Web Economics:
A Case for Agent-Based Digital Libraries

Innes A. Ferguson, Jörg P. Müller, Markus Pischel, and Michael Wooldridge

Agent Systems Group, Zuno Ltd, Fourth Floor, International House
Ealing Broadway Centre, London W5 5DB, United Kingdom

{iaf, jpm, pischel, mjw}@dlib.com

Abstract

We advocate *agent-based digital libraries* as a key technology for solving some of the most important problems associated with the internet today. In particular, the aims of this paper are six-fold: to introduce the state of the art in internet-based information systems and in particular, the World-Wide Web (WWW); to detail some problems associated with the internet/WWW as it currently stands; to discuss the requirements for tools to overcome these problems; to motivate the idea of a digital library (DL) as the kind of system that will meet these requirements; to justify the claim that an appropriate way to build such digital libraries is as an *agent-based information economy*; and finally, to introduce ZUNODL (Zuno Digital Library), a forthcoming commercial framework for building digital libraries as agent-based information economies.

1 Introduction

The internet has been in existence for nearly three decades, and has been in everyday use within academia for at least half of this time. Recently, driven by the widespread availability of cheap but powerful PCs, but perhaps more importantly, the emergence of an easy-to-use front end in the form of the World-Wide Web (WWW), the internet has become increasingly used by everyday office workers and home users. The internet could now be said to have reached a critical mass, in that both the volume and quality of information available on it, and the number of end users, makes it both necessary and useful to academia and commerce.

As the internet has grown, so has the amount of information available on it. The *potential* of the internet as an information resource is enormous, and is immediately obvious to anyone with more than a moment's experience of WWW use. However, *realising* this potential seems to be very hard: building software tools that can help a user to do this is one of the most important challenges facing computer scientists in the late 1990s. We need rich, personalised, added-value services that help us find, retrieve, manage, understand, and interpret the wealth of valuable information that is so obviously present on the internet.

In this paper, we argue that building the tools that provide these services will require techniques developed by AI researchers. Specifically, we argue that the appropriate tools will be *digital libraries*, and that a compelling way to build an internet-based digital library is as an *agent-based information economy*. We begin, in section 2, by describing the state of the art in internet-based information retrieval and management tools, and focus in particular on the inherent limitations of these tools. This leads us, in section 2.3, to identifying some requirements for internet-based information retrieval and management tools, and in section 3, to introducing the notion of a *digital library* as a system that provides added-value services to raw, heterogeneous information sources. In particular, we claim that internet-based digital libraries may usefully be built as *agent-based information economies*; we motivate and explain these ideas at length. To better understand how such a digital library might work, in section 4, we introduce the Zuno Digital Library (ZUNODL), a commercial product that facilitates the development of digital libraries as agent-based information economies. The ZUNODL enables markets that support publishers in vending digital documents to customers, and provides customers with powerful tools for filtering and retrieving documents. A ZUNODL system is a collection of interacting agents, with each agent playing an information consumer, producer, or mediator role. In section 5, we discuss related work, and in section 6, we present some conclusions.

2 The State of the Art

The internet, and in particular its most famous offspring, the WWW, has become *the* icon of the 1990s. No other software product in history has caught the imagination not just of the computing community, but of the wider public in the way that the WWW has. The WWW is variously seen as the convergence of media industries, the nemesis of the book, and even as a tool for fomenting revolution. Ultimately, however, the WWW is simply a framework for information retrieval. What makes the WWW so novel, effective, and exciting is that:

- it allows access to networked, widely distributed information resources;
- it provides a uniform interface to multi-media resources including text, images, sound, video, and so on;
- it is hypertext-based, making it possible to link documents together in novel or interesting ways;
- perhaps most importantly, it has an extraordinarily simple and intuitive user interface, which can be understood and used in seconds.

It is already a cliche, but it is nevertheless worth pointing out that the internet/WWW will fundamentally change many aspects of work. These changes are already happening in academia, but as time passes, we can expect to see similarly radical changes in the way that both commerce and industry work. Eventually, we will see the emergence of an *information industry*, in which digital information is routinely bought and sold across the internet.

