Active: An Essential Nature of the (Future) Web?

Joe Geldart William Song
University of Durham
Department of Computer Science
Durham, DH1 3LE
j.r.c.geldart@dur.ac.uk

ABSTRACT
In this position paper, we describe our vision of the evolution of the Web: the Active Web. We argue for a phased view of Web evolution, describing some trends seen in development so far. We then describe the social, contentious and dynamic nature of human knowledge and show how this may be developed into a principled road map for Web development. We finish by describing a selection of usage scenarios, providing a view of our intended endpoint for this phase of the development of the Web.

Categories and Subject Descriptors
H.3.5 [**Online Information Services**]: Web-based services

General Terms
Human Factors, Standardization, Languages, Theory.

Keywords
World Wide Web, Semantic Web, Pragmatics, Cognitive Science, Multi-Agent Systems, Common-Sense Reasoning.

1. INTRODUCTION
The Semantic Web [1] has seen many exciting developments in the past few years of research and development. As a technology, it is enabling organizations and individuals to consolidate their information in the form of densely linked open data. This enables their processes to work smarter, and harder, for their needs.

The development of the Semantic Web has also highlighted a number of properties of the Web that we feel deserve greater examination. Certainly, the concomitant rise of social networks has provided for a great use-case for the Semantic Web with the Friend-Of-A-Friend (FOAF) project [6][4] fast becoming one of the most commonly used RDF vocabularies (coming high in the ebiquity survey, narrowly behind RDF itself and RDF Schema [5].) Simultaneously though, these networks also suggest that many people are most interested in *social* information; that is, information which arises through social interactions, like Wikipedia.

The purpose of this paper is to act as a rallying call for what we feel is the next stage of the Web's study and evolution. It makes three main contributions to the debate:

- A reanalysis of the key ideas in the development of the Web in terms of phases of research (section 2)
- A proposal to treat information as a socially situated construct, dissolving the traditional boundaries of clients and servers for the purposes of modeling the Web Architecture (section 3)
- A definition of a selection of usage scenarios for the next phase of the Web, intended to inspire development and (heated) discussion (section 4).

This paper should be treated as deliberately contentious and provocative. It does not focus on the technical details of what it discusses, but rather is intended to spur on discussion about what those technical details may turn out to be.

2. A NOMENCLATURE OF PHASES
We feel the marketing-inspired (and insipid) label of 'Web 2.0' is merely an old idea in slightly different packaging. Exploiting an existing model's side aspects when implemented by a particular system (in this case, client-side scripting combined with a dynamic server) demonstrates the flexibility of the model but doesn't significantly advance our understanding of the model's concepts. What other form of classification can we then impose upon models of the Web?

The Web was originally developed as a way of sharing hyperlinked documents with other people whether they be across the world or within the same lab. Documents are put on servers that are always available (or as close as possible) and contain links to other documents allowing a densely linked network of citations and references to form. This has been formalized as the idea (significantly articulated and given a pithy name by Roy Fielding in his dissertation) of Representational State Transfer (REST). This views the Web as a collection of objects which should primarily be manipulated using four key verbs: GET, POST, PUT and DELETE. All manipulation is done via combinations these verbs, which should treat the documents as immutable. In this sense, the Web could be seen as the world's most used functional program. This view of the Web, as a collection of discrete and immutable documents is what we term the Document Web.

The logical next step is to dissolve the boundaries between documents and provide meaning to the structure of the contents. This is the key idea behind the Giant Global Graph. This has been the focus of two mildly competing efforts. One, XML, tries to make document representation mechanical and so maximize the reuse of tools or documents. The other, RDF, tries to provide a formal data-model for the Web. We say mildly competing because there is a difference in direction between the two projects, and some overlap in goals. XML starts with the premise that the 'Document is King' and represents all data in the form of a tree of structured parts. RDF, starts from the idea of the Web, that 'Connection is

King' and works backwards to specific item representations from there. Looking at the Semantic Web project, we can see a large portion of the idea is to liberate data from stifling documents only manipulated as a whole and allow the Web idea to operate right down to the level of a binary assertion. This densely inter-linked network of assertions is what we term the Data Web and represents a true second phase for the Web project.

