

The Next Revolution: Free, Full, Open Person-2-Person (P2P) E-commerce

Prem Devanbu, Dept. of Computer Science University of California Davis, California CA 95616 USA devanbu@cs.ucdavis.edu	Stuart G. Stubblebine CertCo 55 Broad Street – Suite 22 New York, NY 10004 stubblebine@certco.com	Michael F. Uschold Boeing Applied Research PO Box 3707 Seattle, USA mfu@redwood.rt.cs.boeing.com
--	---	--

July 3, 2000

1 Background

Our goal here is to present a long-term perspective on the evolution of a new, fine-grained type of e-commerce on the internet. We begin with a motivating scenario.

Abuelita wants to buy the latest Neal Stephenson cypherpunk thriller. Almost without thinking, she visits the website of riverofbooks.com, a “category killer” e-tailer. She quickly finds what she wants, and with a few easy clicks, has purchased the book for \$28.99, plus \$4.99 shipping and handling. She can’t wait for the book to arrive in 3 business days. Later that day, she rides her bicycle over to the local mall to buy herself a smoothie, and peeks in the window of local used-book seller, where a Lute concert is underway. Imagine her chagrin when she notices that the very same book is on sale in the window for \$14.99, plus \$2.61 in local taxes! She could have had the book, and supported her local book store (and the lute concert!). Bah humbug, she says to herself. What good is the internet, if it can’t tell me that the same book is on sale in my local used bookstore for less?

Abuelita’s problem will be solved with the advent of person-to-person (P2P) e-commerce. She would have simply described what she wanted to buy; her local bookseller would also be continuously updating its on-line description of what it has to sell; the infra-structure would handle the rest, and notify Abuelita that the book she wanted was a short bike-ride away. Our basic position is that the current “.comization” phase of the internet represents an early state in the movement towards making P2P commerce fully electronic. Another revolution awaits, we believe.

2 Current state of e-commerce

Think quickly about how many brandnames of detergent you know about. Most people will probably recall 3 or 4 major national brand names, stocked by all major national department chain-stores. There are many others— at my local organic foods store, there are numerous other smaller brands, which are ecologically sensitive, allergen-free, low-sudsing, perfume-free, and so on. Consumers aware of all these choices might well make better decision. In fact, most consumers simply rely on brand-names to choose a detergent; presumably this leads to lower competition, and inefficiencies in the market for detergents.

The primary reason for this is ignorance, arising from a lack of information. There is no simple, cheap, convenient way for small detergent vendors to get the word out to their customers. This sluggish flow of information, and the resulting ignorance creates the opportunity for big, well-financed manufacturers;

they can advertise on mass media, and over years, build up formidable brand names that dominate all others. Lack of free information flow thus leads to reliance on brandnames, rather than on actual information about products.

One might think, that in cyberspace, where “information flows free”, *there would be a smaller role for brandnames*. Internet users enjoy quick access to a plethora of information, with powerful search engines help guide them towards material they actually read. One might think that ignorance should not be a problem on the internet, and that brand-names would play a limited role. In fact, the current business model of “.coms” is actually reliant strongly on building brand-names! Large portals, e-tailers, etc have precisely the same agenda of building a brand-name and dominating their market segments.

Rather than building a free information-flow “e-topia” where buyers find sellers easily, we are simply re-creating the bricks-and-mortar world of large brand-names in cyberspace. This state of affairs gives rise the type of scenario outlined in Section 1.

3 Barriers to e-commerce

There are several impediments to more open commerce on the cyberspace. What are the reasons why brand-names persist on the internet, and consumers gravitate toward them? We list them one by one.

