Report on Second International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems

Kun-Lung Wu and Philip S. Yu
IBM T.J. Watson Research Center
30 Saw Mill River Road
Hawthorne, NY 10532
{*klwu,psyu*}*@us.ibm.com*

1 Introduction

The Second International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS 2000) was held at the Crowne Plaza San Jose/Silicon Valley in Milpitas, California on June 8-9, 2000. The purpose of this workshop was to bring together leading practitioners, developers and researchers to explore the challenging technical issues and find feasible solutions for advancing the current state of the art in e-commerce and web-based information systems. In particular, the workshop was interested in the infrastructure issues to facilitate e-commerce and Web-based information systems.

WECWIS 2000 was successful. There were three invited talks, one industrial panel discussion and six technical sessions. The keynote speech, "The global trading web: A strategic vision for the Internet economy," was delivered by Dr. Jay M. Tenebaum, VP and Chief Scientist, Commerce One, Inc., on June 8 immediately following the opening remarks by the conference chair. The banquet address, "Business issues in e-commerce," was delivered by Mr. Daniel Druker, General Manager, Hyperion e-Business Division. Finally, a lunch address, "B2C, B2B, N2N, N2M: Why 2 is so instrumental?" was delivered by Mr. Mstafa A. Syed, VP of Technology, VertialNet, Inc.

The industrial panel was moderated by Dr. L. Mason and Dr. Z. Zhang, both of Blue Martini Software. The panelists included J. Becher, Accrue Software; L. Mellot, Business Objects; A. Srivastava, Blue Martini Software; and C. Zhou, IBM. The panel discussion topic was "Can e-business intelligence survive?" Among the many interesting issues being discussed were: Will privacy concerns stunt e-business intelligence utility? Will integrated e-commerce solutions be able to collect and analyze click steams, contents, products and sales data simultaneously? To what extent can out-of-the-box combined e-commerce and e-business intelligence solutions be useful? Is data mining useful in B2B e-commerce? Both positive and negative responses were hotly debated.

There were a total of 30 papers included in the technical presentations, organized into six sessions. They were selected after rigorous reviews by the program committee members. The presented papers cover a wide range of topics, from framework, architecture and protocol issues of e-commerce to various types of e-services to web-based information systems for facilitating e-commerce. The rest of this report provides a brief summary of the technical presentations given in the workshop. The entire workshop proceedings is available from the IEEE Computer Society.

2 B2B E-Commerce and Negotiations

There were five presentations in this session. The first paper, "An architecture for cross-organizational business processes," was written by Y. Hoffner, H. Ludwig, C. Gülcü, and P. Grefen. It described a *CrossFlow* architecture to support the dynamic establishment and enactment of a business relationship between two organizations, based on a contract that specifies this relationship. It was achieved by creating an electronic market where advertising and searching for compatible business partners takes place.

M. Iwaihara presented the second paper, "Supporting dynamic constraints for commerce negotiations." He introduced the notion of dynamic constraints to model complex conditions exchanged during business negotiations. Negotiation conditions were stored in a database and a special query language, called *dynamic constraint algebra*, was designed to access them. Algorithms for processing queries based on dynamic constraint algebra were presented.

The third paper, "An infrastructure for meta-auctions," was coauthored by C. Bornhövd, M. Cilia, C. Liebig and A. Buchmann. It developed the notion

of meta-auction that allows a buyer to roam automatically across auction sites on the Internet. It identified the communication and notification requirements for an Internet-scaled meta-auction system. The authors argued that Internet-scaled business applications, such as meta-auctions, (a) require publish/subscribe as an interaction paradigm, in addition to request/response; and (b) should leverage pro-active information dissemination and caching mechanisms to cope with the scale, distribution and availability issues. They also pointed out the need for metadata-based infrastructures providing common vocabularies for semantically meaningful exchange of data and notifications.

The fourth paper, "A bid evaluation system for internet contract negotiation," was jointly written by Q.-B. Nguyen, M. A. Cohen, and J.-Y. Chung. It presented a bid evaluation system for contract negotiation on the Internet. The system uses a hierarchical decision-criteria framework to handle complex constraint hierarchy with n-ary variables. The system provides a mechanism to analyze and forecast certain events based on seasonal and historical data.

