Computer Science
And Information
Management Program

Asian Institute
Of Technology

AIT-CSIM-TR-26-12-03

Balanced Matching of Buyers and Sellers in E-Marketplaces: The Barter Trade Exchange Model

Peter Haddawy
Namthip Rujikeadkumjorn
CSIM Program
Asian Institute of Technology

Christine Cheng
Dept. of EE&CS
University of Wisconsin-Milwaukee

December 2003

P.O. Box4, Klong Luang, Pathumthani, Thailand, 12120

Balanced Matching of Buyers and Sellers in E-Marketplaces: The Barter Trade Exchange Model

Technical Report #AIT-CSIM-TR-26-12-03

Peter Haddawy
Namthip Rujikeadkumjorn
CSIM Program
Asian Institute of Technology
haddawy@ait.ac.th, Namthip.Rujikeadkumjorn@ait.ac.th

Christine Cheng
Dept. of EE&CS
University of Wisconsin-Milwaukee
ccheng@cs.uwm.edu

December 26, 2003

Abstract

In this paper, we describe the operation of barter trade exchanges by identifying key techniques used by trade brokers to stimulate trade and satisfy member needs, and present algorithms to automate some of these techniques. In particular, we develop algorithms that emulate the practice of trade brokers by matching buyers and sellers in such a way that the volume of goods traded is maximized while the balance of trade is maintained as much as possible. We show that the buyer/seller matching and trade balance problems can be decoupled, permitting efficient solution as well as numerous options for matching strategies.

We model the trade balance problem as a minimum cost circulation problem (MCC) on a network. When the goods have uniform cost or when the goods can be traded in fractional units, we solve the problem exactly using a simplified version of the minimum mean cycle canceling algorithm. Otherwise, we present a novel stochastic rounding algorithm that takes the fractional optimal solution to the trade balance problem and produces a valid integer solution. We then make use of a greedy heuristic that attempts to match buyers and sellers so that the average number of sellers matched to a buyer of a good is minimized.

Finally, we also present results on the empirical evaluation of our algorithms on test problems and simulations. The solutions from our stochastic rounding algorithm are always within 0.7% of the solution obtained from a commercial mixed integer programming package. We evaluate the effectiveness of our algorithm on maintaining balance and on stimulating trade using a simulator built using transaction history data from a trade exchange. The simulation results confirm the barter trade exchange rule of thumb that maximizing single-period trade volume while maintaining balance of trade helps to maximize trade volume over the long run.

1 Introduction

With the movement of business to the Internet, one of the most popular e-commerce models to emerge has been that of the e-marketplace. An e-marketplace is an electronic intermediary that brings buyers and sellers together and provides various support services. These services have traditionally consisted of such things as e-catalogs, search capabilities, and transaction support. More recently researchers have sought to exploit the electronic infrastructure of e-marketplaces and the wealth of information that can be gathered in e-marketplaces to provide sophisticated methods of matching buyers and sellers, using agent-based, auction-based, and broker-based techniques. Of the work in this area that has addressed profitability of the e-marketplace, the overriding concern has been for maximization of single-period revenues [6, 16, 10], with less attention paid to how the techniques fit within a more strategic business model. But as the dramatic down-turn in the e-commerce sector demonstrated, e-business initiatives require solid business models that clearly relate the services provided to the overall profitability of the company [12]. In this paper, we take a particular but quite general e-marketplace business model as our point of departure and use that model to motivate the development of algorithms to support management of trade among buyers and sellers.

The model used in this paper is that of the barter trade exchange, also called retail or commercial barter. A barter trade exchange is a collection of businesses that trade their goods and services, managed by an intermediary. We call the collection of businesses the *barter pool* and call the intermediary the *trade exchange*. The word barter is something of a misnomer since the businesses do not exchange goods as in the synchronous bilateral fashion of traditional barter. Rather, modern barter is multilateral, using a form of private label currency. The trade exchange issues trade dollars to the member businesses and acts as a neutral third party record keeper. When a company sells a good, they receive credit in trade dollars, which they can then use to purchase goods from other members. The value of the trade dollar is tied to the US dollar by not permitting businesses to charge more for their goods in terms of trade dollars than they do in US dollars in the open market, thus preventing devaluation of the currency.