Difficulty of managing information flow Low-overhead, open, commerce requires that matched buyers and sellers find each other without difficulty. Essentially, this is an information retrieval problem. On the current WWW, this matching is achieved by getting customers to visit the right web pages. With a few well-recognized brand-name vendors, the matching of buyers with sellers places modest demands on the WWW infra-structure. The matching is accomplished by a combination of web search, advertising on portals, and direct email appeals. Full P2P e-commerce requires matching between potentially millions of sellers and buyers. We consider two core, inter-related issues: *flow* and *scale*

Flow: On the WWW, information is published “passively” via web servers. To help users find this information, web search engines deploy crawlers that recursively explore all available links from all web pages, and “pull” the content into a repository. This repository is then indexed, and can be searched in response to user queries. For P2P e-commerce, this approach is unsuitable; an individual making an offer on her home page may not be “pulled” in a timely manner (Or not at all! there is some data that indicates that only 15% of web sites are indexed by even the leading search engines). A “push” approach is more likely to succeed.

Scale: With a few vendors and many sellers, the matching problem is quite manageable. With planetary-scale P2P, the infra-structure must potentially match-up millions or even hundreds of millions of offers from buyers and sellers.

Ontology Mismatch Current web-search engines are based on simple keyword searching, using technology based on algorithms invented in the field of information retrieval. Keywords alone are unlikely to support full P2P e-commerce. Pitfalls of keyword based searching include *synonymy*—the tendency of people to use different words to describe the same concept. Word usage also tends to be *polysemic*—the same word is used in different contexts to mean different things. These problems lead to irrelevant and sometimes bizarre information being retrieved by portals and search engines in response to keyword searches; for example, a link to a rubber clothing seller in response to a search for “latex”.

The core problem is *ontology matching*¹. Rather than matching keywords, (which are after all just syntax) we need to match semantics, or the meanings/descriptions of those words. In addition,

¹Ontology is the science of “what is”.

we need to provide people with the assistance they need in constructing proper descriptions of their buy and sell offerings.

Risk Management With planetary-scale P2P e-commerce, users will have good reason to worry about dealing with strangers. Will an unknown seller ship the promised product promptly, and in good condition? Will an unknown buyer claim to have never received a product? Will the price change unpredictably? Does the product conform to safety standards? etc. Risk-averse users, lacking better information, tend to prefer brand-names. The brand-name carries with it a reputation, and thus lower risk. However, some users may also have a risk-cost trade-off strategy; for example, a buyer might be willing to purchase a putative brand-name item from a street vendor if the cost was low enough to justify the risk that it was a fake.

For P2P e-commerce to succeed, we need methods of managing risk. Buyers and sellers must provide in their offerings means of reducing risk, or various approaches to trading off risk for other factors, such as cost. There may be opportunities for third-party businesses to offer risk-reduction services for a fee.

4 Towards a new model

Imagine a new world. Rosalita, in Rancho Cordova, California, is looking for a recipe book on sweets from the southern Indian province of Goa. She first needs to describe this precisely, using standard terminology and logical connectors. So she connects to a public ontology server. Using an easy-to-use GUI, she explores the ontology and constructs a standardized description of what she wants to buy. She deploys this description on her home computer, on which runs her P2P e-commerce server that hosts this description and the description of other items she wants to sell, or buy. She also has described the terms under which she is willing to purchase the item.

Somewhere around that time, Miriam N’Gube, in Bobetti, Ethiopia is selling just such a book, and has deployed her description of the book, and the terms of her sale. Within hours, a world-wide infrastructure for P2P e-commerce has matched Miriam and Rosalita, along with some risk-management, delivery and financial services that will (for a fee) close the deal and complete the fulfillment. Miriam and Rosalita are notified, accept the deal, and by the next day the book and the payment are on the way (in opposite directions).

We believe that several emerging, core technologies will enable the flowering of this type of P2P e-commerce. These technologies will lead to the next business revolution on the web. The resulting “creative disruption” will cause major changes for existng, brand-name e-tailers.

Event Distribution On the current web, an offering from an individual can only be discovered when the containing site is visited by a web crawler. This passive approach will not lead to timely discovery and matching of offerings. We believe that a more active approach is called for: just the publication of an offering will trigger a series of events that will lead to a match. We expect that P2P e-commerce will be based on an infra-structure that supports the publication of and subscription to event streams. While many such systems [6] do exist, they operate within small (LAN-type) networks; however, efforts are underway to scale them up to the internet [9].

