

Preparatory Observations Ubiquitous Knowledge Environments: The Cyberinfrastructure Information Ether

By

Jeffrey Gennari, Maria Harrington, Stephen Hughes, Michael Manojlovich, and Michael Spring
Department of Information Science and Telecommunications
School of Information Sciences
University of Pittsburgh

Contents

1: Introduction.....	2
1.1: Motivation.....	2
1.2: Goals (Intents).....	3
2: Knowledge Environments: A working perspective	4
2.1: Information Stores	4
2.1.1: Granularity, Descriptiveness, and Manipulability.....	5
2.1.2: Audience Scope	6
2.2: Interactive Systems for Communication and Collaboration.....	6
2.2.1: Paradigms.....	6
2.2.2: Functions.....	9
2.2.3: Patterns.....	10
2.3: Infrastructure.....	12
2.4: Toward Complication of the Task: Two Possible Developments.....	13
2.4.1: Surrogates: More than Agents	13
2.4.2: Dynamic Documents: More than Digital	14
3: Conclusions	15
4: Appendices.....	16
4.1: Appendix A: An Information Scientist’s Lament.....	16
4.2: Appendix B: Functions supporting Collaboration.....	19
4.3: Appendix C: Web Surrogates	22
4.4: Appendix D: A Paper on Infrastructure Standards	26

1: Introduction

In the Spring term, 2003, a doctoral seminar¹ was offered in the Department of Information Science and Telecommunications at the University of Pittsburgh on “The Semantic Web: Architectural Patterns for Evolution”. This seminar provided a forum for examination of the concepts and technologies needed to bring about the next-generation vision of Tim Berners-Lee. Starting with the semantic web architecture, the seminar expanded its focus to the more general notion of information marketplaces offered by Dertouzos in his book “What Will Be”². The participants of the seminar explored definitional matters, infrastructure requirements, and technology and standardization needs. In concluding the seminar, the proposal for the NSF workshop on “Ubiquitous Knowledge Environments” to be held in Chatham Massachusetts on June 15-17, 2003, was shared with the members of seminar. What follows are some of the reactions and reflections of the seminar participants related to the proposal. It is our hope that our reflections on the task set before the NSF panel and some of the results of our own work on information marketplaces will be of use to the panel members in their deliberations.

1.1: Motivation

It is our belief that the multitude of efforts and opinions about “Ubiquitous Knowledge Environments” may be aided, like philosophical discussions, by paying attention to what we really mean:

*The results of philosophy are the uncovering of one or another piece of plain nonsense and of bumps that the understanding has got by running its head up against the limits of language. These bumps make us see the value of the discovery.*³

We also hold close Simon’s discussion of the “**Sciences of the Artificial**”⁴ – in particular, his view that artificial phenomena are contingent on the design goals or purposes – in contrast to the sciences of the natural, which seek to uncover “laws” that describe natural phenomena. The sciences of the artificial are based on artifacts which Simon describes as synthesized entities which may imitate natural things and most importantly, as things that may be:

- characterized in terms of functions, goals and adaptation
- discussed in terms of both imperatives and descriptives.

In the context of the current workshop then, we return to Wittgenstein who suggests

*A main source of our failure to understand is that we do not **command a clear** view of the use of our words. -- Our grammar is lacking in just this sort of perspicuity. A perspicuous representation produces just that understanding which consists in 'seeing connections'. Hence the importance of finding and inventing **intermediate cases**.*⁵

Thus, following both Wittgenstein and Simon, it is important for the panel to clarify the grand vision that is to be fleshed out through discussion. “*Ubiquitous Knowledge Environments*” and “*Cyberinfrastructure Information Ether*” are seductive phrases. Our own explorations these past few months convince us that terms such as “knowledge environments”, “cyberinfrastructure”, “information ether”, “semantic web”, “information marketplaces”, or “web services”, are fertile ground for bumps in the head from running up against the limits of language.

¹ The syllabus for the seminar is at: http://www2.sis.pitt.edu/~spring/courses/ds_semantic_web.html

² Dertouzos, Michael, **What Will Be**. Harper, 1997.

³ Wittgenstein, Ludwig, **Philosophical Investigations**, Blackwall, (1998), section 119.

⁴ Simon, Herbert, **The Sciences of the Artificial** (Second Edition). MIT Press, Cambridge, Mass, 1981.

⁵ Wittgenstein, Ludwig, **Philosophical Investigations**, Blackwall, (1998), section 122.

The task set before the panel participants is to develop a proposal for NSF funding of research that will promote the development of Knowledge Environments based on a “cyberinfrastructure.” This is closely tied to the development of a science of “information and knowledge management” along with the tools to make use of the resulting “information ether.” Progress on defining these goals is in part dependent upon solid, clear, and consistent definitions of among other things: work, information, communication, documents, knowledge and collaboration. While each of us “knows” what these terms mean, drilling down leads to a lack of consistent operational definitions. As a simple example, we include a brief digression on the definition of information and knowledge in appendix A. This paper endeavors to enhance the panel discussion by encouraging definitional consistency and clarity. The intent of the effort is to classify and organize some of the possible meanings. We attempt to provide a working perspective on some of the many terms and ideas that serves as a basis for more qualified and clear discussion.

1.2: Goals (Intentions)

This section begins with an effort to define the task of the workshop. We do this by reflecting on, the Atkins Report on Cyberinfrastructure, the RIACS report on Technology Requirements for Information Management, the Ubiquitous Knowledge Environments: Cyberinfrastructure Information Ether proposal submitted to NSF, and our own efforts.

“Revolutionizing Science and Engineering Through Cyberinfrastructure”, or the Atkins Report⁶, was a broad-ranging report that had its roots in the Partnerships for Advanced Computation Infrastructure (PACI) program funded by NSF and related to the various supercomputer centers. Building and extending that initiative – vastly expanding it, the report proposes a new initiative:

*The most fundamental goal is to empower radical new ways of conducting science and engineering through the applications of information technology.*⁷

The specific vision of cyberinfrastructure is captured in the following:

*Applications are enabled and supported by the **cyberinfrastructure**, which incorporates a set of equipment, facilities, tools, software, and services that support a range of applications. Cyberinfrastructure makes applications dramatically easier to develop and deploy.... Cyberinfrastructure also increases efficiency and quality and reliability by capturing commonalities among application needs, and facilitates the efficient sharing of equipment and services.*⁸

The RIACS report looks at technology requirements for information management in support of a selected set of application domains. This report is focused on how the vast stores of information needed in various application domains might be managed. The major result of the study was the recognition of a common thread:

*A common thread is the need to interoperate with many diverse information resources, and hence to assure that future systems will be interoperable not only with past and current standards, but adaptable to resources that are not yet recognized.*⁹

The high level recommendation of the study is a research program that “pursues a science of information management” where that science is defined as follows:

*A science of information management would deal with the underlying principles of information management and how humans deal with it. By contrast, information technology provides the tools and systems to achieve the desired functions and goals.*¹⁰

⁶ Atkins, Daniel E., et. al., “Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue Ribbon Advisory Panel on Cyberinfrastructure,” January 2003, available online at http://www.communitytechnology.org/nsf_ci_report/.

⁷ Ibid. page 35

⁸ Ibid. page 35

⁹ Graves, Sara, Craig A. Knoblock, & Larry Lannom, “Technology Requirements for Information Management,” RIACS Technical Report 2.07, November 2002, available online at <http://www.riacs.edu/trs/>, page 3.

As we see it, this report establishes two goals. The first is to understand the principles by which information is managed and the second is to define how humans use information.

In the proposal for the NSF workshop the focus is most succinctly defined by the following:

The successful realization of [a ubiquitous information infrastructure] will be supported by a science of information management that will yield new generations of knowledge environments?

This suggests three elements: an infrastructure (for communication), a science of knowledge management, and knowledge environments.

The seminar at the University of Pittsburgh on the next generation web concluded that the design goal might be stated as follows:

Systems that support direct and indirect human interaction and collaboration require rich information stores based on rules governing how information should be aggregated, stored and transferred and a supportive and generalized infrastructure.

Synthesizing the various findings, recommendations and objectives suggests three things that need to be addressed:

- development of knowledge environments that support communication and collaboration enabled by
- the construction of an infrastructure that is based on
- a principled set of rules governing how information should be aggregated, transferred, and used.