The purpose of the Web is to allow us, its users, to share information and understanding. The Document Web allows us to post pages detailing the most tedious minutiae of a topic and link them in with other pages so that people can find them and share in our ennui. The Data Web allows us to tear down the artificial silos that divide our knowledge and benefit from emergence; the whole is more than the simple sum of its parts. It seems debatable whether it is possible to go further than that, however it is worth noting that the Data Web is very well-suited to simple aggregation, but questionable in its approach to human-oriented knowledge.

Both the Document Web and the Data Web rely upon a very simple idea; that of providing names to things. In the Document Web, these names are locators; they tell you where to find documents so that you can download them. Appropriately, these names are called Uniform Resource Locators (URLs). The Data Web goes a step further and provides the ability to name things you can't download such as the book you just read, the idea you just had and the action you're about to do. These names are called Uniform Resource Identifiers (URIs) and subsume URLs. What could be simpler? You label the world from top to bottom with these atoms[1] to allow you to make unambiguous statements that may be merged with the statements of others. As a concept, URIs are perfectly named. They provide a uniform, basic model of identity to the Web.

However, very little information of interest to people is completely uncontroversial. Even that statement is controversial with some who take an extreme position of objectivism so allow us to elaborate. We're used to balancing evidences in our everyday lives, and the brain seems to have this ability built in. Our natural mode of expression in language is replete with modifiers, expressions of uncertainty and partial information. There is plenty of evidence to suggest that we think in uncertain terms too [7], not only having degrees of belief about every proposition but also being uncertain about our uncertainty depending on how much evidence has been accounted for [11].

If it proves necessary, we can usually make a binary decision and lump things into the boxes 'yes' and 'no', but to do so loses a lot of the information which was inherent in the statement. Someone receiving a Boolean valuation of a statement has no means to ascertain how sure the producer was of the statement or how much evidence went into it. Worse, the receiver is limited in their choices as to what to do with this statement if they too are acting in a Boolean fashion. If they agree with it, or rather have no reason to disagree, then they can just accept it or reject it based on the source. However, should the information conflict with the receiver's own beliefs and they trust the source then what should they do? It seems prudent therefore to cast off the shackles of Boolean truth-values and consider something more flexible.

Unfortunately, it isn't so easy as to just assign one or more floating-point numbers to each statement made. A re-evaluation of the assumptions of our semantics is required. Let us turn back to the idea of universal names. With an assumption of no-disagreement, this is perfectly fine as we can treat everyone's knowledge as part of the same shared understanding. What happens if we allow disagreement, our motivation for introducing truth-values that are more complex? Well it should be clear that, even if just from a pragmatic perspective, the denotation of a concept should be based upon what we know about it. However, this means that if we disagree about the veracity of various statements made about the concept, we also must disagree about the denotation. The import of this is that Uniform Resource Identifiers cannot *identify* resources universally; different people are able to have different understandings of the same resource[2]. Resources can only be *named* and then understood relative to a given agent's background context [3]. Therefore, in contrast to typical assumptions that people bring to the Semantic Web through the Web Architecture, a URI is insufficient to stand for a denotation on its own.

Seen from a cognitive perspective, this is hardly surprising. However, it is only with the development of the Web that logical modeling has collided with real, uncontrolled information representation on a global scale. There is no opportunity to hide behind toy examples now; we must rise to the challenge and produce a semantic framework that is capable of dealing with real human knowledge representation. In short, we must look carefully at what we know about how humans think.