Content-based Routing As the number of these types of “offering” events grows, the nature of the offerings also can get more complex: incorporating not just the description of the salable item, but also such aspects as risk management, delivery schedules, pricing policies etc.

Managing the forwarding of these very large number of P2P offering events, with complex information content is an important problem. Whereas the current internet infra-structure routes packets to their destination based on addresses, we need to route these events based on the

information they contain. *Content-based routing* is an approach [8] is a technique for routing events based on their semantic content. This approach needs to be developed further, both in terms of scaling up the volume of such routing techniques, and also handling intricate semantics in the content.

Ontology Matching and Authoring An underlying problem in the descriptions of offerings, and matching offerings, is the meanings of the terms used in the descriptions. Divergence and incoherence must be avoided; term usage must be based on agreed-upon semantics. In addition, the descriptions should be constructed using well-founded and computationally manageable logics (*e.g.*, description logics [7]). Well-designed, standardized, freely available public ontologies will provide the basis on which standardized offering descriptions can be authored. Standardized ontologies are already being developed in several specialized domains, such as health (The HL7 Group), manufacturing (STEP Express), *etc.* See www.ontology.org for other examples.

But who will build these general ontologies? We believe that once the revolutionary promise of P2P becomes well understood, there will be a groundswell of interest in constructing broad-based standard ontologies. We expect that the open-source movement will embrace this endeavor; voluntary teams will discuss, debate, and consensually [11] construct *ipso-facto* standard ontologies for different sundry sub-domains of P2P e-commerce, such as antiques, gardening tools, cook-books, and so on.

To facilitate the authoring of offering descriptions in these formal logical languages, we expect that there will be sophisticated authoring systems. These authoring systems will transparently contact free, public ontology servers, and support in the construction of offering descriptions. Research into knowledge base browsing and authoring tools [10] will serve this endeavor.

Risk Description and Management Before matched offerings from parties as widely separated in space and culture as Roasalita and Miriam can lead to a deal, there must be means of managing the perceived risk from both sides. Both parties need assurance that the deal will go through as desired.

Currently, there are companies, such as Escrow.com and i-escrow.com, which offer risk-reduction services. For a fee, Escrow.com will handle the details of the payment, delivery and fulfillment, including such deviances such as the buyer rejecting the deal. For P2P e-commerce to flourish, these and other risk reduction services need to be developed. Other bases for risk reduction might be certified histories, (such as those offered by Better Business Bureaus) liability insurance, guaranteed protection against fraud (such as offered by credit cards), anonymity/privacy mechanisms, e-cash *etc.* These services will be offered on-line protected by appropriate cryptographic mechanisms.

We have given short descriptions of various innovations that we believe will pave the way for full, free and open P2P e-commerce on a planetary scale. This is not meant to be a complete list; there are surely other issues that we have neglected, and still others may yet arise.

5 Related Work

There are many efforts that are attempting to solve related problems.

Current Web-the Massive middle server Currently, several auction sites allow customers to offer items for sale and conduct trade. Unless all potential traders list on a monolithic site, there is a risk of not finding trading partners; potential partners may list on a different site, and/or use an incompatible ontology. With a single business “owning” this market and becoming the “universal middleman”, however, there is tremendous scope of monopolistic abuse of consumers, leading to

escalation of costs and other inefficiencies. There may also be systems issues relating to reliability, scalability, security etc. Our approach proposes a finer-grained approach, to matching traders, based on automated content-based distribution of trading offers. The distribution would be handled by a loose federation of servers, where failure is well-tolerated.

Open Trading Standards There are several related efforts to create open standards for trading on the internet. The *open trading protocol* [3] addresses the actual completion of a transaction. Its scope includes the accompanying details, such as delivery, payment, financing etc. There are various cryptographic schemes for proof of delivery, proof of payment and such variations as multiple payments, return-for-refund, etc. It assumes, however, that buyers and sellers have identified each other already, and does not provide any means of initiating this.