Accepting that work will need to proceed with less than perfect definitions, this paper looks to find ways of thinking about the problem that might help to focus the discussion. The ideas presented below constitute straw solutions that participants may well reject in favor of more specific or accurate ideas. It is our hope that the discussion based on the specifics defined herein will cause unarticulated assumptions to be minimized.

2: Knowledge Environments: A working perspective

Knowledge environments are the result of distilling knowledge, objectifying it, and then embedding it in a collection of shared information artifacts. The goals that we are seeking to achieve are implied in the names given to the spaces that have begun to emerge. They are as concrete as “Digital Libraries” and as abstract as “Information Ether”. Clearly, digital libraries imply goals having to do with organization and access to information and knowledge stores. The semantic web, as put forward by Berners-Lee, is an effort to make the resources of the web more accessible for artificial understanding, shielding humans from needing to surf a tidal wave. Dertouzos’ information marketplaces, on the other hand, imply new kinds of services that are available via the infrastructure that is emerging. These are but a few examples of “knowledge environments” that will likely to materialize in the future. The purpose of this discussion is less focused on the particulars of “What will emerge”, and more focused on “What is needed to support these kinds of things that should emerge”.

Communication and Collaboration can take many forms. Working from the premise that knowledge has been transferred from the human to the environment, we are focused on communication through artifacts. Further, we have great concern that the design of the artifacts might allow a machine to serve as an intermediary to help with semantic level processing. Ultimately, what is being discussed is an infrastructure to support computer-mediated communication and collaboration using intelligent artifacts. In the following sections, knowledge environments are described in terms of information stores, interactive systems, and infrastructure.

2.1: Information Stores

We choose to focus our attention first on managing the information objects through “information stores” – systems that define rules for applying various operations on the artifacts (creation, storage, dissemination, use, etc).

¹⁰ Ibid: page 4.

Frequently, definition of operations on the artifacts are dictated by the representation. There is a growing body of research that supports various kinds of information stores – text documents, images, audio, video, etc. Choosing the representation may be influenced by a number of factors:

- Cognitive/Perceptual: Is there a format that optimizes the transfer of the information?
- Preference: Can the presentation reflect the learning styles of the audience?
- Aesthetics: Is the representation consistent with existing information in the store?
- Cost: How expensive is this representation in terms of production, storage, transmission, etc?
- Authorship: Is the author better equipped to create a certain type of document?

This approach offers a valuable perspective on how the consumer might ultimately receive the information and certainly must be taken into consideration. However, it is not conducive to a discussion of overall operational management; the various representations quickly splinter this discussion into disparate technologies such as latent semantic indexing, edge detection, phoneme analysis, or key-frame indices. While representation is an important characteristic and well understood, there are several other characteristics that need to be considered as well.

2.1.1: Granularity, Descriptiveness, and Manipulability

Decisions made based on representation consider how information is encoded in the document. We are interested in how to decode the document – especially if its meaning is to be interpreted by the machine. In general, it is easy for a person to extract relevant knowledge from a document. Given the current state of artificial intelligence, it is less easy to programmatically extract information from a document. In the web community this is understood by the adage, “the web is machine-readable, but not machine understandable”. Since it is not practical to have the consumers repeatedly inferring knowledge structures from the documents, designers are encouraged to impose a knowledge structure on the document. This can be done by first level structures – the schema for a document or by second level data -- *meta-data*. Almost since the start of the web, some have advocated that pages be marked with “meta” tags. From TEI to PICS to RDF, researchers have been exploring methods for embedding metadata in the artifact, providing commentary that makes it easier to locate and extract knowledge. Characterizing a document’s meta-data should not be a simple binary measure – present or absent. Rather, the degree to which knowledge about a document is made accessible can be expressed as a continuum along three dimensions: granularity, descriptiveness and manipulability.

Meta-data can describe a large aggregate or an atomic element. The resolution at which meta-data can be applied is limited by an object’s *granularity*. A coarse document might have metadata that explains that it is a doctoral thesis about information visualization. Alternatively, a fine-grained approach to the same document might provide meta-data information about various sub-sections, e.g. “this section describes the use of preattentive cues to isolate critical values.” It is important to observe that while granularity is bounded by syntax, it is the semantics associated with the level of granularity that is of import. We are not necessarily interested in decomposing objects into words, pixels or frames, unless the selection of those particular elements offers additional insight. Thus, highly granular objects can be divided into meaningful segments that provide the ability to discriminate knowledge within a portion of the object. This brings us to the second dimension of interest. *Descriptiveness* captures the amount of detail that is provided by the meta-data for each granule. Meta-data may be broad, providing only a few keywords to summarize a lengthy concept, or it may adhere to a robust ontology that provides minute detail on several attributes of the object.

Highly collaborative environments require a wide range of operations on the artifact (revise, annotate, supplement, etc), while other environments, e.g. simple communication, may require a more restricted set of operations, e.g. reading only. *Manipulability* refers to the degree to which the artifact is prepared to be operated upon. Like descriptiveness, manipulability is bounded by the granularity of the object. It is the granularity of a database that allows for DBMS operations. Similarly, granular documents or images might be joined, extracted, summarized, described, etc. In general, analog objects are continuous in nature and resist decomposition make them more difficult to describe or manipulate. Digitized objects, voice or images, are granular, but the granules may not be easily mapped to higher level components for manipulation or description. Simply put, as any programmer knows, inferring that a set of pixels within a given color range constitute a straight line of a given color and length is

difficult. That same line, as a component of a CAD-CAM drawing already has all of the meta-data attached and the routines to manipulate the underlying structure are relatively easy to write.

2.1.2: Audience Scope

Audience scope is another important factor that needs to be considered in the design of an information store. We observe that there are several distinct layers of audience scope that have a profound effect on the management of the document. Specifically, one can identify five levels of audience scope:

- Personal
- Group
- Organizations
- Enterprise
- Archival

This classification suggests an increasing diversity of consumers. A personal memo or “To-Do” list is created specifically by the author, for the author; most of the knowledge remains with the author and the document serves simply as an external cue. Group documents capture the knowledge of a small band of collaborators. Contributors have a collective mindset and shared vocabulary with respect to the focus of the document. Documents prepared at the organizational level can tap into the culture of the institution to facilitate comprehension. Enterprise documents extend to populations outside an organized association and must reach consumers at a societal level. Such documents may rely on ethnic and social literacy to convey meaning. Ultimately we use the term Archival to refer to documents where the author cannot make any assumptions about the consumer (who may be several generations or civilizations removed) and must embed all aspects of the knowledge in order to successfully communicate.

We believe that there is a positive correlation between granularity, descriptiveness and audience scope. The broader the audience, the more explicitly knowledge must to be embedded into the information store. Highly descriptive and granular documents might be needed to characterize a denser information store that comes with enterprise and archival level documents. Conversely, personal documents don’t need to be described because the limited audience is already intimately familiar with the knowledge.

Interoperability between information stores will also play a large role in the level of discourse that is permitted. This issue is often addressed with a wave of the hand and uttering the incantation: “Ontology!” While we believe that ontology processing will likely play a large role in the infrastructure of future knowledge environments, it is too often taken as a solution that already viable, rather than a placeholder for a future technology. Normalization of the granularity, descriptiveness, manipulability and audience scope across several information stores may be an important first step in understanding whether the pursuit of a common ontology will be tenable.

2.2: Interactive Systems for Communication and Collaboration

Systems to support communication and collaboration, in the context of “knowledge environments” and “information ether”, will be more functional, integrated and seamless than the environments we know today. This section looks at a few of the overarching paradigms suggested, the functionality required, and the key design patterns for these new environments. These, collectively, describe the boundaries of the kinds of systems to be built, and infer some of the infrastructure requirements. The discussion is limited to communication and collaboration systems.

2.2.1: Paradigms

Smart rooms and virtual reality have been suggested by some as the future of computing. How they accomplish this goal is less clear. While agents, visualization, augmentation, and other techniques will play a role in the future of computing, just what overall motif or leitmotif will make sense is less clear. We briefly outline several paradigms for interaction design discussed by Preece et al.¹¹:

¹¹ Preece, J., Rogers, Y., and Sharp, H. (2002) *Interaction Design: Beyond human-computer interaction*. Hoboken, NJ: John Wiley & Sons.