3. FROM SEMANICS TO PRAGMATICS

The current models of the Web are very passive and static things. By contrast, humans are active and dynamic. All the information that is on the Web is a product of human action in some form. Whether it be written by hand, or the result of a human-conceived piece of software that automates a task, there is nothing available in the medium which hasn't been touched human thought in some fashion. Now, human conception isn't a static thing; we aren't born with everything we are ever going to know. We learn, we adapt, we make mistakes in our beliefs that we then correct. All this happens instinctively, and without effort.

It should be clear from this that the assimilation of information, its understanding and subsequent dissemination, can be seen as a form of process. We learn by acting. We communicate by acting. Our use of the Web is just a particular form of action, allowing us to find parceled snippets of another's thoughts. As the disagreement problem shows, there is no inherent semantics to the information on the Web, just the meanings we acquire through our readings.

The study of meaning situated within behavior is known as pragmatics, which in its strongest form may be summed up by the phrase 'meaning is doing'. Recent commentary has suggested that the Semantic Web should become, in a sense, the Pragmatic Web focusing research on pragmatic issues such as communication layers rather than knowledge representation languages [10][2]. A pragmatic foundation for the Web would entail viewing all information in the context in which it is communicated. We must seek understanding of the process of understanding.

This then is what we propose as the next phase of the Web; accepting the fundamental role that process, particularly communication, plays in defining meaning. The Web exists to enable communication. This entails a number of important changes in perspective. Rather than treating information on the Web as having meaning in and of itself, it only gains its meaning through the

[1] We should make clear that the idea of URIs as atomic is not without controversy.

[2] It is worth noting that this is different to the so-called 'Superman Paradox', which is based upon multiple names. This issue arises with a single name having different denotations to different agents.

interactions of its users, be they machine or human, and its embodiment in the world that the users occupy. Equally, the desire for openness and decentralization entails that we accept disagreement and provide mechanisms to deal with it that do not require us to discard information unnecessarily. Information is active in that it affects the understanding of other information.

Rather than a fundamental distinction between producers (the servers) and consumers (the clients), we should treat production as merely as one outcome of the consumption process. We propose a Web of *communicating* agents, with no *a priori* distinctions between them, all (at least notionally) communicating on the same footing using symmetric protocols. This proposal, of a Web of active and communicated information produced and consumed by active agents acting within the dynamic real world, we term the Active Web.

Naturally, such a significant change of presumptions leads to a mismatch between architectures. As mentioned before, the Web-as-is is based on the principles of REST. The restriction of the protocol to just a few verbs serves very well. Ideally, we'd like to save the spirit of REST, essentially captured by a minimal communications protocol, whilst adapting it to our current purposes.

The Active Web is based around the idea of loose coalitions of agents communicating in order to further their own understanding and goals. As such, the ideal protocol will be symmetric and asynchronous; that is, every command will come in a pair and won't require an immediate response. The asynchronous nature will benefit a true loose community where given agents may need time to consider their response, or may be only available sporadically.

The communication-focused approach of this proposal shifts the primary research question of the Web from developing a sufficient, singular knowledge representation format to the interactions between agents and the social life of information. Our aim then becomes to support and understand the dynamics of these loose, ad hoc communities of information holding agents. We must ask, how should these communities organize so as to make access to appropriate, relevant information as straightforward as possible for us humans.

A human society is connected by its culture, the shared knowledge which contextualizes all discourse and interaction. A culture contains not only the social mores of a community (ranging from whether or not we burp after a meal to the laws of the land) but also many assumed beliefs, questions and goals. Members of the society may use these knowledge items to guide inference, shape decision-making and form the basis of successful communication. If a community shares a belief, that belief need not be communicated every time. Such knowledge is often called 'common-sense' [9].

We choose to term the shared knowledge of a loose ad-hoc community, the community's *cloud*. The visceral image is that of gossamer, strung between the agents in the community. It has a gradual boundary, is dynamic and constantly shifting and may mingle with other clouds as they come near. A cloud is the fine glue which connects a collection of agents, the matters being communicated and the agents themselves (as each agent has knowledge regarding the other agents connected by the cloud.) We see the Active Web as being glued together by these clouds of knowledge (Figure 1.)

