

國立中山大學八十九學年度碩博士班招生考試試題

科目：資管所（論文評述第一節） 資訊管理類論文

共 頁 第 頁

5

1. Briefly discuss the strength and weakness of the research described in the paper, "An empirical study of adopters and non-adopters of the Internet in Singapore" in the following three categories: theoretical foundation, methodology, and implication. (25%)

5

2. Briefly discuss how each of the following theories can be applied to improve the research in the aforementioned paper. (25%)

1. The Cognitive Fit Theory (Vessey, 1991)
2. The Media Richness Theory (Daft and Lengel, 1984)
3. The Technology Acceptance Model (Davis 1989)

10

Daft, R.L. and Lengel, R.H. "Information Richness: A New Approach to Managerial Behavior and Organizational Design," in *Research in Organizational Behavior*, Staw, B. and Cummings, L. (Eds.), 1984, pp. 191-233.

10

Davis, F. D. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly*, Vol. 13, No. 3, September 1989, pp.319-340

15

Vessey, I. "Cognitive Fit: A Theory-based Analysis of the Graphs Versus Tables Literature," *Decision Sciences*, Vol. 22, 1991, pp. 219-240.

15



Research

An empirical study of adoptors and non-adopters of the Internet in Singapore

Thompson S.H. Teo^{*}, Margaret Tan¹

*Department of Decision Sciences, Faculty of Business Administration, National University of Singapore,
10 Kent Ridge Crescent, Singapore 119260, Singapore*

Received 21 November 1996; accepted 4 August 1998

Abstract

The Internet has been given a great deal of attention around the world and the number of host computers and users on the Internet has been increasing approximately at an exponential rate. However, despite this rapid growth, there are relatively few empirical studies on the Internet. This paper reports on a questionnaire survey of both adopters and non-adopters of the Internet among organizations in Singapore. The organizational characteristics, benefits of adopting the Internet, reasons for not adopting the Internet, criteria for selecting Internet access service providers, etc. were examined, thereby contributing to a better understanding of the Internet phenomenon in Singapore. © 1998 Elsevier Science B.V. All rights reserved

Keywords: Internet; Adoption; Benefits; Internet service providers; Singapore

1. Introduction

The Internet had its origins in the early 1960s, when the US Department of Defense realized the need for a decentralized computer network that would provide the Pentagon with a command and control communications system in the event of contingencies such as nuclear war [1]. This early network, known as the Advanced Research Projects Agency network (ARPANET), gradually appeared in university and government research laboratories and eventually became the

technological underpinnings of the Internet [11, 21]. Programs were written to enable people to exchange electronic mail, tap into remote databases, run supercomputers at a distance, and brainstorm via electronic bulletin boards.

The most powerful innovation was the communications protocol that gave the Internet its name – the Internet Protocol. It allows any number of computer networks to link up and act as one. The full set of ARPANET protocols, known as TCP/IP (Transmission Control Protocol/Internet Protocol), were adopted in the academic and research communities and then into commercial computing. By the late 1980s, millions of computers and thousands of networks were using TCP/IP and, from their interconnections, the modern Internet emerged.

^{*}Corresponding author. Tel.: +65-874-3036; e-mail: teosh@nus.edu.sg

¹Tel.: +65-874-3157; E-mail: fbatann@nus.edu.sg

The Internet will have an expected 200 million users worldwide by the end of this century; a substantial jump from the estimated 56 million users in 1995 and just 2 million users in 1990 [4]. With such a large base of potential users, it is no wonder that an increasing number of organizations worldwide have tapped its potential as a medium to boost their business. Furthermore, virtual markets on the Internet can provide new areas of opportunities for retailers, producers and consumers [22]

1.1. The Internet in Singapore

The Internet was first introduced to the academic community of the National University of Singapore in late 1980. In 1992, after setting up Technet, the National Computer Board, the National University of Singapore, and the Ministry of Education started to explore the possibility of exposing schools to the Internet. The following year, it was introduced into secondary schools and junior colleges.

In line with Singapore's IT2000 plan [14, 15], the first commercial Internet access service provider (IASP), SingNet, was launched in July 1994. It provided the gateway for businesses and the general public to easily access information and electronic resources available on the Internet. This enabled Singapore to move a step closer to achieving its vision of becoming an intelligent island in the twenty-first century.

In September 1995, Technet was privatized and renamed Pacific Internet. As the second IASP, it initiated an aggressive marketing campaign to attract subscribers. As a result of intense competition, both IASPs reduced prices. With subscribers enjoying lower rates, the number of subscribers grew. The launching of the third service provider, Cyberway, on 22nd March 1996, resulted in further acceleration of the growth of the Internet.

Today, there are an estimated 150 000 Internet users in Singapore, including individuals and organizations, a substantial growth from the 52 000 users in early 1995. It has been estimated that the total number of Internet users in Singapore will reach nearly one million by the end of this decade [13, 23]. With the popularity and almost exponential growth, it is surprising that there are relatively few published empirical studies on the Internet.

2. Research method

A questionnaire survey was used to collect data. The instrument was derived from an extensive literature review of published materials in academic and practitioner journals and periodicals pertaining to the Internet and innovation adoption. The survey collected data on (1) the organizational features of the firm; (2) ranking of Internet applications; (3) championship for Internet adoption; (4) objectives of Web sites; (5) benefits of adopting the Internet; (6) reasons for not adopting the Internet; (7) criteria for selecting Internet access service providers; and (8) benefits of intranet.

2.1. Data collection procedures

The development of the questionnaire involved a series of pretests with students and faculty members over a period of about two months. Particular attention was focused on the wording, structure, sequence and overall presentation. It was then pilot-tested with managing directors of two local companies. Based on feedback, amendments were made to improve the clarity of wordings so as to further increase readability and comprehension. The final questionnaire was reviewed by two faculty members before being sent to a random sample of 500 companies in Singapore. A cover letter addressed to the senior executive in the company explained the objective of the study and a stamped return envelope was enclosed for the response. A second mailing was sent to non-respondents about three weeks later. Out of the 500 companies, 30 declined participation; 195 questionnaires were received. However, only 188 responses were usable. Hence the usable response rate was 37.6 percent, which is adequate for the purpose of the study.

3. Results

Out of the 188 firms, 131 (69.7 percent) have an Internet account and are regarded as adopters. Out of the adopters, 55 (42.0 percent) have Web sites. Among the non-adopters, 29 (50.9 percent) intend to adopt the Internet within the next six months.

Table 1 shows the respondents' profile and the results of chi-square tests between organizational

Table 1
Demogra
bb title
Chairman
Managing
General M
JO/IT M
Finance/I
Marketing
Others
Industry
Architect
Education
Finance/B
Insurance
Manufactu
Pharmaceut
Printing/P
Services
Transporta
Travel/Tou
Others
Organizati
Governme
Local comp
Local comp
Subsidiary
Others
Number of
100
11-300
31-600
61-1000
101-2000
2000
Annual reve
10
1-100
1-300
1-600
1-1000
1000
Mode of info
Additional
Involving rol
Integral role
Mode of pro
Production
Health
Security
Crime
In son

Table 1
Demographic profile

data.	Adopters	Non-adopters	Chi-square
<i>Job title of respondents</i>			
Chairman/CEO			
Managing Director	3	0	df=6
General Manager	9	6	chi-square=9.49
IT Manager	15	9	p=0.148
Finance/HR Manager	66	18	
Marketing/Operations Manager	10	9	
Others	8	2	
	15	9	
<i>Industry</i>			
Architecture/Engineering			
Education	13	4	df=10
Finance/Banking	1	0	chi-square=7.84
Insurance	8	2	p=0.644
Manufacturing	2	0	
Pharmaceutical/Biotechnology	59	29	
Printing/Publishing	1	3	
Services	5	1	
Transportation	19	10	
Travel/Tourism/Hotel	8	4	
Others	11	3	
	4	1	
<i>Organizational type</i>			
Government-owned company			
Local company with local ownership	12	1	df=4
Local company with major foreign ownership (joint venture)	54	21	chi-square=7.87
Subsidiary of foreign company	16	4	p=0.096
Others	39	27	
	10	4	
<i>Number of employees</i>			
100			
1-300	2	8	df=5
301-600	21	14	chi-square=23.68
601-1000	29	17	p=0.00025
1001-2000	26	9	
2000	14	5	
	39	4	
<i>Annual revenue (\$ Million)</i>			
30			
<100	4	7	df=5
1-300	40	33	chi-square=25.31
301-600	28	6	p=0.00012
601-1000	18	7	
>1000	13	1	
Others	23	2	
<i>Use of information technology</i>			
Additional role			
Supporting role	32	26	df=2
Integral role	48	17	chi-square=11.08
	51	11	p=0.004
<i>Stage of product life cycle</i>			
Production			
Launch	3	1	df=3
Stability	54	17	chi-square=4.03
Decline	67	34	p=0.258
	2	3	

* In some categories, the total (adopters+non-adopters) is <188 due to missing responses.

variables such as industry sector, firm size, and adoption behavior.

3.1. Hierarchical level of respondents

More than 80 percent of the respondents hold management positions in their companies. There is no significant relationship between respondents' hierarchical level and Internet adoption. The high hierarchical levels of adopters and non-adopters of the Internet provide some assurances as to the validity of responses, since these respondents can generally be expected to be more knowledgeable about their organizations than respondents from low hierarchical levels. Furthermore, respondents have an average working experience of seven years in their companies and 11.5 years in their respective industries.

3.2. Industry sector

No significant relationship was found between industry sector and adoption of the Internet. Apparently, regardless of industry sector, information technology (IT) has proliferated throughout the economy, as companies seek new ways to rationalize their operations and compete more effectively.

3.3. Organizational type

No significant relationship was found between organization type (e.g. government, local companies, or foreign firms) and adoption of the Internet. Interestingly, out of 13 government-owned companies, 12 are adopters, thereby reflecting the leading role that the Singapore government plays in the adoption of the Internet.

3.4. Firm size

Firm size is measured by two components: number of employees [8, 18] and annual revenue [6, 9]. More than 75 percent of the companies are large in terms of size, with more than 300 employees and an average annual revenue ranging from \$10 million to more than \$1 billion. The results showed significant relationship between Internet adoption and these two measures of firm size. A possible explanation is that larger firms, having more resources, may have a

greater need to stay at the forefront of technology than smaller firms.

3.5. Role of information technology (IT)

Three types of roles (traditional, evolving and integral) of IT were examined in this study. The results showed significant relationship between the role of IT and Internet adoption. Generally, the role of IT for adopters of the Internet is mainly an evolving or integral role. Conversely, for non-adopters, IT appears to play a traditional role.

3.6. Phase of product life cycle

We also investigated the product life-cycle phase (introduction, growth, maturity, decline) of the major products/services of both adopters and non-adopters of the Internet. The chi-square test indicates that there is no significant relationship between the stages of the product life cycle and adoption of the Internet.

4. Internet applications

Six Internet applications were identified and the adopters were asked to rank them from 1 (most frequently used) to 6 (least frequently used) in terms of usage. The results of the ranking are shown in Table 2 and the applications are as follows: World Wide Web (WWW), electronic mail (e-mail), usenet news, chat, file transfer protocol (FTP), and telnet.

It is not surprising to note that electronic mail (e-mail) is the most widely used internet application. This is consistent with Cockburn and Wilson's [5]

Table 2
Rank of Internet applications

Rank	E-mail	WWW	FTP	Usenet	Telnet	Chat
1st	64	55	2	0	0	0
2nd	43	46	12	10	3	1
3rd	3	7	55	15	10	4
4th	3	2	18	16	19	6
5th	1	2	2	17	12	5
6th	0	0	0	1	12	4
Total	114	112	89	59	56	41

Note: The total for each Internet application is different because some respondents did not rank all the six applications.

technology work which reported similar results. In contrast, use of e-mail, news, chat, and telnet are used less frequently since they are mainly used for educational and recreational purposes.

5. **Championship**

In the adoption of many IS technologies, a factor that has consistently been found to be important is the existence of champions [3, 7]. They provide the necessary drive and effort to initiate the adoption of new technologies. Our findings are consistent with those of past IS research. Out of the 131 adopters of the Internet, 83 (63.4 percent) have a champion for the Internet adoption process. Among these champions, 40 (48.2 percent) are from top management, 36 (43.4 percent) are from middle management and only seven (8.4 percent) are from lower management. The majority (46.1 percent) of the champions come from the IT/IS departments and ≈28 percent of the champions come from corporate headquarters or administrative departments.

Respondents were also asked to rate their champions in terms of degree of knowledge about business and IT on a five-point Likert scale ranging from 1 (lowest) to 5 (high). In terms of knowledge about business, the mean score was 4.29; while in terms of knowledge about IT, the mean score was 3.95. Champions are generally rated as possessing relatively high levels of business knowledge as well as having the necessary IT knowledge. Thus, champions will be in a better position to decide how IT can be leveraged to support the business, create business value for customers, and enhance competitive advantage.

their products and services on the Internet. The Internet represents an alternative form of advertising and is comparatively cheaper than the traditional means [16]. The Internet is certainly convenient for announcement and promotion of new products and services to customers worldwide.

3. *For direct selling/marketing of products and services:* Only 16 (29.1 percent) use their Web sites for direct selling/marketing of their products and services. This finding is generally consistent with King [12] who found that most Web sites of Fortune 500 firms are used for providing information rather than changing the way of doing business through online commerce. One likely reason is that the security of commercial transactions on the Internet is still being debated [10]. Furthermore, existing electronic shopping systems provide a vendor/customer dialectic that offers low product differentiation and product comparability [2].

Other objectives of Web sites mentioned by respondents include providing corporate presence on the Internet, having hands-on experience for future application of the Internet, reservation systems for taxis, collecting information about visitors, providing search engines and a vehicle for customers to communicate with the firm 24 hours a day, seven days a week.

7. **Benefits of adopting the Internet**

Adopters of the Internet were asked to evaluate its benefits on a five-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree). The results, which are shown in Table 3, are as expected and consistent with previous research [19].

Table 3
Benefits of adopting the Internet

Benefits	Mean (SD)
Convenient access to worldwide information	4.50 (0.59)
Creation of worldwide electronic presence	4.13 (0.70)
Extend global market reach	3.96 (0.71)
New business opportunities	3.85 (0.73)
Improve customer service	3.79 (0.72)
Selling of products through the Internet	3.66 (0.79)
Lower operational costs using electronic commerce	3.58 (0.81)

Objectives of Internet Web sites

The three main objectives for having Web sites are as follows:

To provide information to customers: Out of the 55 adopters with Web sites, 51 (92.7 percent) use their Web sites to inform customers of their products, as well as provide company and other relevant information pertaining to the products.

To advertise a firm's products and services: Thirty-four (61.8 percent) design Web sites to advertise

Table 4
Reasons for not adopting the Internet

Reasons	Mean (SD)
Staff will waste time surfing the Internet	3.20 (0.91)
Do not have expertise	2.79 (1.04)
Intranet irrelevant to business	2.71 (1.04)
Costs too much to adopt the Internet	2.71 (0.85)
Organization is too small	2.45 (0.92)
Staff are computer illiterate	2.31 (0.88)

8. Reasons for not adopting the Internet

Non-adopters were asked to evaluate the reasons for not adopting the Internet on a five-point scale ranging from 1 (strongly disagree) to 5 (strongly agree). The results are shown in Table 4.

One main concern of non-adopters is that their staff may waste too much time surfing the Internet instead of performing their normal duties.

9. Criteria for selecting Internet service providers

Adopters were asked to evaluate their criteria for selecting Internet access service providers (IASPs) on a five-point Likert scale ranging from 1 (not important) to 5 (very important). The results are shown in Table 5.

As expected, access speed appears to be the primary concern. However, it is quite surprising that special promotions and recommendations from family/friends are rated as least important.

Table 5
Criteria for selecting Internet service providers

Criteria	Mean (SD)
Access speed of the Internet	4.66 (0.52)
Technical support services	4.48 (0.65)
Reputation of service provider	4.27 (0.69)
Usage package that suits needs	4.09 (0.76)
Costs	4.06 (0.80)
Easy to apply	3.77 (0.88)
Special promotion	3.19 (0.90)
Recommendation from family/friends	2.79 (1.01)

10. Intranet

This study also obtained information about the usefulness of intranet as perceived by organizations in Singapore. Essentially, intranet is a descriptive term being used for the implementation of Internet technologies (e.g. NETSCAPE browser) within a corporate organization, rather than for external connection to the global Internet. Intranet provides access and sharing of information among various departments within the organization [20].

Out of 131 adopters of the Internet, 28 (21.4 per cent) have intranet in their organizations. The reason for such a small percentage of intranet users could be that intranet is a technology which is still relatively new, and firms may prefer to adopt a 'wait-and-see' attitude. Quelch and Klein [17] proposed that creating Internet networks (intranet) to facilitate electronic communications and transactions among employees, suppliers, distributors, etc. may be the Internet's principal value to multi-national corporations.

Table 6 shows how adopters of intranet evaluated the benefits derived or that can be derived from using intranet, on a five-point scale ranging from 1 (strongly disagree) to 5 (strongly agree).

Generally, all respondents perceive the facilitating of communication within the organization and providing convenient access to company information for employees and customers as the most obvious benefits.

11. Concluding comments

Despite the rapid growth of the Internet in Singapore, about one-third of the firms sampled do not have an Internet account. Of the remaining two-thirds, only ca. 40 percent have WWW homepages for their firms.

Table 6
Evaluation of Intranet benefits

Item	Mean (SD)
Facilitate communication	4.43 (0.61)
Easy access to company information	4.32 (0.62)
Eliminate or reduce paperwork	3.86 (1.01)
Enhance coordination of organizational tasks	3.82 (0.88)
Enhance user/employee productivity	3.82 (0.91)

For potential adopters of the Internet, the findings are instructive as to the common objectives of Web sites, potential benefits of Internet adoption as well as the need for a champion to promote the adoption of the Internet. For IASPs and policy makers, this study provides insights as to why some firms do not adopt the Internet. The study also examines the criteria for selecting IASPs. Policy makers can perhaps use similar criteria to evaluate current and potential IASPs as to their suitability as providers of Internet services. Existing IASPs should continue to provide competitive price, high access reliability, excellent customer service and technical support if they want to retain existing customers as well as to attract new customers.

Acknowledgements

The authors wish to thank Wong Kok Buk for his assistance in data collection.

