Decentralized Marketplaces and Web Services: Challenges and Opportunities

Itai Yarom, Zvi Topol, Jeffrey S. Rosenschein

School of Engineering and Computer Science
Safra Campus, Givat Ram
Jerusalem, Israel 91904
{jarom,zvit,jeff}@cs.huji.ac.il

Abstract. Centralized marketplaces have a variety of shortcomings: central authorities may possess excessive power in defining and designing the rules of trade, and current developments in peer-to-peer (P2P) and other distributed architectures lacking a central point of control render centralized markets less appropriate. We argue that decentralized marketplaces, although having their own disadvantages, are more suitable to these kinds of distributed environments.
The transition from centralized to decentralized marketplaces will be more likely to occur if participants find themselves *economically* better off in the latter. We envisage that following this transition, market makers will either disappear, or alternatively provide their specialized infrastructure services (e.g., matching between potential transaction partners or providing reputation mechanisms) as standard products in the marketplace.
This paper presents work in progress, which is oriented towards verifying the above arguments with respect to *Web service marketplaces* — marketplaces where Web services are the subject of trade. We believe that the emergence of such marketplaces is imminent, and poses many challenges and opportunities that have not yet been addressed.

1 Introduction

Decentralized marketplaces, where trade is carried out in a distributed manner, have become increasingly popular. Yet most electronic marketplaces today are still centralized, i.e., based on a central market authority that is responsible for designing and controlling the trading environment, setting the underlying rules of the market. Such marketplaces suffer from a variety of shortcomings.

First, traders are constrained to present their products or services and to interact according to rules set by the market maker; negotiation mechanisms and protocols, matching between buyers and sellers, and even additional services such as reputation information about the traders are all defined and controlled by the central market authority. Such control and enforcement mechanisms, of course, do have desirable properties. They provide a well-defined environment, in which commerce takes place in an orderly manner. Yet, in these settings there are real constraints on how traders choose their trading partners and trading strategies (which are limited in the sense that they have to be compliant with the rules of trade). Furthermore, honest traders might not even be allowed to trade when they do not fully satisfy the market's rules (a typical example is

NASDAQ delisting the stock of a company whose market value falls below a predefined threshold).

Second, many centralized marketplaces tend to overlap when it comes to trader audience and products, e.g., eBay, Yahoo Auctions, and Bidz are all general-purpose B2C and C2C marketplaces. Traders who wish to get an overall picture of an entire market are required to participate in separate centralized marketplaces, trading in parallel in all markets. The fragmentation of the marketplace is solely a function of separate, competing central market authorities.

Third, the increasing popularity of highly decentralized and ubiquitous computing environments, e.g., grid computing, Web services, and mobile devices, along with the development of robust and dynamic P2P architectures, leads to the creation of a decentralized, heterogeneous, inter-connected environment lacking a central point of control. Using the model of centralized marketplaces for such environments is an uncomfortable fit.

Fourth, any central point of control will tend to be less fault tolerant than a well-designed, decentralized architecture. Centralized marketplaces may be more vulnerable to malicious behavior, and traders are forced to trust the centralized market maker to be protective and honest. This setting can lead to situations where the centralized authority exploits its excessive power, and behaves dishonestly.

The intrinsic problems of centralized marketplaces, however, are not automatically solved by decentralization. Instead, a whole new series of fundamental questions are raised, as to how to effectively exploit the potential of decentralized markets. How can we mitigate the risks entailed by a distributed environment (for example, how can participants refrain from transactions with dishonest traders)? And what would be an appropriate market infrastructure (e.g., which communication protocols and product specification languages should be used)?

Ongoing research regarding trust and reputation mechanisms, e.g., [28, 29] might provide answers to some of the above questions, while current work on overlay networks, standardization of Web services, and the Semantic Web certainly provides other pieces of the puzzle.

In addition, despite the variety of disadvantages that centralized marketplaces might have in comparison to decentralized markets, it is still not clear that traders themselves would indeed prefer the decentralized alternative. If participants are *economically* better off trading in centralized marketplaces, traders might be reluctant to move to decentralized platforms. If, on the other hand, traders have an incentive to opt for decentralized alternatives, we envisage that market makers will either disappear or provide their infrastructure services, e.g., matching between potential transaction partners or providing reputation mechanisms, as standard products or services in the marketplace.