The *CommerceNet* [4] initiative defines the *eCo* architecture for open electronic commerce in the B2B arena. *eCo* includes mechanisms for sellers and buyers to dynamically find a business partner, and then execute a transaction. The matching mechanisms work within the context of a *market* domain, which includes a registry mechanism. Services provided by participants can be queried dynamically for available functionality, much like some component object models (*e.g.*, COM) allow dynamic identification of supported interfaces. The market and registry mechanisms of *eCo* are ill-suited to P2P, however, where thousands or millions of potential trading partners may enter and leave different markets.

Ontology-based efforts There are several projects underway to construct open, free and public ontologies. While the specific goals of these projects differ, they all have a common need to integrate information from different sources, and hence the need for common terminology. Projects include the Open Dictionary Projects [2], which is a browseable, searchable directory site, that is ontology-based and open for editing and creation. The Standard Upper Ontology effort [1] which aims to provide standard definitions for general-purpose terms. There are several others, which we omit here for brevity.

There are also a number of projects which aim to exploit large ontologies for knowledge integration. For brevity, we just mention information broker project at KSL [5], which aims to integrate information from heterogeneous sources, using a common domain model, to answer users' queries. The domain model also uses a "query-by-reformulation" approach to help users formulate queries.

6 Conclusion

The efficient operation of free markets depends on the free flow of information between buyers and sellers. Prior to the internet revolution, this information flow was sluggish, and dominated by mass media; small-scale buyers and sellers often operated with incomplete information. This led to the creation of large centralized organizations that capitalized on brand-names. At the moment, we are just simply re-injecting the bricks-and-mortar world of monolithic, brand-name based enterprises into the internet. This is based on the traditional assumption that information flow will continue to be restricted, and that customers will continue to operate with a paucity of information, and preserve their reliance on brand-names.

The development of new technologies such as standardized ontologies, content-based routing, large-scale event distribution mechanisms, and risk management protocols will streamline information flow on the internet, and lead to more efficient markets. The emphasis on large brand-names and B2B e-commerce will be tempered with a new emphasis on person-to-person e-commerce. This will lead to much more fine-grained, democratic type of commerce that will no longer be focused large scale B2B or B2C, but also on individuals buying and selling to meet their needs on a smaller scales. Great

efficiencies will result, as end-users make more informed decisions, and the need for middlemen is reduced.

We believe that this revolution will largely be driven by publicly supported, or voluntary open-source based efforts, and will result in the “creative disruption” of many current internet enterprises.

References

- [1] IEEE Computer Society, Standards Activity Board, Standard Upper Ontology Study Group. <http://itsc.ieee.org/suo/>
- [2] The Open Dictionary Project <http://www.opendict.org>
- [3] Open Trading Protocol <http://www.otp.org>
- [4] The eCo Framework Project, Commerce Net project <http://www.commercenet>
- [5] Fikes, R, Englemore, R., Farquhar, A., Pratt, W., Network-based Information Brokers, Knowledge Systems Lab Technical Report KSL-95-13.
- [6] Balachander Krishnamurthy and David S. Rosenblum. "Yeast: A General Purpose Event-Action System," *IEEE Transactions on Software Engineering*, Volume 21, Number 10, Pages 845-857, October 1995.
- [7] Diego Calvanese, Giuseppe De Giacomo, and Maurizio Lenzerini Representing and Reasoning on XML Documents: A Description Logic Approach *Jnl. of Logic and Computation*, 1999
- [8] Antonio Carzaniga, David S. Rosenblum and Alexander L. Wolf, Content-Based Addressing and Routing: A General Model and its Application, Department of Computer Science, University of Colorado at Boulder, Technical Report CU-CS-902-00, Jan. 2000.
- [9] David S. Rosenblum and Alexander L. Wolf. "A Design Framework for Internet-Scale Event Observation and Notification," *Proceedings of the Sixth European Software Engineering Conference / ACM SIGSOFT Fifth Symposium on the Foundations of Software Engineering*, 1997.
- [10] L. G. Terveen and David A. Wroblewski, A Collaborative Interface for Browsing and Editing Large Knowledge Bases, in *Proceedings of AAAI'90*.
- [11] L. G. Terveen and David A. Wroblewski, A Tool for Achieving Consensus in Knowledge Editing, in *Proceedings of AAAI'91*