Each user can be considered to have a personal cloud which contains knowledge relevant to him or her, and which is organized in a way that is relevant to them. This knowledge should be virtually encapsulated to ensure appropriate privacy, containing the personal agents that make up the *society of mind* [8] of the user's existence on the Web. All interactions with other clouds are mediated by this personal cloud so that it is said to *surround* them. The cloud may be thought of as behaving like an emergent 'super-agent' so that clouds can form themselves into higher-level clouds surrounding emergent social groups such as teams, societies, etc. through the communication of items of knowledge. The clouds may also contain information derived from the existing Web.

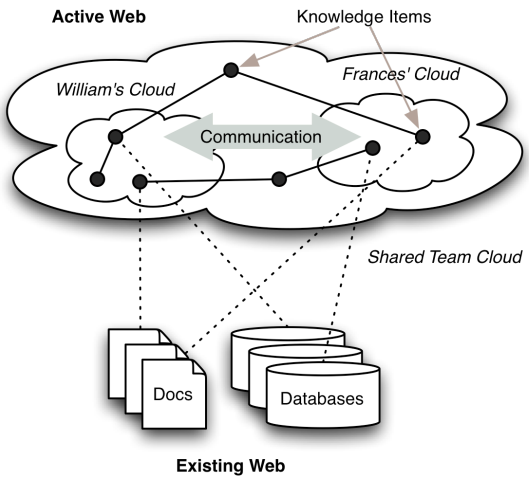

Figure 1. The structure of clouds in the Active Web.

The main issues of the Active Web are then; the dynamics and interactions of the clouds of knowledge from which it emerges, the appropriate representation of the knowledge itself (including a means for incrementally growing a cloud for a given community through the communications between the community's members) and the interaction between these clouds and the outside, non-Web world (including the humans it is intended to support.) It is vital for the 'boot-strapping' of the Active Web idea that, at the very least, one can re-use the information from the Document and Data Webs.

4. USAGE SCENARIOS

In order to provide an idea of what we see the Active Web evolving into, we provide three usage scenarios. These scenarios each address a particular level of interaction; from the personal up to the inter-organizational. They are presented in the form of user stories, describing how each of three characters uses the Active Web's concepts to further their aims. It should be noted that these levels are merely here for explanatory purposes; the vision of the Active Web is that the clouds are emergent and compositional allowing arbitrary and dynamic scales of interaction.

4.1 Personal Use

William is a project manager involved in a successful trans-European telecommunications company. The project teams he supervises frequently require guidance whilst he is away from the office.

Whilst he is away, his personal information cloud still surrounds him, supplying answers to the questions he asks. His cloud also *interacts* with those of his employees and supervisors in order to give him a holistic picture of the work that his teams are doing.

Frequently the information William gets must be actively *updated* and *consolidated* to give a coherent perspective. The Active Web allows these processes to be automated, and conflicting information to be integrated.

If one of his teams' members needs to flag something for his attention, their cloud can actively communicate with his cloud.

Rather than waiting for him to see the message, and then ask for the relevant information he requires to form a response, the Active Web can anticipate his need by *adjusting* and *adapting* the information and processing of the cloud to have as much of the necessary information ready for his input. If the request is urgent, his cloud can talk to pervasive UI services available in his locality and immediately call his attention to the matter wherever he may be.

4.2 Intra-Organizational Use

Frances is a requirements engineer in one of William's teams. She has to regularly communicate with her colleagues in order to ensure her requirements documents are consistent and correct with respect to the conceptual models of the project stakeholders.

With the Active Web, her cloud is constantly *seeking* and consolidating information. It is also actively *growing* new knowledge through *conversations* with the clouds of her colleagues, constructing a coherent view of the domain. These discourses may happen on their own, as a continuous knowledge seeking and inquiry process, or 'piggy-back' on the inter-personal communications (via IM, email, etc.) she has with the other members of the team. Her cloud may then extract discourse-context relevant information from the conversation, which is used to guide the inferences it makes as well as the meaning of the inter-machine communications being made.