The reality of WWW use in the late 1990s is, however, still somewhat removed from this goal. There are several hurdles that must be overcome if the internet is to be used as the backbone for this future industry. We can divide these problems into two categories: *human* and *organisational*.

2.1 Human Factors

The most obvious difficulty from the point of view of human users of the World-Wide Web is the well-known "information overload" problem [Maes, 1994]. People get overwhelmed by the sheer amount of information available, making it hard for them to filter out the junk and irrelevancies and focus on what is important, and also to actively search for the right information. Search engines such as Lycos and Yahoo attempt to alleviate this problem by indexing largely unstructured and unmanaged information on the WWW. While these tools are useful, they tend to lack functionality: most search engines provide only simple search features, not tailored to a user's particular demands. In addition, current search engine functionality is directed at textual (typically HTML) content — despite the fact that one of the main selling features of the WWW is its support for heterogeneous, multi-media content. Finally, it is not at all certain that the brute-force indexing techniques used by current search engines will scale to the size of the internet in the next century. So *finding* and *managing* information on the internet is, despite tools such as Lycos, still a problem.

In addition, people easily get bored or confused while browsing the WWW. The hypertext nature of the WWW, while making it easy to link related documents together, can also be disorienting — the "back" and "forward" buttons provided by most browsers are better suited to linear structures than the highly-connected graph-like structures that underpin the WWW. This can make it hard to understand the topology of a collection of linked WWW pages; indeed, such structures are inherently difficult for humans to visualise and comprehend. In short, it is all too easy to become lost in hyperspace. When searching for a particular item of information, it is also easy for people to either miss or misunderstand things.

Finally, the WWW was not really designed to be used in a methodical way. Most WWW pages attempt to be attractive and highly animated, in the hope that people will find them interesting. But there is some tension between the goal of making a WWW page animated and diverting and the goal of conveying information. Of course, it *is* possible for a well-designed WWW page to effectively convey information, but sadly, most WWW pages emphasise appearance, rather than content. It is telling that the process of using the WWW is known as "browsing" rather than "reading". Browsing is a useful activity in many circumstances, but is not generally appropriate when attempting to answer a complex, important query.

2.2 Organisational Factors

In addition, there are many organisational factors that make the WWW difficult to use. Perhaps most importantly, apart from the (very broad) HTML standard, there are no standards for how a WWW page should look. In particular, there are no standards for metadata, or semantic markup, which would allow content providers to annotate their pages with information defining the content of the page. There are some good reasons for this, chief among them being that beyond the obvious (author, title, abstract, date), there is no real consensus on what is useful or possible; there are also significant technical problems with formalisms for defining document content.

Another problem is the cost of providing online content. Unless significant information owners can see that they are making money from the provision of their content, they will simply cease to provide it. How this money is to be made is probably the dominant issue in the development of the WWW today. As it currently stands, the WWW has a number of features that limit it as an "information market". Many of these stem from the fact that the WWW has academic origins, and as such, it was designed for free, open access. The WWW was thus not designed to be used for commercial purposes, and in particular, no consideration at design-time was given to issues such as:

- *privacy and security* — anyone wishing to use the WWW for commerce must implement their own privacy/security mechanisms;
- *billing/revenue* — no built-in billing mechanisms are provided by the WWW: they must be implemented over the basic WWW structure; in addition, the WWW was not designed with any particular revenue model in mind;
- *reliability* — the internet, and hence the WWW, is inherently unreliable, in that data and connections are frequently lost, and it thus has unpredictable performance.

> These limitations may be accepted by home/hobby users, but they represent a very real obstacle in the way of the wider business use of the WWW.

We stress that these are not criticisms of the WWW — its designers could hardly have anticipated the uses to which it would be put, nor that they were developing one of the most important computer systems to date. But these are all obstacles that need to be overcome if the potential of the internet/WWW is to be realised. The obvious question is then: what more do we need?