References

- [1] P. Baran, On distributed communications networks. Rand Corporation, 1962.
- [2] J.B. Baty II, R.M. Lee, InterShop: Enhancing the vendor/customer dialectic in electronic shopping, *Journal of Management Information Systems* 11(4), 1995, pp. 9-31.
- [3] C.M. Beath, Supporting the information technology champion, *MIS Quarterly* 15(3), 1991, pp. 355-374.
- [4] C. Bourmellis, Internet '95, *Internet World*, 1995, 47-52.
- [5] C. Cockburn, T.D. Wilson, Business use of the World Wide Web, *International Journal of Information Management* 16(2), 1996, pp. 83-102.
- [6] W.J. Doll, G. Torkzadeh, The relationship of MIS steering committees to size of firm and formalization of MIS planning, *Communications of the ACM* 30(11), 1987, pp. 972-977.
- [7] V. Grover, An empirically derived model for the adoption of customer-based interorganizational systems, *Decision Sciences* 24(3), 1993, pp. 603-640.
- [8] T. Guimaraes, Y.P. Gupta, Measuring top management satisfaction with the MIS department, *OMEGA* 16(1), 1988, pp. 17-24.
- [9] S.E. Harris, J.L. Katz, Firm size and the information technology investment intensity of life insurers, *MIS Quarterly* (1991) 333-352.
- [10] D.L. Hoffman, T.P. Novak, Commercializing the information superhighway: Are we in for a smooth ride? *Owen Manager* 15(2), 1994 pp. 2-7.
- [11] W.J. Kettinger, National infrastructure diffusion and the U.S.

- information superhighway, *Information and Management* 27, 1994, pp. 357-368.
- [12] D. King, Fortune 500 on the web: The road to second-level effects, *Proceedings of the 29th Annual International Conference on Systems Sciences*, 1996, pp. 463-470.
- [13] T. Lee, Internet pioneer here calls for laws to protect online shoppers, *The Straits Times*, January 18, 1997.
- [14] T.Y. Liang, Organized and strategic utilization of information technology, *Information and Management* 24, 1993, pp. 329-337.
- [15] National Computer Board. The IT2000 Report. NCB, Singapore, 1992.
- [16] M. Perton, Profiting from the net, *Asia, Inc.* (1995) 50-57.
- [17] J.A. Quelch, L.R. Klein, The Internet and international marketing, *Sloan Management Review* (1996) 60-75.
- [18] L. Raymond, Organizational context and IS success, *Journal of Management Information Systems* 6(4), 1990, pp. 5-20.
- [19] C. Soh, Q.Y. Mah, F.J. Gan, D. Chew, E. Reid, The use of the Internet for business: The experience of early adopters in Singapore, *Internet Research* 7(3), 1997, pp. 217-228.
- [20] A.L. Sprout, The Internet inside your company, *Fortune* (1995) 91-94.
- [21] The Economist., The Internet: The accidental superhighway (1995) 5-20.
- [22] H. Wang, A conceptual model for virtual markets, *Information and Management* 32, 1997, pp. 147-161.
- [23] J. Yap, 1995: The year Internet became household word, *The Straits Times* (1996).

Thompson S.H. Teo is a Senior Lecturer in the Department of Decision Sciences at the National University of Singapore. He received his Ph.D. in Management Information Systems from the University of Pittsburgh. His current research interests include IS planning, teleworking, IT personnel, IT adoption and diffusion, and electronic commerce. He has published more than 30 papers in international journals such as *Behavior and Information Technology*, *Communications of the ACM*, *Decision Sciences*, *European Journal of Information Systems*, *Information and Management*, *Information Technology and People*, *International Journal of Information Management*, *Journal of Economic Psychology*, *Journal of Management Information Systems*, and *Omega International Journal of Management Sciences*.

Margaret Tan is an Associate Professor in the Department of Decision Sciences and the Graduate School of Business at the National University of Singapore. She received her Ph.D. in Management information systems from the University of Queensland, Australia. Her research interests include information technology productivity and governance, national information infrastructure in Asia Pacific, and the diffusion and adoption of electronic commerce. Currently, she is the Chair of the Executive Committee of Pacific Asia Conference for Information Systems (PACIS) and the Chair of the Asian Assembly for the Decision Analysis Society (DAS) of the Institute for Operations Research and the Management Sciences (INFORMS). She is also the council member of the Association for Information Systems (AIS) for the Asia Pacific and an associate of the ORBICOM Network of UNESCO Chairs in Communications.

國立中山大學八十九學年度碩博士班招生考試試題

科目：資管所（論文評述第二節）資訊管理類論文

共 頁 第 頁

下面問題主要是根據“Competing on the Internet: The Case of Amazon.com”這篇個案文章，你的答案應以中文為主，除了引用該文內容及特殊專有名詞。答案請儘量簡單扼要。總分五十分。

- 一、有關個案研究的文章，通常是尋找與主題相關的文獻，建立理論基礎，指導個案研究的進行與文章的寫作。你認為本篇文章的主題為何？並說明相關文獻與理論基礎。(15分)
- 二、個案研究的問題與當時研究的社會及歷史背景有關，而個案研究的主要貢獻在於產生理論或是提出某些意涵(implication)。你認為本文的主要貢獻為何？(15分)
- 三、個案研究所得結論，經由邏輯分析與辯證，可應用於類似情境。本文結論與網路書店的經營模式有關。你認為若在台灣經營網路書店是否可利用本文之研究成果？若是在台灣要做電子商務之個案研究(請自選一主題)，應要如何進行？(20分)

Competing on the Internet: The Case of Amazon.com

SURESH KOTHA, *University of Washington*

Rapid technological change and the growth of the Internet have enabled firms to rewrite the rules of the competition in many sectors of the US economy. Although many academics recognize the importance of the Internet, research highlighting successful business models that some firms have developed to compete via the Internet have been slow in coming. Using an in-depth case study on Amazon, this paper highlights how this exemplar firm is exploiting this emerging technology-driven media to rewrite the rules of competition in the book retailing industry. © 1998 Elsevier Science Ltd. All rights reserved

Rapid technological change and the growth of the Internet, and the subsequent emergence of the World Wide Web (WWW) have enabled firms to rewrite the rules of the competition in many sectors of the US economy. Although strategy researchers recognize the growing importance and the significance of the WWW, scholarly research highlighting the business models that firms are using to compete successfully, and detailed discussions of the *modus operandi* used by 'pure' Web-based firms have been slow in coming.

Also, losses by major corporations are so widespread on the Internet that Don Logan, CEO of Time Warner, declared publicly that his firm's Web site, 'Pathfinder, gave a new definition to the term black hole.' It appears that the only firms currently making money on the Internet are those selling services and products to enable others to develop their Web sites (Armstrong and Hagel, 1996; Tapscott, 1996). However, there are a successful few (e.g., Amazon, Online Auction, Yahoo!) who have managed to buck the trend and create successful business models that other firms are desperately trying to emulate. *How have these firms managed to find a recipe for success when others have failed to find a business model that works?*

This paper is an attempt to address this question. I do so by examining the dynamics of implementing

a Web-based strategy using a detailed case study of Amazon.com, a firm that is solely dependent upon the WWW for its revenues and survival. Amazon provides a singular case in which the frequently hyped WWW is actually changing how consumers buy products and services. Not content to just transplant the physical book retailing format to the WWW, the firm's founder, Jeff Bezos, is attempting to transform the process of buying books through technology that tap the interactive capability of the Internet. Over the past two years, he has built a fast-growing business using his firm's web site (<http://www.amazon.com>) (see Figure 1). The firm's Website has become an underground sensation for thousands of book lovers around the world who spend hours perusing its vast electronic library, reading other customers' amusing on-line reviews and, more importantly, ordering books.




Figure 1 Amazon's Rapid Growth in Revenues. Source: Company Records.

Using an in-depth case study on Amazon, I highlight how this firm is exploiting the electronic medium to rewrite the rules of competition in the global book retailing industry. I highlight the dynamics of pursuing a Web-based strategy with a focus on operational and competitive aspects. The Amazon case explores how a new entrant to the book retailing industry is creating a prototype of how book retailing, and on-line retailing in general, might evolve. In that sense, it is about how one firm (*in casu* Amazon) is 'creating the future' (Hamel and Prahalad, 1994) in an industry dominated by well established firms, such as Barnes and Noble, and Borders, Books and Music.

Understanding the World Wide Web

It is now estimated that about 50 million people in 140 countries are connected to the Internet through 16 million servers. Also, the typical Internet user is young, affluent, and well educated. From a commercial perspective, the demographics of Internet and World Wide Web users make them part of an extremely attractive market segment. The average age of computer users is 39, while the average age of a typical Internet user is 32. About one in 10 Internet users (more than 3 million) is under 18 and uses the Internet from home or school. About 64% of Internet users have at least a college degree with a median household income of \$60,000 (Internet World, 1995). Also, every day approximately 150 new businesses come onto the Internet, and their total number is estimated to be over 40,000 and growing. A majority of these companies use the Internet as a public relations tool to promote their products and services (The Economist, 1997a). However, a significant number of on-line sites are also becoming electronic merchants.

Although much of the scholarly work on the Internet is in its infancy, there is an emerging body of practitioner-oriented literature that examines the many facets of the Internet. For instance, many authors have chronicled the rapid evolution and growth of the Internet (Gilder, 1996; Weintraut, 1997; Tapscott, 1996; Negroponte, 1995). Still others have focused on the people who were responsible for its early evolution and its continued growth (Reid, 1997; Wallace, 1997). Additionally, many articles in the popular-business press have highlighted and discussed the Internet's growing universality, and its impact on society in general, and commerce in particular (The Economist, 1997a; Tapscott, 1996).

It is oft-noted that the Internet signals a fundamental shift in the nature of competition in certain, if not all, industries (Armstrong and Hagel, 1996; Evans and Wurster, 1997). Notes Reid:

The Internet's speed, scope, and scale — plus its uncanny intrinsic nature to warp or obviate experience-built knowledge, principles, and concepts — is unlike anything

mankind has experienced in this century, if not in all time.' (Reid, 1997, p. xiii)

Statements such as the ones by Reid (1997, p. xiii) are not uncommon (see also Tapscott, 1996). To many observers, the Internet presents both an opportunity and threat for commerce.

On the one hand, it is perceived as a threat because it enables people and businesses to connect *directly*, thereby, sidestepping intermediaries such as distributors and retailers (Lohr, 1997). Also, since the Internet fundamentally alters the 'economics of information' (Evans and Wurster, 1997), and the way information is communicated, it threatens incumbents in many industries such as the newspaper, travel, real estate, and book retailing industries, to name a few. In the past, many firms in these industries exploited the information asymmetries between buyers and sellers to make a profit. The growing popularity of the Internet as an universal communications medium is reducing information asymmetries between buyers and sellers, thereby increasing market efficiency (Kambil, 1997). More specifically (Evans and Wurster, 1997, p. 74):

The rapid emergence of universal technical standards for communication, allowing everybody to communicate with everybody else at essentially zero cost, is a sea change.... Those emerging open standards and the explosion in the number of people and organizations connected by networks are freeing information from the channels that have been required to exchange it, making those channels unnecessary or uneconomical.

In other words, falling 'transaction costs' (Williamson, 1979), and the ability to communicate the required information via the Internet, are both eroding the profit margins and the competitive advantages of certain businesses (Kambil, 1997).

On the other hand, the Internet provides new entrants, in certain industries, with opportunities to fundamentally alter the dynamics of competition. Using a two-dimensional framework *The Economist* suggests that certain businesses are better suited (than others) for electronic commerce (see Figure 2). Since shopping for a mortgage can be difficult and tedious, a Website offering straightforward and easy-to-understand comparisons can prove to be successful. In contrast, because buying a CD is a relatively straightforward and easy task, on-line merchants who sell CDs need to offer far more than a physical music store to draw shoppers to their on-line stores. Based on this framework, the authors of the framework speculate that 'the more tiresome a purchase is in the physical world, the more likely consumers are to try an on-line alternative.'

Although this framework is helpful to entrepreneurs (and existing businesses) contemplating entering the Internet, it does little to explain how existing electronic merchants should successfully exploit the technological capabilities of the Internet.




Figure 2 A Framework of Competing On-line. Source: Adapted from *The Economist*, 1997a.

Hagel and Armstrong (1997) address this issue by enumerating ways by which on-line businesses can leverage the emerging medium's capabilities to develop and sustain sources of competitive advantage. They focus on the unique capabilities of the electronic medium to form communities. Noting that the notion of community has always been at the heart of the Internet, they argue that many customers join one or more on-line communities (e.g., AOL, MSN) because these communities serve the need for communication, information and entertainment. Specifically, these authors highlight four needs that on-line merchants should recognize in order to be successful. They include the needs for: transaction, interest, fantasy, and relationship.

Communities of transaction emphasize transactions (i.e., buying and selling of products and services) and provide the information related to carrying out those transactions. In such communities, the participants 'are encouraged to interact with one another in order to engage in a specific transaction that can be informed by the input of other members of the community' (Armstrong and Hagel, 1996). *Communities of interest* bring together participants who interact more extensively with one another on specific topics. These communities, relative to communities of transaction, emphasize a higher degree of interpersonal communications. *Communities of fantasy* focus on groups of people that like to indulge in fantasy and make-believe. These are participants who like to create new environments, personalities, or stories. Finally, *communities of relationship* attempt to bring together participants who have had certain life experiences that often are intense and are seeking others with similar

experiences to share their experiences. Communities of relationship enable the formation of deep personal connections because the participants are often aware of one another's actual identities and experiences.

According to Hagel and Armstrong (1997) these communities are not mutually exclusive. Moreover, many commercial Websites rarely exploit the community-building capabilities of the electronic medium. They simply advertise their products on the WWW and hope that somebody will buy something. The main purpose of on-line communities is to develop a critical mass of participants by being the first mover and, thereby, making it difficult for other new entrants to draw customers away from existing communities.

More specifically, Armstrong and Hagel (1996, p. 135) posit that:

'By adapting to the culture of the Internet, however, and providing customers with the ability to interact with one another in addition to the company, businesses can build new and deeper relationships with customers. We believe that commercial success in the on-line arena will belong to those businesses that organize electronic communities to meet multiple social and commercial needs. By creating strong on-line communities, businesses will be able to build customer loyalty to a degree that today's marketers can only dream of and, in turn, generate strong economic returns.'

Internet usage was growing at 2300 per cent a year: Anything that's growing that fast is going to be ubiquitous very fast

In other words, the normative proposition that emerges from this discussion is that firms that are able to develop on-line communities effectively are more likely to succeed relative to those that simply advertise their products on the Internet.

Although the rationale for developing communities is clear, the mechanisms used by on-line firms to develop and sustain such communities needs elaboration. Also, the normative proposition that firms that establish on-line communities outperform others (firms without communities) needs to be examined more critically. Finally, despite the growing literature on the Internet there has been little detailed discussion on the dynamics of implementing a 'pure' Web-based strategy in a firm that has been acknowledged by many to be successful. The intent of this paper is to address these gaps.

The Case of Amazon.com

In 1994, Jeffrey Bezos, a computer science and electrical engineering graduate from Princeton University, was the youngest senior vice-president in the history of D.E. Shaw, a Wall Street-based investment bank.

Dur
the
was
that
quit
dre
sold
pect
adv
sing
ven
Yor
that
ven

Ren
Bez
seve
on-l
offic
in J

Initi
2.4
atio
ble
Aug
moi
\$8 n
a ve
fun
and
had
180,
24
state
late
in J

In e
thes
seve
a fo
Wel
and
maj
dev
net.

To
boo
ove
its i

Op
Unl
elvt
to s
bus
Cus

Euro

During the summer of that year, one statistic about the Internet caught his imagination — Internet usage was growing at 2300% a year. His reaction: 'Anything that's growing that fast is going to be ubiquitous very quickly. It was my wake-up call.' Bezos left his job, drew up a list of 20 possible products that could be sold on the Internet and quickly narrowed the prospects to music and books. Both had a potential advantage for on-line sale: far too many titles for any single store to stock. He chose books. To start his new venture, on-line retailing of books, Bezos left New York City and moved to Seattle because he concluded that Seattle was *the* ideal location to start his on-line venture (for reasons explained later).

Renting a house in Bellevue, a suburb of Seattle, Bezos started work out of his garage. He soon raised several million dollars from private investors for his on-line venture. Operating from a 400-square-foot office in Bellevue, he launched Amazon on the WWW in July 1995.

Initially, the firm's revenues doubled in size every 2.4 months and revenues for the first year of operations were \$5 million. (These revenues are comparable to a large Barnes and Noble superstore.) By August 1996, book sales were growing at 34% a month. During this time, the firm was able to attract \$8 million from Kleiner, Perkins, Caufield and Byers, a venture-capital firm based in Silicon Valley that has funded well-known firms such as Sun Microsystems and Netscape. Through 31 December 1996, Amazon had sales of more than \$16 million to approximately 180,000 customer accounts in over 100 countries.¹ On 24 March 1997, the firm filed an S1 (the registration statement) application with the SEC and a few weeks later, it went public. Sales for the six months ending in June 1997 were \$43 million (see Figure 1).

In early 1997, Amazon had over 250 employees. Of these, 14 employees managed customer support and seven employees attended to marketing. In addition, a few employees managed 'content' on the firm's Web site, including such tasks as Web page updating and formatting book reviews for display. The vast majority of the remaining employees work on developing software tools for operating on the Internet.

To highlight Amazon's pioneering role in on-line book retailing, it is necessary to discuss the firm's overall approach to competitive positioning within its industry.

Operating a 'virtual' Bookstore

Unlike traditional bookstores, there are no bookshelves to browse at Amazon nor are there salespeople to service customers. Moreover, the firm is open for business 24 hours a day and has a global presence. Customers from 100 countries have purchased books

from the firm. The firm is also devoid of expensive furnishings, and money is spent sparingly. All contact with the company is done either through its Website or by e-mail.

At the firm's Website, customers can search for a specific book, topic or author, or they can browse their way through a book catalog featuring 40 subjects. Visitors can also read book reviews from other customers, the *New York Times*, the *Atlantic Monthly*, and Amazon's staff. Customers can browse, fill up a virtual shopping basket, and then complete the sale by entering their credit card information or by placing their order on-line and then phoning in their credit card information. Customer orders are processed immediately. Books in stock (mostly best sellers) are packaged and mailed the same day. When their order has been shipped customers are notified by e-mail. Orders for non-bestsellers are placed with the appropriate book publisher by Amazon immediately.

Shunning the elaborate graphics that clutter many Websites, Amazon instead loads up its customers with information. (*Time* magazine rated Amazon one of the 10 'Best Websites of 1996.')