- Ubiquitous computing
- Pervasive computing
- Wearable computing
- Augmented Reality and Physical/Virtual Integration
- Attentive Environments and Transparent Computing

Ubiquitous computing describes an environment where computers disappear onto the environment. Mark Weiser of Xerox PARC¹² suggested that the most effective computer interface is one that is invisible. It would enhance the world that already exists—in contrast to multimedia representations on screen operate require the user to concentrate on the virtual objects as they move about the screen. Weiser’s goal was to create technology that would meld into the physical world and extend human capabilities. His original design comprised the use of computer “tabs”, “pads”, and “boards” which would be easy to use, as they would be seen as a metaphor of a post-it note, sheet of paper, and blackboard. These devices, proposed several years ago, function similarly to our present personal digital assistants, tablet computers, and large screen terminals. The envisioned, and prototyped, devices were more connected than what we use today and automated the transfer of information permitting users to focus on the task allowing the system to manage data gathering, transfer, and display tasks.

Pervasive computing is in some ways simply another name for ubiquitous computing in that technology is employed to create a seamless integration of information gathering and presenting devices. Unlike the ubiquitous computing environment, pervasive computing involves a wide variety of devices from telephones to automobiles to appliances. This explosion in the number of intelligent devices has resulted in a changing view of just what kind of device is needed to access distributed information and in what visualization. A person moving around in the world could demand far more of their wireless and mobile system to the point of providing an augmented experience for the user that would add information. Hoffnagle¹³ gives an example of a user and a portable device that would not only allow a person to watch the sky; it would also overlay the real view with a projection of the constellations synchronized with the location of the observer and the time of year. A next generation of interactive systems will be required to obtain such information without regard to location or display device.

Both ubiquitous and pervasive paradigms lead to smart rooms or smart spaces. These are work environments that use embedded computers, information appliances and multi-modal sensors to allow the user to interact with computer systems in a far for efficient manner. The National Institute of Standards and Technology in the Smart Space Laboratory has been conducting research in this area and how to improve the work environment by the use of embedded computers, information appliances, and multi-modal sensors allow task to be performed far more efficiently than having to search for information via traditional computer interfaces¹⁴. Besides the traditional distributed multimedia distributed databases, the multi-modal character of smart rooms will require spoken as well as visual document retrieval and indexing.

Wearable computing is focused on those components of the ubiquitous and pervasive computing environments that are most intimately associated with humans. Ark and Selker¹⁵ suggest that embedding computers appeals to the general population for four reasons: (1) computing is spread throughout the environment; (2) users are allowed to be mobile; (3) information appliances are readily available and; (4) communication is easier between machine and human. When these requirements are met, the user begins to see information retrieval as less a mechanical operation that requires unique interface understanding and more of a natural flow of information akin to that of human conversation or observation of their surroundings. The creation of widespread wireless communication networks is becoming a reality allowing access to remote information in databases and other information

¹² Weiser, M. (1991) The computer for the 21st Century. *Scientific American*, 265 (3), 94-104

¹³ Hoffnagle, G. (1999) Preface to special pervasive computing issue. *IBM Systems Journal*, 38 (4), 502-503.

¹⁴ Rosenthal, L. and Stanford, V. (2000) NIST Information Technology Laboratory Pervasive Computing Initiative. *IEEE Ninth International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, June 14-16, 2000, Gaithersburg, Maryland.

¹⁵ Ark, W.S. and Selker, T. (1999) A look at human interaction with pervasive computers. *IBM Systems Journal*, 38 (4), 504-507.

repositories while remaining mobile even to the point of incorporating computing services into wearable computers such as jackets, glasses, etc. Rhodes et al. describe one potential use for this form of interactive system as a tour guide that would provide relevant information to users as they randomly walk about an exhibition and as they move to various attractions throughout a city.¹⁶

Augmented reality is possible with intimately associated computing devices. Ishii and Ulmer¹⁷ have proposed that ubiquitous computing will extend into what they describe as tangible user interfaces. The focus is the actual integration of computation augmentations into the physical environment. There will be a growing desire to incorporate digital information with physical objects and surfaces such as buildings. This would allow people to carry out everyday activities without any specific attention to a computer interface. For our current discussion, the most obvious example would be the development of dynamic books, physical books that are embedded with digital information. The first of these are already available in the form of customizable greeting cards that display an animation or even user provided photos to the recipient. This static application will most likely evolve into books having a dynamic digital framework permitting both user customization as well as display interfaces that would accept content updates from central repositories or publishers.

Closely allied with tangible user interfaces is augmented reality where a virtual representation is superimposed upon physical objects either by the use of wearable computers or in smart rooms. By blending the real and the virtual, these interface allow users to see each other as well as the virtual objects. Such visualization will change communication behaviors toward that of face-to-face rather than screen collaboration. In such an interactive system, the physical objects and their interactions become as important as the virtual images. For collaborative augmented reality environments, the need for cooperation, independence of control per user and individuality in the displayed data will be required to create an effective augmented reality experience^{18,19}.

Transparent computing and attentive environments, as the name suggests, involve the computer attending to the user's needs and anticipating what the user wants to do. This implies a major shifting of the burden of interactive functions onto the system. This makes the computer interaction implicit by responding to the user's gestures and expressions. Gaze-tracking technology will permit the computer to determine what aspect of the environment the user is wishing to activate, such as the television or computer interface. By monitoring eye movement, the system will be capable of directing access to the desired web page, document, etc. Some research has focused on using Grid computing powered by both sensor arrays and distributed processing to try to anticipate the needs of the researcher and create workflow models for accessing computing resources. Intelligent middleware will create new job workflow requests based upon high-level specifications of the desired results.

Trying to distill all of these efforts and to understand the underlying themes is not easy. As we see it, the future of computing is ubiquitous, aware, embedded, and distributed. While there is still some overlap in these terms, we think they capture and articulate the four essential qualities of knowledge environments. We define the terms as follows:

- **Ubiquitous** refers to the accessibility of the computing resources. They follow the user from place to place and thus the location of the user is not constrained.
- **Distributed** refers to the interaction that takes place between components to carry out some task. Thus, while ubiquity refers to the availability of access, distribution refers to the internal processing. A

¹⁶ Rhodes, B., Minar N., and Weaver, J. (1999) Wearable computing meets ubiquitous computing; reaping the best of both worlds. In *Proceedings of the Third International Symposium on Wearable Computers (ISWC '99)*, San Francisco, 141-149.

¹⁷ Ishii, H. and Ullmer, B. (1997) Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of Computer Human Interface '97*, 234-241.

¹⁸ Billingham, M. and Hirokazu, K. (2002). Collaborative augmented reality. *Communications of the ACM*, 45 (7), 64-70.

¹⁹ Schmalsteig, D., Fhurmann, A., Szalavri, Z., and Gervautz, M. (1996) Studierstube: an environment for collaboration in augmented reality. In *Proceedings of the Collaborative Virtual Environments (CVE '96) Workshop*, Nottingham, Great Britain.

distributed system carries out the functions needed by the users without regard to where the actual resources or processes take place. One might suggest that distribution is the internal view and ubiquity is the external view.

- **Aware** refers to the fact that the interaction between the human and the system is informed by context. The extent to which the awareness is "inferred" is a matter of some significant research and can be as simple as the system being informed of who the user is and gathering data for that user to things as complex as inferring the nature of the task the user is engaged in and taking action based on that.
- **Embedded** refers to the movement of computational elements into everyday appliances. Like Anti-lock Braking Systems, the future will see "smart" doors, thermostats, refrigerators, etc. Like ABS, many of these systems will be important single function systems. Others will share information with other subsystems. At what point we can take user aware displays and input devices and conclude that the aggregate of devices at a location provides for a "smart room" is not clear. This suggests that there is a close relationship between being aware and being embedded. At the same time, we see that each brings a separate research focus.