In this paper, we describe ongoing research focusing on the comparison between centralized and decentralized marketplaces with respect to *Web services*. Web services are of much interest to both academia and industry, and a great deal of work is being carried out in this field in a variety of directions. However, relatively little research has been conducted on *markets of Web services*, where Web services are the objects of trade. Pricing models and trading strategies in Web service marketplaces have not yet been explored, despite the fact that they entail many new challenges and opportunities.

2 Web Service Marketplaces

We here provide a brief review of Web services, and argue that Web service marketplaces are related to information marketplaces, yet pose a set of complex questions and challenges that have not been fully answered by work on information marketplaces.

Web services are machine-friendly software components designed to work interoperably when deployed over heterogeneous computing environments. Interoperability is achieved through a set of standards based on the XML markup language, such as UDDI (Universal Description, Discovery and Integration language) for the discovery of Web services, WSDL (Web Services Description Language) for the definition and description of Web services, and SOAP (Simple Object Access Protocol) for the underlying messaging protocol.

The Semantic Web augments the current Web through the use of *ontologies* that provide meaningful definitions for different information structures on the Web. OWL (Ontology Web Language), an extension of RDF (Resource Description Framework), enables the creation and use of such ontologies, and OWL-S (Ontology Web Language for Services) is a set of ontologies used to describe Web services for the Semantic Web.

On top of these standards, much work has been carried out in the field of Web service composition (e.g., [24, 21, 3]), selection (e.g., [22]), and verification (e.g., [7, 8]). The aim of that research is to ensure the integrity and functionality of Web services, and to provide automatic means for the selection and composition of basic Web services into meaningful, sophisticated Web services capable of executing complex tasks. That work establishes building blocks for the development of rich marketplaces for Web services.

Web service markets are related to information markets [26]. Web services, described with OWL-S for the Semantic Web, could provide, inter alia, meaningful information through XML documents that they output in response to queries. Web service providers could thus be considered as service providers in the sense of information markets, while Web service consumers could be considered as information consumers. Composition of Web services in this context could be seen as providing richer information through manipulation of information supplied by simpler Web services.[1]

Yet, the fact that every Web user is a potential provider and/or consumer of Web services introduces many new issues. There may be large differences between various Web services in terms of reliability and trustworthiness, both in the quality of information provided and in the availability and fault tolerance of the service. How could a service consumer verify that the information provided to him is reliable? What should she do when a Web service fails to deliver? How would trustworthiness and reputation measures affect the price of the service?

Composite Web services pose additional challenges. How should a composite service be priced? What happens if one or more of the services making up the composite service fail? How will this interdependence influence the price of the composite service? Based on which parameters should a service consumer select a composite service?

[1] It should be noted, however, that there is more to Web services than information-supplying, and as part of our future work we intend to explore these issues.

Other important questions center on the nature of relationships between Web service providers and Web service consumers. What types of relations are worthwhile to establish between a consumer and a provider (e.g., ad-hoc or long-term)? Should there be coalitions of consumers?

It is also not clear what the structure of the market would be. Would Web service traders prefer to meet in a centralized marketplace like eBay, or would they prefer to trade in a decentralized manner? Which type of traders will participate in the marketplace? How are different costs to be allocated among various Web service providers that offer parts of a composite Web service? Will composite Web service providers have a legitimate right to exist in the marketplace? What about basic Web service providers?

Clearly, there are many avenues of research in this field, investigating the challenges and questions mentioned above. In the next section, we present our model for a marketplace for Web services, extending the work of Yarom et al. on information marketplaces [26].