2.3 Requirements

In order to realise the potential of the internet, and overcome the limitations discussed above, we require a framework that:

- gives a single coherent view of distributed, heterogeneous information resources;
- gives rich, personalised, user-oriented services, in order to overcome the "information overload" problem — it must enable users to find information they really want to find, and shield them from information they do not want;
- supports electronic commerce, with secure, trusted mechanisms for buying, selling, negotiating, cooperating, and so on;
- is scalable, distributed, and modular, to support the expected growth of the internet and WWW;
- is adaptive and self-optimising, to ensure that services are flexible and efficient; and finally,
- integrates seamlessly with the WWW.

3 Digital Libraries

In this section, we will argue that *digital libraries* provide the functionality to satisfy the requirements stated above, and further, that an appropriate infrastructure to underpin internet based digital libraries is that of the *agent-based information economy*. We begin, in the following section, by discussing the question of exactly what a digital library is.

3.1 What is a Digital Library?

The internet today is a highly dynamic, unstructured repository for information. It is not a library. To see why, consider that a library is not simply a roomful of books; it is a roomful of books together with *services*. The minimal service we would expect of a library is indexing, so that we could find desirable information relatively easily. Better libraries provide richer services. For example, a good librarian would know about the library's users, and use this knowledge in their interests. Such a librarian would, for example, not simply search for books or articles when specifically asked, but would proactively offer relevant content to a user as it became available. With the exception of some (fairly crude) search facilities, the WWW offers none of these services. The goal of digital library research is to develop software tools and techniques to provide them. Lynch *et al* define a digital library as:

> "An electronic information access system that offers the user a coherent view of an organized, selected, and managed body of information" [Lynch *et al.*, 1995]

The internet has some features that make it a difficult environment to deal with from a software development point of view. The size, dynamic nature, and lack of central control in particular are very troubling; they together conspire to make the domain extraordinarily complex. Traditional software development techniques and models are simply not up to the task of dealing with this complexity: we need new ways of thinking about and developing software. We believe that some recent AI developments provide a solution. In particular, we believe that *agent-based information economies* will enable us to build digital libraries that meet these requirements. In the subsections that follow, we will explain what we mean by the term agent-based information economy, and justify our claim that such an approach is appropriate.

3.2 Agent-Based Systems

Agent-based systems are a new paradigm for software development. By an *agent*, we mean a computer system that is capable of *flexible autonomous action* in some environment in order to meet its design objectives [Wooldridge and Jennings, 1995]. By *flexible* we mean:

- *reactive* — an agent should be capable of responding to environmental changes and modifying its behaviour accordingly, in time for these changes to be useful;
- *pro-active* — an agent should be capable of exhibiting goal-directed behaviour, generating and attempting to achieve new goals as opportunities arise; and
- *social* — an agent should be capable of interacting with other agents in order to achieve its goals.

An *agent-based system* is one that is designed or implemented as a collection of such interacting agents. We claim that an agent-based approach is highly appropriate for a digital library. To see why, consider that the usefulness of such an approach is indicated by many domain attributes, but in particular: distribution of data or control, the lack of any global viewpoint, highly dynamic environments in which central control is undesirable or infeasible, or a system that is usefully conceived as a collection of interacting semi-autonomous components [Bond and Gasser, 1988b]. The WWW has exactly these properties: data and control are obviously distributed, there is no central control or viewpoint, and it is highly dynamic. Finally, a digital library must offer *personalised* services, automatically tailored to each individual user's needs. Agents acting as "expert assistants" have for some time been recognised as a useful medium through which to deliver such individual services [Maes, 1994].

3.3 Information Economies

In designing a digital library as a collection of agents, a key issue to consider is what organisational form one should impose on the agents in order to support the requirements of the domain. Given the growing influence of internet commerce, one can argue that it should be an organisational form that incorporates the concepts of utility, monetary value, negotiation, and brokering. In addition, given the uncertainty in the future development of the internet, both in terms of the services that might be offered and the services that might be needed, it also needs to be highly adaptive.