4.3 Inter-Organizational Use

Hector works in a university involved in a joint project with the company at which William and Frances work. The two groups, have a meeting once a week to discuss the progress of the project. The university's campus is fully integrated with the Active Web, providing a *pervasive* semantic environment for interaction with each user's personal information cloud and *discourse-context awareness*.

At the review meetings, each participant receives constant feedback as to relevant information for their position from their personal cloud, through the environment (via screens, mobile phones, auditory feedback, etc.) Each person's cloud understands its user's circumstances best and so can provide the most appropriate information and feedback to support their participation. The clouds interact with each other as a matter of course and so provide an emergent community cloud for the project as a whole that supports the aims of each level of interaction.

5. CONCLUSIONS

The Web has grown to become one of the most important inventions of the 20th century. This has been in part because of fate, but also because its architecture facilitated many shared aims in human nature. The subsequent developments of the Web that break down the barriers between documents (the Semantic Web) are allowing even greater, automated leverage of the vast quantities of information now available through the World Wide Web.

We have argued that the next stage for the Web must be the earnest analysis of the fundamental nature of human interaction; both action and communication. Projects like FOAF, and social networking in general, have succeeded because they allow humans to express the most relevant information of all: social information. All information may be seen as being social in nature, being developed and given meaning by the interactions within societies of humans, and between humans and their actions.

We believe that the Web is fundamentally active; i.e. that its meaning comes from its behavior. As such, we further believe that the most significant questions for research must be those to do with this active nature.

If we take on board the arguments presented here, then the main questions become: how should information (and behavior) be organized, how should we handle disagreement and how should the Web be embodied in the worlds we humans occupy. Some starting points to answer these questions have been proposed in this paper, but it falls to posterity to determine the ultimate solutions.

6. ACKNOWLEDGMENTS

We would like to thank Sean Palmer, Danny Ayers, Kevin Reid, Robin Berjon and John Cowan for insightful discussions about the topics in this paper.

7. REFERENCES

[1] Berners-Lee, T., Hendler, J. and Lasilia, O. The Semantic Web. Scientific American, 284 (May 2001), 28–37.

[2] Bussler, C. Is Semantic Web Technology Taking the Wrong Turn? IEEE Internet Computing, 12 (January/February 2008), 75–79.

[3] Connolly, D. A Pragmatic Theory of Reference for the Web in Proceedings of the 15th International World Wide Web Conference (Edinburgh, UK, 23rd – 26th May 2006).

[4] Ding, L., Zhou, L., Finin, T. and Joshi, A. How the Semantic Web is Being Used: An Analysis of FOAF Documents in Proceedings of the 38th Annual Hawaii International Conference on System Sciences (Hawaii, 3rd–6th January 2005).

[5] eBiquity Group, Top RDF Namespaces. http://ebiquity.umbc.edu/blogger/2007/09/23/top-rdf-namespaces/ (Accessed on 14th February, 2008.)

[6] The Friend of a Friend (FOAF) Project. http://www.foaf-project.org/ (Accessed on 13th February, 2008.)

[7] Lakoff, G. Women, Fire and Dangerous Things: What Categories Reveal About the Human Mind. Chicago, 1987.

[8] Minsky, M. The Society of Mind. Simon and Schuster, New York, 1986.

[9] Mueller, E.T. A Calendar With Common Sense in Proceedings of the 5th International Conference on Intelligent User Interfaces (New Orleans, Louisiana, 2000), 198–201.

[10] Schoop, M., de Moor, A. and Dietz, J. The Pragmatic Web: A Manifesto. Communications of the ACM, 49:5(May 2006), 75–76.

[11] Wang, P. Confidence as Higher-Order Uncertainty in The Second International Symposium on Imprecise Probabilities and Their Applications (Ithaca, New York, June 2001), 352–361.