3 The Market Model

The infrastructure of the market model used in our research is based on a framework of Web services. This framework handles the technical issues of describing and posting services or products, and locating relevant agents that offer a given service or product. In our model, products or services possess two attributes: Price, and Quality of Service (QoS) of the agent offering the product or service. At each market period, each agent chooses with probability ρ whether to take an action on the following turn or not. Our market model comprises three types of agents:

- *Service Consumer (SC)* — The Service Consumer agent buys information commodities from the Service Provider and from the Composite Service Provider.
- *Service Provider (SP)* — The Service Provider agent sells services to the Service Consumer and to the Composite Service Provider agents.
- *Composite Service Provider (CSP)* — The Composite Service Provider agent buys services from the Service Provider agent and then sells the newly created composite service to the Service Consumer agents.

The Utility of the agents in the market is defined as $U = V - C$, where V is the value function and C is the cost function of the agent. Different types of agents have different types of V and C functions. For example, the V function of the Service Consumer agent is the value of the information it purchased, while the C function is the price it pays for it. On the other hand, the V function of the Service Provider agent is the price it gets for the information it offers, while the C function is the cost of producing this information. For the Service Consumer agent the costs can include payment for information products, payment for information services, the costs of agent discovery (e.g., UDDI directory), and the expected costs of fraudulent transactions.

In our simulations, we used a 'normalized' version of the utility function: $U = \frac{V-C}{\#Trans}$, where $\#Trans$ is the number of transactions of all agents: ($\#Trans = \Sigma_{i \in Agents} Trans_i$).

The following simple example illustrates the concepts and definitions introduced above. Consider an agent interested in receiving a comprehensive review of financial information regarding the companies included in the S&P500 index. This agent might approach a composite Web service in order to purchase such information. The composite Web service will locate a Web service providing formatted data regarding stock exchange rates, and a Web service providing analysis reports regarding the relevant companies, using the UDDI directory or other service location directories. It will then purchase the information from the Web services, and thereafter process the data and return the desired information to the agent. The Web service giving stock exchange rates provides its formatted data to the composite Service through the manipulation of raw data that it acquires from a service provider, e.g., Yahoo finance.

4 Simulations

In this section, we describe simulations comparing centralized and decentralized markets. We first look at a marketplace where the number of Service Consumer agents is significantly larger than the number of Service Provider and Composite Service Provider agents. This setting represents a decentralized marketplace where trade is conducted within clusters of agents. In this configuration we will have two basic information commodities, three types of Service Provider agents (analyzing, translation, and information source services), five Composite Service Provider agents, and one hundred Service Consumer agents.

The second configuration is of a market where all agents can buy, sell, or offer new services. The centralized version of this market is similar to eBay, while the decentralized version of it is similar to Gnutella. In this configuration, the number of sellers and the number of buyers is the same, since all agents can sell and buy services. We have one hundred agents, and each one of them can offer the three types of services mentioned in the first configuration.

4.1 The Agent Pricing and Shopping Algorithms

Each of the Service Consumer, Service Provider, and Composite Service Provider agents will use the pricing and shopping strategies that were introduced by Kephart et al. [12], and which were also used by Yarom et al. [25]. The pricing algorithms are as follows:

1. GT (Game Theory) — Kephart et al. have shown that there is not a single pure strategy that is in Nash equilibrium for sellers to establish the price of a commodity. There is, instead, a mixed strategy that is in Nash equilibrium. This mixed strategy instructs each agent to choose prices randomly using a function $p(F)$, where F is a random value between the cost c of the commodity and its value v (in our case $F \in [0,1]$). $p(F)$ is given by $p(F) = c + \frac{w_1 * (v-c)}{\Sigma_{i=1}^{S} i * w_i * (1-F)^{i-1}}$, where S denotes the number of information suppliers in the market (i.e., both Service Provider and Composite Service Provider agents in our case), and w_i is the fraction of buyers that compare i prices.
2. MY (Myoptimal) — The agent sets the price of the commodity in the market to maximize its short-term profit (i.e., it assumes that the current known market conditions do not change, which is a reasonable assumption in the short-term). This

method requires knowledge about the Service Consumers' population, the number of competing agents, and all of the agents' prices.

3. DF (Deviate Follower) — The agent keeps increasing the price of a commodity as long as its profit increases. The agent decreases the price when its profit drops a certain level. The agent continues decreasing the price as long as its profit increases. When the profit starts to decrease and passes a certain level, then the agent begins to increase the price.