For many featured books, it offers capsule descriptions, snippets of reviews and 'self-administered' interviews posted by authors. Additionally, the firm offers space for readers to post their own reviews and then steps out of the way and lets its customers sell to each other.

When asked why people come to his firm's site, Bezos responds (Fast Company, 1996):

'Bill Gates laid it out in a magazine interview. He said, 'I buy all my books at Amazon.com because I'm busy and it's convenient. They have a big selection, and they've been reliable.' Those are three of our four core value propositions: convenience, selection, service. The only one he left out is price: we are the broadest discounters in the world in any product category.... These value propositions are interrelated, and they all relate to the Web.'

At Amazon all books are discounted. Bestsellers are sold at a 40% discount and the other books at a 10% discount. The firm is able to provide deep discounts because they have a lower cost structure than a physical bookstore. According to Bezos: 'We turn our inventory 150 times a year. That's like selling bread in a supermarket. Physical bookstores turn their inventory only 3 or 4 times a year.' The firm's small warehouse is used only to stock mainly bestsellers and consolidate and re-pack customer orders. Moreover, only after the firm receives a paid-customer order does it request that a wholesaler ship the book to Amazon. The firm then ships the book to the customer. The firm owns no expensive retail real estate and its operations are largely automated.

Although the firm advertises in print, it spends a substantial amount on Web advertising. According to Jupiter Communications, the firm spent over \$340,000 for the first half of 1996 and ranked 34th

in Web ad spending. These expenses have gone up significantly in recent months, primarily because the firm entered into multi-year, \$1.8 million advertising agreements with three Internet aggregators, Yahoo!, Excite and AOL. For the six months ending June 1997, marketing and sales related expenses amounted to \$7.7 million.

Since Amazon is an Internet-only retailer, Web advertising gives it a unique opportunity to track the success of an ad by the number of click-throughs to the store's Web site and the number of Internet surfers who actually purchase something. However, industry analysts estimate that only between 2 and 3% of people who see an ad on the Web will actually click-through to see more.

Additionally, the firm advertises in large circulation newspapers such as *The Wall Street Journal*, *The New York Times*, and *San Jose Mercury News*, and on Internet search-engine sites such as Yahoo!, Excite, the Microsoft Network, and Microsoft's Slate magazine. Amazon keeps its banner ads simple, with just a few words and a Web address.² More recently, the firm has started to advertise on CNN, the cable news channel.

According to Bezos, Amazon differentiates itself from potential rivals in many ways, besides just marketing and aggressive brand promotion. He observes (Fast Company, 1996):

'People who just scratch the surface of Amazon.com say — 'oh, you sell books on the Web' — they don't understand how hard it is to actually be an electronic merchant. We're not just putting up a Web site. We do 90% of our customer service by e-mail rather than by telephone. Fourteen of our 110 employees do nothing but answer e-mail from customers. There are very few off-the-shelf tools that help do what we're doing. We've had to develop lots of our own technologies. There are no companies selling software to manage e-mail centers. So we had to develop our own tools. In a way this is good news. There are lots of barriers to entry.'

How does Amazon's approach to competitive positioning in the book retailing industry correspond to the descriptions provided in the emerging literature on on-line electronic commerce?

Notion of Disintermediation

There is a growing consensus that playing the role of 'middleman' on the Internet can be tricky, albeit not impossible. As noted, this is because the 'essence of the Internet is that it enables people and businesses to connect *directly* — sidestepping intermediaries like distributors and retailers' (Lohr, 1997). However, in general, on-line consumers are more interested in making better-informed purchases quickly, rather than necessarily getting the lowest price (The Economist, 1997a; Kambil, 1997). Therefore, on-line alter-

natives are more likely to attract customers if these alternative are less tiresome than conducting business in the physical world (see Figure 2).

Earlier, I noted that Bezos narrowed his choice of potential on-line business to music CDs and books. He chose books because:

'There are so many of them! There are 1.5 million English language books in print, 3 million books in all languages worldwide. This volume defined the opportunity. Consumers keep demonstrating that they value authoritative selection. The biggest phenomenon in retailing is the big-format store — the 'category killer' — whether it's selling books, toys, or music. But the largest physical bookstore in the world has only 175,000 titles.... With some 4200 US publishers and the two biggest booksellers, Barnes and Noble and Borders Group Inc., accounting for less than 12% of total sales, there aren't any 800-pound gorillas in book selling.'

In contrast, the music industry had only six major record companies. These companies controlled the distribution of records and CDs and, therefore, had the potential to lock out a new business threatening the traditional record-store format. Thus, by choosing books, and not CDs, Bezos had implicitly recognized that being an intermediary on the Internet is fraught with uncertainty. However, the Amazon case illustrates that being an intermediary on the Internet is a distinctive possibility, provided that an on-line venture adds value beyond that of its physical counterpart.

Competing as an Information Broker

The firm competes as 'information broker' for books and not just as a retailer of books as do most physical bookstores. It has found a way to use the emerging medium to offer services that a physical bookstore cannot match. According to Bezos:

'At Amazon customers can romp through a database of 2.5 million titles (over five times the largest superstore's inventory), searching by subject, title or author. When one selects a book, the site is programmed to flash other related titles that they may also want to buy.'

To keep customers interested in Amazon, the firm offers two forms of e-mail-based services to its registered customers. They include the 'Eyes' program and 'Editor's Service.' The 'Eyes' program is a personal notification service in which customers can register their interests in a particular author or topic. Once registered, they are notified each time a new book by their favorite author or in their topic is published. The 'Editor's Service' program provides editorial comments about featured books via e-mail. Many full-time editors at Amazon read book reviews, pore over customer orders, and survey current events to select the featured books. These and other freelance editors employed by the firm provide registered users with e-mail updates on the latest books

they have been reading. Both of these 'virtual' services are automated and, more importantly, are available free-of-charge to all registered users.

Such services, claims the firm's CEO, are vital for the firm's success. This is because 'If you make customers unhappy in the physical world, they might each tell six friends. If you make customers unhappy on the Internet, they can each tell 6000 friends with one message to a newsgroup. If you make them really happy, they can tell 6000 people about that. You want every customer to become an evangelist for you.' But the firm does not limit its efforts to being just an information broker, it attempts to create a sense of community through other mechanisms.

Earlier I noted that firms that create 'communities' that satisfy relational and transactional needs will reap the benefits of greater customer loyalty and gain important insights into the nature and needs of their customer base (Hagel and Armstrong, 1997). *In the context of book retailing, what mechanisms has Amazon instituted that enable the formation of a community? Also, what are the different communities that the firm's have attempted to develop?*

Virtual Community of Customers and Retailers

Community of Customers. There are four types of needs that electronic communities should attempt to satisfy in order to be successful (Armstrong and Hagel, 1996). They include the needs for: (1) transaction, (2) interest, (3) fantasy, and (4) relationships. I observed that the purpose of building communities is to develop a critical mass of participants by being the first mover and, thereby, making it difficult for new entrants to draw customers away from existing communities (Armstrong and Hagel, 1996). In contrast to other on-line businesses, Amazon is among the few pioneers attempting to leverage the medium's capability to develop on-line communities. The firm has instituted mechanisms that address at least three of the four needs identified in the extant literature.

For instance, the firm offers space for readers to post their 'own' reviews. It then steps out of the way and lets its customers sell to each other. Thus, a large part of the editorial content on the firm's Website is created by customers themselves (along with the firm's editors). As this content grows, it attracts others to add to the richness of the mix, thus creating a virtuous cycle (The Economist, 1997b). In essence, this is an explicit attempt to create a community around the *needs for transaction*. In other words, customers to the

firm's Website have the option to interact with others and be informed before they buy books.

But the Amazon case illustrates that they go beyond this and attempt to address the needs of communities of interest. Notes Bezos [quoted in Fast Company, 1996]:

'We want customers who enter Amazon.com to indicate whether they want to be 'visible' or 'invisible.' If they choose 'visible,' then when they're in the science fiction section, other people will know they're there. People can ask for recommendations — 'read any good books lately?' — or recommend books to others. I'm an outgoing person, but I'd never go into a bookstore and ask a complete stranger to recommend a book. The semi-anonymity of the on-line environment makes people less inhibited.'

This ability to be 'visible' or 'invisible' suggests that the firm is attempting to develop mechanisms that have the potential to promote greater interaction among customers. Customers can leverage the semi-anonymity that the electronic medium provides, and seek greater interaction with others based on their specific interests. In other words, these mechanisms are an attempt to enable interaction that participants might undertake in the 'virtual' context more so than in the 'physical' context of a traditional bookstore.

Recently, the firm has embarked on an innovative venture with the Pulitzer Prize winning author, John Updike. This author is co-authoring a 'collaborative' original on-line story (on a real time basis) entitled 'Murder Makes the Magazine.' Visitors to Amazon's Website are invited to participate in collaborating with Updike over 46 days. Then, Updike will conclude the on-line novel. The writers will contribute paragraphs and a few of these contributions will be selected by the firm's editorial staff to become part of the novel, and those chosen will each receive \$1000 for their effort. Notes Bezos:

'This is all about fun....We are committed to giving our customers rich and unique experiences. Forty-four talented people will get to collaborate in real-time with John Updike, the greatest living writer. We will all watch this collaboration unfold every day for 46 days.'

The innovative venture is an attempt to focus on groups of people who like to indulge in fantasy. More likely, these are the participants who like to create new environments, personalities, or stories (Armstrong and Hagel, 1996).

In sum, the approach employed by Amazon is consistent with Armstrong and Hagel's (1996) arguments that 'by adapting to the culture of the Internet, however, and providing customers with the ability to

The firm offers space for readers to post their 'own' reviews. It then steps out of the way and lets its customers sell to each other

interact with one another in addition to the company, businesses can build new and deeper relationships with customers.' It appears these activities have promoted more loyalty among the firm's customers. Currently, about 40% of the book orders come from repeat customers (SEC, 1997).

Although the extant literature has highlighted the notion of communities based on the needs of customers (i.e., emphasizing the demand side of the business), it has yet to address how firms are establishing communities of retailers to make their presence more ubiquitous on the Internet. Creating an on-line community of retailers is an attempt to attract a greater number of customers to the firm's Website. In other words, this is an attempt to emphasize the supply side. In this regard Amazon has taken the concept of 'community' building one step further.

Community of Retailers. Using an approach called the 'Associates Program,' the firm is attempting to generate more customer traffic to its Website. This program is a 'referral' service from other Websites to Amazon's 2.5 million item catalog. For instance, an associates Website, such as Starchefs — which features cookbook authors — recommends books and makes a link from its Web page to Amazon's catalog page for the books. According to *The Economist* (1997a, p. 10):

'Amazon.com knows that it will probably never be the very best site for rock climbing information or quantum physics discussions, but that the sites specialising in such subjects could be great places to sell books. A link to Amazon is an easy, and potentially lucrative, way for such specialist sites to do that at one remove; a click on the link takes a viewer to Amazon's relevant page.'

The associated Website then earns referral fees for sales generated by these links. Partners receive weekly referral fee statements and a check for the referral fees earned in that quarter. More than 15,000 sites (as of August, 1997) have already signed up under this program, and earn a commission of 8% of the value of books bought by the referred customer.

In contrast to physical bookstores that require a large outlay of financial capital for expansion, the approach used by Amazon leverages the capabilities of electronic medium without any additional overhead costs. This is clear from Bezos's statement that, 'The technology of the Web has made it possible to set up micro-franchises, and with zero overhead.'

Unlike the emphasis on promoting interactivity and communication among customers on Amazon's site, the focus of the 'Associates Program' is to gain greater name and brand recognition on the Internet where over 40,000 store fronts are currently operating. In such a noisy and fragmented environment capturing the consumer's attention is a critical

precursor to product sales (Kambil, 1997). Also, given that analysts estimate that only between 2 and 3% of people who see an ad on the Web actually click-through to see more, promoting the firm's ubiquitous presence via such programs greatly increases the probability of drawing customers to Amazon's Website.' Thus, by being a pioneer in promoting the notion of micro-franchises, the firm has built a substantial presence on the WWW.

Amazon is currently growing at the rate of 20–30% a month and part of the reason for this rapid growth is the firm's Associates Program (Bezos, 1996). Also, many of the important special interest sites have already become part of the Amazon electronic micro-franchising network. The more this network grows, the more it attracts other partners, creating a virtuous cycle of sites wanting to be associates. Moreover, as noted by the *Economist* (1997a):

'[T]he community of book-buyers Amazon has built, which provides an enviable amount of positive feedback. People come to the site because it has the most reviews written by readers, and they often stop to write some of their own, attracting even more people. The more books you buy and the more information you give Amazon.com about your tastes, the better it will become at finding things you might like. It will send you an e-mail message when a book you are looking for (or might just be interested in) is coming out, hooking you ever more firmly.'

As this network increases in size, it makes it difficult for other competitors to be ubiquitous on the WWW. More importantly, as the scale of Amazon's operations increase, it has the potential of 'locking out' new entrants to the on-line segment of the book retailing industry altogether. In this sense, the firm is attempting to harness the notion of 'increasing returns to scale' (Arthur, 1988) and competition here is beginning to resemble approaches used by some firms in 'network' industries.

Taken together the mechanisms instituted by the firm in both the supply and demand sides appear to be generating a large number of repeat customers to the firm's Website.

Changing Nature of the Firm's Competitive Advantage

It appears that Bezos recognizes that many of the competitive advantages that Amazon has developed are transient in nature. With a few exceptions (e.g., proprietary software, reputation, the sense of community, and first mover advantages), many of the mechanisms that are at this Website can be imitated by competitors, albeit not easily. To defend itself against competitors, for the past year Amazon has been developing a detailed purchasing history and profile of its customers. It now has a vast and unique database of customers' preferences and buying pat-

terns, tied to their e-mail and postal addresses (The Economist, 1997a). Notes Alberto Vitale, chairman of Random House Inc. (The Wall Street Journal, 1996): 'Amazon is creating a database that doesn't exist anywhere else. Book publishers have never had much market data about readers, and some are already salivating for a peek into Amazon's files.'

Amazon is still experimenting with ways to leverage and exploit this database. For instance, the firm is working on models that will customize information based on the user characteristics and preferences. Such an approach could open up further possibilities for 'mass-customizing' services without significant additional costs (cf. Kotha, 1996). Doing so may enable the firm to further entice people to return to its Website. However, despite the entry of a plethora of well-known physical bookstore firms (e.g., Barnes and Noble) and other copy-cat on-line book retailers, the firm's revenues continue to increase rapidly. *Now that large competitors such as Barnes and Noble and Borders Books have set-up stores on the Internet, how has Amazon fared?*

Performance Implications

During mid-1996, a plethora of virtual bookstores began sprouting on the Internet. Not surprisingly, many of these on-line bookstores are modeled after Amazon. As noted, despite this entry, the firm's revenues continue to increase rapidly.

Along with mechanisms instituted by the firm to form on-line communities, this growth can be attributed to 'first mover advantages' (cf. Lieberman and Montgomery, 1988). This observation is consistent with Harrington and Reed's (1996) arguments that 'Early movers in this rapidly evolving world can get ahead by attracting businesses to new on-line marketplaces, by determining the range of products and services these marketplaces offer, by setting the rules of interaction, and by reserving the value-added positions for themselves.' Therefore, overcoming first-mover advantages is perhaps an important consideration for businesses contemplating entry into book retailing.

This rapid growth can also be attributed to the firm's Seattle location. Seattle is a thriving area for software talent and, more importantly, is adjacent to the world's biggest book warehouses operated by Ingram Industries in Oregon. Observes Bezos [quoted in Fast Company, 1996, October–November]:

It sounds counterintuitive, but physical location is very important for the success of a virtual business. We could have started Amazon.com anywhere. We chose Seattle because it met a rigorous set of criteria. It had to be a place with lots of technical talent. It had to be near a place with a large numbers of books. It had to be a nice place to live — great people won't work in places they don't want to live.

Finally, it had to be in a small State. In the mail-order business, you have to charge sales tax to customers who live in any State where you have a business presence. It made no sense for us to be in California or New York.... Obviously Seattle has a great programming culture. And it's close to Roseburg, Oregon, which has one of the biggest book warehouses in the world.'

This counterintuitive argument becomes more salient when one compares Amazon with the Internet Bookshop, an on-line bookstore founded at approximately the same time by Darryl Mattocks in England. Although both Bezos and Mattocks had independently conceptualized their business around books, Amazon has grown to over \$40 million in revenues (as of September 1997), whereas Mattocks' venture is still under a million dollars.

The comparison between the two firms becomes more dramatic when one examines their respective market capitalization. Amazon was valued around \$500 million during its IPO (and its value has increased significantly since), and the Internet Bookshop, in contrast, was valued at \$10 million on Britain's Ofex, a small pseudo-market that matches trader's orders (The Economist, 1997b). According to The Economist (1997b, p.69):

'It is tempting to conclude that Mr. Mattock's is just a victim of Britain's less vibrant venture-capital industry. But he also helped himself to come second. He did not research the industry in the same systematic way as Mr. Bezos. He started the firm in Oxford because that was where he lived, and he liked the town. But Oxford is not a book-distribution centre....Worse, the Internet Book Shop does not sell titles until their British release date, which can be months after the American publishing date for many books.... Significantly, Amazon's sales outside America are ten times the Internet Book Shop's sales outside Britain.'

In contrast to the argument that on-line commerce is (1) making retailing a global industry, and (2) geography (location) is not a consideration for businesses on the Internet, my study of Amazon suggests that location, at least for on-line book retailing, is still an important consideration.

Some critics contend that despite all of the attention paid to Amazon's strategy and approach on the Internet, the firm is yet to turn a profit (see Figure 1). Acknowledges Bezos:

'[Although] we have seen no profits and our strategy is to invest very heavily right now because we believe so much in the future of this business.... We are spending significant amounts of money investing in marketing and building our brand name and customer acquisition and so on, and that's one of factors that holds down profitability.'

Perhaps, it is too premature to judge the 'performance' implications of the approach used by Amazon

to gain a foothold in the emerging on-line book retailing industry. Based on cybershare (and revenues), the firm is currently acknowledged to be the largest on-line bookstore on the WWW.

Implications

This paper has explored Amazon's unique strategy in the on-line US book retailing industry which is characterized by change and growing intense competition. Despite the limitations inherent in drawing conclusions from a single study, it is possible to draw some useful inferences regarding Amazon's approach to electronic commerce.