The barter industry is interesting as a test bed for market design because a barter pool is a relatively closed economy about which we have very detailed information due to the book keeping function of the trade exchange. The trade exchange maintains a general profile for every member business, as well as complete records of all transactions between members. A barter pool has many similarities with a traditional economy, with the trade exchange playing a role analogous to that of the federal government in regulating the economy. The exchange controls such variables as monetary supply, interest rate, rate of commission (analogous to revenue tax), and even supply and demand through its ability to selectively recruit new member businesses. Interestingly, although it has control over all these parameters, the trade exchange works to stimulate the barter pool economy primarily by making recommendations to member businesses through trade brokers.

The success and survivability of the barter business add to its attractiveness as a model to study. The barter trade exchange industry has existed for over forty years, surviving numerous changes in the economic landscape. The International Reciprocal Trade Association [3] estimated that the total value of goods and services bartered by businesses through barter companies reached USD 7.87 billion in 2001. This number was an increase from USD 6.92 billion in 1999 and was the third consecutive year the industry saw over 12% growth. There were an estimated 719 trade companies active in North America in 1999 with some 471,000 client businesses [2]. Examples of active barter trade exchanges with a Web presence include BizXchange, ITEX, BarterCard, and Continental Trade Exchange.

The rest of this paper is organized as follows. In Section 2, we provide a description of the operation of barter trade exchanges, identifying key techniques used by trade brokers in order to stimulate trade and satisfy member needs. In this paper we focus on implementing techniques for maximizing trade and maintaining balance of trade within the barter pool. In Section 3 we present a formalization of this problem and in Sections 4 and 5 we present novel efficient algorithms for its solution, using minimum cost circulations on networks and stochastic rounding techniques. In Section 6, we present empirical evaluation of our algorithms. In Section 7 we discuss related work and show that the balance problem is closely related to the multi-unit combinatoral auction problem. In Section 8 we present directions for future work.

2 Barter Trade Exchange Model

Given its important role in B2B commerce, there is a surprising lack of literature on the barter trade exchange industry. An exception is the work of Cresti, which examines theoretical economic rationale for development of the barter industry in industrialized countries [4], as well as investigating the macroeconomic variables influencing the industry in the United States [5]. But there exists no formal literature describing the barter trade exchange industry on an operational level. Our interest lies in understanding how managers and brokers in a trade exchange manage the operations of the exchange in order to maximize their company's profits. Thus our first step in conducting this work was to gather information through extensive interviews with industry experts. We also communicated with them periodically to verify the assumptions behind our models. We interviewed two executives at BizXchange (www.bizx.bz), a relatively new but rapidly growing trade exchange located in the San Francisco Bay and Seattle areas. Since its inception in January 2002, BizXchange has grown to include over 600 member businesses. The two executives we interviewed have over 28 years of combined industry experience, have founded and built several successful barter networks, and have served on the Boards of the International Reciprocal Trade Association and the National Association of Trade Exchanges.

A barter pool can be viewed as a carefully managed micro economy. Managers of trade exchanges attempt to recruit member businesses in such a way that the mixture of supply and demand in the pool is balanced. Member businesses are typically small to medium size enterprises that offer goods and/or services. They fall into the broad categories of operating expenses, employee benefits, and travel and entertainment. Examples of typical businesses in barter trade exchanges include car rental, catering, advertising, office equipment and furniture, office supplies, dental services, health clubs, restaurants, and hotels. Henceforth we will use the term *goods* to refer to goods and services.