2.2.2: Functions

There are a variety of different functional components that might be used to provide support for high level collaboration and communication. Work at the University of Pittsburgh over the last several years has focused on adding additional levels of functionality to collaboration systems. The research test-bed, known as CASCADE²⁰, was developed to explore how various communications streams and information streams could augment collaboration related to document creation. Part of the CASCADE research program involved analysis of both the collaboration literature and collaboration systems.²¹ The analyses cataloged functions included in collaboration systems. The functions were related to:

- e-mail
- database management
- document management
- calendaring
- conferencing
- management information systems and decision support services
- network and administrative services.²²

Classification of functions into topical areas, led to the following organization:

- Message transaction support and management
- Activity coordination and meeting support and management
- Shared information store support and management
- Workgroup support and management

The augmentation group in the department has attempted to classify and describe this development a number of different ways. As a part of this current effort, several of our taxonomies were reexamined to see if a more predictive classification of system functionality were possible. We are satisfied that the four categories of functionality developed early in the CASCADE effort continue to hold as a valid guide for examining functionality. However, it is clear that there are various degrees of integration and sophistication of that functionality. It appears that systems tend to grow and evolve with increasing "intelligence" present in the system. It appears that this "intelligence" may be operationally defined at four levels:

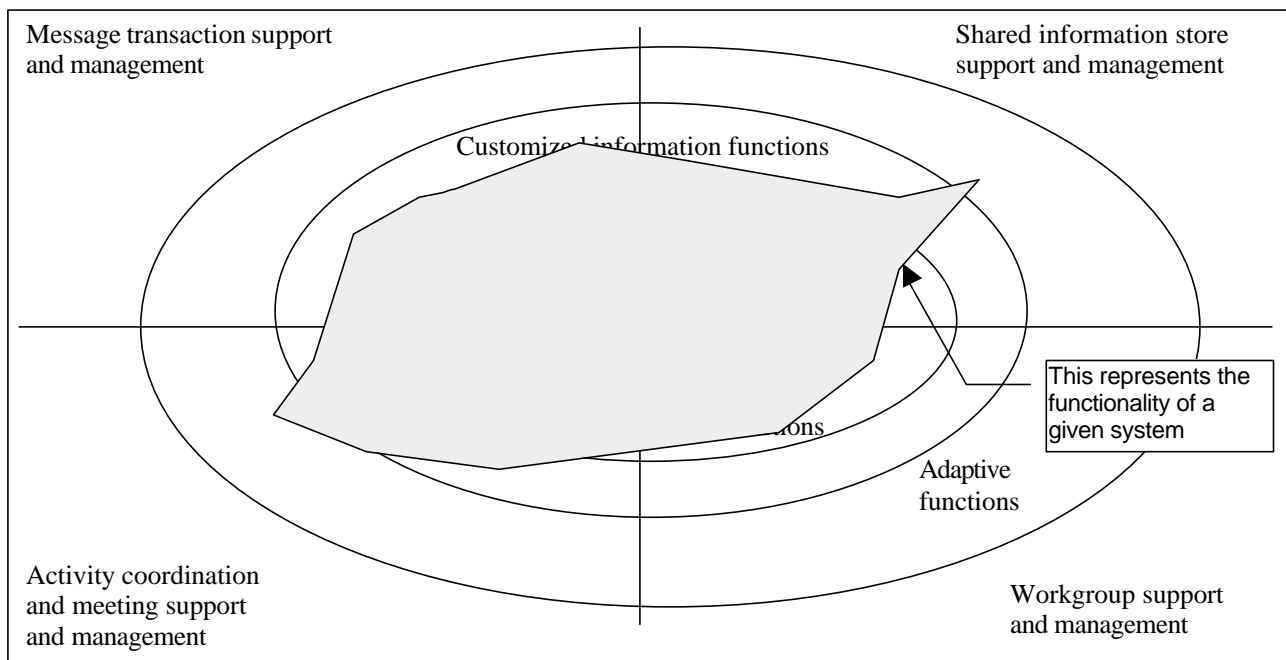
²⁰ CASCADE – Computer Augmented Support for Collaborative Authoring and Document Editing – was funded by NIST and resulted in both a prototype system for standards development and a research testbed.

²¹ Sapsomboon, B. Andriati, R. Roberts, L. and Spring, M.B. Software to Aid Collaboration: Focus on Collaborative Authoring, NIST Research Project Technical Report 1, January 30, 1997, DIST, University of Pittsburgh.

²² A more detailed list is included in appendix 3.

- **Rule based functions** constitute the lowest level of functionality. Classifying systems at this level does not imply that the systems are simple. The rule based functions can be exceptionally sophisticated. It simply implies that the focus is on the algorithms for the process.
- **Information enabled functions** have rules and information stores which can inform the rules. Although this information may be initially configured, it generally remains static once it has been initialized. These information stores can be as simple as a dictionary of words and as complex as a set of rules for fire wall construction.
- **Customized information functions** have information stores that are continually tailored to the user, either explicitly or implicitly. An example of this would be address books that grow with use of the system.
- **Adaptive functions** have all of the advantages of customized information functions with the added functionality that rules are also modified based on time and experience. An example here would be a system that employs a Bayesian or neural net in which weights and rules change with use.

Combining function areas with function sophistication yields the classification shown below. While this classification is still far from perfect it does provide a single matrix to which the various functions can be mapped and it does have some predictive power in that it postulates that function development starts at the center and moves toward the periphery. In part, this movement is controlled by the state of the art in computing. That is, the functions in the outermost rings demonstrate the highest level of adaptive intelligence. A given system might be described by outlining the sophistication and focus of its functionality as shown by the sample shading.



2.2.3: Patterns

The upper levels of function sophistication are based on the ability of the system to gather peripheral information that supports the task as well as information about the preferred methodology. Two important patterns emerge based on what is done with that information. On one hand, the system may augment the viewer's ability to make decisions by affecting the *presentation*. On the other hand, this information might be used to power automation of the task, affording *substitution* of human judgment with the system's action.

Presentation is concerned with aligning the flow of information to the viewer with the formats and arrangements that optimize understanding. There are a number of presentation operations that may be taken to influence the decisions of the viewer including:

- **Ordering** – The sequence in which a collection of information is viewed is can have a significant influence a person’s comprehension and judgment. Order may reflect importance, priority or an inherent sequence of knowledge. Moreover, the order may be restricted to a linear progression, or it could be bound by a loose hierarchy of prerequisite knowledge.
- **Filtering and Aggregating** – the deliberate omission of non-essential information can be used to focus attention to relevant aspects of the information. Likewise, insight may be better achieved if component information is collected from several sources and presented as a unit.
- **Marking** - introducing annotations to the existing display can be used to call attention to a feature (e.g. highlighter strokes), emphasize a point (e.g. underlining) or elaborate the existing information with additional text or iconography (e.g. post-it notes)
- **Providing Alternate Representations** – As mentioned earlier, the learning styles of the viewer can play a role in understanding. Generating visualizations or other representations from existing data sets could allow the system to align the presentation with the appropriate learning style.

Further discussion of some of these techniques can be found in Brusilovsky’s taxonomy for adaptive hypertext navigation²³.

While influencing the presentation is driven largely by the peripheral information, the nature of substitution is shaped more by methodological information. As such, the upper limits of substitution are bounded only by our imagination; Turing envisioned systems that could complete activities in a manner that is indistinguishable from humans. Clearly, this level of substitution has not been attained. However, simple substitution has found its way into our everyday lives. For example, word processors are able to parse our documents, infer that we accidentally typed the preposition ‘from’ when we meant to use the noun, ‘form’, and automatically make the correction.

Substitution might be further characterized under three design conditions.

- **Repetitive, Detail-oriented actions** – Conditioned responses that are unnecessarily subjected to human error are ideal candidates for substitution. Simple examples are abundant, such as completing a login protocol – storing and procuring the correct username and password. However substitution can also handle detail-oriented actions that are associated with more complex conditions and variability in judgment. For example, reacting to DBMS triggers, the system may initiate a meeting with appropriate members of the team, handling the details of coordinating schedules, reserving resources, sending reminders, etc.
- **Validation** – Authenticating that a document comes from a trusted source, or corroborating questionable information are time-consuming activities. Criterion for assessing the authority of a site can be transferred to the system
- **Time-sensitive Operations** – Systems can easily monitor the time and remind a user that action needs to be taken on a certain task. Going a step further, responsibility to complete these tasks can be transferred to the monitor, especially when there are negative consequences for tardiness. Dependence on the system would likely be commensurate with the consequences and faith in the system’s ability to complete the task without error. For example, one would probably readily accept a system that automatically renewed library books. Further, a system that automatically filed a tax return might be declined in favor of a system that automatically filed for an extension to file.

