Trust on the Web: Some Web Science Research Challenges

In *UoC Papers: E-Journal on the Knowledge Society*, 7, 2008

Kieron O'Hara & Wendy Hall

Intelligence, Agents, Multimedia Group
School of Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
{kmo,wh}@ecs.soton.ac.uk

Abstract: Web Science is the interdisciplinary study of the World Wide Web as a first-order object in order to understand its relationship with the wider societies in which it is embedded, and in order to facilitate its future engineering as a beneficial object. In this paper, research issues and challenges relating to the vital topic of trust are reviewed, showing how the Web Science agenda requires trust to be addressed, and how addressing the challenges requires a range of disciplinary skills applied in an integrated manner.

Key words: Web Science, trust, privacy

Introduction: Web Science

The Web is one of the most ubiquitous and transformative technologies humankind has ever produced, and in a remarkably short space of time has been embedded into a very large number of social activities, ranging from e-science to e-commerce, from e-government to entertainment, from citizen journalism to cybercrime. Yet we still remain in remarkable ignorance of the trajectory of the Web's development (and, conversely, of what risks the Web faces given its major exposure to the world); now-mainstream activities, such as blogging, file-sharing or social networking were unheard-of just a few years ago, and their appearance and sudden blossoming took most scientists and pundits by surprise.

There are doubtless many reasons for this, not least being the rapid inflation of such activities from niche practices to mainstream behaviour. But one of the most important is that we are lacking the conceptual tools and focused effort required to understand the Web. Of course it is a piece of computer engineering, but it is not simply the sum of TCP/IP, HTML, HTTP, PageRank, Ajax, URIs and whatever else. It is also created, written, linked and read by hundreds of millions of people.

Hence the Web is beyond the purview of any individual discipline, even computer science. Google's search algorithm PageRank is an impressive piece of work, but understanding the algorithm does not tell you about its place in the Web; for that you would need to understand the function of search, the complex real-world environment plagued by bad behaviour such as Google-spoofing, the economics of Google's click-based advertising business model, the engineering of Google's indexing and caching methods and so on. Web Science transcends faculties, requiring theoretical science,

empirical science, engineering, social science and humanities. In 2006, a group of computer scientists launched the Web Science Research Initiative (WSRI - http://webscience.org/) to help create the interdisciplinary study area that would support our understanding of the Web. As the founding directors of WSRI put it:

> Web science is about more than modeling the current Web. It is about engineering new infrastructure protocols and understanding the society that uses them, and it is about the creation of beneficial new systems. It has its own ethos: decentralization to avoid social and technical bottlenecks, openness to the reuse of information in unexpected ways, and fairness. It uses powerful scientific and mathematical techniques from many disciplines to consider at once microscopic Web properties, macroscopic Web phenomena, and the relationships between them. Web science is about making powerful new tools for humanity, and doing it with our eyes open. (Berners-Lee et al 2006).

In this paper we wish to present an example of a particular problem whose solution(s) require the full breadth of scale and disciplinary experience to which Berners-Lee et al allude: the problem of trust on the Web.

There are a number of reasons why trust is an issue online: the decentralised environment, the lack of supporting contextual factors, the artificiality of many agents, the fluidity of identity, the highly heterogeneous user base. But most importantly trust is essential. Should you trust the content you download, given the lack of a central moderator? If you use automated services, should you rely on their effectiveness, and should you give them access to sensitive information? How do you know that people are representing themselves accurately in social networks? Any Web user must learn to place trust in content, in services and in people wisely and safely. Jennifer Golbeck (Golbeck 2006) lists three major challenges to applying trust online. *Trust management* is the process of determining who has access to which information or resources. *Trust computation* is the method of deriving a level of trust in a resource on the basis of the available data. *Integrating trust into applications* involves building applications that can function by placing trust accurately enough for its purposes.