First, I have identified several industry conditions (e.g., location, proximity to suppliers) that may be important for the successful pursuit of an on-line book retailing strategy. Specifically, it appears that first mover advantages may be critical for the success of on-line retailing. Despite the imitation of this firm's strategy by its rivals, Amazon has been able to establish its reputation as the most recognized 'Internet' on-line book retailer. The reputation effects derived from first mover advantages can be important. This is consistent with Hilts' (1997) observations that, 'Their [Amazon.com] timing was exquisite, and they promoted the heck out of it.' Hence, *timing the entry* into on-line commerce may prove critical to its success ultimately. In other words, researchers can examine whether a firm that is a first mover into on-line retailing reaps superior returns, relative to competitors that are late entrants.

Second, although the emerging literature discounts the location as being not so relevant to on-line commerce (cf. Reid, 1997), the study of Amazon suggests that for on-line merchants, especially those who sell physical products and are intermediaries in the value

chain, location can be a source of competitive advantage. Location can be critical because of the need to rapidly acquire the product for distribution to the customer, and the cost saving resulting from not having to carry large inventories. Location can also determine the availability and the ability to attract top talent needed to compete on this emerging medium.

Third, the mechanisms instituted by the firm were an outgrowth of in-house programming. Although some of these mechanisms are becoming widely available, many still are proprietary. With the increased access to financial capital, the firm is becoming more sophisticated in its application of the Internet technology. In other words, the firm's organizational capabilities to exploit the new medium have increased over time and continue to increase. Thus, many of its capabilities are not readily available to new entrants without these same resources. Ultimately, the firm's ability to envision and build innovative on-line mechanisms using software tools developed internally depends upon the human capital it is able to attract.

Fourth, Barnes and Noble's (and other competitors) entry into on-line book retailing does not seem to have negatively impacted Amazon's continued rapid growth. Further, an examination of the firm's stock price (since its IPO offering) indicates that many market analysts perceive its approach to electronic commerce to be sustainable (Figure 3). However, at this stage, one can only speculate that Amazon's continued growth suggests that brand names such as Barnes and Noble, established over other media, do not automatically transfer over to the Internet. This is clearly an interesting question for future research. Perhaps this is due to the fact that the opinion leaders for the Internet are likely to be different from those of the other media (Kambil, 1997).

Finally, for a firm competing in the on-line environment, the ability to continually attract customers, and

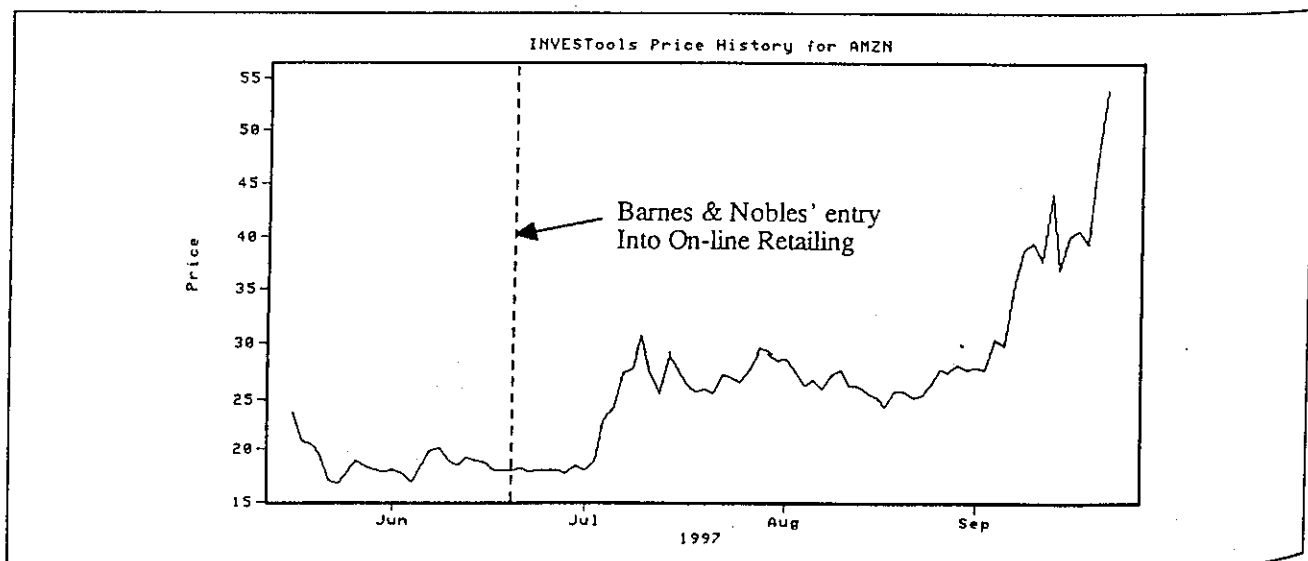


Figure 3 History of Amazon.com's Stock Price since its IPO.

get them to return to their on-line store is critical for a competitive advantage. One generalizable approach to attract and retain customers is through building communities (Hagel and Armstrong, 1997). The Amazon case highlights many mechanisms that the firm has instituted to foster a sense of community. It is this sense of community that the firm hopes will promote customer loyalty and fend off imitators from taking cybershare away from Amazon (The Economist, 1997a). Although, customer loyalty is important for any business, it can be vitally important for firms competing on-line. This is because of the lack of built-in switching costs between on-line merchants offering identical merchandise. Also, according to Armstrong and Hagel (1996: 141):

The value of successful electronic communities will be the intense loyalty they generate in their participants, which is what favors first movers into this area. The organization of successful electronic communities will depend on skills and the right iconoclastic mind-set, not capital. As a consequence, this arena may favor bold entrepreneurs with constrained resources over established corporations.'

Although the mechanisms that Amazon has created may be unique to book retailing, the concepts (e.g., Eyes program, and Editors' choice, and Associates Program) that the firm has instituted are perhaps generalizable to other firms operating on the WWW.

Concluding Remarks

The intent of this paper was to explore the dynamics and issues pertaining to implementing the on-line book retailing strategy employed by Amazon. By envisioning and developing an innovative on-line approach to book retailing Amazon, a relative new entrant, has successfully challenged the established firms (e.g., Barnes and Noble) in this industry. In an effort to stay competitive, and understand the transformation unleashed by Amazon's entry, many established firms are beginning to emulate the business model pursued by Amazon. However, given the rapidity of change, the uncertainty surrounding the dynamics of the Web, and the likely moves and counter moves by players in this industry, it is hard to predict the exact future state of this industry. However, in its effort to 'create the future,' it appears that Amazon is well ahead of its competitors. Also many of the mechanisms it has instituted are prototypes of what one might expect to observe in on-line retailing. In that sense, these mechanisms are recipes for superior competitive advantage on the Internet.

Acknowledgements

I thank Anil Nair and Dick Moxon for their comments on an earlier version of the document.

Notes

1. The firm's self-reported measures indicate that average daily visits (not 'hits') have grown from approximately 2200 in December 1995 to approximately 50,000 in December 1996, and repeat customers currently account for over 40% of orders (SEC, 1997).
2. Web advertising is gaining increasing legitimacy. Revenue for advertising on the World Wide Web rose 83% to \$46.4 million in the second quarter of 1996. The figure is expected to reach \$312 million by the end of the year. This amount is still quite small in comparison to the \$30 billion spent on television advertising each year.
3. Moreover, Bezos adds: There are a couple of objectives with a promotion like this. One is to bring more people to the Website and some fraction of those people will convert into customers. It also helps branding. The second objective with a contest like this is create a fun environment on the Website that sort of replaces some of the fun components that you can do in a physical bookstore that you can't do on-line. You know you never have soft sofas and tasty lattes online, but you can do a lot of fun things online that can't be done in the physical environment.'
4. This approach is not without its risks (The Economist, 1997a). First, customers may be more loyal to the referring site and, second, the referring sites also may transfer their allegiance to some other bookseller if they offer a more 'lucrative' deal. Therefore, it appears managing the network of micro-franchises is as important as maintaining the interests of the community of customers. In a sense both the supply-side and the demand are highly related.

References

- Armstrong, A. and Hagel, J. (1996) The real value of on-line communities. *Harvard Business Review* May/June, 134-141.
- Arthur, B. (1988) Self-reinforcing mechanisms in economics. In *The Economy as an Evolving Complex System*, eds P. Anderson et al. Addison-Wesley, Reading, MA.
- Bezos, J. (1996) Interview transcripts.
- Evans, P. B. and Wurster, T. S. (1997) Strategy and the new economics of information. *Harvard Business Review* September/October, 70-83.
- Fast Company* (1996) Who's writing the book on Web business? October-November, pp. 132-133.
- Gilder, G. (1996) *Telecosm*. Simon and Schuster, New York.
- Hagel, J. and Armstrong, A. (1997) *Net Gain: Expanding Markets Through Virtual Communities*. Harvard Business School Press, Boston.
- Hamel, G. and C.K. Prahalad (1994) *Competing for the Future*. Harvard Business School Press, Boston.
- Harrington, L. and Reed, G. (1996) Electronic commerce (finally) comes of age. *The McKinsey Quarterly* 2, 69-83.
- Hilts, P. (1997) Quoted in the *Seattle Times*, 5 January.
- Kambil, A. (1997) Doing business in the wired world. *IEEE* May, 56-61.
- Kotha, S. (1996) From mass production to mass customization: the case of the national industrial bicycle company of Japan. *European Management Journal* 14, 442-450.
- Lieberman, M. B. and Montgomery, D. B. (1988) First mover advantages. *Strategic Management Journal* 9, 41-58.
- Lohr, (1997) Digital commerce. *The New York Times*. 19 May, p. C5.
- Luhmann, N. (1990) *Essays in Self-reference*. Columbia University, New York.
- Negroponte, N. (1995) *Being Digital*. Knopf, New York.
- Philadelphia Business Journal* (1996) 27 September.
- Publishers Weekly* (1996a) 1 January.
- Publishers Weekly* (1996b) 11 March.
- Reid, R.H. (1997) *Architects of the Web: 1000 Days that Build the Future of Business*. John Wiley and Sons Inc., New York.



SEC (1997) *Amazon.com, 10Q Filing with the Securities and Exchange Commission*. Edgar Database. US Department of Commerce.

Tapscott, D. (1996) *The Digital Economy*. McGraw-Hill, New York.

The Economist (1997a) A survey of electronic commerce, 10 May, pp. 1-18.

The Economist (1997b) A fable concerning ambition, 21 June, p. 69.

The Financial Times, 7 October.

The New York Times (1996) A nonchain bookstore bucks the tide, 8 September, p. 4.

The Wall Street Journal (1996) Reading the market: how a Wall Street whiz found a niche selling books on the Internet, 16 May, p. 1.

US Bureau of the Census (1992) US Department of Commerce Government Publications, Washington, DC.

Wallance, J. (1997) *Overdrive: Bill Gates and the Race to Control Cyberspace*. J. Wiley and Sons Inc., New York.

Williamson, O. E. (1979) Transaction-cost economics: the governance of contractual relations. *Journal of Law and Economics* 22, 233-261.

Weintraut, J.N. (1997) Introduction. In *Architects of the Web*, ed. R.H. Reid. John Wiley and Son Inc., New York.



SURESH KOTHA,
School of Business
Administration, Uni-
versity of Washington,
Box 353200, Seattle,
Washington 98195,
USA.

Suresh Kotha is an Assistant Professor of Strategic Management at the School of Business Administration, University of Washington. His teaching and research interests are in the areas of competitive strategy, manufacturing strategy, and US/Japan comparative practices. Dr. Kotha is currently conducting a study that focuses on identifying and highlighting various mechanisms that exemplar on-line firms are employing to create and sustain their competitive advantage on the World Wide Web.

L
A
A

ANI
Man

The c
for a
the c
and r
vides
orgar
renew
exam
tiona
know
sition
legic
text.
provi
a us
unity
their
with
partn
ing
with
and
bases
ate v
oppo
partn
cle fi
ent
that
of al
conn
learn
consi
effect
alliar
All ri

Please read the paper "Integrating Information from Multiple Independently Developed Data Sources" and answer the following questions.

1. What are the "conflicting information" problem and the "object identification" problem discussed in the paper? Please use examples to illustrate these two problems. (10%)
2. The paper proposed the use of MATCH clauses to specify how to identify the same objects from various sources. Please clearly identify the limitations of this representation. (16%)
3. (a) Assume the global ID table (for recording matching instances) be materialized before the query processing. Please modify the merge query processing algorithm (as shown in page 245-246) accordingly. (10%)
(b) What are the advantages and disadvantages of the materialized global ID table approach as compared to the approach that dynamically constructs the global ID table during the query processing phase (as proposed in the paper). (14%)

Integrating Information from Multiple Independently Developed Data Sources

I-Min A. Chen*

Bioinformatics Systems, Gene Logic, Inc.
2001 Center Street #600, Berkeley, CA 94704
Email: ichen@genelogic.com

Doron Rotem

Lawrence Berkeley National Laboratory, Berkeley, CA 94720
Email: D_Rotem@lbl.gov

Abstract

We propose in this paper an approach to integrating information from multiple independently developed data sources, in which incomplete or conflicting information prevails. We do not assume that equivalent objects in different sources have the same identity or key values, or there exists a global ID table uniquely identifying these objects. Instead, a global user view is constructed upon these data sources, and the view definition specifies how objects from various information sources are identified, matched and merged. Two query processing strategies are presented for processing queries against the global user views: (1) a merge-based approach which first merges all objects from various sources and then performs query evaluation on the merged result, and (2) a condition push-down approach that passes as many query conditions as possible in subqueries to local sources and then processes the returned subquery results. To validate our algorithms and to conduct performance studies, a prototype that reads the global view definition and processes input queries using both merge and push-down strategies was implemented. Our performance studies show that the push-down approach outperforms the merge approach when large portions of non-qualified instances can be removed from all sources by condition push-down. Otherwise, the merge approach has better performance.

1 Introduction

When integrating information from multiple independently developed data sources, one will inevitably encounter problems such as: *name differences* (e.g., synonyms or homonyms), *format differences* or *mismatched attribute domains* (e.g., attribute values are represented in different formats or different units, or are encoded into different code values), and *missing and conflicting data* [2]. With synonyms and format differences in key values it is difficult to identify equivalent

*Part of the work was done when the author worked at Lawrence Berkeley National Laboratory, and was supported by the Office of Health and Environmental Research Program of the Office of Energy Research, U.S. Department of Energy under Contract DE-AC03-76SF00098.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIKM 98 Bethesda MD USA

Copyright ACM 1998 1-58113-061-9/98/11...\$5.00

objects from various data sources. Two kinds of integrations are of special importance: (1) *join integration* that interconnects related objects in different sources, and (2) *union integration* that combines similar or equivalent local objects from different sources to form global objects [6]. This paper focuses on the latter.

Many approaches have been proposed to address the information integration problem (among others, [1, 2, 3, 4, 5, 6, 7, 9, 10, 11, 12, 13, 14]). In these approaches, a global schema or view is defined to integrate local sources. Look-up tables, functions, virtual attributes or dynamic attributes, and external procedures may be used to resolve the domain mismatch problem. All approaches are subject to at least one of the following restrictions: (1) They either do not identify equivalent objects from different information sources, or assume the existence of global ID values or global ID tables. (2) They assume each local attribute value is mapped to at most one global (view) attribute value. (3) They do not handle missing attribute values (or unknown values) at query processing. (4) Users are limited to certain "fixed" object integration approach, or are required to learn a complicated language to define how similar or equivalent objects are to be merged. (See section 7 for more detailed discussion of related work.)

In this paper, we propose an approach to defining one or more global user views for integrating information from multiple independently developed data sources. Instead of assuming uniform and consistent local data sources, we give the users different options of specifying conflict resolutions, attribute value conversions and object identification methods in the global view definitions. Users can issue queries against the global views to retrieve objects that are known to be true or are considered as *maybe* true.

Retrieval queries against a global view can be processed by (1) first "materializing" global objects by *merging* corresponding information from local sources, and then evaluating query conditions on the global objects [1, 5], or by (2) *pushing down* applicable selection conditions into subqueries to local sources, and then combining the results returned by these subqueries [1, 4, 6, 10, 12]. The first alternative requires processing a large amount of data locally, which may take lots of space and can be slow. Condition push-down is very complicated if one does not know whether an object exists only in one local data source or in many. Let us illustrate this using a restaurant example: suppose restaurant *A* is returned as an inexpensive restaurant by a data source S_1 , but not by another data source S_2 . It is not clear whether this is because *A* is not listed in source S_2 , or because *A* is listed in source S_2 as an expensive restaurant (a

conflict). If it is the first case, then A should be returned to the user. If it is the second case, and source S_2 is considered as more reliable than S_1 , then A should not be returned. In the case that one only considers objects co-existing in both sources [12] or that there is a global ID table, this is not a problem. However, when considering integrating information from multiple independent data sources in general, this is a problem that has to be resolved. Besides, not all selection conditions can be pushed down to local sources. With all the complexity in condition push down, the push-down approach does not necessarily guarantee better performance.

The rest of this paper is organized as follows. In section 2, we present a restaurant example that is used to illustrate our approach throughout the paper. The construction of global views for information integration is described in section 3. A *merge* query processing strategy that matches objects from two sources, and then evaluates query condition on merged objects is discussed in section 4. In section 5 we describe problems with condition push-down, and present an algorithm for *push-down* query processing. A performance comparison of the two strategies using a prototype is given in section 6. Section 7 presents related work, and section 8 gives concluding remarks.

2 A Restaurant Example

2.1 Source Description

We consider two information sources for the Bay Area restaurants: Source S_1 contains Fodor's restaurant guide of San Francisco restaurants.¹ There are 11 attributes: *name*, *address*, *city*, *location*, *price*, *cuisine*, *star*, *phone*, *lunch*, *dinner* and *description*. Attributes *lunch* and *dinner* indicate in which days a restaurant is open for lunch and dinner, respectively. Other attributes are self-descriptive. Assume *fodors* is stored in a Unix *ndbm* database, and is accessed via a local search engine. (We called the combined system *ndbm+*.) Source S_2 contains San Francisco Focus restaurant guide of the Bay Area restaurants.² There are 9 attributes: *name*, *address*, *city*, *phone*, *price*, *cuisine*, *lunch*, *dinner* and *description*. Assume *focus* is stored as a table in a relational database.

These two sources are described in Figure 1. Each source has a name, a database type, a data source, and a set of attributes. A source attribute has a name (which is unique in this source) and a data type. If a source attribute is part of a key, then it is specified as *key*. A multi-valued attribute is specified as *multi-valued*. A source attribute that allows null values is defined as *can be null*.