The shopping algorithms used by the agents in the market are the following (the numbers in parentheses represent the percentage of each type of agent in the Service Consumer agent population in our tested market):

1. Compare-All (70%) — Service Consumer agents compare all of the prices requested for the commodity of interest. Then, the Service Consumer agents choose the agents that ask for the lowest price. This algorithm is similar to the implementation of the ShopBot in [12].
2. Compare-None (10%) — Each Service Consumer agent randomly chooses an agent that offers the requested commodity.
3. Compare-two (20%) — Each Service Consumer agent randomly chooses two agents that offer the requested commodity, and then buys from the cheaper one.

4.2 Results

Agents in a decentralized marketplace have incomplete information regarding the other agents. The DF algorithm does not use any data on other agents, and therefore it could be left unchanged in the decentralized case. The MY pricing algorithm, on the other hand, uses data on all other agents for setting a new price. In a decentralized market, this data will be incomplete and available only about several (k) agents at a given moment. We denote MYk as the MY pricing algorithm applied to k agents. In our simulation, sniffing is executed on two agents ($k = 2$).

We first explore the effect of incomplete knowledge in decentralized markets on agents' profit. We expected that the incomplete data of the decentralized market would reduce the agents' profit. However, as can be seen in Table 1, the profit of the Composite Service Provider agents increases in a decentralized market. Incomplete information reduces the agent's awareness of other agents' behavior, and therefore reduces price competition in the market, resulting in higher profits.

Agents' algorithm	Agents' Profit
DF	0.972
GT	1.394
MY	1.412
MY_k	1.476

Table 1. The Composite Service Provider agents' profit

Agents' algorithm	Agents' Profit
DF	0.33
MY	0.37
MY_k	0.41

Table 2. The agents' profit when each agent is SC and CSP agent, in a market with 100 agents.

In the pure decentralized market, there are one hundred agents that can each have the characteristics of Service Consumer, Service Provider, and Composite Service Provider. In this case, the calculations involved in the GT algorithm are too complex, and therefore impractical. Table 2 shows that the various pricing algorithms result in similar profits, while the MYk algorithm gains the highest profit.

The profit of agents in the second scenario are lower, compared to the profit of the Composite Service Provider agents in the first scenario. This is because in the second scenario the ratio between agents that offer information to agents that want to purchase it is 1 (compared to $\frac{1}{20}$ in the first scenario). The increase in competition over customers reduces prices and profits.

5 Related Work

5.1 Decentralized Marketplaces

Some of the first research concerned with P2P agent-based marketplaces is that of Youll [27]. In this work, a framework for an agent-mediated P2P marketplace was suggested and implemented in Java using a centralized trader registry. Searching for potential traders is done using an adaption of the Contract Net protocol [5]. The use of Web Services that are based on standard protocols, e.g., WSDL and SOAP, along with a decentralized registry implemented over a P2P network such as Gnutella (as suggested in [17]) seems a more suitable choice for P2P marketplace implementations. The important issue of traders' trustworthiness was briefly mentioned in Youll's work, but not really handled. Coalitions between traders were not checked. In addition, no simulations or analysis of experimental results were provided to justify why decentralized marketplaces might be better than centralized ones.

Another justification for decentralized marketplaces can be derived from [16], where a P2P-agent double continuous auction was presented. It was shown that such a decentralized auction displays price convergence behavior similar to that of the centralized auction. However, the P2P auction outperforms the centralized auction in the sense that it has a constant cost in the number of message rounds needed to find the market equilibrium price when the number of traders increases, whereas a central auctioneer incurs a linear cost in this case. It would be interesting to extend this work and introduce quality of products or trustworthiness of agents into the simulation.

Turner et al. describe in [23] a P2P resource market where buyers and sellers trade their surplus resources, e.g., CPU cycles, storage, and bandwidth. They suggest that each entity in the market might issue its own currency, and propose the Lightweight Currency Protocol as a standard protocol for the interaction between users and currency issuers. This protocol might be a building block on top of which P2P marketplaces could be implemented.