In the remainder of this paper, we will survey research challenges relating to online trust from the point of view of Web Science, showing that a road map for Web Science research into the problem of trust will among other things have to cross varying scales and disciplines. We begin philosophically and psychologically, examining the features of trust and its management that would ideally be replicated online. Next, we switch to sociology, and examine some of the attitudes to trust that prevail among Web users. The next section takes the position of politics and security, looking at the relationship between trust and the related concept of privacy. Finally, we bring in technology, examining developments to promote trust on the Semantic Web. In the conclusion, we discuss how ‘the problem of trust’ has, from a point of view consistent with a traditional division of disciplines, decomposed into several research challenges. We take this as *de facto* evidence that the inclusive approach of Web Science is essential.

Trust and trustworthiness

Without trust the Web wouldn’t function; exchanges of resources or information require all sorts of risk-taking. But what do we get in return for the risk? One clear function of trust is as a method of complexity reduction (Luhmann 1980). Trust

enables us to ask or pay others to act on our behalf. If we trust them, we do not have to carry out the subtask ourselves, and neither do we have to monitor or micromanage their performance. It enables us also to ask advice or receive instruction; if we do not trust our advisors, we will have to acquire their expertise before we can act confidently. So, although there is a considerable body of evidence that trust has a large moral component (Uslaner 2002), trust performs the important social function of increasing the efficiency of social interaction. Online, the moral dimension to online trust is less well established, and may not be very relevant at all (social networking may be changing this). Online trust is generally reduced to an evidence-based cost/benefit/risk analysis of expectations of whether performance will live up to our expectations (see the survey of methods in Golbeck 2006).

Connecting trust and trustworthiness

Nevertheless, although trust is a good thing, and commentators have argued that high-trust societies have advantages over low-trust societies (Fukuyama 1995), we should note that increasing trust is not a solution to all social problems. There is an important distinction that is often blurred in discussion between trust and trustworthiness (cf. Hardin 1996). Trust is an attitude of one agent, X to another, Y. Trustworthiness is a property of an agent. These are usually relativised to a task – one person trusts another *to do something particular* (we will ignore this caveat in the rest of our discussion). So, X trusts Y just in case X *believes* Y is trustworthy.

X trusting Y will certainly reduce the complexity of X's life, but if X's belief is actually false, then the cost to X in resources risked in the transaction will be correspondingly high. What is just as important to X is that her belief about Y is actually true (viz, that Y actually is trustworthy). Conversely, Y's trustworthiness does him no good as long as no-one believes he is. There is an opportunity cost to being trustworthy without trust.

The incentives are skewed in a difficult way. X benefits from Y's trustworthiness, but controls only her trust; Y benefits from X's trust, but controls only his trustworthiness. What is essential – what a functioning society does – is to link trust and trustworthiness so that ideally all and only trustworthy people are trusted. The challenge, then, is not how to increase trust, but rather how to create a causal link between trustworthiness and trust. This presents us with the first set of Web Science research challenges.

- How should we maintain a causal link between trust and trustworthiness? What incentives and economic models are available to promote trust and trustworthiness together?

Local and global trust

Trust can be decomposed into *local* and *global* trust (O'Hara 2004), depending on the evidence one uses for supporting one's beliefs about the trustworthiness of others. Local trust relies on personal acquaintance mediated through high-bandwidth interactions. We take note of varying signals given out by agents – facial expressions, clothing, language and so on, and these signals provide the causal link that connects trust and trustworthiness. The downside of the system is that once a signalling convention (wearing a suit, say) has been adopted and codified in a society, then it can be faked by untrustworthy people. Hence trustworthiness signalling systems need to be updated constantly in an arms race with untrustworthy counterfeiters.

Global trust involves outsourcing our trust decisions to trusted institutions, so X would trust Y on the basis of a certificate from an institution that Y was trustworthy. Such institutions do not 'solve' the problem of trust – they merely shift it, because X must decide whether she trusts the certifying institution, not Y. But the institutions do have the important effect of globalising trust, because X can take informed trust decisions about people she has never met. An institution has all sorts of economies of scale that enable it to perform more thorough investigations of its subjects, and so it can establish a very strong connection between trust and trustworthiness. However, the downside is systemic risk – one mistaken certification could lead to all the institution's output being ignored, and its clients withdrawing their trust from everyone it has certified.