Each "selectable" attribute is specified with keyword *select*. It is possible that some attributes allow range selections, while others allow only value matching using "*=*". Since *source description* [8] is not the main focus of this paper, we simply assume that legal attribute conditions are determined by the database type of the sources. Source type *ndbm+* allows value matching, range selection, set comparison using *in*, *not in* and *contains*, and null value checking; only conjunctive query conditions are supported. Source type *rdbms* supports all the above subcondition types except for *contains*, and disjunctive conditions are allowed.

2.2 Problems

Problems in integrating information from the two sources include:

¹<http://www.fodors.com/>

²<http://sfbay.yahoo.com/external/focus/index.html>

```
SOURCE fodors
DB ndbm+ DATA_IN "fodors"
ATTRIBUTE name: varchar(40) key, select
ATTRIBUTE address: varchar(80) key, select
ATTRIBUTE city: varchar(20) select
ATTRIBUTE location: varchar(40) select
ATTRIBUTE price: char(10) select
ATTRIBUTE cuisine: varchar(40) multi-valued,
                select
ATTRIBUTE star: char(1) select
ATTRIBUTE phone: char(12) select
ATTRIBUTE lunch: varchar(20) select, can be null
ATTRIBUTE dinner: varchar(20) select, can be null
ATTRIBUTE description: text
```

```
SOURCE focus
DB rdbms DATA_IN "focus"
ATTRIBUTE name: varchar(40) key, select
ATTRIBUTE address: varchar(80) key, select
ATTRIBUTE city: varchar(20) key, select
ATTRIBUTE phone: varchar(40) select
ATTRIBUTE price: char(4) select
ATTRIBUTE cuisine: varchar(40) select
ATTRIBUTE lunch: varchar(22) select, can be null
ATTRIBUTE dinner: varchar(22) select, can be null
ATTRIBUTE description: text
```

Figure 1: Two Restaurant Guide Databases

- **Name Difference:** Restaurant names are not unique; restaurant chains have several restaurants with the same names. The same restaurants may be listed with slightly different names in two sources (e.g., "Helmand" versus "Helmand Restaurant"). Similar name differences occur in other attributes (e.g., address) as well.
- **Mismatched Domains:** Information such as price, cuisine, lunch and dinner is described differently in both sources. For example, one restaurant guide lists a restaurant as "closed weekends," while the other lists the same restaurant as "open Mon-Fri." Attribute *cuisine* in source *fodors* is multi-valued, while *cuisine* in source *focus* is single-valued.
- **Missing Data:** No single guides provide complete information. Some restaurants are listed in one guide, but not in the other. Some properties such as restaurant rating are available only in one source. Information regarding restaurant opening time may be missing in either source.
- **Conflicting Information:** The two sources list conflicting opening time, price ranges and/or cuisine types for the same restaurant.
- **Object Identification:** There are no global ID values or inter-database ID tables for identifying the same restaurants from both sources.

3 Integrating Information from Multiple Sources

3.1 Constructing Global Views

Global views are constructed to facilitate multiple source information integration. A view has a unique view name, is defined on one or more information sources, and is associated with one or more view attributes. (In this paper we only integrate information from two sources; however,

the proposed approach can be extended to more sources.) View `Restaurant` in Figure 2 is a global view constructed on top of the two restaurant guides described in the previous section.

View attribute values can be obtained from merging multiple source attributes. (Attribute merging will be discussed in Section 3.3.) Each source attribute value can be taken "as is," or can be converted using external procedures or conversion functions.

External procedures are a set of system or user-provided procedures that perform complicated value transformations. For example, `p.GetAddress` is an external procedure that examines and extracts only street number and street name of an address.

Conversion functions convert source attribute values to one or many view attribute values. For example, `t.ConvertDay1` converts a local source value "No weekends" to destination view attribute values { "Mon" "Tues" "Wed" "Thur" "Fri" }. Although conversion functions can also be implemented as external procedures that perform the same value conversion tasks, the definition of conversion functions provides explicit value conversion information in the view definition, and also facilitates condition push-down during query processing (to be discussed in later sections). For a multi-valued attribute A , $t(A) = \cup \{ t(a) \mid a \in A \}$.

Although it is not shown in the `Restaurant` view definition, a view attribute can take values from multiple local attributes of the same source. Similarly, two view attributes can take values from the same local source attribute possibly using different external procedures or conversion functions.

3.2 Identifying Objects

View attributes can be associated with `MATCH` clauses to specify how to identify the same objects from various sources. Only source attributes of a view attribute are allowed in the `MATCH` clause of this view attribute. External procedures or conversion functions can be used in matching. For example, view attribute `Address` has two source attributes `fodors.address` and `focus.address`, and therefore it can be associated with a `MATCH` clause specifying how values of the two source attributes are to be matched. An external procedure named `p.GetAddress`, which extracts street number and street name from an address, is applied to match addresses from both sources. For example, both `p.GetAddress("1 Maple St.")` and `p.GetAddress("One Maple Street")` are evaluated to "1 Maple," and therefore the two addresses are considered as the same.

Instances in local sources are identified and matched to form global view instances. In our restaurant example, a restaurant from `fodors` and a restaurant from `focus` are considered as the same restaurant if they have matching name, address and city attribute values. If no attributes in a view are associated with `MATCH` clauses, then we consider all objects from different sources as distinct.

3.3 Merging Attribute Values

When merging information from multiple sources, one has to take into account that maybe no single source provides complete information, and that two sources can provide conflicting information. We consider null values as *unknown*, and two sources providing conflicting information if they return different non-null values on the same attribute of the same instance. (One can easily extend this definition to consider a null attribute value and a non-null attribute value to be

conflicting.) In the case that inconsistency occurs, a user may prefer information from a source that is considered as more reliable.

Let S_1 and S_2 be two local sources for a view V . Suppose i is a global view instance with corresponding local instance i_1 in S_1 . If source S_2 provides complete information, then there should exist an equivalent instance of i_1 , say i_2 , in S_2 . Instances i_1 and i_2 are to be matched and merged to form the global instance i . If such i_2 does not exist in source S_2 , then we consider that it is because S_2 is incomplete.³ This is as if there were an instance i_2 in S_2 with all attribute values unknown. Based on the above assumption, we will be able to define attribute merging operations. Let Va be a view attribute of V which is defined on local attributes or attribute conversions A of S_1 and B of S_2 . Four merging operations are provided for defining Va :

No-Merge - $Va: A$

View attribute values are obtained from single source attribute or attribute conversion A . For example, attribute `Star` is defined solely on `fodors.star`.

- if i_1 in S_1 exists, then $Va(i) = A(i_1)$;
- otherwise, $Va(i) = \text{NULL}$.

Union-Merge - $Va: A ; B$

View attribute values are defined as the union of multiple source attribute values. We use symbol `;` to denote union-merge. For example, view attribute `Cuisine` is defined as: `fodors.cuisine ; focus.cuisine`.

- if both i_1 in S_1 and i_2 in S_2 exist: then $Va(i) = A(i_1) \cup B(i_2)$;
- if only i_1 exists, then $Va(i) = A(i_1)$;
- if only i_2 exists, then $Va(i) = B(i_2)$.

Intersect-Merge - $Va: A , B$

View attribute values are defined as the intersection of multiple source attribute values. We use symbol `,` to denote intersect-merge. For example, view attribute `Lunch` is defined as: `t.ConvertDay1(fodors.lunch), t.ConvertDay2(focus.lunch)`.

- if i_1 in S_1 and i_2 in S_2 both exist, then: (i) if $A(i_1)$ is null, then $Va(i) = B(i_2)$; (ii) if $B(i_2)$ is null, then $Va(i) = A(i_1)$; (iii) otherwise, $Va(i) = A(i_1) \cap B(i_2)$;
- if only i_1 exists, then $Va(i) = A(i_1)$;
- if only i_2 exists, then $Va(i) = B(i_2)$.

Prefer-Merge - $Va: A > B$

Attribute values from one source are preferred (if values are not unknown). We use symbol `>` to denote prefer-merge. For example, view attribute `Price` is defined as: `t.ConvertPrice1(fodors.price) > t.ConvertPrice2(focus.price)`.

- if i_1 in S_1 exists and $A(i_1)$ is not null, then $Va(i) = A(i_1)$;
- else if i_2 exists, then $Va(i) = B(i_2)$;
- otherwise, $Va(i) = \text{NULL}$.

³We neglect the case that S_1 mistakenly includes an incorrect instance i_1 .


```

VIEW Restaurant
ON fodors, focus
ATTRIBUTE Name: fodors.name > focus.name MATCH p_GetName(fodors.name) = p_GetName(focus.name)
ATTRIBUTE Address: focus.address > fodors.address
MATCH p_GetAddress(fodors.address) = p_GetAddress(focus.address)
ATTRIBUTE Location: fodors.location
ATTRIBUTE City: fodors.city > focus.city MATCH fodors.city = focus.city
ATTRIBUTE Phone: p_GetPhone(fodors.phone); focus.phone
ATTRIBUTE Price: t_ConvertPrice1(fodors.price) > t_ConvertPrice2(focus.price)
ATTRIBUTE Cuisine: fodors.cuisine ; focus.cuisine
ATTRIBUTE Lunch: t_ConvertDay1(fodors.lunch), t_ConvertDay2(focus.lunch)
ATTRIBUTE Dinner: t_ConvertDay1(fodors.dinner), t_ConvertDay2(focus.dinner)
ATTRIBUTE Star: fodors.star
ATTRIBUTE Description: fodors.description; focus.description

CONVERSION t_ConvertDay1
{ ("No", { "No" } ), ("No weekends", { "Mon" "Tues" "Wed" "Thur" "Fri" } ),
  ("No Sun", { "Mon" "Tues" "Wed" "Thur" "Fri" "Sat" } ),
  ("No Mon", { "Sun" "Tues" "Wed" "Thur" "Fri" "Sat" } ),
  ("No Sun Mon", { "Tues" "Wed" "Thur" "Fri" "Sat" } ),
  ("No Tues", { "Sun" "Mon" "Wed" "Thur" "Fri" "Sat" } ),
  ("No Mon Tues", { "Sun" "Wed" "Thur" "Fri" "Sat" } ) }

CONVERSION t_ConvertPrice1
{ ("Under $20", "$"), ("20--$30", "$$"), ("30--$50", { "$$$" "$$$$" } ), ("Over $50", "$$$$$" ) }

CONVERSION t_ConvertDay2
{ ("No", { "No" } ), ("Sun-Fri", { "Sun" "Mon" "Tues" "Wed" "Thur" "Fri" } ),
  ("Mon-Fri", { "Mon" "Tues" "Wed" "Thur" "Fri" } ),
  ("Mon-Sat", { "Mon" "Tues" "Wed" "Thur" "Fri" "Sat" } ),
  ("Tues-Sun", { "Sun" "Tues" "Wed" "Thur" "Fri" "Sat" } ),
  ..... ("daily", { "Sun" "Mon" "Tues" "Wed" "Thur" "Fri" "Sat" } ) }

CONVERSION t_ConvertPrice2
{ ("$", "$"), ("$$", { "$" "$$" "$$$" } ), ("$$$$", { "$$$$$" "$$$$$" } ) }

```

Figure 2: The Global View

3.4 Queries

Users can query a global view V using conjunctive select queries with format:

```

select  $A_1, A_2, \dots, A_m$  from  $V$ 
where  $cond(B_1)$  and ... and  $cond(B_n)$ 

```

where $A_1, \dots, A_m, B_1, \dots, B_n$ are all attributes of V , and $cond(B_i)$ is an attribute condition with format: (1) B_i is (not) null, (2) $B_i = (\neq, >, \geq, <, \leq, \text{contains}) v$, or (3) B_i in (not in, contains) S , where v is a constant value and S is a set of values. If B_i is multi-valued, then condition $B_i > (\geq, <, \leq) v$ is interpreted as: any value of $B_i > (\geq, <, \leq) v$.

We consider missing attribute values as *unknown*. For example, an instance that is listed only in focus will have unknown Star attribute values. Additional selection operation *select+* is provided for retrieving also instances that have unknown values on the condition attributes. For example, the following query retrieves all restaurants that either have Star = "Y" or have unknown Star values: *select+ Name from Restaurant where Star = "Y"*;

4 Merge Query Processing

In order to process a retrieval query against a global view, we can first identify the same objects in various information sources, merge proper attribute values of these objects, and

then evaluate the query on the merged results. We present in this section a *merge query processing approach* based on this strategy.

Let Q be a conjunctive query against global view V as described in the previous section:

```

select (or select+)  $A_1, A_2, \dots, A_m$  from  $V$ 
where  $cond(B_1)$  and ... and  $cond(B_n)$ 

```

The basic steps of this query processing strategy are as follows:

1. Modify query Q by adding to the *select* clause all the view attributes that are associated with match clauses or are in $\{B_1, \dots, B_n\}$, but not selected by Q . Let Q' be the new query with modified *select* clause.
2. For the two sources of V , S_1 and S_2 , prepare two subqueries Q_{s1} and Q_{s2} : the *select* clause of Q_{s1} (Q_{s2}) contains all attributes of S_1 (S_2) that serve as local source attributes (with or without conversions) of any view attributes in the *select* clause of Q' . Q_{s1} (Q_{s2}) has no associated *where* clause.
3. Let S_2 be the source with a slower retrieval system, and S_1 be the source with a faster retrieval system. For each instance o_1 returned by subquery Q_{s2} , if o_1 matches an instance o_2 returned by subquery Q_{s1} , then record o_1 and o_2 as matching instances in a *global*

ID table. (The global ID table records only *key* attributes of o_1 and o_2 , if there are such key attributes; otherwise, the entire instances are recorded.)

- For each o_1 that has matching instance o_2 , merge o_1 and o_2 to form a global object o based on the view definition.
- For each o_1 that has no matching instances, construct a global object o based on the view definition and o_1 attribute values alone.

Evaluate query condition of Q (i.e., *where* clause of Q) on o :

- If Q uses *select*, then the standard 2-valued logic is applied, and all objects that are evaluated to *true* are returned with a status value *true*.
- If Q uses *select+*, then 3-valued logic is applied in condition evaluation, and all objects that are not evaluated to *false* are returned with a status value *true* or *maybe*.

If o is evaluated to *true* (or *maybe*), then project attributes of o that are selected by the original query Q , include the selected and projected instances in the query answers.

4. For each instance o_2 returned by subquery Q_{S_2} , if o_2 is not included in the global ID table, then construct a global object o based on the view definition and o_2 attribute values. Evaluate query condition of Q on o (as described above), and include the projected instance o in the query answer if o is evaluated to *true* or if o is evaluated to *maybe* in a *select+* query. \diamond

(Example 1.) Suppose the following query (Q0) is issued to select all good San Francisco restaurants that serve Italian food and are open for Sunday dinners:

```
select+ Name, Address, City, Location, Phone,
        Price, Cuisine, Lunch, Dinner, Star
from Restaurant
where City = "San Francisco"
and Cuisine contains "Italian"
and Dinner contains "Sun" and Star = "Y";
```

Since all view attributes associated with match clauses (i.e., Name, Address and City) or in query condition are selected, we do not need to modify the *select* clause of (Q0).

Two subqueries (Q0.1) and (Q0.2) are issued against sources *fodors* and *focus*, respectively:

```
(Q0.1) select name, address, city, location, phone,
           price, cuisine, lunch, dinner, star
from fodors
```

```
(Q0.2) select name, address, city, phone, price,
           cuisine, lunch, dinner
from focus
```

Identify the same restaurants in both sources based on name, address and city values. External procedures *p.GetName* and *p.GetAddress* are used to convert name and address values, respectively. A global ID table is established to record matching instances.

For all the matching pairs recorded in the ID table, a global object o is generated by obtaining (1) Name value from *fodors.name*, (2) Address value from *focus.address*, (3) Location value from *fodors.location*, (4) City value from *fodors.city*, (5) Phone values from unioning *fodors.phone* (after converted by external procedure *p.GetPhone*) and *focus.phone*, etc.

For an object o_1 (or o_2) that is listed only in *fodors* (or *focus*), a global object o is generated from o_1 (or o_2) attribute values alone using proper external procedures or conversion functions. For objects that exist only in *focus*, Location and Star values are considered as *unknown*. (Therefore, subcondition Star = "Y" will be evaluated to *maybe*.)

Evaluate the query condition on o using 3-valued logic, report o to the user if it is evaluated to *true* or *maybe*. Report the status as well. \square

Because query conditions are evaluated on the merged results, queries are not limited to conjunctive queries. This approach is also easy to implement. Moreover, if several retrieval queries are issued in the same session, then the global ID table only needs to be constructed once. However, this approach has several disadvantages. First, it may not always be possible to store a global ID table due to resource limitations. Also, this approach cannot utilize the existing query optimization techniques, and therefore can be very slow when processing large amount of data. Moreover, condition evaluation is performed after merging, and it is required to perform object identification on all instances instead of only on qualified instances.

5 Push-Down Query Processing

In a multidatabase environment with mismatched domains, incomplete information and conflicting data, query optimization by pushing down applicable selection conditions can have unexpected side-effects [1]. We first describe in this section problems with condition push-down, and then present an algorithm for push-down query processing.

5.1 Problems with Condition Push-Down

Assume a user is interested in all the inexpensive restaurants in the Bay Area. Suppose the given query is processed by "pushing down" the selection condition (i.e., price is inexpensive) to subqueries against the two information sources *fodors* and *focus*, and then merging results returned by the two subqueries. Further suppose a restaurant named "Moonlight Cafe" is returned by *focus* but not by *fodors*. It seems that Moonlight Cafe should be included in the query answer. However, if we examine this problem more closely, then we will find out there are actually two cases: (i) Moonlight Cafe is not listed in *fodors*, or (ii) Moonlight Cafe is listed in *fodors* as an expensive restaurant. Since prefer-merge operation is to be applied on Price, Moonlight Cafe should be included in the answer in case (i), and should not be included in the answer in case (ii). It is the same if intersect-merge is to be applied.

This problem occurs because in general only objects satisfying query conditions are returned to the users. "Maybe" instances can also be returned to the users if *select+* is used. But instances that are known to be false are not returned. Therefore, a user cannot distinguish the above two cases by simply looking at the returned results. This problem will not occur if the merge query processing strategy is adopted, because all instances from local sources are returned for condition evaluation, and therefore case (ii) will

not happen.