In this respect, economic markets provide a useful metaphor for designing complex distributed systems [Wellman, 1994; Ferguson and J.Karakoulas, 1996]. In particular, market-based economies are, by design, decentralized, capable of organising themselves without central control, and capable of adapting themselves to changing supply and demand conditions. The market, or *information economy*, metaphor, therefore, seems highly appropriate for the digital library domain: producers (sellers) and consumers (buyers) of information interact in a market-like environment. Brokers, which mediate interactions between consumers and producers and which have specific knowledge of the supply and demand for information products, can help increase the overall efficiency of the market.

The digital library function is also comparatively expensive in terms of resource consumption when compared to, say, the search engines we find on the internet today. For example, a digital library may need to store substantial amounts of user profiling information; in addition, it may need to manage large numbers of regularly scheduled alerting or persistent query services. Efficient resource management within a digital library is thus particularly important. Fortunately, markets have the capability to manage resources very effectively: efficiency and prudence tend to be rewarded, inefficiency and imprudence penalised. Through their feedback mechanisms, market systems are capable of evolution, with better — more efficient, more powerful, cheaper — products and services tending to dominate over those of lower quality. In fact, because it is difficult — if not actually impossible — for people (users, designers, system administrators) to manually optimise the performance of any large-scale distributed information system, a system such as an internet-based digital library must be capable of *self*-optimisation. A market-based design makes this possible.

Finally, by starting from the point of view of an information economy, with concepts like "producer", "consumer", "price", and "profit" built in, we are much better positioned to support real commerce, of the type we ultimately intend to facilitate. Indeed, as the internet and WWW become increasingly commercialised, the need for effective profit-oriented agents which can act on behalf of their owners and seek payment for services rendered is likely to increase dramatically [Genesereth and Ketchpel, 1994]. However, in order for such sophisticated agents to materialize, key work remains to be done in defining and deploying techniques for brokering of information products and for linking requests of demand (consumer) agents to the ability of supply (producer) agents to provide the required information items [Brown *et al.*, 1995; Kuokka and Harada, 1995]. A number of these issues have been addressed directly in our agent-mediated digital library framework. In the following section we describe this framework and the various economic roles played by its constituent agents.

4 A Case Study: The Zuno Digital Library

In this section, we present a case study: the Zuno Digital Library (ZUNODL). This system is in fact not a single digital library, but a collection of tools and techniques for building digital libraries as agent-based information economies. The main features of ZUNODL may be summarised as follows:

- *Agent-based*: The ZUNODL framework allows a digital library to be implemented as a collection of cooperating agents. Agents are the key to providing the robust, proactive, personalised services that we seek.
- *Information economy*: The agents in a ZUNODL system take on the role of producer, consumer, or mediator, acting in an information economy.
- *Decentralised framework*: The ZUNODL framework is not centrally managed: information owners own, manage, and provide access to their content. They retain ultimate control.
- *Domain independence*: ZUNODL is not directed at any particular type of domain — ultimately, it can be used to build digital libraries for both internet and intranet applications, with many different types of information.
- *WWW front-end*: the ZUNODL framework does not require end-users to have access to powerful hardware or software: the user interface to the system is the WWW.

4.1 Main Concepts

A ZUNODL system is realised as a collection of agents, each of which plays one of three key roles:

- *Producers*: Producers correspond to information owners — these are organisations or individuals that have content they wish to sell.
- *Consumers*: Consumers correspond to the library end users, who obtain access to the library services via a web interface.
- *Facilitators*: Facilitators are rather like brokers — they map between producers and consumers.

Agents can be physically located in one of three places: on the producer's computers, on the user's computers, or on computers run by the ZUNODL system service provider. The service provider will typically correspond to an institution (such as a university library).