M. R. Lee presented the final paper in this session, "Context-dependent semantic values for e-negotiation." The author presented the use of semantic value as the unit of exchange and a context-agent to facilitate the exchange during e-negotiation. An XML framework is used to declare the values of the properties of the semantic values. An ontology server specifies the concepts and terminology mapping. The explicit representation of semantic value reduces errors and frees applications from being concerned with conversions. It also makes possible to understand the impact of semantic changes to the data among different heterogeneous information systems.

3 Infrastructure Issues

The first paper, "A scalable approach for subscription-based information commerce," was coauthored by A. Celik and A. Datta. The authors considered an information environment where clients subscribe to information objects sent via a broadcast onto a shared network (e.g., wireless, LAN, the Internet). They proposed a scalable system, called *Drop Groups*, to allow a client to access its objects of interest only for its subscription period.

R. Behrens presented the second paper, "MONTANA: towards a Web-based infrastructure to improve lecture and research in a university environment." This paper reported on MONTANA, which is a digital media archive for teaching and research. It showed how to model structures and describe metadata by the use of XML. It combined RDF with Dublin Core to provide broad interoperability and data reuse. Several problems due to the use of XML as a data model were highlighted.

The third paper, "Supporting global replicated services by a routing-metric-aware DNS," was jointly written by W. Tang, F. Du, M. W. Mutka, L. M. Ni and A.-H. Esfahanian. The authors proposed DNS extensions to support global replicated services. The DNS server collects routing metrics from routers for cached replicated servers. Upon receiving a DNS request from a host, the DNS automatically returns the IP address of a server available from an IP address pool that has the best routing metrics preferred by the host. Such approach makes it easier for content providers to support global replicated services flexibly and transparently.

H. Yu and A. Vahdat coauthored the next paper, "Building replicated Internet services using TACT: a toolkit for tunable availability and consistency tradeoffs." It is well known that replication is an important approach to achieving scalable, high-performance and high-available systems. The authors argued that Internet services can benefit from dynamically choosing availability/consistency tradeoffs. Three consistency metrics were used to show how consistency can be meaningfully quantified for many Internet services. They presented the design of TACT toolkit that allows Internet services to flexibly and dynamically choose availability/consistency tradeoffs, enabling differentiated availability/consistency quality of services.

The final paper, "DWINS: A dynamically configurable Web-based information system," was coauthored by X. Dong and F. Du and L. M. Ni. This paper described a dynamically configurable Web-based information system for inventory management. The database structure can be adjusted any time, even after the system has been configured and starts running. All existing functions are applicable to all new database objects once they are created. The functions of the system are not fixed and key operations can easily be extended and configured.

4 E-Services, E-Advertising, and E-Markets

The first paper, "Surveying the E-services technical landscape," was presented by H. Kuno. He presented an e-service vision of Hewlett Packard for e-commerce in which a rich array of nimble modular e-services are accessible by virtually anyone and any device. This paper surveyed the current efforts to meet the technical challenges presented by e-services. It summarized

the efforts of major vendors, including HP, Microsoft, Sun, IBM, Lucent, and Hitachi. It compared and contrasted the functionality, characteristics, and limitations of those vendor's products.

The second paper, "Agent-mediated Internet advertising," was coauthored by D. Siew and X. Yi. The authors considered the issues of using software agents to automate several of the time consuming stages of Internet advertising. They surveyed the application trends of software agents in Internet advertising and proposed some solutions to these trends.

The third paper, "Advanced dynamic property evaluation for CORBA-based electronic markets," was jointly written by A. Schade, C. Facciorusso, S. Field and Y. Hoffner. Dynamic property updates are important in electronic markets and can be supported by the CORBA Trading Service. However, specifying, executing, and managing the algorithms used to compute the values of the dynamic properties are outside the scope of the related CORBA standard. The authors presented the concept and implementation of a generic engine which can be used to create, edit, manage, and execute dynamic property evaluation algorithms. The approach has many advantages, such as an improved development environment, a simulation test bed, improved management and maintenance capability and a structured way of linking the trading service to legacy systems.

The next paper, "Information kiosk system by cooperation between agents and experts using situation adaptive scenarios," was coauthored by Y. Taniguchi, M. Hirotoshi, H. Yajima, and N. Komoda. It described an information kiosk system where service providers can control telecommunication channels by the cooperation between agents and experts. The system was motivated by certain specific problems encountered by earlier kiosk systems installed in many of the convenience stores in Japan. The authors described a prototype implementation and some of the encountered problems.