Also, set operations such as in or contains cannot always be pushed down. Suppose a user is interested in restaurants satisfying the following condition: Cuisine contains { "Spanish", "Mexican" }. Assume a restaurant Sunrise Taqueria is listed as a Spanish restaurant in one source, and is listed as a Mexican restaurant in the other. By pushing down the selection condition, Sunrise Taqueria will not be selected by either source, although it should be included in the query answer because Cuisine is defined as the union of cuisine types in both sources.

Moreover, if union-merge operation is to be applied on any attributes, even when we are sure that some objects should be included in the query answer, it is still necessary to check corresponding information for these objects in both sources in order to perform union-merge on certain attributes.

We can learn from the above examples:

1. Certain subconditions (e.g., conditions that involve set operations) may not be pushed down to local information sources. This depends on the attribute definitions.
2. An object should be (partially) "materialized" if it is returned by at least one subqueries. For example, since "Moonlight Cafe" is returned by a subquery against focus, instances in fodors should be examined to see whether there are matching instances. Proper attribute value merging should be conducted if matching instances are found.
3. Because of the various problems (e.g., queries with set comparators, information sources that cannot handle certain types of condition evaluation), condition evaluation should be repeated on all the "materialized" objects.

5.2 Query Processing Algorithm

Using the same query Q as described in the previous section, the basic steps of this query processing strategy are as follows:

1. Add all the view attributes that are associated with match clauses or are in $\{ B_1, \dots, B_n \}$ to the select clause of Q if they are not selected by Q . Let the new query with modified selection clause be Q' . The where clause of Q' is exactly the same as where clause of Q .
2. For a source of V , S_1 (or S_2), prepare a subquery $Q_{s,1}$ (or $Q_{s,2}$):
 - The select clause of $Q_{s,1}$ (or $Q_{s,2}$) contains all attributes of S_1 (or S_2) that serve as local source attributes (with or without conversions) of view attributes in the select clause of Q' .
 - The where clause of $Q_{s,1}$ (or $Q_{s,2}$) is constructed as follows. For each subcondition $cond(B_i)$ ($1 \leq i \leq n$):
 - (a) If attribute C_j of S_1 (or S_2) is a local source attribute of view attribute B_i without using conversions, and C_j is "selectable," then generate a subcondition $cond'$ as listed in Figure 3. If $cond'$ is not generated, then neglect this subcondition; otherwise, if Q' uses select+, C_j allows null values, and $cond'$ is not a nullness testing condition, then append: or C_j

is null to $cond'$. If $cond'$ can be processed by the source database, then add $cond'$ to the subquery; otherwise, neglect this subcondition.

- (b) If $t(C_j)$ is a local attribute source of B_i (where t is a conversion function), and C_j is "selectable," then generate a subcondition: $cond'$: C_j in D_j , where D_j is a set of source values in t whose destination view attribute values satisfy $cond(B_i)$; if Q uses select+, and C_j allows null values, then append: or C_j is null to $cond'$. If $cond'$ can be processed by the source database, then add $cond'$ to where clause of $Q_{s,1}$ (or $Q_{s,2}$); otherwise, neglect this subcondition.
- (c) If $p(C_j)$ is a local attribute source of B_i (where p is an external procedure), then neglect this subcondition.
- (d) If B_i has no local components in S_1 (or S_2), and Q uses select (instead of select+), then stop constructing $Q_{s,1}$ (or $Q_{s,2}$) since subquery $Q_{s,1}$ (or $Q_{s,2}$) should always return null results.

3. If subquery $Q_{s,1}$ is constructed, then execute this subquery. For each object o_1 returned by subquery $Q_{s,1}$, match o_1 with each object of S_2 . If o_1 matches an object o_2 of S_2 , then record o_1 and o_2 as matching objects in a global ID table.

- For each o_1 that has matching instance o_2 , merge o_1 and o_2 to form a global object o based on the view definition.
- For each o_1 that has no matching instances, construct a global object o based on view definition and o_1 attribute values alone.

Evaluate query condition of Q (i.e., where clause of Q) on o :

- If Q uses select, then the standard 2-valued logic is applied, and all objects that are evaluated to true are returned with a status value true.
- If Q uses select+, then 3-valued logic is applied in condition evaluation, and all objects that are not evaluated to false are returned with a status value true or maybe.

If o is evaluated to true (or maybe), then project attributes of o that are included in select clause of the original query Q , include the selected and projected instances in the query answers.

4. If subquery $Q_{s,2}$ is constructed, then execute this subquery. For each object o_2 returned by subquery $Q_{s,2}$, if o_2 is not included in the global ID table, then match o_2 with instances of S_1 using similar procedures as described above for o_1 of $Q_{s,1}$. \diamond

(Example 2) Using the same query (Q0) as in Example 1, two subqueries (Q0".1) and (Q0".2) are issued against sources fodors and focus, respectively:

```
(Q0".1) select name, address, city, location, phone,
          price, cuisine, lunch, dinner, star
          from fodors
          where city = "San Francisco"
```

$cond(B_i)$	$cond'$		
	no-merge or prefer-merge	union-merge	intersect-merge
B_i is null	C_j is null	C_j is null	(none)
B_i is not null	C_j is not null	C_j is not null	C_j is not null
$B_i = v$	$C_j = v$	$C_j = v$	(none)
$B_i \neq v$	$C_j \neq v$	$C_j \neq v$	$C_j \neq v$
$B_i > (\geq, <, \leq) v$	$C_j > (\geq, <, \leq) v$	$C_j > (\geq, <, \leq) v$	$C_j > (\geq, <, \leq) v$
B_i in S	C_j in S	C_j in S	(none)
B_i not in S	C_j not in S	C_j not in S	C_j not in S
B_i contains S	C_j contains S	(none)	C_j contains S
B_i contains v	C_j contains v (C_j multi-valued) / $C_j = v$ (C_j single-valued)	C_j contains v (C_j multi-valued) / $C_j = v$ (C_j single-valued)	C_j contains v (C_j multi-valued) / $C_j = v$ (C_j single-valued)

Figure 3: Condition Push-Down

and cuisine contains "Italian"
and star = "Y";

```
(Q0".2) select name, address, city, phone, price,
         cuisine, lunch, dinner
from focus
where city = "San Francisco"
and cuisine = "Italian"
and ( dinner in { "Sun-Fri", "Tues-Sun", ...,
                 "daily" }
    or dinner is null );
```

Execute the first subquery,⁴ and for each object o_1 returned by this query, match o_1 with instances in focus. If there is a matching object o_2 , then add (keys of) o_1 and o_2 to the global ID table, create a new global instance o by merging o_1 and o_2 based on view definition, and evaluate the condition of Q0 on o using 3-valued logic. If there are no matching instances, then create a new global instance o' from o_1 alone based on view definition, and evaluate query condition on o' .

Also execute subquery (Q0".2). For each returned object o_2 , if o_2 is not recorded in the global ID table, repeat a similar procedure for o_1 on o_2 . \square

6 A Comparison of the Two Approaches

In order to study the performance and validate the query processing algorithms, a prototype was implemented in C++. Currently this prototype can integrate information from Oracle, Sybase and ndbm databases. In this section we only present a comparison between the Oracle and ndbm data sources. Let us now examine the performance of the two query processing approaches presented above. Suppose source S_1 has m instances, and source S_2 has n instances. Assume S_1 is the source with a slower retrieval system.

If the *merge* query processing strategy is adopted, then the m instances in S_1 should be matched with the n instances in S_2 first to identify the same instances. (Recall that the source with a faster retrieval system is used in the "inner loop" for object matching.) An ID table is established, and global objects are constructed for query condition evaluation.

If the *push-down* query processing approach is adopted, and p instances of S_1 and q instances of S_2 are returned by

the subqueries Q_{s1} and Q_{s2} with pushed down conditions, respectively. The p instances returned by Q_{s1} are matched with the n instances in S_2 , and the q instances returned by Q_{s2} are matched with the m instances in S_1 for identifying objects and creating the global ID table. p global objects are generated from instances of S_1 , and are evaluated. The q instances returned by S_2 are checked to see if they are already in the ID table. A new global object is constructed for each instance of S_2 that is not already in the ID table for query evaluation. This approach will require a total of $(pn + qm)$ instance matches. The values p and q are dependent on the query conditions that can be pushed down. However, it is easy to see that if $p \approx m$ or if $q \approx n$, then the push-down query processing strategy is likely to perform more instance matches, and the constructed ID table will have almost the same numbers of entries. Moreover, since the results returned by either subquery need to be matched with objects from the other source, the push-down query processing strategy cannot utilize the optimization technique of using a faster source in the "inner loop" of object matching. In this case, the push-down query processing strategy actually has longer execution time when compared with the merge query processing strategy. There are many factors that can result in $p \approx m$ (or $q \approx n$). One reason is that no query conditions can be pushed down to source S_1 (or to S_2). Another reason is that a data source cannot process all or certain types of query conditions.

The above arguments are backed up by our performance studies. We used two data sources: one is a Unix ndbm database with a local query condition processor (called ndbm+), and the other is an Oracle relational database. The prototype communicates with the two data sources through the application program interfaces. In the following tests, we issued 8 queries with increasing numbers of returned instances in the query results to evaluate the two approaches. Query (Q1) does not return any instances, while query (Q8) returns all instances. The performance time (in *second*) was obtained using Unix function `getrusage`. We identified multiple factors such as: data size, size of overlap, data consistency, database system and buffer sizes, as potential parameters that may affect query processing costs. We conducted extensive experiments to evaluate the relative contributions of each of these parameters. Due to lack of space we present here only a subset of these tests.

data size: Figure 4 shows the comparison of the merge and the push-down approaches with 10,000 objects in each

⁴Note that subcondition on dinner is not included because `fedora` on `ndbm+` cannot process disjunctive conditions.

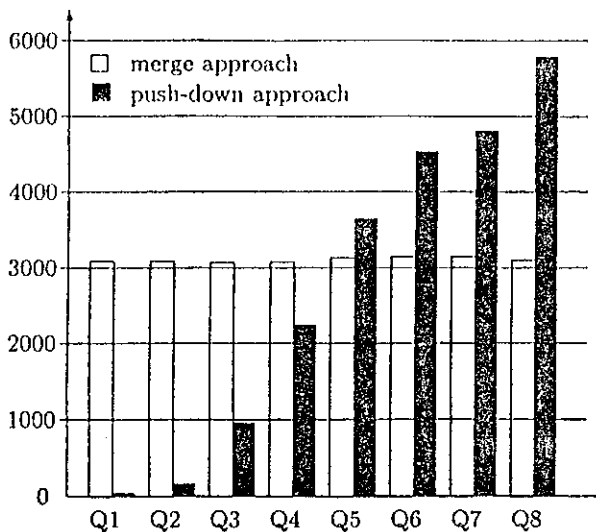


Figure 4: Source 1 (ndbm+ 10,000 objects), Source 2 (Oracle 10,000 objects), 20% consistent overlapping, buffer size = 100




Figure 5: Source 1 (ndbm 10,000 objects), Source 2 (Oracle 10,000 objects), 20% consistent overlapping, buffer size = 100

data sources and with 20% of data are overlapping. In this case, $m = n = 10,000$, and the result shows that the push-down approach performs better when condition push-down removes a large numbers of unqualified instances (i.e., when $p + q < 10,000$ in the first 4 queries), and the merge approach performs better when large numbers of instances satisfying query conditions. The same results hold for smaller data sources (e.g., 1000 objects in each).

overlapping size: In the second test, we use the two data sources with 80% of overlapping data. In this case, we experience shorter execution times because the total numbers of instances are reduced. However, the same results still hold.

data consistency: In another test, we use inconsistent overlapping data in two sources. This test gives similar results as in Figure 4.

database system: We repeat the tests using different database systems (e.g., Sybase, and a file system with local condition processor), and get similar results. However, if we use a pure ndbm without local query processor (i.e., Source S_1 cannot process query conditions, and $p = m$ in all cases), then the push-down approach loses its advantage of removing unqualified instances in S_1 , and thus always performs worse than the merge approach (see Figure 5).

buffer size: Both merge and push-down algorithms can be revised to use buffers to store subquery results. Increasing buffer sizes can reduce the query execution time to a certain extent, but the same conclusions still hold.

7 Related Work

Multidatabase system MRDSM [9] allows defining *dynamic attributes* to resolve unit or semantic differences in local attributes. Each local attribute value is mapped to no more

than one dynamic attribute value. It takes into account that related attributes in two or more local databases may contain inconsistent values. Since MRDSM is a multidatabase language system [2] that does not support a global schema, related attributes are considered as distinct attributes instead of being integrated into one global attribute. Therefore, conflicting resolution is irrelevant.

Inconsistent data in multidatabase systems are handled by specifying *conditional functions* in advance in [4] (and in a later extension [10]). A conditional function of an attribute A defines how A value of an object o is retrieved or computed depending on subtypes to which o belongs. A global query is then partitioned into subqueries for different subsets of objects, and the global query result is the union of subquery results. This approach relies on the existence of global ID values or global ID tables to determine to which subtypes an object belongs.

Unknown attribute values (called *partial values*) are handled in [5]. In order to support partial attribute values, traditional relational operations are extended. This approach again assumes global ID values, and supports only intersect-merge for merging attribute values of the same objects. Conflict resolution is neglected.

Reference [1] extends the work [5] to handle inconsistent non-key attributes. Key attribute values are still required to be consistent. This work introduces the concept of *flexible relations* to capture all possible inconsistent attribute values. A query q against flexible relations is first transformed into related subquery q' to each data source with all conjunctive conditions in q converted into disjunctive conditions in q' . Subqueries results are then merged to evaluate the original query condition in q to determine the set of identities for the result. Identities are then utilized to retrieve information from individual data sources in the second round to get the final query result of q . Our push-down approach differs from this approach in the following aspects: (1) We do not require consistent key or ID values. (2) We have an algorithm to determine which condition can be and should be pushed down to each individual (heterogeneous)

data source. (3) We "materialize" all candidate instances before the final query evaluation. The reasons why we need a different condition push down strategy are because we do not assume consistent identities, and because we have richer object-oriented view constructs and allow set comparisons in query conditions.

An interesting reasoning approach for inter-database instance identification has been proposed in [13]. This approach is then utilized to create a global ID table in [14]. The ID table contains all instances in local sources, which is different from our approach of generating ID tables to contain only matching objects. Conflicting information is resolved using generalization. This approach only focuses on the integration of particular instances, and does not address the query processing issue.

Object-oriented views are used to integrate heterogeneous databases in the ViewSystem [7]. This paper addresses only structure conflicts but not value conflicts. It is proposed that if there are common instances in multiple sources, then these common instances should be materialized and integrated before query processing. But in the case of disjoint instances, query decomposition with condition push-down should be used. No performance studies are conducted.

The TSIMMIS project [6, 11, 12] provides a mediator mechanism for information integration. When integrating similar or equivalent objects from two sources, this approach does not require global ID values or global ID tables. However, when considering objects that appear in single source, it still requires *specified object ids* to identify exported objects [11]. Global views are constructed using a powerful but complicated logic language. Because of the complexity, it is more challenging to optimize queries using this approach when compared with ours. Partial values are not handled by this approach.

8 Conclusions

The contributions of this paper include:

- We proposed an easy and flexible view mechanism for union information integration of multiple independently developed data sources. The approach does not assume the existence of global ID values or ID tables.
- We presented a merge strategy and a condition push-down strategy for query processing in such environment.
- We implemented a prototype and conducted performance studies of the two proposed query processing strategies, and identified factors that affect query performance time. The push-down approach was shown to outperform the merge approach in the case that all data sources can perform certain condition evaluation to filter out a large portions of non-qualified instances.

Future studies include: (1) to incorporate our view mechanism to existing work in object-oriented views, (2) to extend the subset of retrieval queries that can be supported, and (3) to investigate further improvements to the condition push-down query strategy.

References

- [1] Agarwal, S., Keller, A.M., Wiederhold, G., and Saraswat, K. Flexible Relation: An Approach for Integrating Data from Multiple, Possibly Inconsistent Databases, *Proc. 12th Inter. Conf. on Data Engineering*, 1995.
- [2] Bright, M.W., Hurson, A.R., and Pakzad, S.H. A Taxonomy and Current Issues in Multidatabase Systems, *IEEE Computer*, 25, 3, March 1992.
- [3] Connors, T., and Lyngbaek, P. Providing Uniform Access to Heterogeneous Information Bases, *Proc. 2nd Inter. Workshop on Object-Oriented Database Systems*, September 1988. (Also available as *Lecture Notes in Computer Science 334*, Springer-Verlag.)
- [4] Dayal, U., and Hwang, H.-Y. View Definition and Generalization for Database Integration in a Multidatabase System, *IEEE Trans. on Software Engineering*, SE-10, 6, November 1984.
- [5] DeMichiel, L.G. Performing Operations over Mismatched Domains, *Proc. 6th Inter. Conf. on Data Engineering*, 1989.
- [6] Garcia-Molina, H. et al., The TSIMMIS Approach to Mediation: Data Models and Languages (Extended Abstract), NGITS Workshop, 1995.
- [7] Kaul, M., Drosten, K., and Neuhold, E.J., ViewSystem: Integrating Heterogeneous Information Bases by Object-Oriented Views, *Proc. 7th Inter. Conf. on Data Engineering*, 1990.
- [8] Levy, A.Y., Rajaraman, A., and Ordille, J.J. Querying Heterogeneous Information Sources Using Source Descriptions, *Proc. 22nd VLDB Conference*, 1996.
- [9] Litwin, W., and Vigier, Ph. Dynamic Attributes in the Multidatabase System MRDSM, *Proc. 3rd Inter. Conf. on Data Engineering*, 1986.
- [10] Meng, W., and Yu, C. Query Processing in Multidatabase Systems, in *Modern Database Systems: The Object Model, Interoperability, and Beyond*, W. Kim (ed.), ACM Press, New York, New York, 1995.
- [11] Papakonstantinou, Y., and Garcia-Molina, H. Object Fusion in Mediator Systems (Extended Version). Available by anonymous ftp at db.stanford.edu as the file `/pub/papakonstantinou/1995/fusion-extended.ps`.
- [12] Papakonstantinou, Y., Garcia-Molina, H., and Ullman, J. MedMaker: A Mediation System Based on Declarative Specifications, *Proc. 12th Inter. Conf. on Data Engineering*, 1996.
- [13] Wang, Y.R., and Madnick, S.E. The Inter-Database Instance Identification Problem in Integrating Autonomous Systems, *Proc. 6th Inter. Conf. on Data Engineering*, 1989.
- [14] Wang, R., and Madnick, S.E. Facilitating Connectivity in Composite Information Systems, *DATA BASE*, Fall 1989.