4.2 Producer Agents

For each producer (information owner) in a ZUNODL system, there are two agents. One of these, the *Library Service Agent*, represents the digital library within the information owner's network, the other, the *Catalogue Agent*, represents the information owner within the digital library itself:

Library Service Agents (LSAs): These agents represent the digital library within the information owner's network; they are under the control of the information owner, and in essence provide the digital library with an interface to the owner's content. LSAs have the ability to veto any operation requested by the library. The key functions of LSAs are managing access to content, billing, and managing the security of content.

Catalogue Agent (CAs): These components represent the interests of the information owner within the digital library itself. Each CA is twinned with an LSA; CAs will thus "mirror" LSAs. Each CA will ultimately have a cached "model" of its associated content, and will be able to rapidly respond to queries about this content.

4.3 Facilitator Agents

Facilitator agents are rather like brokers; their goals are to map a user's requirements onto the producers that are best able to satisfy these requirements, and also to allow for a certain amount of "producer push", by mapping producer content to the consumers that are most likely to have an interest in the content. Unlike brokers, however, facilitators are not third parties. They are "owned" by whoever runs the digital library.

Search Agents (SAs): Perhaps the key components in ZUNODL are search agents. SAs, as their name suggests, provide the basic search engine functionality that underpins ZUNODL. However, they provide additional functionality, over and above that provided by current internet search engines. First, SAs can make use of thesauri and taxonomies in order to manipulate queries to obtain better results for the user. For example, suppose a user issues a very specific query that results in a small number of results. Then an SA can use a thesaurus to look for related words, in the hope that these will yield better results. Taxonomic hierarchies of words can also be used to broaden queries; put crudely, if a search for "chair" fails to yield useful results, then an SA can use a taxonomic hierarchy to see that chairs are a type of furniture, and do a search for "furniture" instead. The query is thus generalised, with the result that more hits are obtained. SAs can also do enriched searching. Thus:

- searches may exact, may be broadened if insufficient matches are found, or alternatively, they may be based on related terms;
- searches need not be just keyword-based — they can be based on keywords, subject, title, author, and potentially other metadata.

4.4 Consumer Agents

User Interface Agents (UIAs): UIAs represent the interests of the user within the ZUNODL system. Each user that is logged on to the system will be associated with exactly one UIA. Each UIA will request services of the digital library in order to satisfy a user's goals. In order to do this, UIAs maintain user profiles — literally, models of the user. In contrast to current search engines, UIAs will be persistent, in that they will remember about a user from one transaction to the next. If a user logs out of ZUNODLand then logs back in some time later, their profile is retained. In the prototype implementation of ZUNODL (see below), UIAs are all resident on the ZUNODL system service provider's machines; however, this has obvious implications for network traffic use and server hardware support, and for this reason, in the forthcoming product version of ZUNODL, support is likely to be provided for client-side UIAs.

4.5 ZUNODL System Structure

Figure 1 illustrates the way that the ZUNODL system is put together. The main point to note is that the various components of the system are "owned" by different groups. Thus content is owned by information owners, as are the LSAs that give the library access to this content; the user interacts with the system via a web browser running on their own machine; and finally, the other components of the system (essentially UIAs and SAs) are located within the ZUNODL service network.

To see how the system works, it is worth walking through what happens when a user issues a query. The process starts with the user logging on to the ZUNODL system using their WWW browser. When the user has successfully logged in, their UIA is (re-)started. Suppose the user then issues a search query. The UIA will forward this query to one or more SAs, depending on the parameters of the search. For example, if the user has requested a comprehensive search, then the UIA can direct the query at SAs that provide a comprehensive but potentially slow search facility, as opposed to SAs that offer a shallow but fast search. SAs use knowledge of different repositories to direct the query at CAs. CAs may be able to respond to the query directly; if, for example, the query is outside their domain, they can respond immediately, indicating this. If they have sufficient information to answer the query completely, then they can do so. However, they may wish to pass the query on to the corresponding LSA (in the case of a comprehensive full-text search, for example). LSAs are located on information owner machines, however, and the added communication delay in talking to them, in addition to the added requirement for server support on the owner's machines means that they are only used when really required. In