The two design patterns, presentation and substitution have been presented as two distinct options. It is probably more accurate to think of them as extremes on a continuum of augmentation. The system’s confidence in the user model and task methodology may determine whether or not it attempts to influence the presentation or substitute its

²³ Brusilovsky, P. (1997). Efficient techniques for adaptive hypermedia. Intelligent Hypertext: Advanced Techniques for the World Wide Web. C. Nicholas and J. Mayfield. Berlin, Springer-Verlag. **1326**: 12-30

own judgment, or perhaps take some intermediate action. Consider the grammar check in Word; grammatical mistakes induce an annotation in the display. Upon investigating, the system often has a single suggestion for correcting the problem, but the system is unwilling to completely impose its judgment.

2.3: Infrastructure

From our point of view, the infrastructure to support collaboration is pretty simple to design conceptually. A conceptual design of an infrastructure to support collaboration might consist of a hierarchy of services that meet the criteria of being scalable and extensible and that meet the design goals of being gracefully degradable, distributed, and decentralized. The services themselves may be organized in terms of the placement and scope of their operation. We identify a three-tier hierarchy of network, system, or application services. Network services exist as a part of the fabric of the infrastructure itself and are accessible to all connected devices. For instance, DNS is a service that is independent of a particular device but provides services to multiple processes across multiple devices. System services are those services that exist at the operating system level of every device that participates in the infrastructure and are available to all of the programs or processes on a given device. A well-known example of this type of service is the TCP/IP protocol stack. Finally, Application services provide functions specific to a single program or process through such mechanisms as frameworks and APIs.

What core services are needed for the system as a whole to work? Using the classification scheme above as an analytical tool, we have looked at the nature of infrastructures, in general looking for trends and design features that contributed to their success or failure. Drawing on this research we feel the services outlined below are strong candidates for inclusion in this emerging class of distributed collaborative systems. We are confident that there will be services at each level and think the services outlined below are reasonable candidates.

At the network level three services would seem to be logical: resource location, vocabulary, and storage. Resource location is critical for a networked information system. Regardless of the form of the information there needs to be a way to locate resources. In the tradition of DNS, resource location should live in the network and be, by design, simple, extensible, and generalized. Every resource must have a description to serve as input for the resource location service. We envision as a corollary of the resource location service, a vocabulary service that would serve to relate descriptions. (A service that would operate deductively on RDF schema might be an example.) The semantic and pragmatic issues associated with this service are challenging, and the ability to describe resources in non-locational terms adds to the complexity of this service. However, we are confident that a comprehensive way to describe resources is essential to the success of this new class of systems and will evolve in one way or another. Finally, in an environment where users move freely, selected resources required by users must be reliably stored and easily retrieved from the system. A storage service may be implemented to guarantee this availability.

At the system level services two services stand out: security and transaction. Much in the tradition of RSA encryption, there should be common methods to assure the resources are secured. Whether or not this is an essential service or not is unclear. However, it is a great example of a system service as it is positioned at the application's interface to the network universally throughout the network. Classifying a transaction service as a system level service is a matter of debate, but the possibility of a base protocol to manage transactions between parties (e.g. services, objects, etc.) is far from unheard of.

Application level services are the most difficult to illustrate because they will evolve with the software deployed in the system. API's and application frameworks are positioned here as they will simplify, expedite, and standardize development of software for use within the system. We believe that API's will develop around the functional areas defined in section 2.2.2.

The services suggested above and their positioning is far from the only solution. Indeed as these new classes of large scale systems are developed our suggestions may be completely off. However, our suggestions are meant to serve as a catalyst for discussion of how the infrastructure's composition and how it will operate.

Many of the visions of the web hold dearly to the notion that the services not only be distributed, but also decentralized. We distinguish these two terms as follows: A system is distributed when the components that make up the system exist at more than one point allowing for replication and graceful degradation. A distributed system

may be highly centralized in terms of control, or control of the system may be decentralized. Thus, a distributed DBMS system may be controlled centrally. DNS on the other hand is not only distributed, but selected functions are also decentralized. A step beyond this is a federated system. From our point of view, federations are a collection of systems that voluntarily adhere to a common design or ad-hoc standard. From an infrastructure point of view, this means that service cannot be guaranteed. Thus, the current vision of the web itself is distributed and decentralized, and supports numerous federations. Beyond this issue of control, it will be important to decide at what cost extensibility of the system will be a primary focus of the architecture. No doubt, the more extensible functionality is, the more likely it is that the system will be more expensive to develop and deploy. Similar costs are associated with engineering for graceful degradation and scalability. (Appendix D provides more on the infrastructure standards that are required.)

2.4: Toward Complication of the Task: Two Possible Developments

There are any number of developments that could have a significant impact on the kinds of collaborative environments that are imagined here. The progress being made on agents has already attracted a lot of attention. As appliances become more intelligent, and in general as intelligence moves more and more into the periphery of the network, the infrastructure will need to change. The simple, and already evident growth of wireless technology is about to migrate past the era of being a “less”²⁴ technology and will surely have a significant impact on how we think about “environments”. Two developments more closely associated with knowledge environments are likely to have a direct and important impact on the design of the information infrastructure. Agents, which we believe are better thought of as surrogates, are likely to have important and significant impacts. A second change that will have an impact on knowledge environments will be the emergence of new document forms. We briefly highlight the kinds of changes we imagine in this section.

2.4.1: Surrogates: More than Agents

There are a large number of projects that make use of agents to manipulate information stores. Indeed the Berners-Lee view of the semantic web is one that is focused on making the web as an information store more machine understandable. It is our view that agents, while important, may invoke too narrow a view for the current effort. Generally speaking agents are viewed as representing human actors. Robots, spiders, delegates, etc. all share this view. As we see it, surrogates can also represent human actors, but surrogates can be postulated for other objects as well. A document could have a surrogate, as could a calendar, or an organization. (Appendix C provides more on the meaning and import of thinking in terms of surrogates.)

A world of surrogates may be created in which the surrogates all interact with differing levels of autonomy and with functions appropriate to the objects they represent. Some surrogates might only respond, while other surrogates might be more proactive. By way of example, consider the development of an information marketplace for physicians. We might imagine the following kinds of surrogates in such an environment:

- The physician regularly interacts with data about patients. In the case where a physician is dealing with a referral, he or she has to make a decision about whether to trust test results in the patients record. These tests can have metadata associated with them indicating who did the test, under what conditions and with what confidence. The physician could assess this data, but this is a tedious and time-consuming task. One might imagine a surrogate for the test interacting with an agent(surrogate) for the doctor to determine if the test results met the doctors standards for this patients situation. (A doctor might accept a test related to one area of diagnosis but require their own tests in another.) All of this interaction might be trusted to surrogates.
- Doctors often have patients referred to them. These referrals are based on a web of trust that might be recreated using surrogates for doctors and certified appraisals of the doctors by the referring physicians. In

²⁴ Interesting technologies often begin in the shadow of the more important technologists they replace and are initially constrained by their parent technology. “Wireless” networks are simply the next in the line of technologies such as “horseless” carriages and “cordless” appliances.

this way, a physician could ask for a list of doctors meeting certain criteria who had been referred by other doctors known to the doctor making the query.

The key idea here is that a test might have a surrogate that speaks for its reliability and that exposes at an appropriate level information about itself. Similarly, a physician might have a referral surrogate that represents the assessments of that physician's capabilities in previous referrals. The surrogate for these information stores could be proactive as well as reactive making the whole marketplace more dynamic.

2.4.2: Dynamic Documents: More than Digital

In the early days of system design, the task was relatively easy. Analyze the existing system and translate the appropriate processes into algorithms. In this way, many processes such as calculating paychecks moved from manual to computerized processes. In the late 1980's and early 1990's, it became clear that formalization of processes often lead to an understanding of weaknesses in the design of business practices, and "reengineering" became a part of the design process. It was no longer enough to understand the current process. Once that goal was achieved, it was sometimes necessary for information scientists to engage the organization as change agent. All through this process, but clearly through the latter 1990's, it was clear that the system could serve as the locus of more and more intelligent processes off loading low-level intellectual activity from the human. It is debatable just where the line between sophisticated algorithm and machine intelligence is to be drawn, but what is clear is that "best practices" were increasingly purchased with software. At the current time, we note an increasing preference to distribute intelligence in the network and to develop more human independent data acquisition systems – sensor networks.

We believe documents will (are) following a somewhat similar development path. Consider two scenarios:

Web documents started as static documents that were minimally structured. With time, the documents were imbued with scripts and style information that made them more active on the client side. Server side enhancements included client side enhancements and the use of cookies and other information to create increasingly personal documents. XML tagging and metadata were included to make the documents more descriptive.