The variables used for gathering evidence will affect the type, and therefore the scope, of trust. For instance, Huynh et al (2006) present a composite model where four types of mechanism are combined to produce a decision to trust or not to trust in an online agent context.

- Interaction trust. Based on past experience of direct interactions.
- Role-based trust. Defined by role-based relations between the parties.
- Witness reputation. Based on reports from witnesses of past behaviour of the agent.
- Certified reputation. Based on third-party references provided by the agent.

Broadly, local trust is provided by the first of these, the actual experiences of the trusting party in his/her history, and current experiences in looking at e.g. a website. The other mechanisms supply global trust (e.g. the roles the two parties play such as citizen/government, customer/retailer, customer/bank, ISP/customer allow the parties to go beyond personal experience to place their trust).

The decentralised nature of the Web, plus its rapid changes, mean that there are few trusted institutions online. Many institutions trade on their offline reputations (for instance, many banks, universities and government institutions), while very few have developed a reputation entirely online (PayPal would be one example). So despite the global nature of the Web online trust is often local, in that it relies on a person's personal experience in dealing with a website. A user interacts with a website, and assesses the signals given out by that website him- or herself.

There are two obvious problems with this. First, online the user is deprived of the complexity of signal available in the offline world (which include quite unconscious signals of (un)trustworthiness such as a shaking hand or unconfident expression). Online, the signals are basically the visual ones described by the HTML source file of the page, augmented possibly by the roles played by the parties in the transaction (cf. role-based trust), which in general is not a secure source of trust on its own. And second, the designer of the website is in total control of the signals that it gives out; the user has little or no opportunity to engage the website in 'conversation', to see how it 'performs', to 'size it up', as we do offline when we are judging people. When trust is local, based on personal acquaintance, the dice are loaded in favour of the website in these two ways. This presents us with a second set of research challenges.

- How should trust be represented, maintained and repaired on the Web? What variables are important? Will these change as we move from human to

artificial agents? What sort of institutions and methods will globalise online trust?

Bootstrapping

Sometimes the trust/trustworthiness link needs to be started on the basis of little evidence – the so-called bootstrapping problem. Y sets up in business – should X trust him? She cannot without evidence that Y conducts his business in a trustworthy fashion, and when Y has just started there is no such evidence. But unless someone trusts Y to work for them, there will never be such evidence.

Offline, we have a number of strategies – many exploiting the moral nature of trust. There has been a lot of debate between the Weberian idea that trustworthiness causes trust, and the Durkheimian position that trusting people inclines them to behave in a trustworthy manner. In fact, we muddle by with a combination of the two, breaking into the circle and bootstrapping the relationship. We rely on moral notions such as duty, and inclusion of people into our moral community.

But online, bootstrapping is a problem, partly because of the relative unimportance of the moral dimension (with the caveat that social networking may be bringing a more Durkheimian model of trust with it). Most models of online trust are evidence-based, and it is evidence that is lacking in the first place. In terms of the model of Hyunh et al, *ex hypothesi* there is (a) no past experience of direct interactions, (b) a lack of experience in any meaningful role, (c) no witness reports, and (d) no evidence-based third-party references. But trust has to begin, somehow. This gives us another set of research challenges.

- How can effective and accurate trust be bootstrapped?

Web users

Bootstrapping is also important when we consider the Web as a social system, rather than a collection of linked hypertext documents and data. Trust in Internet transactions is higher, unsurprisingly, with Internet users than non-users (Dutton & Helsper 2007); the growth of the Web does depend on those users achieving a level of familiarity in order to reduce uncertainty and increase their confidence. Even ex-users of the Internet trust it more than non-users.