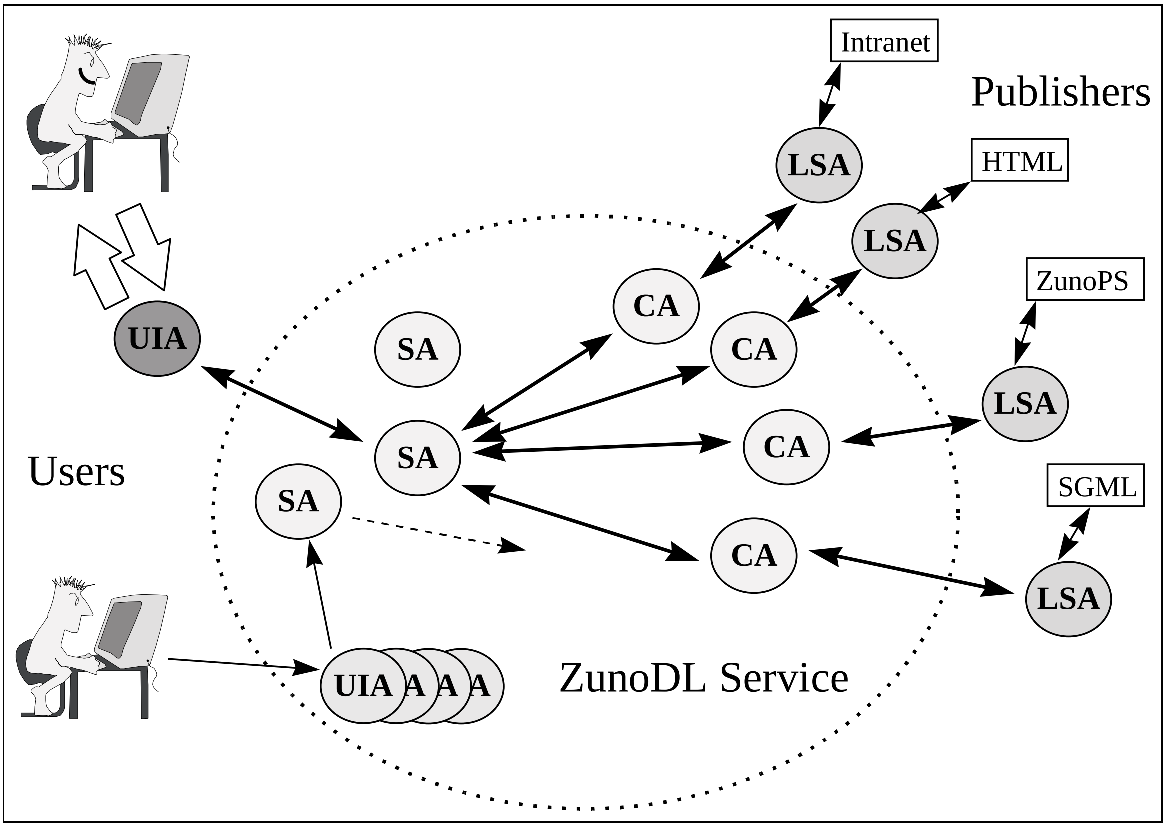

Figure 1: ZUNODL System Structure

any case, the LSA can do searching of its local repository, and can respond to the CA's request, which can in turn pass results back to the SA. Once the SA has begun to receive results, it can start to decide what to do. If the search results are satisfactory, (in volume and quantity), then they can be ordered, sorted, and passed back to the UIA, which can present them to the user. Once the UIA receives the results of the query, it can choose to sort and prioritise them according to the user's interests. For example, the UIA might delete all entries about "James Bond" that appeared in reply to a user's query about "agents" if the user's interests were in software agents rather than fictional spies. If the results to the query are unsatisfactory, the UIA can make use of knowledge about the user to modify the query. For example, it might use "softbot" to replace the query term "agent". The user can then browse the results of the query, viewing the metadata associated with a result item if it is available, and selecting it for purchase if required.

