper (ed.) *Plato: Complete Works*. Indianapolis: Hackett. 506-556.
Quine W.V. & Ullian J.S. 1970. *The Web of Belief*. New York: Random House.
Sunstein C. 2001. *Republic.com*. Princeton: Princeton University Press.
Vargas-Vera M., Motta E. & Domingue J. 2003. 'AQUA: an ontology-driven question answering system' *AAAI Symposium on New Directions of Question Answering*, Stanford, March 2003.
Warwick K. 2002. *I, Cyborg*. London: Century.