The session's last paper, "An open, flexible and configurable system for service composition," was coauthored by F. Casati, S. Ilnicki, L.-J. Jin and M.-C. Shan. The authors argued that, in contrast to e-services delivered point-to-point, there are opportunities for providing value-added, integrated services that are delivered by composing existing e-services, possibly offered by different companies. The paper presented *eFlow*, a model and system that provides a flexible, configurable and open approach to service composition. Moreover, eFlow has an adaptive and dynamic process model that allows processes to transparently adapt to changes in the environment and to changes in needs.

5 XML-Related Issues

Two of the most important data handling technologies are XML documents and relational databases. While relational databases have been used to store a significant amount of today's business data, XML documents have rapidly gaining momentum in e-commerce and Internet-based information exchanges. Hence, an essential part of many XML-based applications is data exchange between relational databases and XML documents. R. Bourret, C. Bornhövd, and A. Buchmann presented the first paper, "A generic load/extract utility for data transfer between XML documents and relational databases." They identified and solved four specific problems: (1) loading data from XML documents into relational tables with a known schema; (2) creating XML documents according to a known DTD from data extracted from a database; (3) generating relational schemas from XML DTDs for on-the-fly storage of XML documents; and (4) generating XML DTDs from relational schemas for on-the-fly extraction of relational data.

The second paper, "DTD-Miner: a tool for mining DTD from XML documents," was coauthored by C.-H. Moh, E.-P. Lim and W.-K. Ng. Although XML documents can be accompanied by a DTD that defines the structure of the documents, the presence of a DTD is not mandatory. The difficulty in deriving the DTDs for XML Documents lies in the fact that DTDs are of different syntax from XML documents and that prior knowledge of the structure of the documents is required. With DTD-Miner, a user can submit a set of similarly structured XML documents and the system will automatically suggest a DTD.

D. Cheung, S. D. Lee, T. Lee W. Song. and C. J. Tan coauthored the next paper, "Distributed and scalable XML document processing architecture for e-commerce systems." They presented a generic architecture for processing XML. It is flexible, scalable, and suitable for distributed network environments. The application was aimed at supporting data exchange in e-commerce systems.

The last paper, "Multimedia presentation components in E-commerce," was jointly written by S. S. Y. Shim, J. Z. Gao and Y. Wang. It presented different ways to include multimedia presentations to Web-based catalog. Instead of using proprietary multimedia document types, the authors argued that Synchronized Multimedia Integration Language (SMIL) elements should be used because SMIL is a standard and it is based on XML. Hence, multimedia presentations with SMIL can be included in XML-based catalog and data exchange.

6 Architecture, Protocols, and Analyses

Five presentations were included in this session. The first one, "E-representative: a scalability scheme for e-commerce," was coauthored by W. Meira, Jr., D. Menascé, V. Almeida, and R. Fonseca. It presented the concept of *e-commerce representatives*, a means to scaling the performance of e-commerce services. E-representatives are programs that execute on a cache server or at nearby machines. E-representatives can be implemented using redirection. The authors presented performance analyses to show the potential performance gains by using e-representatives.

The second paper, "Failure analysis of an E-commerce protocol using model checking," was coauthored by Indrakshi Ray and Indrajit Ray. This paper dealt with the issues of ensuring certain desirable properties be satisfied for e-commerce protocols. In general, e-commerce protocols ensure, among others, confidentiality and integrity of information exchanged, money atomicity, goods atomicity, and validated receipt. The paper showed that model checking can be used to obtain assurance about the existence of those properties in an e-commerce protocol.

The next paper, "Two approaches for pay-per-use software construction," was jointly authored by L. C. Ferreira, R. Dahab, M. P. Aragão and J. A. P. Magalhães. It described two architectures for building pay-per-run software systems. These are systems that allow the user to pay for each execution of an application, instead of buying a more expensive user license. It is a new model for software distribution, where developers charge small fees for each execution in order to increase their user market. The authors detailed each component of the architectures and compared both approaches.