People are remarkably accepting of bad experiences online, such as spam or obscene email, at least up to a point. Such ‘anti-social behaviour’ is assumed to reduce levels of trust, so this might seem a surprising result given the prevalence of spam etc as reported in the media and academic work. One explanation is that those providing such reports tend to be heavy users of the Internet, and therefore much more likely to have bad experiences; general patterns of use must also include a large number of people who use the Internet relatively rarely, but whose voices are seldom heard in discussions about the Internet or the Web (Dutton & Helsper 2007).

As one would perhaps expect from the previous section, people’s online trust mirrors the experience of those placing local trust, i.e. unmediated and based on personal acquaintance (in the online case, with a website). For instance, one review of online trust discerned three *perceptual* factors that were particularly relevant. *Perception of credibility* is to do with honesty, expertise, predictability and reputation. *Ease of use* relates to the simplicity and design of the website. *Risk* is the perceived likelihood of an undesirable outcome (Corritore et al 2003). Risk is of course a pervasive issue with

trust, but the other two factors are strongly connected to the gathering and evaluation of signals of trustworthiness. Credibility signals are designed to display the trustee's expertise, but ease-of-use signals, which include having a well-designed site, avoiding such pitfalls as bad spelling and dangling links, are strictly unconnected with credibility and easy to fake, yet are still important signals. This confirms the findings of an earlier study which found six major features that encouraged trust in e-commerce sites – the site's brand, seals of approval, ease of navigation, a fulfilling ordering experience, the site's presentation and the technologies used to create the website – which again are strongly connected with the signalling systems characteristic of local trust (Cheskin Research 1999).

However, Web users are not particularly efficient at picking up the right signals that provide the causal connection between trust and trustworthiness. Dhamija and colleagues (2006) investigated the reasons why bogus sites work, and discovered to that existing anti-phishing browser cues – the 'signals' which users are supposed to pick up, and which connect trust and trustworthiness – are ineffective. A participant group in that experiment made mistakes 40% of the time (even though primed to look out for phishing sites), and surprisingly neither age, gender nor computing experience were significant variables. The study showed that people are unaware of the sorts of signalling systems that have been developed to ensure trustworthiness (e.g. the padlock symbol to show that the page was delivered securely by SSL), or of the typical strategies of counterfeiters (e.g. using images to mask underlying text, or placing an SSL-padlock in the body of a webpage). Furthermore, users often fail to notice the *lack* of expected signals of trustworthiness. Attention to the needs of actual Web users leads to a further set of Web Science research challenges.

- How can secure systems be made usable and effective for consumers, given the bounded knowledge and rationality of Web users?

Privacy

One of the biggest obstacles to trust on the Web is the threat that digital information provides to the user's privacy (O'Hara & Shadbolt 2008). Thanks to the Internet and the Web, information is very easily copied and transferred. These technologies were founded on a liberal ideology of free-flowing information, conceived in the context of fast-moving scientific research where access to data was limited and the publishing cycle was slowing down the research cycle. The Web provided a means for information-sharing that has clearly boosted, indeed transformed, research.

However, outside the realm of science, information, particularly personal data and intellectual property, has a value which may be threatened by copy and transfer. The trade-off between the benefits and costs of freedom of information is an acute problem on the Web (though not unique to it). It is clear that problems with online privacy are perceived as the major risks by users (Dutton & Helsper 2007). Furthermore, as in general wealthier and better-educated people are more likely to be experienced online users, there is an unequal distribution of privacy whereby privacy-aware individuals have the means to protect privacy (at some cost), while people with a strong Web presence who remain unaware of the risks or unable to understand or afford protections are potential targets (O'Hara & Stevens 2006).

This is a classic Web Science problem, at the interface between group behaviour, individual perceptions, the politics of information and Web engineering. Furthermore, many issues pertaining to security pop up here too, such as the need to retain usability,

and to integrate solutions into standard workflow and Web use. One multi-layered attempt to address this problem is the *Policy-Aware Web* (PAW – Weitzner et al 2005), which suggests a mechanism whereby, if a browser requests information (for example, images to be displayed within a Web page), it receives a modified HTML 401 error which includes the URI of the site's privacy policy (a statement of restrictions of access to the information). The browser would then construct a proof, based on its own credentials, that the policy was satisfied. The site would then check the proof, and release the information if the proof was valid.