The theoretical work of Neeman et al. [15] shows that when buyers and sellers are given the opportunity to choose between trading through a centralized market or through a decentralized one, they would both prefer the centralized option. However, they assume in their work that a homogeneous product is traded and that the centralized market incurs no costs. We are interested (by contrast) in the case where the product traded is not necessarily homogeneous; products from different sellers may have different qualities. Furthermore, not all trading partners can be assumed to be trustworthy. This might lead to the emergence of coalitions based on trading relationships, explored for example in [2].

5.2 Information Markets

Information markets can emerge in different contexts, such as in Digital Libraries (e.g., the Stanford Digital Library Project [13] and the University of Michigan Digital Library (UDML) [18]), and in markets for exchanging expert advice (e.g., Kamoon [11]). Kephart et al. [12, 19] explored the dynamic of agents in information marketplaces, while looking only at centralized markets. Our ongoing research removes this assumption and explores the benefits of decentralized markets.

Buyers and sellers in information markets can be represented by autonomous agents or by humans. Das et al. [4] compared the behavior of agents and humans in a CDA e-market. Their experiments reveal that software agents can outperform humans in markets composed of both trading agents and humans. Additional roles that agents can perform in e-markets include pricing agents (e.g., Pricebots [10]), price comparison agents that help buyers to find the seller with the lowest price (e.g., ShopBot [10]), auction bidder agents (e.g., [14]), recommendation agents (such as those at Amazon.com), and broker agents (e.g., [6]).

5.3 Web Services

Much work is being carried out in the field of Web service composition. Web service flow specification languages such as BPEL4WS (Business Process Execution Language for Web Services) enable manual composition of Web services, whereas the Semantic Web community extends classical AI planning techniques so as to tackle the problem of automatic composition of Web services (see, for example, [24, 21, 3]).

Web service selection and composition, based on reliability, reputation, and trust is also a subject for investigation (see, for example, [22, 20]). Current work towards computation and propagation of trust metrics over the Semantic Web, as in [9], will establish the basis for such selection.

Formal verification of Web services, assuring that a given Web service upholds specified properties and that compositions of Web services satisfy the required functionality, is also being studied, using various classic formal verification methods, such as model checking and process calculi (see, for example, [7, 8]).

6 Summary and Future work

In this paper we challenge the traditional structure of centralized markets, and argue that those markets may well evolve into decentralized ones. Our preliminary results show economic benefits for market players to actually prefer this transition. However, decentralized markets have their own disadvantages. We plan to explore those issues and to identify solutions that will resolve them.

One of the important issues that should be dealt with in decentralized markets is the provision of QoS and trust metrics to the market participants, and the influence of these factors on pricing mechanisms and the volume of trade. These issues exist in centralized markets as well, but they will become even more critical in decentralized markets.

In addition, decentralized markets will have to define pricing mechanisms that will handle decentralized transactions better. Will the same pricing mechanisms used in centralized markets be suitable for decentralized markets? What pricing algorithms and strategies will gain the best profits in such markets? We have started to explore those directions with respect to Web service marketplaces, which pose a wide set of challenges and opportunities of their own.

In the transition to electronic marketplaces from classical ones, new roles emerge for the market parties, while old roles disappear [1]. What kind of roles will emerge in the transition from centralized markets to decentralized ones? The centralized market makers of today will have to identify those roles in order to survive in decentralized markets of the future.

References

1. Joseph P. Bialey and Yannis Bakos. An exploratory study of the emerging role of electronic intermediaries. *International Journal of Electronic Commerce*, 1(3):7–20, Spring 1997.
2. Silvia Breban and Julita Vassileva. A coalition formation mechanism based on inter-agent trust relationships. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, Bologna, Italy, July.
3. Mark Carman, Luciano Serafini, and Paolo Traverso. Web service composition as planning. In *Proceedings of the ICAPS'03 Workshop on Planning for Web Services*, 2003.
4. R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 1169–1176, Seattle, Washington, August 2001.
5. Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63–109, 1983.
6. K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan, 1997.
7. Howard Foster, Sebastian Uchitel, Jeff Magee, and Jeff Kramer. Model-based verification of web service compositions. In *Proceedings of the 18th IEEE International Conference on Automated Software Engineering*, pages 152–163, 2003.
8. Howard Foster, Sebastian Uchitel, Jeff Magee, and Jeff Kramer. Compatibility verification for web service choreography. In *Proceedings of the IEEE International Conference on Web Services, ICWS '04*, 2004. To appear.
9. Jennifer Golbeck, Bijan Parsia, and James A. Hendler. Trust networks on the semantic web. In *Proceedings of Cooperative Information Agents VII, 7th International Workshop, CIA 2003*, pages 238–249, 2003.