"Speeding up electronic commerce activities using CapBaseED-AMS," was coauthored by E. Kafeza and K. Karlapalem. CapBaseED-AMS is a capability-based and event-driven activity management system that supports specification and execution of activities using cooperative information systems paradigm. An activity consists of interdependent tasks that are executed by an agent. When there are many activity instances executing at the same time, many tasks are queued up at agents for execution. The authors proposed a multi-agent system model, introduced a measurement of speed-up/slow-down and developed different scheduling policies that allows for faster executions based on user's requests.

The session's final paper, "Scheduling algorithms for the broadcast delivery of multimedia products," was coauthored by V. Vinay and K. Ramamritham. The paper presented several heuristic algorithms that attempt to maximize the revenue earned by an e-commerce merchant providing broadcasting multimedia product services via the Vertical Blanking Interval (VBI) in TV broadcasting. VBI can be used to broadcast data and products like video clips, software, multimedia packages, etc., thereby satisfying several customers with a single transmission at high speed. The paper studied various algorithms tailored to deal with different customer service policies.

7 E-Commerce Applications

The last technical session contained six presentations. The first paper, "Development of inter-college e-commerce system through extending a legacy system," was coauthored by H. Oiso and N. Komoda. It was about an inter-college e-commerce system that enables colleges to provide extension courses to each other. The system streamlines the internal operations, such as registration, and automates data flow between colleges. The authors implemented the system by extending a legacy system into an end-to-end integrated e-commerce system using Internet, EDI, triggers, and intermediate DB in a cost-effective manner.

The second paper, "How e-commerce can benefit from visualization," was presented by M. Lux. The paper described a system, called *PROSECO*, for distributed visualization of abstract data. The author argued that, in order to make the full utilization of the Internet, a user-friendly access to information is needed for the optimal distribution of information and knowledge. User-friendliness can be achieved with the support of information visualization. The author demonstrated the concept by providing a user-friendly access to economic information.

The third paper, "Construction of online catalog topologies using decision trees," was coauthored by D. Yang, W. Sung, S.-M. Yiu, D. Cheung, W. Ho, T.-W. Lam and S.-D. Lee. The paper focused on the problem of organizing information in a Web site. It introduced a method to automatically construct a web site using decision-tree algorithms. The authors demonstrated the idea in automatic construction of product catalogs.

The fourth paper, "Aggregate path index for incremental web view maintenance," was coauthored by L. Chen and E. Rundensteiner. The authors argued that, as web data keeps growing in an explosive way, web queries need to be materialized to achieve fast query response time. However, compared with traditional databases, it is more challenging to develop efficient maintenance techniques for materialized views over dy-

namic web data. The paper presented an aggregation path index structure to hold a collection of "qualified" objects with respect to the query pattern. By using the pre-computed auxiliary information stored in such an index, the authors showed that web view maintenance can be done more efficiently.

The fifth paper, "An architecture to support distributed data mining services in e-commerce environments," was jointly written by S. Krishnaswamy, A. Zaslavsky and S. W. Loke. It presented a hybrid architectural model for distributed data mining (DDM), which is tailored to the needs of e-business where application service providers sell DDM services to e-commerce users and systems. The system integrates the client-server and the mobile agent technologies. It focuses on the issues of billing users for data mining services.

The last paper, "Dynamic 3D visualization of database-defined tree structures on the WWW by using rewriting systems," was coauthored by H. Noser and P. Stucki. It presented a new Web-based technique for visualizing dynamic large hierarchies, tree or category structures in 3D worlds. It uses CGI-scripting, SQL queries and Applet-based production systems to visualize in an efficient and versatile manner. The authors demonstrated their techniques by visualizing the complete shopping mall structure according to different visualization models.

8 Conclusion

The Internet has revolutionized the way commerce will be conducted in the future, whether it is business to consumer, consumer to consumer or business to business. As shown in the above summary, WECWIS 2000 provided a forum for many practitioners, developers and researchers to exchange their ideas. The topics covered a lot of the important issues in e-commerce and web-based information systems.

However, there are still many more issues need to be solved before e-commerce can work seamlessly and become truly pervasive. For example, the arrival of wireless Internet and broadband networks creates many issues, such as content adaptation to various devices, security, personalization, content hosting and infrastructure supports for wireless devices and multimedia contents.

Acknowledgment

WECWIS 2000 organizers would like to thank all those who took the time and effort to prepare papers and presentations and participated in the workshop.