It is a common misconception that the primary benefit of barter is to avoid taxes. In fact, the US Tax Equity and Fiscal Responsibility Act, passed in 1982, legislated that barter income be treated as equivalent to cash income and taxed on the same basis. Cresti [5] argues that fact the growth of the US barter industry did not slow down since 1982 is evidence that firms do not engage in barter to avoid taxes. She shows empirical evidence that barter is adopted to increase profits and gain a competitive edge and that barter is, in fact, complementary to the cash economy. The executives we interviewed at BizXchange cited as the main benefits of barter: Maximizing productivity by keeping the company busy during slow periods; Leveraging excess inventory; Increasing purchasing power by buying with incremental sales; Increasing the customer base due to the incentives for other members to engage in barter, as well as the marketing efforts of the brokers; Developing customer loyalty since barter relationships form close bonds that are not easily broken; and Participating in a managed economy.

When a business joins a trade exchange, it typically pays a membership fee. This represents a small fraction of the revenues of the trade exchange. The primary revenue is made by charging a fee to the buyer and seller on each transaction. The fee is typically in the range of 6 - 7.5% and is payable in US dollars. When a business joins the trade exchange, they are issued a line of credit in trade dollars, which permits them to make purchases without first having to sell and also gives them flexibility in conducting transactions. The trade exchange charges interest on negative balances, usually at the same rate as major credit cards. In order to give a company some control over how much of their profits are accrued in terms of trade dollars, the trade exchange permits the member to set an upper limit on the amount of trade dollars they are willing to accumulate. The credit line and upper limit define the financial operating range of the business within the barter pool.

Each member is assigned to a trade broker. A broker typically represents a set of 150 - 200 client businesses. The broker's job from the standpoint of the client is to help the client sell his goods and to inform him of goods he might like to buy. The broker's job from the standpoint of the trade exchange is to stimulate trade, since the exchange's revenues are directly tied to trade volume. The broker stimulates trade by working to help clients spend their trade dollars when they have positive balance and generate sales when they have negative balance. The broker's primary tool is the referral, referring potential buyers to suppliers. Note that member businesses are under no obligation to follow the broker's referrals, but experience from trade exchanges shows that they generally do. While the goods a business has to sell are stated explicitly, those that the business wants to buy may be explicitly stated or may be predicted by the broker based on things like the type of business and other goods that the business has purchased in the past.

In carrying out his job, the broker attempts to maximize single-period trade volume while maintaining fairness and balance of trade. Trade is *fair* when each client gets a share of trade proportional to the amount of goods he has to sell and wants to buy. Trade is *balanced* when the amount each member buys equals what it sells. In this paper we focus on the balance problem. For any supply/demand of a business there will typically be several different buyers/sellers to choose from. The question is then which of those buyers/sellers the broker should refer the business to. Most trade exchanges do not make this decision based on price; they leave price negotiation up to the members. Other factors such as convenience of location being equal, the broker maintains balance by basing his decision on the balances of the buyers/sellers. For example, if we have one supplier who has highly positive balance and one who has highly negative balance, the broker will refer the client to the supplier with highly negative balance. Similarly, if we have one buyer with highly positive balance and one with highly negative balance, the broker will refer the buyer with the highly positive balance to the client. One elaboration on the notion of maintaining balance is that brokers often know that certain members will not spend their trade dollars until they reach some positive accumulation, so they set postivive target balances for those companies.

This concludes our description of the operation of barter trade exchanges. In the remainder of the paper we present a mathematical formalization of the balance problem and develop efficient algorithms for its solution.[1] We assume that trade occurs in business cycles: first businesses' supplies and demands are determined, a matching is found, the businesses act on the resulting referrals, and the cycle repeats. The matching problem can be represented by a requirements matrix in which each row represents a member business, each column represents a category of goods, and matrix entries represent quantities to buy or sell. Positive values represent amounts to sell; negative values represent amounts to buy; and zero values indicate either no ability or no interest to trade in that product category. We assume a uniform unit cost for each product category. Since prices are naturally not uniform across suppliers, this unit cost can be an estimate based on the average retail price over suppliers in that category. The internal state of each business is characterized by its current balance, its credit line, and an upper limit on allowed balance. Once a business reaches its credit limit, then it can no longer buy without selling, so its row in the matrix will show no demand. Similarly, when a company reaches its upper limit, it can no longer sell without buying, so its

[1] We have developed a similar technique for solving the fairness problem but omit it from this paper due to space limitations.

row will show no supply.