10. Amy Greenwald and Jeffrey O. Kephart. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 506–511, Stockholm, Sweden, August 1999.
11. Kamoon. http://www.kamoon.com.
12. Jeffrey O. Kephart, James E. Hanson, and Amy R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, May 2000.
13. Steven Ketchpel, Hector Garcia-Molina, and Andreas Paepcke. Shopping models: A flexible architecture for information commerce. In *Proceedings of the Second ACM International Conference on Digital Libraries*, pages 65–74, Philadelphia, Pennsylvania, July 1997.
14. J. Morris, P. Ree, and P. Maes. Sardine: Dynamic seller strategies in an auction marketplace. In *Proceedings of the Conference on Electronic Commerce (EC '00)*, pages 128–134, Minneapolis, Minnesota, October 2000.
15. Zvika Neeman and Nir Vulkan. Markets versus negotiations: the predominance of centralized markets. Working Paper, 2002.
16. Elth Ogston and Stamatis Vassiliadis. A peer-to-peer agent auction. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, Bologna, Italy, July.
17. Massimo Paolucci, Katia P. Sycara, Takuya Nishimura, and Srinivasan Naveen. Using DAML-S for P2P discovery. In *Proceedings of the International Conference on Web Services, ICWS'03*, pages 203–207, June 2003.
18. S. Park, E. H. Durfee, and W. P. Birmingham. Emergent properties of a market-based digital library with strategic agents. *Autonomous Agents and Multi-Agent Systems*, 3(1):33–51, 2000.
19. Jakka Sairamesh and Jeffrey O. Kephart. Price dynamics and quality in information markets. *Decision Support Systems*, 28(1–2):35–47, March 2000.
20. Munindar P. Singh. Trustworthy service composition: Challenges and research questions. In *Proceedings of the Autonomous Agents and Multi-Agent Systems Workshop on Deception, Fraud and Trust in Agent Societies*, Bologna, Italy.
21. Evren Sirin, Bijan Parsia, , and James Hendler. Composition-driven filtering and selection of semantic web services. In *AAAI Spring Symposium on Semantic Web Services*, 2004.
22. Raghuram M. Sreenath and Munindar P. Singh. Agent-based service selection. *Journal on Web Semantics (JWS)*, 1(3):261–279, April 2004.
23. David A. Turner and Keith W. Ross. A lightweight currency paradigm for the P2P resource market. In *Proceedings of the Seventh International Conference on Electronic Commerce Research*, 2004. To appear.
24. Dan Wu, Bijan Parsia, Evren Sirin, James Hendler, and Dana Nau. Automating DAML-S web services composition using SHOP2. In *Proceedings of the 2nd International Semantic Web Conference (ISWC2003)*, pages 195–210, Sanibel Island, Florida, October.
25. Itai Yarom, Claudia V. Goldman, and Jeffrey S. Rosenschein. The impact of infocenters on e-marketplaces. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1290–1291, Bologna, Italy, July 2002.
26. Itai Yarom, Claudia V. Goldman, and Jeffrey S. Rosenschein. The role of middle-agents in electronic commerce. *IEEE Intelligent Systems*, 18(6):15–21, November/December 2003.
27. James E. Youll. Peer to peer transactions in agent-mediated electronic commerce. Master's thesis, Massachusetts Institute of Technology — MIT Media Lab, Cambridge, Massachusetts, September 2001.
28. B. Yu and M. Singh. Distributed reputation management for electronic commerce. *Computational Intelligence*, 18(4):535–549, November 2002.
29. B. Yu and M. Singh. Detecting deception in reputation management. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 73–80, Melbourne, Australia, July 2003.