Word documents began, like other word processing documents as static collections of characters. With time, change tracking, automatic spell and grammar checking, styles, and intelligent agents were added to help users – e.g. "Are you writing a letter?"

Within a constrained infrastructure developed for collaborative authoring at the University of Pittsburgh²⁵, we were able to keep users aware of changes in documents, develop ad-hoc hypertext structures from information about the documents, track dozens of pieces of information about the document that has historically not been tracked – e.g., number accesses, number of minutes open, number of comments, etc. These data were used to assesses user attitude toward the document, construct ballots, inform editing, etc.

In line with these tentative evolutionary changes, we suggest that a likely future for documents, and other media types, is increased dynamism. Not only will documents become digital and structured, but they will become active. By this we mean to suggest that rather than users having to locate documents, an appropriate infrastructure could allow for the development of dynamic documents that could locate human who might be interested in them. As just one simple example, imagine a web based document where scroll bar movement was recorded and fed back to the information store responsible for the document. Imagine that 30% of the viewers of the document scrolled back multiple times from a page where a concept was mentioned to the location in the document where the concept was defined. Further imagine that the information store noted that 50% of those users than used an online source to get more information about the concept. It is possible that this activity might infer the need for more explanation of the concept and different placement of that information.

²⁵ The CASCADE system – Computer Aided Support for Collaborative Authoring and Document Editing. See <http://www.sis.pitt.edu/~CASCADE>

Dynamic documents might know a lot of things and be able to take actions based on what they know. Consider for example the following kinds of information and capability:

- Bookmarks, annotations, and highlighting might be preserved and interpreted by a document
- Documents might assist in defining words, finding relevant explanations, revealing data underlying some graphic or statistic, carry out or demonstrate some process or calculation.
- Documents could keep track of who has read or is reading the document, how much time is spent in various sections.
- Documents could translate a term or phrase, pass on a question about the content someone who could answer it, find remedial or advanced treatments of a given topic. The document could keep track of questions and answers and try to answer questions based on prior answers
- Documents could provide summaries of the comments made, analyses of reading patterns, help the author revise the structure for the book or add new sections when the topic area changes.

3: Conclusions

We are more concerned at this point in time with clarifying definitions and directions than we are with drawing conclusions. At the same time, we would be remiss if we did not draw some conclusions from our investigations. By way of furthering the discussion, we draw two conclusions:

First, we are very taken with the diversity of discussions that are taking place today and with the inability of the various groups to understand the underlying motivations and visions of the future. Indeed, beginning with the premise that there will be a next generation of the web, the shape of that web differs greatly depending on who you talk to. While this is healthy, it is also wasteful to the extent that the discussions are unfocused and unnecessarily confrontational based on a lack of understanding. Clarification of our positions in contrast one with the other is an important goal and one we see as appropriate for the workshop.

Second, while we have funded research to advance the state of the art, the widespread adoption of new technological capabilities is lacking. We believe that in part, this lack of adoption reflects a lack of embedded easy to use technological infrastructure. Put another way, a simple and easy to adopt infrastructure is critical to the adoption and use of new technologies and new capabilities. It is the standardized and well-documented infrastructure that supports the advanced features. From ASCII to the DARPA supported internet standards to the ISO OSI model, reference standards have supported the development of more complex systems by providing a clear and well defined infrastructure model. The current situation is one in which business demands have caused an under provision of infrastructure standards and a stranding of advanced technologies. The relative roles of industry, NIST, and NSF in providing for standards are a matter of historical record. We urge the panel to consider an NSF funded effort to understand and solidify the infrastructure needed by committing resources to the standardization effort.

4: Appendices

4.1: Appendix A: An Information Scientist's Lament

4.1.1: Information and Knowledge

As information scientists, we would be remiss if we did not address the core of our field – the term information. We teach our students that we may distinguish five related terms – signal, data, information, knowledge, and wisdom. There are disagreements about meaning as we move down this list of terms, but we can generally agree that each is more and more intimately tied to human or personal activity. That is to say, it is easier to imagine signals and data existing independent of human activity than it is knowledge and wisdom. The organizers of this workshop have chosen to compare and contrast the words information and knowledge in the convening topic. We propose that the term information be reserved to describing atomic units – whether in isolation or aggregation. Knowledge on the other hand may be reserved to refer to both the organizing principles and the resulting organized aggregate of information. Thus, from this point of view, the semantic web would have more to do with the knowledge that structures the organization of information.

We might go further to suggest that a “*knowledge environment*” is an environment that has been structured in accord with some semantics. As we will discuss below, these semantics may take the form of “*metadata*”, which presents an interesting, and we would suggest potentially oxymoronic combination of terms. Coincidentally, we think that the conveners choice of the term “*information ether*” is consistent with this nascent view of the use of the terms. That is, we would suggest that the implication of an information ether is of a “space of rarefied (atomic) elements that provides the basis for the permeation and transmission of something”, i.e. the information ether is the basis for the transmission of knowledge. We would not care to take this definitional issue too much further as we have already reached or entered the domain of religious debates about the meaning of information and knowledge. It is simply our internet to suggest that information is more atomic in nature – not the same as facts, but close – and that knowledge is more related to the organizing principles that serve to codify, organize, and make information useful.

One final thought. Before moving on, we have found some thinking about information to be particularly useful. The academic definitions of information go on and on. Three authors have taken what we find to be rather productive approaches to understanding information and that serve in accord with all of the other definitions to clarify the technology of information. Bob Lucky, in *Silicon Dreams*, reviews and expounds upon the communication theory of information. He states:

*Shannon's information theory is a philosophy of information from the point of view of communications....It gives us a mathematical measure of information and of the information capacity of a communications channel.*²⁶

He goes on to work with the very traditional definitions of information and entropy and to address issues such as when a fact “is not information”, i.e., when it does not reduce uncertainty, when nothing new is learned. He goes on in his inimical way to suggest that “*the purpose of writing is information storage.*”²⁷

Lucky is clear and informative in his analysis. His focus is on information as artifact. In parallel, Shoshanna Zuboff, studying the application of information technology to the workplace observes:

*Information technology not only produces action but also produces a voice that symbolically renders events, objects, and processes so that they become visible, knowable, and shareable in a new way....The word that I have coined to describe this unique capacity is *informate*. Activities, events, and objects are translated into and made visible by information when a technology *informates* as well as automates.*²⁸

²⁶ Lucky, Robert, *Silicon Dreams*, New York, Saint Martin's Press, 1989, p38.

²⁷ *Ibid.*, p92.

²⁸ Zuboff, Shoshanna, ***In the age of the Smart Machine***, New York, Basic Books, 1984, pp9-10

This use of information as verb as well as noun is echoed by Dertouzos several years later:

...information can be a noun or a verb. Text, sounds, images, videos are information nouns with names like the Bible, Marseillaise, and Star Trek. Computer programs that transform text and images and perform work are information verbs...Humans produce information as both a noun (speech, writing, gestures) and verb (processing of office work using their brains).²⁹

Dertouzos goes on his book to make a compelling case for “information marketplaces” as places where this information as verb is exchanged. He also makes a compelling case for suggesting that the hard won analysis of the characteristics of information (noun) as commodity distinct from physical commodities in many ways is flawed when applied to information (verb) in the information marketplace.

So we now conclude with this final caveat that we need to take care in our assumptions that we are talking about the same thing when we talk about information, knowledge, wisdom, etc. Put most simply, and without recourse to jargon, information may be some thing – a fact, but it is also true that the exchange of information between humans involves the action of informing or not informing the recipient. Indeed, we arrive at an apparent conundrum which says that information is relative – it only exists when it changes the state of the recipient – when they are informed.

4.1.2: Semantics and Ontologies

Related to this discussion of information is the whole matter of semantics. For better or for worse, Berners-Lee chose the term “semantic web” to describe his vision of the next generation. Put most simply, semantics is the study of meaning and semantic relates to the meaning of language. The one can infer that the semantic web is one in which the meaning of the language is understood, and it would appear safe to assume that this is exactly what Berners-Lee had in mind. A web of information resources where the meaning of the resources were somehow exposed. At the heart of this proposal is a recognition that as humans traverse the world wide web, they are able to infer the intent or meaning of the words they see on the pages. This use of the web, where semantics are inferred by humans, is less than optimal given the shear size of the web and the growing use of programs to sort and organize the web.