4.6 Discussion

By designing ZUNODL as an agent-based information economy, we are able to deliver a system that satisfies the requirements we set out in section 2.3. First, Library Service Agents and Catalogue Agents provide a uniform and coherent view of heterogeneous information resources. Second, the information overload problem can be overcome by providing and combining different agent services such as personalization, resource discovery, recommendation, and collaborative filtering. Third, by fixing on the market model from the very beginning, we design in support for electronic commerce. Fourth, by using agents, we make ZUNODL modular (new agents can be introduced to represent new information owners) and scalable (agents apply local load balancing strategies). Fifth, by building ZUNODL as an information economy, we are able to ensure that it will be self-optimising and adaptive: efficient resources will tend to be used more than inefficient resources, and will thus receive positive feedback. In the same way, a market system can automatically evolve to meet the requirements of a changing environment. Sixth, ZUNODL provides a web-based graphical user interface that provides seamless integration with the WWW.

From the perspective of a designer of agents and agent-based systems, building a digital library such as ZUNODL comes with a number of interesting challenges and requirements for AI research and software engineering:

- it must support the robust handling of failure due to the (temporal) unavailability of resources or communication failures;
- interoperability is crucial: agents must be able to communicate and cooperate with a wide range of other system entities;
- participating in a market-like organization requires agents to support higher-level cooperation protocols and built-in models of utility;
- agents as well as organizations must be adaptive: at the agent level, this can be achieved by providing feedback-based learning mechanisms; at the agent system level, explicit representations of and reasoning about organizational structures seem necessary;
- the digital library architecture as well as the infrastructure underlying the multi-agent system need to cater for the uncertain: scalability is a *must*;
- the design and the operation of real-world multi-agent systems is impossible without appropriate design-time and run-time tools, that support building, running, and monitoring of agents and their interactions; and finally
- tools need to be grounded in a *design methodology* that provides guidelines for an agent-oriented analysis, design, and implementation process.

5 Related Work

The architecture most similar to ours in spirit is the University of Michigan Digital Library (UMDL) [Birmingham, 1995]. UMDL is designed as an agent-based distributed system in which agents encapsulate different highly specialized library tasks, made accessible by a standardized communication interface. The absence of a general-purpose agent architecture for UMDL agents is one of the most prominent differences to ZUNODL. Furthermore, ZUNODL being a commercial system, personalization is crucial; thus, we focus on adaptability of the user interface agent acting on behalf of the consumer in the information economy, whereas the focus on adaptivity in UMDL is on the producers' end. Also, functionalities such as secure purchasing and billing are currently not considered in UMDL.

In [Kahn and Wilensky, 1995], Kahn and Wilensky have proposed a framework for a class of distributed digital information services. The basic building blocks in their model are digital objects that are stored in decentralized repositories. A digital object consists of two parts: content and metadata. Repositories are accessed using a *repository access protocol*. Like ZUNODL, Kahn and Wilensky's model uses the notion of a service in order to describe encapsulated library functionalities. Zuno facilitator agents can be seen as an embodiment of their *value-added reference services*. Thus, this architecture can be regarded a basic framework based on which an agent-based digital library can be built. However, it does not address issues of personalization and adaptability which are crucial to our approach.

Similar arguments can be put forward for DIENST [Davis and Lagoze, 1994], an HTTP-based protocol with an object-oriented interface to a document model, which was used to build an online digital library of technical reports from a group of US computer science universities. This work does not address issues of billing, IPR, and security. Resource discovery and searching are realized in a rather *ad-hoc* way, and no clean architectural way is described for fitting in added-value services that correspond to those provided by Zuno facilitator agents.