Such an approach has a lot of promise, but there are still tensions within society as a whole between privacy and transparency. New types of interaction including social networking, lifelogging and the maintenance of virtual identities will demand new assessments of the risks to privacy (Bailey & Kerr, 2007; O'Hara et al 2008) and consequent re-examination of the law (Allen 2008). It is in this space that Web Science faces some of its most daunting challenges.

- What privacy issues arise in a Web environment of increasingly sophisticated information sharing? Can traditional forms of regulatory regime cope? To what extent are the service providers going to become the legal gatekeepers for public authorities in terms of delivering their public policy objectives, e.g. Web policing for what is judged to be "illegal and harmful content"?

Privacy, trust and security issues loom large in a number of specific areas of Web activity, for instance e-health. The use and ownership of information set up another set of research challenges.

- What legal regimes are appropriate where users are heterogeneous and often inexperienced? For example, in e-health scenarios, where it is important that professionals have effective and timely access to information, who should own information about patients, where patients have an interest in privacy but many others have an interest in the information?

The Semantic Web, provenance and social networks

The use of the Web's own technology (URIs, theorem proving, error handling) to address privacy with the PAW points to the possibility of doing the same thing for trust. This is indeed essential, as projected developments to the Web, particularly including the Semantic Web (Shadbolt et al 2006) are premised on the inclusion of a technological layer to address the issue of trust. If we look at a layered view of the various formalisms and protocols which are currently seen as making up the Semantic Web (SW), trust has a vital place at the top of the edifice (Figure 1).