An informal assessment of the number of web server hits by spiders harvesting and analyzing pages suggests that a growing amount of bandwidth is devoted to less than optimal efforts to use full text indexing of web pages to aid in this process of finding. Further, given the increased use of programs to produce pages based on cookies or other personal or state information used in the interchange, spiders become less effective at gathering information that might be used to infer the meaning or intent of the pages. Thus the semantic web is one in which the meaning or the intent of the resources is apparent at some objective level. How is this to be accomplished? Ontologies! If we say ontology, as we often do just to get a bump in the head ala Wittgenstein, we have roamed back into the realm of philosophy where ontology is theory about the nature of being and epistemology is the study of knowledge and knowing. (In this context, we may have wished to understand the epistemology of the web!) In any case, the AI community in computer and cognitive science uses ontologies as a *specification of a conceptualization*. An ontology is a formal description of concepts and relationships among them. Ontologies are built as a meta-framework for the purpose of enabling knowledge and information sharing. Pragmatically, ontologies are normally written as a set of definitions of formal vocabulary.

If the world were a simple place, we might have one single ontology that would reflect the concepts in the real world and all of the relationships among them. To some extent we do already all share some concepts and relationships, but our personal ontologies vary in degree of detail and scope. But again, if we had one, and if it could be used to classify appropriately granular atoms of information, it would be possible to classify and relate all the information on the web. Without delving too deeply into descriptive logic, it is fairly easy to understand that there are any number of relationships that might be specified between concepts – from a simple as x is a subclass of y to the more complex anything that is an x is not a y. The issue of making sense of the web using ontologies begins with the

²⁹ Dertouzos, Michael, **What Will Be**. Harper, 1997, pp53.

notion of a shared set of conceptualizations and rapidly gets more complex as increasingly complex relationships are allowed.

At one end of the continuum, the semantic web would be enabled by a global shared ontology that could be used to classify resources. In this sense, it would be little different than a thesaurus used to classify items. At the other end of the continuum, the semantic web may be imaged as a set of resources describes by hundreds of partially related ontologies with extensive and varied relationship language. Each end of the continuum and all the points in between make sense for different visions of the future of the semantic web. Unfortunately, these different visions and their mapping to the kind of ontological sophistication imagined, are seldom articulated and related.

4.2: Appendix B: Functions supporting Collaboration

Collaboration systems may be differentiated on the array of services they provide. The tables below are from the analysis done preparatory to the development of the CASCADE system. While dated, the services are still generally descriptive of those required. The first two tables organize the features by topical areas. The last table provides a functional organization consistent with the taxonomic system described in section 2.2.2.

Services	Level	Features
Email	Basic	creation-postpone-delivery of messages
		folders
		filters
		threading
		address book
		aliases
		form response
		multimedia
		spelling checking
		automatic quoted reply
	Extended	process initiation, i.e. task automation via email
		forwarding
		auto-reply
		speech act-based dialogue
other full-fledged editor features		
Database Management	Basic	creation-storage-retrieval-removal of items
		multiple data presentation/views
		association (linking) mechanism
		concurrent access controls
		indexing and sorting
		versioning
	Extended	conflict resolution
		scripting and programming language
		automatic task processing
		distributed database or client/server support
Document Management	Basic	electronic forms (creation process)
		document conversion
		library functions
		commenting
		linking
		link management
		SGML supports
	Extended	document-management-specific user interface e.g. navigation tools, color coding
		writer-commenter management tools
		linguistic analysis
		visualization tools for intra- inter-document analysis

Preparatory Observations

Calendaring - Scheduling	Basic	creation-removal and rescheduling of items (items include meeting, single event, multiday event, repeated event, call, and to-do
		Individual
		Group
		categorical scheduling
		filter
		prioritizing
		follow up on changes
	Extended	reminders
		calendaring-specific user interface e.g. color-coding, banner
		automatic time allocation for both individuals and group
		automatic conflict resolution
Conferencing-Asynchronous	Basic	process initiation e.g. automatic group notification/broadcasting
		integration with other information e.g. linking from calendar items to database records
		basic email features
Conferencing-Synchronous	Basic	discussion forum (threaded messages)
		bulletin board
Conferencing-Asynchronous	Extended	extended email features
	Conferencing-Synchronous	Basic
computer based audio		
computer based video		
mixed public private windows		
Management Information Systems	Basic	brainstorming tools
		idea organizing tools
		idea prioritizing tools
		voting tools
		anonymous posting/voting
		access to various information resources e.g. on-line repository
		asynchronous tools (may need access control features)
editor		
Workgroup Utilities	Basic	network management, e.g. group membership
		administrative services,
		e.g. document access and document movement
	Extended	aliases
		overlapping groups
		versioning
		conversion
		printing

Preparatory Observations

Service Group	Service	Feature	Details
Message transaction support and management	Email, voice mail, fax	multimedia	
		filters	
		threading	
		addressing & aliases	
		folders	
	Message transaction management	bulletin boards	
	Structured messages	processing of incoming	calendar
		edi	
		forms	
generation of outgoing		e.g. ballots, meeting setup, meeting reminders	
	integration with other mail systems		
Activity coordination & meeting management	Task/process oriented	routing	
		approval	
		annotation	
		conversion	
	Meeting oriented	discussion for a (asynch. bulletin board)	
		synchronous "whiteboards" -symbolic, audio, video	
		structured	brainstorming decision making
		real-time chat	
	Scheduler	e.g. group calendaring reminder	
	Shared Information objects- documents, databases, designs, etc.	Document Management	structuring
conversion			
version control			
library function			
access control			
linking & link management			
commenting			
conflict resolution			
electronic forms			
DBMS		structuring	
		conversion	
		access	
		automatic data replication	
		database synchronization	
		hooks to external databases	
Other (e.g. CAD/CAM)		structuring	
		conversion	
	access		
Workgroup Utilities	Scripting language		
	Server Management	documents databases	
	Administrative Services		
Other	Built-in Tools	text editor	
		spread-sheet	
		remote \& local OLE	
		searching facility	
	Integration capability	with Internet	
		with external gateways	
	with other applications		

4.3: Appendix C: Web Surrogates

4.3.1: Introduction

Several visions of the future of the web are closely associated with agents and programmatic processing of web information. In the current view, we have observed a dichotomy between “active” agents and “static” information stores. As information scientists, we would like to imagine that information stores may be dynamic as well. Thus, we might imagine that active information stores could find users just as easily as user agents could find information stores. To help think in this frame, we substitute the term surrogates for agents. We don’t find any objective reason why the term surrogate is better than agent. It simply has less meaning associated with it, and can be used as an aid in communicating about information stores in a broader and more dynamic way. With this in mind, and hoping to foster some out of the box thinking, we paint a picture in the following sections of a world populated by surrogates, for both **providing** information stores and **requesting** information stores. Thus, as a functional definition of web surrogates, one might think of them as information objects that broadcast descriptive, semantic information about themselves and at the same time receive information from other web surrogates that are related, needed, or appropriate given the context.

4.3.2: Web Surrogates Defined

When are web surrogates appropriate, instead of an agent? How are they different? We think of web surrogates as being appropriate to assist people and agents in three main categories; 1) information objects, 2) tools and appliances and 3) as guardians for the living.

Information objects could have web surrogates. A document, for example could have a web surrogate to broadcast descriptive semantic data about its’ content and to receive information that it may need, such as citations or data. Other information objects, such as books, museum collections, and art are commonly considered artifacts. These artifacts represent the interaction of thought, creativity and storage in some physical object. These objects have meaning locked in the medium, and currently the substance of that meaning is very quiet on the web.

Furthermore, since a document may incorporate different types of data, qualitative, quantitative and visual, the surrogate could make use of web services³⁰ for specific presentation tasks. Such tasks could be in finding and using the “best” presentation of the information to facilitate user decision support. Thus, tools are needed by the information objects web surrogates.

Tools and appliances could have web surrogates as well. Databases and applications, agents and web services may also be considered as tools that may benefit from having web surrogates. These tool web surrogates could broadcast and receive necessary information from other web surrogates, and thus interact with information object web surrogates. These tools may range from software programs, statistical applications, visualization systems, GIS and virtual reality tools, to devices such as my PDA, cell phone, or laptop.