國立中山大學八十九學年度碩博士班招生考試試題


科 目：資管所（論文評述第二節）資訊科技類論文

共 頁 第 頁

Please answer the following questions by referencing the attached article, Negotiation Supports in a Commodity Trading Market.

- 1 (10%) What are the major contributions of this paper? Any application limitations?
- 2 (5%) Please describe the differences between the proposed TELCOT system and the other systems mentioned in the paper.
- 3 (15%) Please explain the rationales behind the design of the estimated offer price and the quote change index.

4. If the buyer's target price is lower than a seller's bid price which is lower than the current market quote, the estimated offer price is between the buyer's target price and the quote price. That is, in the case of



the estimated offer price is:

$$offer_{jt} = \min\{target_j, + q_change_t^{1/\alpha} \cdot (quote_t - target_t), bid_i\}.$$

The quote change index is

$$q_change_t = \min\left\{\frac{\max\{e_quote_{t+1} - target_t, 0\}}{quote_t - target_t}, 1\right\}.$$

That is, if the estimated quote at $t + 1$ is higher than the quote at t , then the index is 1; if the estimated quote at $t + 1$ is lower than the buyer's target price, then it is 0; otherwise, it is a ratio of the estimated quote-target difference to the quote-target difference.

- 4 (10%) In section 3.2, a hierarchical constraint satisfaction approach is used for the elicitation of preference on sellers. Please explain the implication denoted at the end of section 3.2, which asserts that the buyer's strategy to lower the target price and negotiation offer prices as an effect does not guarantee better deals.
- 5 (10%) A buyer's target price is the desired price of the buyer want to pay for the product. It is subjective sometimes. Please propose a proper method to assist the buyer in the determination of a reasonable target price and evaluate the proposed approach in terms of its strengths and constraints.

Negotiation Supports in a Commodity Trading Market

Young U. Ryu
Sung-Chul Hong
Decision Sciences Department
The University of Texas at Dallas
JO 4.4, P.O.Box 830688
Richardson, Texas 75083-0688, USA
ryoung@utdallas.edu

Abstract

Advanced electronic trading systems, such as the third phase of the TELCOT system, provide a negotiation mechanism for traders. The purpose of such a negotiation mechanism is to make the trading procedure more flexible. In principle, an electronic trading system establishes an environment in which traders select their partners based on posted trading terms such as product bid/offer prices, product attributes, and trading terms. Most trading systems, including the first and the second phases of the TELCOT system, assume that trading terms are fixed once they are posted by traders. The third phase of the TELCOT system, on the other hand, allows buyers to negotiate prices of sellers' products. As a result, the third phase of the TELCOT system increases the total volume of products traded in the market. However, its negotiation capacity is restricted in that a buyer can extend only one counter-offer to a seller and it does not provide any support functions such as the negotiation price estimation. Assuming the improvement of the third phase of the TELCOT system, we propose negotiation support functions of the elicitation of a buyer's preference on sellers and the estimation of negotiation prices. The preference elicitation is based on hierarchical constraint satisfaction, which is a technique of multi-constraint decision making. The negotiation price estimation function is proposed as a function of a buyer's and seller's price requirements, and the market quote.

Keywords: Commodity Trading, Electronic Marketplace, Hierarchical Constraint Satisfaction, Negotiation Support

1. Introduction

An automated system serving as a trading intermediary in an electronic marketplace reduces traders' costs of searching for trading partners, and this reduction leads to increases in the efficiency of trading transactions (Bakos, 1991). Such cost reduction and increases in the trading

transaction efficiency brought many forms of automated trading systems (Malone, Yates, & Benjamin, 1987), such as the TELCOT system for cotton trading (Lindsey, Cheney, Kasper, & Ives, 1990).

In commodity trading, a trader acquires potential trading partners' information of the commodity and decides trading partners in a certain time limit. In traditional trading of the commodity with heterogeneous attributes, traders negotiate over variable trading terms such as prices, delivery dates, and payment terms. During a negotiation, a buyer and a seller consider attributes of commodity, and bid/offer trading terms to the other. An electronic commodity trading system, replicating the traditional commodity market, should facilitate certain negotiation supporting mechanisms. For instance, the third phase of the TELCOT system provides a primitive form of negotiation environment, through which a buyer can examine a seller's posted information of the commodity price and other attributes, and offer a trading price of the commodity.

An improved form of negotiation environment should provide at least the market information and traders' preference analyses functions (Robinson, 1997). That is, it should help a trader retrieve information on potential trading partners' products and establish a preference order on them, so that the trader can start negotiation with the most preferable trading partner. In addition, a price determination function would be desirable in such a negotiation environment. Yen, Lee, and Bui (1996) propose such a negotiation support environment in a trading system called the intelligent clearinghouse. After considering traders' product and trading constraints, it establishes a trader's preference order on his potential trading partners and proposes a negotiation price in the middle point between the bid and the offer prices. This negotiation procedure is performed from the market-maker's point of view. That is, the market-maker acts as an arbiter reducing the gap of sellers and buyers' conflicting trading requirements.

We propose, in this paper, a negotiation support environment for buyers, which can improve the third phase of the TELCOT system. The negotiation price suggestion mechanism is more sophisticated, which is based on not only buyers' and sellers' price constraints but also the current market quote and the quote changes. A buyer's preference ordering on sellers is considered as a hierarchical constraint satisfaction problem (Wilson & Borning, 1993). The goal of the proposed negotiation support environment is to facilitate buyers with a flexible trading market and to improve the effectiveness of the market.

Oliver (1996) proposes sixteen types of electronic commerce systems, which are classified by factors of negotiation vs. non-negotiation, system autonomy, the number of parties under consideration in trading partner selection, and the number of issues, i.e., attributes considered in trading partner selection. The negotiation support environment of this paper assumes that the trading market allows negotiation of trading terms of price, product delivery terms, etc. The system itself does not make binding offers; it supports traders to conduct negotiation during the process of trading partner selection. That is, the system does not have an automated negotiation capability. Each trader evaluates more than one potential trading partners. Thus, the trading partner selection is a multilateral matching process. Even though we focus on the negotiation of price, we do not exclude the possibility of negotiation of other trading terms. The proposed negotiation support environment, therefore, corresponds to Oliver's category of multilateral and multi-issue negotiation systems without autonomous negotiation decision capabilities.

2. TELCOT and negotiation support systems

TELCOT is a computer-based real-time cotton trading system which is developed by Plains Cotton Cooperative Association (Lindsey et al., 1990). There are three phases of trading provided by the TELCOT system. The first phase with regular offers collects buyers' bids for 15 minutes and selects a highest bidder whose price is not lower than the lowest price requested by sellers. This phase is a replica of floor trading of stocks. The second phase with firm offers improves the first phase by allowing a seller to offer a firm ask price, which is analogous to a sell offer in stock markets. The firm ask price remains until the seller's products are sold or the seller withdraws it. A buyer who accepts the seller's firm ask price first is selected as the trading partner.

The third phase of the TELCOT system supports a counter-offer function, through which a buyer can notify a seller an offer price based on the seller's requested price. Thanks to its negotiation capacity, this counter-offer function has gained popularity among buyers. Its main advantage is the flexibility in trading agreement that increases the

trading volume significantly; that is, a counter-offer may result in trading that otherwise would not occur in the second phase with firm offers (Lindsey et al., 1990). However, its negotiation capacity is limited in that upon a buyer's counter-offer, the seller must either accept or reject, but he is not allowed to propose a further counter-offer to the buyer.

Even though the third phase of TELCOT significantly improves the previous phases by facilitating a real-time negotiation environment, its supporting capacity is very primitive. It provides an interactive system through which buyers can collect market information and propose counter-offer prices to sellers. However, it does not provide any function for buyers to establish preference orders on sellers. Further, it does not support buyers in the determination of counter-offer prices of sellers' products. It is up to individual buyers to analyze market information, select sellers to negotiate with, and determine negotiation prices.

Yen et al.'s (1996) intelligent clearinghouse is a brokerage system that facilitates negotiation functions from the market-maker's point of view. It assumes that both sellers and buyers post their buy and sell orders. With a certain protocol, it matches traders based on their posted trading requirements such as bid and offer prices. When there are traders who fail to find their trading partners, the clearinghouse determines pairs of traders if some minimal changes in their trading requirements would satisfy each other. Then, it suggests specific trading terms to both parties so that their trading requirements would agree. For instance, when the discrepancy between the sell and buy orders is due to their bid and ask price difference, it proposes a negotiation price at the middle point between the two prices.

The intelligent clearinghouse has a market structure that is different from the third phase of the TELCOT system. The TELCOT system acts as a marketplace-provider. It is mainly a buyer who finds his trading partners. If the buyer does not agree on a specific seller's posted price, the system provides a channel for the buyer's counter-offer and the seller's response. However, the intelligent clearinghouse acts as a match-maker who determines trading partners of buyers and sellers if their trading terms are agreeable. Otherwise, it arbitrates between buyers and sellers in their discrepancy in trading terms.

The third phase of the TELCOT system and the intelligent clearinghouse provide negotiation support environments that establish on-line channels for negotiation or supply and analyze negotiation information. Their main goals are not the automation of negotiation, but the increase of trading effectiveness or productivity by supporting humans in contracting and trading situations.

There is another stream of research on negotiation support systems in electronic commerce, which aims at automated negotiation, corresponding to Oliver's (1996) elec-

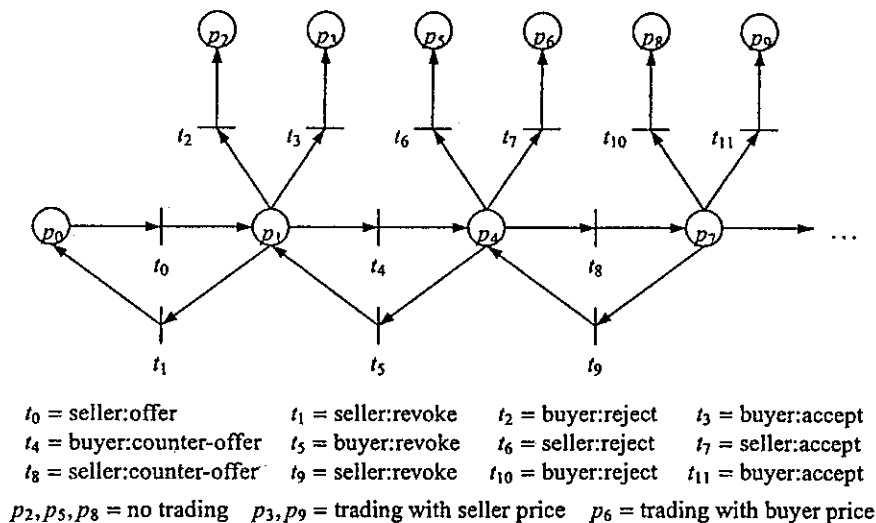


Figure 1. Trading negotiation procedure

tronic commerce systems with intelligent autonomous negotiation agents. Beam and Segev (1997) further classify them into two types: one with learning capabilities and the other without learning capabilities. An automated negotiation support system without learning capabilities has intelligent artificial negotiation agents which have a knowledge-base with negotiation strategies programmed by human negotiation experts. In an automated negotiation support system with learning capabilities, intelligent artificial negotiation agents have machining learning capabilities, often without initial knowledge of negotiation strategies. They build negotiation strategy knowledge through training or actual negotiation sessions.

Chavez and Maes's (1996) Kasbah is an agent-based negotiation environment for trading without learning capabilities, which autonomously performs negotiation and determines trading partners based on buyers' and sellers' negotiation strategies programmed to it. Dworman, Kimbrough, and Laing (1996) and Oliver (1997) present more advanced intelligent autonomous negotiation agents based on genetic algorithms with learning capabilities.

In the following section, we propose a negotiation support environment in a TELCOT-like trading market. We do not purport to develop an automated negotiation support system like Dworman et al.'s (1996) and Oliver's (1997) systems; instead, we propose a few functions, e.g., negotiation price determination and evaluation of potential trading partners, crucial for negotiations in electronic marketplaces, which are similar to two of the three main features of the Kasbah negotiation agent system. The proposed negotiation support environment corresponds to a multilateral negotiation support system without an autonomous negotiation

decision capability in Oliver's (1996) framework for electronic commerce systems. Nevertheless, we expect that such additions to systems like TELCOT will contribute to the increase in trading effectiveness.

3. Negotiation support functions

In traditional trading, negotiation starts with a seller's offer to sell his products under certain trading terms such as a price. Upon a seller's offer, a buyer either accepts the offer, rejects the offer, or proposes a counter-offer, unless the seller revokes the offer; upon the buyer's counter-offer, the seller either accepts the counter-offer, rejects the counter-offer, or proposes a further counter-offer, unless the buyer revokes the counter-offer; and so on. This typical trading negotiation is illustrated in the Petri net of Figure 1 (in which circle nodes labeled by p_k denote states of affairs, or trading situations at specific moments, bar nodes labeled by t_l denote traders' actions and responses that change trading situations, and directed arcs denote pre- and post-conditions of actions.)

In a commodity trading market, sellers' initial offers are made by posting their price requirements and commodity attributes. If we consider the trading negotiation procedure of Figure 1 in a commodity trading market, p_0 is the initial market situation to sellers who enter the market. From p_0 , a new state p_1 is reached when a seller posts his product and trading attributes (t_0). The state p_1 is the initial market situation to a buyer who enters the market. At that situation, the public information on sellers' products is available to the buyer. The negotiation support environment, then, estimates the buyer's counter-offer prices (i.e., negotiation

offer prices) of outstanding sellers and establishes a preference order on sellers' products. Note that the estimation of negotiation offer prices is performed before the preference ordering on sellers' products. That is, in the preference elicitation, the negotiation support environment uses the estimated negotiation offer prices, instead of sellers' posted prices of products.

As the result of the preference ordering, the buyer selects sellers for trading or negotiations. At p_1 , the buyer can take one of three actions for each of the outstanding sellers, unless the seller has revoked his initial offer; e.g., to reject a specific seller's terms (t_2) by not responding to the seller if the seller is not selected, to accept a specific seller's terms (t_3) if the seller is selected and his offer price is acceptable without further negotiation, or to extend an offer to buy to the seller (t_4) with terms and an offer price that the buyer determines based on the suggested negotiation offer price, if the seller is selected and his offer price is to be negotiated. When the buyer extends an offer to the seller (p_4), now, it is up to the seller who can either reject the buyer's counter-offer (t_6), accept it (t_7), or extend a further counter-offer to the buyer (t_8), unless the buyer has revoked his offer to the seller. Upon the seller's counter-offer (p_7), the buyer has the same options to respond.

In the course of the trading negotiation procedure in a trading market, the proposed negotiation environment provides two functions: (1) the estimation or suggestion of negotiation offer prices and (2) the elicitation of preference on sellers' products from sellers' product and trading attributes and a buyer's trading constraints. In the following sections, these two functions are described in detail.

3.1. Estimation of negotiation offer prices


Yen et al.'s (1996) intelligent clearinghouse system estimates a negotiation price in the middle point between a seller's bid price and a buyer's offer price, when the offer price is lower than the bid price. Chavez and Maes's (1996) Kasbah negotiation agent system has a more sophisticated negotiation price estimation function, which is based on buyers' and sellers' range of trading prices, that is, those between a desired price and a highest (or lowest) acceptable price. Still, the negotiation price estimation is a function of buyers' and sellers' prices. In this paper, however, we view that a negotiation price ($offer_{ij}$) offered by buyer j to seller i at time t depends not only on the seller i 's bid price (bid_i) and the buyer j 's target price ($target_j$) but also the market quote at t ($quote_t$) and the quote change (q_change_t).

Buyer j 's target price is a price of commodity that j subjectively considers is reasonable. It corresponds to a buyer's desired price (i.e., "what the user wants to pay for the good") of the Kasbah negotiation agent system (Chavez & Maes, 1996). On the other hand, a buyer's price constraint, which


is used in the elicitation of preference on sellers, contains a price corresponding to the buyer's highest acceptable price (i.e., "the highest price the user is willing to pay for the good") of the Kasbah negotiation agent system. The major difference between the target price and the desired price of Kasbah is: the target price is one of factors for the estimation of the negotiation offer price, which is to be close to the target price, on the other hand, the desired price is the lower boundary (and the highest acceptable price the upper boundary) of the negotiation offer prices. That is, the target price is not necessarily the lower boundary of the negotiation offer prices.

The quote change is expressed by an index between 0 and 1, measuring the historical movement of quotes in the market. It is determined from each seller's bid price, the buyer's target price, and the estimated quote in the future (e_quote_{t+1}) obtained through a time series analysis on past quotes. The details of the estimation of negotiation offer prices are as follows.

1. If a seller's bid price is lower than both the buyer's target price and the current market quote, then the estimated negotiation price is the seller's bid price. That is, in the case of



or




where the 'x' mark indicates an approximate position for the estimated price, the estimated offer price is:

$$offer_{ij} = bid_i.$$

Note that a negotiation offer price must not be suggested higher than the seller's bid price, because we can hardly imagine any trading situation in which a buyer pays more than what a seller is asking for. If the suggested negotiation offer price equals to the seller's bid price, there is no need for negotiation; the buyer can buy the seller's product at the seller's bid price.

2. If a seller's bid price is lower than a buyer's target price but higher than the current market quote, then the estimated offer price is between the current quote and the seller's bid price. That is, in the case of



the estimated offer price is:


$$offer_{ij} = quote_t + q_change_t^{1/\alpha} \cdot (bid_i - quote_t).$$

The quote change index is

$$q_change_t = \min \left\{ \frac{\max\{e_quote_{t+1} - quote_t, 0\}}{bid_i - quote_t}, 1 \right\}.$$

That is, if the estimated quote at t_1 is higher than the seller's bid price, then the index is 1; if the estimated quote at t_1 is lower than the quote at t then it is 0; otherwise, it is a ratio of the quote difference to the quote-bid difference. The parameter $\alpha > 0$ measures the buyer's urgency.

3. If the buyer's target price is higher than the current quote but lower than a seller's bid, then the estimated offer price is between the current quote and the buyer's target price. That is, in the case of



the estimated offer price is:


$$offer_{jt} = quote_t + q_change_t^{1/\alpha} \cdot (target_j - quote_t).$$

The quote change index is

$$q_change_t = \min \left\{ \frac{\max\{e_quote_{t+1} - quote_t, 0\}}{target_j - quote_t}, 1 \right\}.$$

That is, if the estimated quote at $t + 1$ is higher than the buyer's target price, then the index is 1; if the estimated quote at $t + 1$ is lower than the quote at t then it is 0; otherwise, it is a ratio of the quote difference to the quote-target difference.