3 The Barter Universe

In the barter universe $\mathcal{B} = (C, P, \alpha, R)$, $C = \{c_i, i = 1, \ldots, m\}$ denotes the set of companies, and $P = \{p_j, j = 1, \ldots, n\}$ denotes the set of goods. For each j, the unit cost of good p_j is α_j barter dollars. The $m \times n$ *requirements matrix* R specifies the number of units each company is willing to supply or purchase for each good. We shall assume that R is an integer matrix, and use the convention that if c_i is a *seller* of p_j then $R_{ij} > 0$, and if c_i is a *buyer* of p_j then $R_{ij} < 0$, and if c_i is not interested in p_j then $R_{ij} = 0$.

A *trade set* in the barter universe is specified by an $m \times n$ matrix T, where T_{ij} indicates the number of units of p_j that c_i sold or bought in the trade. There are two conditions T must satisfy for each j:

i. If c_i is a seller of p_j then $0 \leq T_{ij} \leq R_{ij}$, and if c_i is a buyer of p_j then $0 \geq T_{ij} \geq R_{ij}$, and if c_i is not interested in p_j then $T_{ij} = 0$.

ii. $\sum_{i=1}^{m} T_{ij} = 0$; i.e., for each p_j, the total number of units sold equals the total number of units bought.

The *volume* of trade set T is $vol(T) = \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{m} |T_{ij}|$, the total number of units of all the goods exchanged in the trade.

Throughout this paper, we will consider *maximal trade sets*, which are trade sets with the largest volumes. For each j, let S_j and N_j denote, respectively, the *sellers* and *buyers* of p_j. If the supply of p_j exceeds its demand then a maximal trade set T will have $T_{ij} = R_{ij}$ for all $c_i \in N_j$. Similarly, if the supply of p_j is no more than its demand then $T_{ij} = R_{ij}$ for all $c_i \in S_j$. Hence, the number of units traded for each p_j is maximized, and every maximal trade set has volume equal to $\sum_{j=1}^{n} \min\{\sum_{i \in S_j} R_{ij}, \sum_{i \in N_j} |R_{ij}|\}$. We note that it is straightforward to find a maximal trade set for $\mathcal{B}$ in time $O(|P||C|)$.

4 Finding the Most Balanced Maximal Trade Set

After the trade specified by T takes place, we define the *balance of* c_i as $b_i = \sum_{j=1}^{n} \alpha_j T_{ij}$, and the *absolute balance due to* T as $ab_T = \sum_{i=1}^{m} |b_i|$. Our goal is to solve for the maximal trade set T^* that is *most balanced*; i.e., $ab_{T^*} \leq ab_T$ for every other maximal trade set T. The following lemma follows directly from the fact that $\sum_{i=1}^{m} T_{ij} = 0$ for each j. It will be useful later on.

Lemma 4.1. *For any trade set T, $\sum_{i=1}^{m} b_i = 0$.*

The problem of finding the most balanced maximal trade set can be formulated as an integer program. Thus, one approach we can take is to first relax the condition that the entries of a trade set T be integers. The integer program reduces to a linear program, and the latter can be solved optimally in polynomial time. We can then use some procedure to transform a fractional solution to an integer one. In this section, we shall take a similar approach. However, instead of solving the relaxed problem as a linear program, we shall show in the next subsection that it can be reduced to a minimum-cost circulation (MCC) problem on an appropriate network, and that it can be solved optimally by running a bounded number of breadth-first searches because of the network's simple structure. In the subsequent section, we will present a novel randomized rounding procedure that produces an integer maximal trade set T' from a fractional trade set T so that the expected balance of a company in T' is equal to its balance in T. By enforcing this constraint, we hope that $ab_{T'}$ will be close to ab_T. In Section 6, our simulations indicate that the MCC solutions combined with the randomized rounding procedure are as competitive as the solutions obtained from a commercial mixed-integer programming (MIP) package. Furthermore, in all but one instance, our approach is significantly faster than the MIP package.