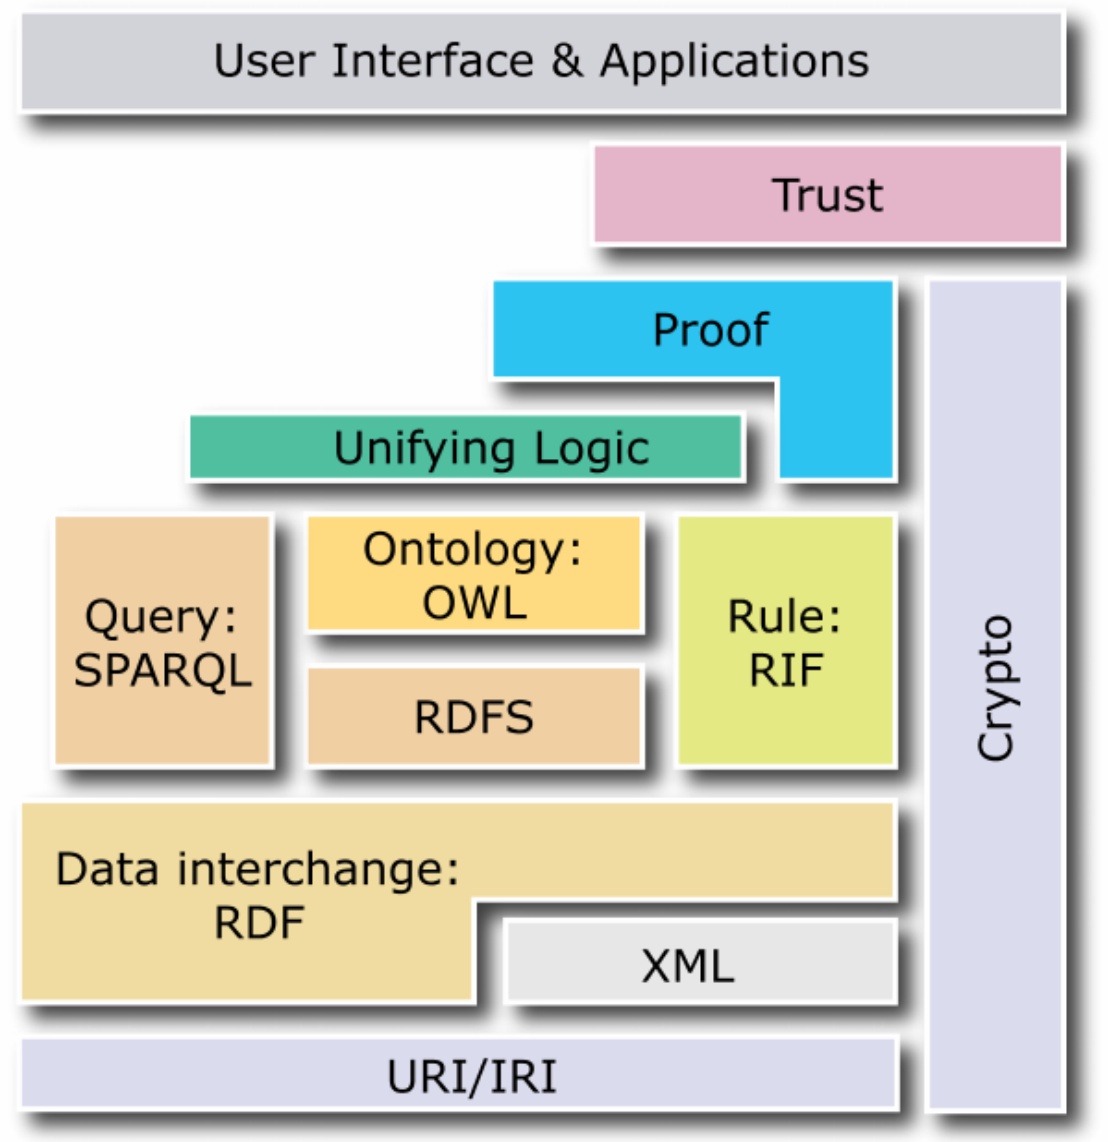

Figure 1: The Layered View of the Semantic Web

Work on the trust layer is at an early stage, and is currently very fragmented, partly because of a range of opinion on what is likely to establish trust in the data which the SW is intended to make available (Golbeck 2006). It is certainly the case that without a trust layer, the amount of data which will be made available may well be less than the SW's designers hope – another example of the interaction between social pressures and technological change.

There are a number of approaches under research, including the proof-based idea implicit in the PAW. Another aspect is the creation of metadata about the provenance of an information resource, detailing where it came from and what methods were used to generate it; this metadata can be used to inform trust decisions. Moreau et al (2008) describe a method of metadata provision that is crucially sensitive to the lifecycle of provenance, using an open data model for documentation of a resource which serves user queries over a representation of what processes were involved in generating that resource. The key point is that this idea is consistent with a number of computing trends towards open applications, composed dynamically which derive results on the fly.

A third method exploits a quirk about trust, which, although not transitive (if A trusts B, and B trusts C, it does not follow that A trusts C, or even that A should rationally trust C), has transitive qualities (A's trust of B might lead A to give B's opinion of C extra weight). So, for example, Richardson et al (2003) use social networks with trust to calculate the belief a user might have in a statement by finding paths through a social network from the user to a node that represents the opinion in question. Trust values along the paths are concatenated and aggregated to provide a global trust value.

One influential SW project is Friend-of-a-Friend (FOAF), a framework for representing information about people and their social connections, enabling the iterative creation of a social network via connective predicates such as "A knows B" (Brickley & Miller 2007). Such a network can make it possible to make judgements about someone you don't know, purely by seeing their place in a network. For instance, one could decide to allow access to your information to people within two steps of you in the network. Or you could use a weighting algorithm to determine how trustworthy someone is likely to be, given that he or she is known in some role by certain people you trust. Golbeck et al have added a trust module to FOAF allowing people to rate how much they trust each other, either with respect to a particular topic, or in general (Golbeck et al 2003). FOAF is beginning to be applied not just to the Web as a whole, but to the important sub-world of social networks (Golbeck & Rothstein 2008; O'Hara et al 2008).