The Stanford University Digital Library Initiative proposes a so-called *information bus architecture* for a digital library. The different components of the digital library (information sources, library services, and users) interact with the bus. Intermediate *protocol machines* handle translations between different languages and formats used by different participants. The approach suggests that specific services will perform facilitator functions such as search, filtering, and notification. However, the focus of this approach clearly is on *interoperability*; issues such as adaptability and personalization which are directly supported by the Zuno agent-based architecture, are not explicit in the information bus architecture.

Finally, the ZUNODL framework builds upon an increasing body of work in agent-based systems [Bond and Gasser, 1988a] and market-based systems [Wellman, 1994]. Space limitations prevent a deeper discussion of the impact of this work on the ZUNODL framework.

6 Conclusions

In this paper, we have argued that the technology of digital libraries will help us to overcome some of the most important problems with the internet today, and that autonomous, cooperating agents which provide rich, added-value services are the appropriate building blocks of such a digital library. In particular, organizing these agents into an information economy offers us the adaptability, efficiency, and scalability that is needed to deal with the ever growing demands of the information producers and consumers that populate the internet.

In [Etzioni, 1996], Etzioni has presented the metaphor of an information food chain for the web, where high-level information services “intelligently hunt and feast on herbivores”, i.e., on standard web search engines. While autonomous agents seem ideal candidates to populate the upper end of the information food chain, we believe that the food chain metaphor is somewhat limited as it does not account for the interaction mechanisms between peers on the information food chain required to provide not only useful, but also intelligent services. We believe that the market mechanisms described in this paper have the potential to support this interaction.

Acknowledgments: The views expressed in this paper are those of the authors, and do not necessarily reflect those of Zuno Ltd.

References

[Birmingham, 1995] W. P. Birmingham. An agent-based architecture for digital libraries. *D-Lib*, July 1995.

[Bond and Gasser, 1988a] A. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann, Los Angeles, CA, 1988.

[Bond and Gasser, 1988b] A. H. Bond and L. Gasser, editors. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann Publishers: San Mateo, CA, 1988.

[Brown *et al.*, 1995] C. Brown, L. Gasser, D. E. O’Leary, and A. Sangster. AI on the WWW: Supply and demand agents. *IEEE Expert*, 10(4), 1995.

[Davis and Lagoze, 1994] J. R. Davis and C. Lagoze. A protocol and server for a distributed digital technical report library. Technical Report CS:TR94-1418, Department of Computer Science, Cornell University, 1994.

[Etzioni, 1996] O. Etzioni. Moving up the information food chain: Deploying softbots on the world wide web. In *Proceedings of AAAI-96 (Abstract of invited talk)*, 1996.

[Ferguson and J.Karakoulas, 1996] I. A. Ferguson and G. J.Karakoulas. Multiagent learning and adaptation in an information filtering market. In *Proceedings of the AAAI Spring Symposium on Adaptation, Coevolution and Learning in Multiagent Systems*, Stanford University, 1996.

[Genesereth and Ketchpel, 1994] M. R. Genesereth and S. P. Ketchpel. Software agents. *Communications of the ACM*, 37(7):48–53, 1994.

[Kahn and Wilensky, 1995] R. Kahn and R. Wilensky. A framework for distributed digital object services. Technical Report TN95-01, Corporation for National Research Initiative (CNRI), Reston, VA, 1995.

[Kuokka and Harada, 1995] D. Kuokka and L. Harada. On using KQML for matchmaking. In *Proceedings of the First International Conference on Multiagent Systems (ICMAS-95)*, pages 239–245. AAAI Press / MIT Press, 1995.

[Lynch *et al.*, 1995] C. Lynch, A. Michelson, C. Preston, and C. A. Summerhill. CNI white paper on networked information discovery and retrieval. http://www.cni.org/projects/nidr/www/toc.html, 1995.

[Maes, 1994] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):31–40, July 1994.

[Wellman, 1994] M. P. Wellman. A computational market model for distributed configuration design. In *Proceedings Conference of the American Association for Artificial Intelligence*, pages 401–407, 1994.

[Wooldridge and Jennings, 1995] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.