Lastly and perhaps the most importantly, web surrogates could be used to act as guardians for the living. Avatars, Guides, Automaton are all just a few terms being used to approach this concept. Web surrogates for the living; people, animals, trees could also broadcast relevant information about them selves and receive relevant information about other web surrogates. Far more than a dating service, far more than a real time GIS location device, far more than a personalized shopper on the internet, the individual’s web surrogate could offer life critical medical service, and could be used to revolutionize the medical industry globally.

4.3.3: One example

What would happen if all individuals in the world had access to an easy to use tool that they could use to build their web surrogate? Envision an all inclusive tool, for both **providing** information and **requesting** information at the

³⁰ Kuno Harumi, Sahai Akhil, “My Agent Wants to Talk to Your Service: Personalizing Web Services through Agents”

individual level – one that can scale eloquently to aggregates of individuals, family units, communities, companies, or governments. The tool would provide a way for creating the requests on the web – a lot of RFPs, or requests for products/services - in a standard, semantic representation. A key assumption is that demand proceeds supply, so make it possible to express the demand. Yet with an open standard, responders - or the providers of that information - will have an incentive to find those consumers requesting information.

4.3.4: Capture

Since the tool has to first provide a way for creating the requests on the web, the web surrogate would have to first capture all variables that are important to an information objects, tools and appliances and for individuals or the living, for the purpose of conducting both broadcasting and receiving of information for interaction. The form of interaction could be e-commerce, both for the buy and sell sides of the transactions, or the form could be interactions on many doctor's recommendations of how to improve a patient's health. Ultimately, the information in the web surrogate can provide a way to optimize an individual's life.

Furthermore, the defining attributes of an individual's web surrogate is that they can be user controlled, semi-autonomous, or autonomous. That it first must be able to capture data about the individual's choices, needs and variables. That these parameters are complicated equations that involve the quantification of all preferences of all products and services in terms of cost, quantity, quality, and time, risk, opportunities and aesthetics, and perhaps other variables, essentially what economists refer to as utility functions. One main variable could be to differentiate between information needed and information wanted. "Need vs. Want" definitions will be needed to help the optimization process.

The data about an individual behavior could be gathered from historical data stores, such as medical records, tax records, credit card companies, utility companies, and bank records. All choices and consequentially historical transactions for all buy and sell activity could be gathered and integrated into the web surrogate. Explicit instruction and data entry from the individual with an easy to use web based GUI could be used to gather data on the individual preferences. Interfaces could be expanded to include complicated gaming and simulation applications of situations that could be used to build a model of the individual through interaction. Furthermore, "what-if" scenarios could be used for exploring behavior under new circumstances. Such a GUI could help people to express their preferences and needs, even if they can't articulate them. Data capture could be augmented modules that observe choices, actions and purchases and learn the individual's preferences over time. These modules would enable the surrogate to dynamically change over time as preferences change. The surrogate could also dynamically adjust to new exogenous factors such as business cycles and economic trends.

4.3.5: Interactions

Once data is captured, it has to be expressed for both *providing* information and *requesting* information. CMU's WEbMate³¹ and MIT's MyAdvocate were similar ideas, however with slightly different focus. The web surrogate will have to publish data and information about itself. Providing and requesting information on the web is a lot like buying and selling products and services. The simplest idea is to have standard interfaces for expressing. Providing and requesting information, as are currently available when buying and selling products and services in the form of RFPs, and the corresponding bid responses.

4.3.6: Required Framework

Standards are essential for communication and for multiple simultaneous development efforts and efficiencies. Additionally if the web surrogates are all created in a standard way, a standard process, and a standard technological interface, their proliferation could help to boot strap a de facto Semantic Web into existence. One method is to use

³¹ Chen, L. and Sycara, K. (1998). "WebMate: A personal agent for browsing and searching", In Proceedings of the 2nd International Conference on Autonomous Agents, Minneapolis, MN, May 10-13.

existing and emerging standards such as XML and UDDI³², and other XML meta-languages in appropriate domains. The web surrogate must be expressed with standardized ontologies.

When standardized ontologies do not exist, collaboration has to be supported, once again in a standard process, to create or extend an ontology. Such an ontology editor, as OntoEdit³³ would be required. Once expressed on the semantic web, the web surrogate may accept responses to the advertised requests.

4.3.7: Decision Making

Once the requested information is received, the web surrogate must express the received information in terms of the parameterized expression of the individual, so as to find matches or close matches. Activity by the web surrogate could be autonomous, semi-autonomous or under user's control. Filters could be used to filter incoming messages and to rank messages relative to the parameters. Filters and E-mail notification are common practice even today, however constrained to specific applications, such as a job search web site like monster.com. Mitigation of cognitive overload may become essential, and that will require filtration and presentation inferences. If the user opts for control, then decision support tools become essential. The web surrogate could make use of web services for specific presentation tasks, finding and using the "best" presentation of the information facilitating user's decision support. Basically, presentation depends on data type and user style, and qualitative data presented in documents and presentations, quantitative data presented in spreadsheet formats and or statistical and visual data presented in charts, visualizations, geographical information systems and virtual reality.

Web surrogates requesting information – the demand side - will operate in a variety of different ways. Autonomous transactions may be good for steady state activities and leverage existing e-commerce infrastructure. Such repetitive tasks commonly found in existing electronic fund transfer (EFT) interfaces of the individual's bank's system, tax payment systems, accounting systems, lending, credit card payments, and others could be leveraged. The payment to the current provider is steady state, but not the contract. For example, a utility company that provides electrical or gas service, has a contract for the service, which could be managed by an agent that seeks a better price and service agreement continuously and in real time.

Semi-autonomous transactions that require a user decision to execute could also be built into the web surrogate.³⁴ An example of such a hybrid may be if the user receives some proposals from utility companies that are more competitive than the current provider. The user may want to reserve the actual decision for themselves and not the web surrogate. However, once the decision is made, web surrogates may be used to update the EFT detail, thus canceling the EFT to one utility company and setting up the new EFT. Or the reverse may be true; a user may want the selection of the new provider to be automated along maximizing a variable, or ranking of variables, and the managing of the EFT to be under user control. Also in some situations, complete and total user control may be required for certain transactions, as well.

4.3.8: Auctions

The actual transaction of e-business with auctions has been a reality for several years. Most notably, eBay.com for the retail sector and Freemarkets, for B2B, all use software that allows for dynamic, real-time auctions. The web surrogate that participates in auctions could instantaneously and simultaneously allow for participation in multiple negotiations of requests for products and services. Furthermore, if instructed, it could execute selection and transaction of products or services, and contract settlements and so they may also automatically participate in auctions, for competitive pricing. Web surrogates could, if allowed leverage existing e-commerce infrastructure.

³² Akkiraju Rama, Flaxer David, Chang Henry, Chao Tian, Zhang Liang-Jie, Wu Frederick, Jeng Jun-Jang "A Framework for Facilitating Dynamic e-Business Via Web Services"

³³ Sure York, Erdmann Michael, Angele Juegen, Staab Steffen, Studer Rudi, and Wenke Dirk "OntoEdit: Collaborative Ontology Development for the Semantic Web."

³⁴ Sycara, K. and Lewis, M. (2002). From data to actionable knowledge and decision. Proceedings of the Fifth International Conference on Information Fusion, Annapolis, MD, July 7-11, 2002.

Furthermore, trading blocks, team work and coalitions may form. These web surrogates may, if they find other web surrogates where partnerships are mutually advantageous, dynamically aggregate to form coalitions and groups for increased purchasing power and negotiation. So, for efficiency gains from an economic standpoint and optimize the individual's utility function with teamwork. Existing agent infrastructure such as RETSINA³⁵ could possibly be adapted to such activities.

4.3.9: Conclusion

Lastly, optimization of an expressed utility function of information objects, tools and appliances and guardians for the living, only becomes possible when they are expressed in a standard way and as mathematical equations. Where along any variable dimension, auctions and aggregations may dynamically form. Total optimization and recommendations to optimize could be constructed.

³⁵ Dedker Keith S., Sycara Katia, "Intelligent Adaptive Information Agents", Kluwer Academic Publishers, Boston.

4.4: Appendix D: A Paper on Infrastructure Standards

A paper looking at e-business issues and the next generation web will be presented in July at a conference in Portugal. It provides additional background on infrastructure issues that need to be addressed for collaborative knowledge environments.