The fact that there is a wide variance in methods of computing trust, and a similar range of contexts in which that has to be done, means that a single means of dealing with the problem is unlikely – and consequently that there will not be a purely technological solution, and that the interdisciplinary range of Web Science will need to be leveraged. For instance, analysis of the costs and benefits of different strategies in the manner of the framework (O'Hara et al 2004) will be required; this framework does not draw conclusions but rather details the costs of various strategies of applying trust, and also subdivides the costs into operational costs, opportunity costs, risk, costs of betrayals and service payments.

The question of bootstrapping trust also has to be taken into account – how should a new entrant 'break into' a social network? Supplying provenance metadata has to be part of the answer, as does a policy-based approach. And one advantage of the social networking approach is that once one has established social relations with one or two of the network members, one does have some sort of network presence, however minimal, which provides an opportunity for more transactions potentially leading to more trust. (O'Hara et al 2004) also discuss which strategies for placing trust can help circumvent the bootstrapping problem. In general, the more optimistic the strategy, the better for bootstrapping. Centralised strategies also can work, but not only are they hard to scale, but they work against the Web's decentralised ethos.

Hence the technology presents us with more research challenges.

- What languages and ontologies are appropriate for expressing the requirements for online trust? At the moment, the work on trust on the Semantic Web is relatively sparse and unfocused – how should it be focused?
- How can accountability and transparency be engineered into information use? Given the ability to describe information policies, how can they be enforced? How can we ensure the quality of provenance metadata?

- How should the trust layer of the Semantic Web be integrated with the layers lower down to create a seamless interaction for the user?

What is fascinating is that this set of research challenges brings us right back to the start, and our first set of challenges about establishing and maintaining the causal link between trust and trustworthiness.

Conclusion

Trust cannot be engineered, but mechanisms can be put into place that aid standard mechanisms to create trust. As noted above, the Web has aspects which put trust at risk; it is by going with the grain of society that the Web can promote trust of itself and also support the globalisation of trust in wider society. To understand the grain of society, it is essential to understand the interactions of online trust at a number of levels ranging from the micro (the protocols governing the transfer of information) to the macro (the social effects of information being passed around) and all stages in between (for example, the individual psychology of online trust).

There are many specific research challenges with which Web Scientists must wrestle, many of which we have highlighted and bulleted in this paper. Any or all of these would make fascinating research projects or PhD topics – but none of them can be properly approached from the perspective of a single discipline, even computer science. Our aim has been to indicate the broad range of disciplines required to understand the problem; it is doubtless not an exhaustive list, but in its breadth it is perhaps a very strong argument for the importance of Web Science for the future not only of the Web, but our Web-enabled society.

References

ALLEN, A.L. (2008). “Dredging up the past: lifelogging, memory and surveillance”. *University of Chicago Law Review*. Vol.75, pp.47-74.

BAILEY, J., KERR, I. (2007). “Seizing control? The experience capture experiments of Ringley & Mann”. *Ethics and Information Technology*. Vol.9, pp.129-139.

BERNERS-LEE, T., HALL, W., HENDLER, J., SHADBOLT, N., WEITZNER, D. (2006). “Creating a science of the Web”. *Science*. Vol.313, pp.769-771.

BRICKLEY, D., MILLER, L. (2007). *FOAF Vocabulary Specification 0.91*, http://xmlns.com/foaf/spec/ [Date of consultation 11/7/08].

CHESKIN RESEARCH AND STUDIO ARCHETYPE/SAPIENT (1999). *eCommerce Trust Study*, http://www.cheskin.com/cms/files/i/articles//17__report-eComm%20Trust1999.pdf [Date of consultation 10/7/08].

CORRITORE, C.L., KRACHER, B., WIEDENBECK, S. (2003). “On-line trust: concepts, evolving themes, a model”. *International Journal of Human-Computer Studies*. Vol.58, pp.737-758.

DHAMIJA, R., TYGAR, J.D., HEARST, M. (2006). “Why phishing works”. *Conference on Human Factors in Computing Systems (CHI 2006)*, http://people.seas.harvard.edu/~rachna/papers/why_phishing_works.pdf [Date of consultation 10/7/08].