4. If the buyer's target price is lower than a seller's bid price which is lower than the current market quote, the estimated offer price is between the buyer's target price and the quote price. That is, in the case of



the estimated offer price is:


$$offer_{jt} = \min\{target_j + q_change_t^{1/\alpha} \cdot (quote_t - target_j), bid_i\}.$$

The quote change index is

$$q_change_t = \min \left\{ \frac{\max\{e_quote_{t+1} - target_j, 0\}}{quote_t - target_j}, 1 \right\}.$$

That is, if the estimated quote at $t + 1$ is higher than the quote at t , then the index is 1; if the estimated quote at $t + 1$ is lower than the buyer's target price, then it is 0; otherwise, it is a ratio of the estimated quote-target difference to the quote-target difference.

5. If the buyer's target price is lower than the current quote which is lower than a seller's bid price, then the estimated offer price is between the buyer's target price and the seller's bid price. More specifically, if the current quote is rising, it is between the quote and the seller's bid price; if the current quote is dropping, it is between the target price and the quote. That is, in the case of



the estimated offer price is

$$offer_{jt} = target_j + (2 \cdot q_change_t)^{1/\alpha} \cdot (quote_t - target_j),$$

if $q_change_t < 0.5$; otherwise,

$$offer_{jt} = quote_t + [2 \cdot (q_change_t - 0.5)]^{1/\alpha} \cdot (bid_i - quote_t).$$

The quote change index is

$$q_change_t = 0.5 \cdot \frac{\max\{e_quote_{t+1} - target_j, 0\}}{quote_t - target_j},$$

if $e_quote_{t+1} \leq quote_t$; otherwise,

$$q_change_t = 0.5 \cdot \min \left\{ \frac{e_quote_{t+1} - quote_t}{bid_i - quote_t}, 1 \right\} + 0.5.$$

When the estimated quote at $t + 1$ is higher than the quote at t : if the estimated quote at $t + 1$ is lower than the buyer's target price, then the index is 0; otherwise, it is a ratio of the estimated quote-target difference to the quote-target difference. When the estimated quote at $t + 1$ is higher than the quote at t : if the estimated quote at $t + 1$ is higher than the seller's bid price, then the index is 1; otherwise, it is a ratio of the quote difference to the quote-target difference.

As noted previously, the future quote, a key factor in the determination of negotiation prices, is to be estimated through a time series analysis. Since we are estimating the quote in the near future, it is desirable to apply techniques of exponential smoothing in order to place more weights on recent quote data. Among the components of the time series analysis, the irregular component may contain not only the well-behaved white noise but also the ill-behaved pink noise. In the time series analysis, the pink noise, which typically has big deviations, must be excluded. However, due to the effect of the pink noise, the current quote may be severely overstated or understated. Thus, for $quote_t$ in

the above methods to determine negotiation prices, instead of using the actual quote at t , we may use the quote value at t determined by the curve obtained as the result of the time series analysis. When the deviation of the white noise is substantially big, it will also eliminate the effect of the white noise.

3.2. Elicitation of preference on sellers

Assume a set of sellers, in which seller i is characterized by n product attributes $a_{i,1}, a_{i,2}, \dots, a_{i,n}$ posted in the trading market or determined by the previous negotiation price estimation methods. Let x_k be a variable for a product attribute over the domain of D_k . A buyer's arithmetic constraint is an expression of ' $x_k = d$ ', ' $x_k \leq d$ ', or ' $x_k \geq d$ ', where x_k is over a numeric domain D_k . The function $e(c, i)$ for a constraint c and a seller i measures the degree of the satisfaction of the constraint c by attributes of the seller i 's product. That is, $e(x_k = d, i) = |d - a_{i,k}|$, $e(x_k \leq d, i) = \max\{a_{i,k} - d, 0\}$, and $e(x_k \geq d, i) = \max\{d - a_{i,k}, 0\}$. A Boolean constraint is either $x_k = d$ or $x_k \neq d$, where x_k is over a non-numeric domain D_k . A Boolean constraint is evaluated in such a way that $e(d_k = d, i) = 0$ if $a_{i,j} = d$; otherwise, $e(d_k = d, i) = 1$; and $e(d_k \neq d, i) = 1 - e(d_k = d, i)$.

Suppose the buyer has a set C of constraints that are hierarchically organized by the importance of constraint satisfaction (Wilson & Borning, 1993). That is, the set C of constraint is partitioned to $\Pi_C = \{\pi_1, \pi_2, \dots, \pi_m\}$, such that the satisfaction of constraints in π_k is more important than the satisfaction of constraints in π_l if and only if $k < l$.

It is required that the buyer specifies a price constraint. The product constraint has an upper-bound price, which the buyer is willing to pay for products. In general, it is not lower than the target price that was used to estimate the negotiation offer price in the previous section. The price constraint is on a negotiation offer price, not on a seller's bid price, because the actual trading price is negotiable. Thus, a buyer's preference ordering on sellers is in part based on the potential trading prices, not on sellers' bids.

If the buyer's constraints are linearly ordered, that is, every π_k is a singleton set, we can use a simple pairwise lexicographic comparison (Wright, 1975) to determine the preference on sellers. That is, a seller i is preferred to another seller i' if for some $k \leq m$,

$$e(c_k, i) < e(c_k, i') \text{ and} \\ e(c_l, i) = e(c_l, i') \text{ for every } l < k.$$

In general, however, we must have an aggregation operator (Wilson & Borning, 1993) or utility function (French, 1988; Roberts, 1979) σ_k for each π_k in order to define:

$$E(\pi_k, i) = \sigma_k(\{e(c_1, i), e(c_2, i), \dots, e(c_p, i)\}),$$

where $\pi_k = \{c_1, c_2, \dots, c_p\}$. Then, a seller i is preferred to another seller i' if for some $k \leq m$,

$$E(\pi_k, i) < E(\pi_k, i') \text{ and} \\ E(\pi_l, i) = E(\pi_l, i') \text{ for every } l < k.$$

In the initial elicitation of the preference on sellers, the estimated negotiation offer price for each seller (or an offer price that the buyer determines based on the estimated negotiation offer price) is evaluated against the buyer's price constraint. Upon a seller's counter-offer, the buyer has two options. First, he can consider the seller's counter-offer price as a firm price and re-elicite the preference on sellers based on the evaluation of the counter-offer price against the buyer's price constraint. Second, he can increase the target price for the seller, re-estimate the negotiation offer price for the seller, and then re-elicite the preference on sellers. When the negotiation offer price was higher than the original target price, the new target price for the seller would be between the negotiation offer price and the original target price.

The above negotiation offer price estimation and preference elicitation methods restrict the use of a seller's simple strategy of bidding as high as possible and a buyer's corresponding strategy of setting the target price and the highest acceptable price as low as possible. With all other factors remaining the same, the estimated negotiation offer price for a seller will not be lower than those for sellers with lower bid prices. This does not support sellers' strategy of inflating bid prices. However, if the negotiation offer price for a seller is estimated higher due to his high bid price, he will be evaluated lower in a buyer's preference ordering on sellers. Thus, he has smaller change to receive offers from buyers.

With all other factors remaining the same, a buyer with a lower target price will estimated negotiation offer prices that are not higher than those with a higher target price. Since it is a multilateral trading situation and thus there is a competition among buyers, however, with low negotiation offer prices, his chance to succeed in transaction becomes low. This implies that the buyer's strategy to lower the target prices and negotiation offer prices as an effect does not guarantee better deals.

3.3. An example

Assume that sellers are to post the product price, fiber length, and color in a cotton trading market. The color attribute is measured by a real number between 0 and 1 denoting the gray scale based on the market's color specification standards. Suppose there exist three sell orders: $s_1 = (\$30.20, 12, 0.5)$, $s_2 = (\$31.50, 13, 0.7)$, and $s_3 = (\$31.30, 11, 0.6)$. Let the buyer's target price be

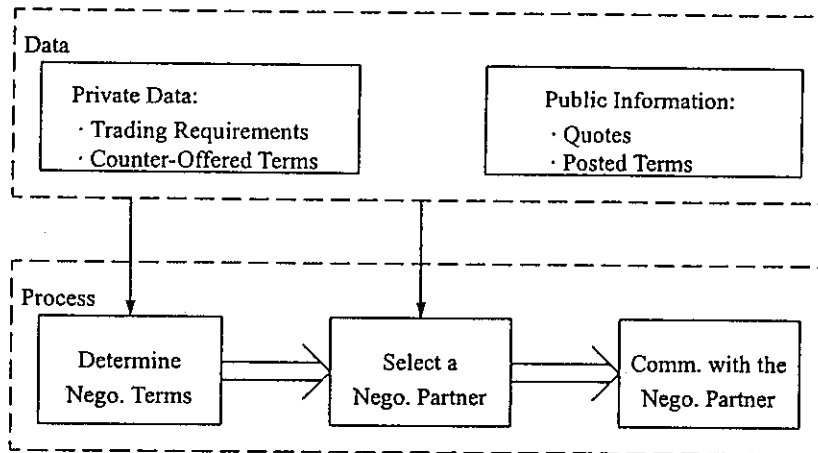


Figure 2. Negotiation Support Functions

\$29.15, the current quote \$31.10, and the quote change index 0.6. Further assume the buyer urgency index is 1. Then, the estimated negotiation offer prices are

$$\begin{aligned} o_1 &= \min\{29.15 + 0.6 \cdot (31.10 - 30.20), 30.20\} \\ &= \min\{29.42, 30.20\} \\ &= 29.69, \end{aligned}$$

$$\begin{aligned} o_2 &= 31.10 + 2 \cdot (0.6 - 0.5) \cdot (31.50 - 31.10) \\ &= 31.18, \end{aligned}$$

and

$$\begin{aligned} o_3 &= 31.10 + 2 \cdot (0.6 - 0.5) \cdot (31.30 - 31.10) \\ &= 31.14, \end{aligned}$$

respectively.

Suppose the buyer has a price constraint of ' $x_1 \leq 31.15$ ', which is the most important, a color constraint of ' $x_3 \geq 0.6$ ', and a fiber length constraint of ' $x_2 \geq 13$ ', which is the least important. By the price constraint, s_1 and s_3 are preferred to s_2 . By the color constraint, s_3 is preferred to s_1 . Thus, s_3 is selected as the seller to be negotiated with. Upon the buyer's offer to buy s_3 's product at the price of \$31.14, suppose s_3 extends a counter-offer with the price of \$31.16. The buyer considers it as the seller's firm price. Now, among the three sellers, s_1 is the most preferable. Thus, the buyer extends an offer to buy s_1 's product at the price of \$29.69. Upon the seller's acceptance of the price, the buyer's search for his trading partner succeeds.

3.4. Extensions

So far, we have presented a negotiation support environment for buyers with functions of negotiation offer price estimation and elicitation of preference on sellers. In this

section, we sketch a few on-going research subjects, which extend the proposed negotiation support environment.

The buyer's target price is one of factors governing the estimation of negotiation offer prices. We viewed that the target price is a price that the buyer subjectively considers as a fair price of sellers' products, corresponding to a buyer's desired price of the Kasbah negotiation agent system (Chavez & Maes, 1996). If the target price is unreasonably too low, the estimated negotiation offer price may be also very low and the buyer has a bigger chance to fail the negotiation. Thus, it would be desirable for the negotiation support function to assist the buyer in the determination of a reasonable target price. This is an issue that will be addressed in the future extension of the paper.

In the elicitation of preference on sellers, a buyer's negotiation offer prices, instead of sellers' posted or counter-offered prices, are considered. One drawback is that, in principle, negotiation prices for all sellers must be estimated, and it will cause a computational overhead. However, certain heuristics can reduce the number of sellers for whom the negotiation offer prices are to be estimated. For instance, if a seller has non-price requirements or constraints that must be satisfied, their satisfaction will reduce the number of sellers under negotiation considerations. We are currently investigating such heuristics.

In addition to the estimation of negotiation offer prices, there is another important issue, which buyers might expect the negotiation support environment to help determine. Intuitively, if the market quotes are expected to decrease continuously, a buyer may postpone negotiation. However, since the trading market under consideration is a multi-lateral, instead of bilateral, commodity market, buyers are competing. This implies that a buyer looking for some specific buyers' products cannot wait for a long time. That is, if

a buyer waits more, he may have a better deal; but his chance to have the deal decreases. Another reason supporting our conjecture of the decreasing chance of the buyer's successful trading is that sellers may want to sell their products as quickly as possible when the market quotes are expected to decrease continuously.

If the buyer knows the probabilities of successful trading negotiation with specific negotiation offer prices at specific negotiation time points, he may determine an optimal negotiation time. As such information is not easily available, we might build a model with the probabilities that sellers are available at specific time points. That is, an approximate model to determine the negotiation time point t^* for the most preferable seller would be $p_t \cdot (\bar{\pi} - \pi_t) = \max_{t \in T} \{p_t \cdot (\bar{\pi} - \pi_t)\}$, where p_t is the probability that the seller is available at t , $\bar{\pi}$ is the upper-bound price that the buyer is willing to pay, π_t is the offer price for the negotiation at t , and T is a time span under consideration.

4. Concluding remarks

Negotiation support functions in a trading market can increase the total volume of the traded products by reducing the discrepancy in trading requirements of buyers and sellers who otherwise would not successfully find their trading partners. The third phase of the TELCOT system, aiming at this purpose, introduced a form of negotiation environment. However, its negotiation capacity is restricted in that a buyer can extend only one counter-offer to a seller and it does not provide any support functions such as the negotiation price estimation. Assuming the improvement of the third phase of the TELCOT system, we proposed a negotiation support environment, which corresponds to a multilateral and multi-issue negotiation support system without an autonomous or automated negotiation decision capability in Oliver's (1996) framework for electronic commerce systems, even though only the issue of price was explicitly addressed in the paper. Features the proposed negotiation support environment are summarized in Figure 2.

It provides process functions of the estimation of negotiation prices, the selection of a negotiation party through the elicitation of a buyer's preference on sellers, and the on-line communication with the selected negotiation party, corresponding to the key negotiation activities of the Kasbah negotiation agent system (Chavez & Maes, 1996). Among these three functions, first two are addressed in the paper. The negotiation support functions utilize the publicly available information such as the price quotes in the market and sellers' posted trading terms including product attributes and prices, and individual buyers' private data. Individual buyers' private data include their trading requirements or constraints and trading terms counter-offered by sellers, which override the posted terms. The second type of pri-

ate data occurs in the course of a buyer's negotiation with a seller. For instance, suppose that upon a seller's posted price of \$31.30, a buyer negotiates with the seller with a price of \$31.14; next, the seller counter-offers with a price of \$31.20. Then, the seller's counter-offer price of \$31.20 becomes the buyer's private datum, which overrides the seller's originally posed price of \$31.30 in the further negotiation.

The negotiation price estimation function is proposed as a function of a buyer's and a seller's price requirements, and the market quote. In the estimation of negotiation offer prices, the market quote change plays an important role. For instance, when the buyer's target price is lower than the quote but the seller's bid price is higher than the quote, if the market quote is decreasing, the negotiation offer price is estimated to be between the buyer's target price and the current quote; if the market quote is increasing, it is estimated to be between the current quote and the seller's bid price. The preference elicitation is viewed as a problem of hierarchical constraint satisfaction, which is a technique of multi-constraint decision making.

In order to verify the usability of the proposed negotiation support environment, we plan to perform simulations on three different settings: the second phase of the TELCOT system with firm (i.e., non-negotiable) offers of sellers, the third phase of the TELCOT system with one round of buyers' counter-offers, and the proposed negotiation support environment allowing interactive negotiation sessions between buyers and sellers. The goal of the simulations is to show that the proposed negotiation support environment increases the total volume of products traded in the market. That is, the use of the negotiation support environment would help traders find their trading partners; but without it they would fail to trade.

References

- Bakos, J. Y. (1991). A strategic analysis of electronic marketplaces. *MIS Quarterly*, 15(3), 295-310.
- Beam, C., & Segev, A. (1997). Automated negotiations: A survey of the state of the art. Tech. rep. 97-WP-1022, Fisher Center for Information Technology & Management, Walter A. Haas School of Business, University of California, Berkeley.
- Chavez, A., & Maes, P. (1996). Kasbah: An agent marketplace for buying and selling goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London, UK.
- Dworkin, G. O., Kimbrough, S. O., & Laing, J. D. (1996). On automated discovery of models using genetic programming: Bargaining in a three-agent coalition game. *Journal of Management Information Systems*, 12(3), 97-125.
- French, S. F. (1988). *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood, West Sussex.
- Lindsey, D., Cheney, P. H., Kasper, G. M., & Ives, B. (1990). TELCOT: An application of information technology for com-

- petitive advantage in the cotton industry. *MIS Quarterly*, 14(4), 347-357.
- Malone, T. W., Yates, J., & Benjamin, R. I. (1987). Electronic markets and electronic hierarchies. *Communications of the ACM*, 30(6), 484-497.
- Oliver, J. R. (1996). *On Artificial Agents for Negotiation in Electronic Commerce*. Ph.D. thesis, Operations and Information Management Department, The Wharton School of the University of Pennsylvania.
- Oliver, J. R. (1997). A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3), 83-112.
- Roberts, F. S. (1979). *Measurement Theory with Applications to Decision Making, Utility, and the Social Sciences*, Vol. 7 of *Encyclopedia of Mathematics and Its Applications*. Addison-Wesley, Reading, MA.
- Robinson, W. N. (1997). Electronic brokering for assisted contracting of software applets. In *Proceedings of the Thirtieth Annual Hawaii International Conference on System Sciences*, Vol. 4, pp. 449-458.
- Wilson, M., & Borning, A. (1993). Hierarchical constraint logic programming. *Journal of Logic Programming*, 16(3/4), 277-318.
- Wright, P. (1975). Consumer choice strategies: Simplifying vs. optimizing. *Journal of Marketing Research*, 12(1), 60-67.
- Yen, J., Lee, H. G., & Bui, T. (1996). Intelligent clearinghouse: Electronic marketplace with computer-mediated negotiation supports. In *Proceedings of the Twenty-Ninth Annual Hawaii International Conference on System Sciences*, Vol. 3, pp. 219-227.