DUTTON, W.H., HELSPER, E.J. (2007). *The Internet in Britain 2007*. Oxford: Oxford Internet Institute.

FUKUYAMA, F. (1995). *Trust: The Social Virtues and the Creation of Prosperity*. New York: Free Press.

GOLBECK, J. (2006). "Trust on the World Wide Web: a survey". *Foundations and Trends in Web Science*. Vol.1, Iss.2, pp.1-72.

GOLBECK, J., PARSIA, B., HENDLER, J. (2003). "Trust networks on the Semantic Web". *Proceedings of Cooperative Intelligent Agents*, Helsinki, 2003, http://www.mindswap.org/papers/CIA03.pdf [Date of consultation 11/7/08].

GOLBECK, J., ROTHSTEIN, M. (2008). "Linking social networks on the Web with FOAF: a Semantic Web case study". *Proceedings of AAAI '08*. http://www.cs.umd.edu/~golbeck/downloads/foaf.pdf [Date of consultation 11/7/08].

HARDIN, R. (1996). "Trustworthiness". *Ethics*. Vol.107, Iss.1, pp.26-42.

HYUNH, T.D., JENNINGS, N.R., SHADBOLT, N.R. (2006). "An integrated trust and reputation model for open multi-agent systems". *Journal of Autonomous Agents and Multi-Agent Systems*. Vol.13, pp.119-154, http://eprints.ecs.soton.ac.uk/12593/1/jaamas-dong.pdf [Date of consultation 10/7/08].

LUHMANN, N. (1980). *Trust and Power*. Chichester: Wiley.

MOREAU, L., GROTH, P., MILES, S., VÁZQUEZ-SALCEDA, J., IBBOTSON, J., JIANG, S., MUNROE, S., RANA, O., SCHREIBER, A., TAN, V., VARGA, L. (2008). "The provenance of electronic data". *Communications of the ACM*. Vol.51, Iss.4, pp.52-58.

O'HARA, K. (2004). *Trust: From Socrates to Spin*. Duxford: Icon.

O'HARA, K., ALANI, H., KALFOGLOU, Y., SHADBOLT, N. (2004). "Trust strategies for the Semantic Web". *Workshop on Trust, Security and Reputation on the Semantic Web*, Hiroshima, Japan, http://eprints.ecs.soton.ac.uk/10029/ [Date of consultation 11/7/08].

O'HARA, K, SHADBOLT, N. (2008). *The Spy in the Coffee Machine: The End of Privacy As We Know It*. Oxford: Oneworld.

O'HARA, K., STEVENS, D. (2006). *inequality.com: Power, Poverty and the Digital Divide*. Oxford: Oneworld.

O'HARA, K., TUFFIELD, M.M., SHADBOLT, N. (2008). "Lifelogging: issues of privacy and identity with Memories for Life". *Workshop on Identity in the Information Society*, Arona, Italy, May 2008, http://eprints.ecs.soton.ac.uk/15993/ [Date of consultation 11/7/08].

RICHARDSON, M., AGRAWAL, R., DOMINGOS, P. (2003). "Trust management for the Semantic Web". In D. Fensel, K. Sycara, J. Mylopoulous (ed.), *The Semantic Web – ISWC 2003*. Berlin: Springer, pp.351-368.

SHADBOLT, N., HALL, W., BERNERS-LEE, T. (2006). "The Semantic Web revisited". *IEEE Intelligent Systems*. Vol.21, Iss.3, pp.96-101.

USLANER, E.M. (2002). *The Moral Foundations of Trust*. Cambridge: Cambridge University Press.

WEITZNER, D.J., HENDLER, J., BERNERS-LEE, T., CONNOLLY, D. (2005). "Creating a Policy-Aware Web: discretionary, rule-based access for the World Wide Web". In E. Ferrari, B. Thuraisingham (ed.), *Web and Information Security*. Hershey,

PA: Idea Group Inc, http://www.w3.org/2004/09/Policy-Aware-Web-acl.pdf [Date of consultation 11/7/08].