4.1 When the maximal trade sets can be fractional

We can solve for the most balanced fractional maximal trade by first constructing a network, and then finding the minimum cost circulation on the network. (We refer the readers to [1] for an excellent introduction to the field of network flows.) Let us start with the following example. Suppose there are four companies and two goods which costs \$2 and \$3 respectively. Below is a requirements matrix R:

	p_1(\$2)	p_2(\$3)
c_1	50	20
c_2	70	-30
c_3	-20	15
c_4	-10	-40

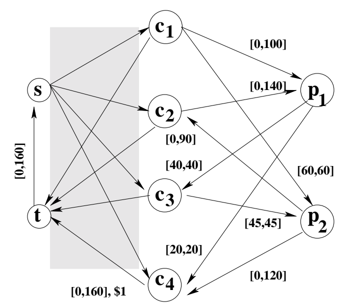

Figure 1: The network that corresponds to the barter universe in the example.

The corresponding network is shown in Figure 1. We only indicated the edge costs when they are non-zero. There is a directed edge from c_i to p_j (or p_j to c_i) if c_i is a seller (or buyer) of p_j. The lower and upper bound capacities of an edge between c_i and p_j indicate the minimum and maximum *total cost* of the units of p_j that c_i can sell or buy. The edges (s, c_i) and (c_i, t) will be used to indicate the value of b_i if $b_i > 0$ and if $b_i < 0$ respectively. The upper bound capacities on all edges between $\{s, t\}$ and nodes in C, and (s, t) are set to 160 because this is the largest balance any of the companies can have at the end of a trade set.

For an arbitrary $\mathcal{B}$, we shall create the network as follows. Let $G = (V, E)$ be a directed graph where

$$\begin{aligned} V &= C \cup P \cup \{s, t\}, \\ E &= \{(p_j, c_i), \forall c_i \in N_j, \forall j\} \cup \{(c_i, p_j), \forall c_i \in S_j, \forall j\} \cup \{(s, c_i), \forall c_i \in C\} \cup \{(c_i, t), \forall c_i \in C\} \cup \{(t, s)\}. \end{aligned}$$

We shall use functions $l : E \to \mathbf{R}$ and $u : E \to \mathbf{R}$ to denote the lower and upper bound capacities on the edges of G. For each j, if the supply of p_j exceeds its demand then let all edges $e = (p_j, c_i)$, where $c_i \in N_j$ have $l(e) = u(e) = \alpha_j |R_{ij}|$, but let all edges $e' = (c_{i'}, p_j)$ where $c_{i'} \in S_j$ have $l(e) = 0$ and $u(e) = \alpha_j R_{ij}$. The capacities are chosen because, in a fractional maximal trade set, all buyers of p_j will have to purchase all their demands for p_j, but the sellers of p_j need not sell all their supplies of p_j. When the supply for p_j is no more than its demand then we do the opposite. Let all edges $e' = (c_{i'}, p_j)$ where $c_{i'} \in S_j$ have $l(e) = u(e) = \alpha_j R_{ij}$, and let all edges $e = (p_j, c_i)$, where $c_i \in N_j$ have $l(e) = 0$ and $u(e) = \alpha_j |R_{ij}|$. Let the edges (s, c_i) and (c_i, t), for all c_i, and (t, s) have lower and upper bound capacities of 0 and ∞ (or a large number) respectively.

Finally, we shall use $\gamma : E \to \mathbf{R}$ as the cost function for the edges of G. For each c_i, let $\gamma(s, c_i) = \gamma(c_i, t) = \$1$, and let all other edges have a cost of \$0. Let us denote this network as (G, l, u, γ).

A *circulation* $f : E \to \mathbf{R}$ on the network assigns a *flow value* to each edge. We say that the circulation f is *feasible* if (a) for each $e \in E$, $l(e) \le f(e) \le u(e)$, and (b) for each $v \in V$, the flow into v equals the flow out of v. The *cost of* f is $\sum_{e \in E} \gamma(e) f(e)$. We shall now show that there is a correspondence between the fractional maximal trade sets for $\mathcal{B}$ and the feasible integer circulations for (G, l, u, γ).

• *Fractional maximal trade sets to feasible integer circulations.* Let T be a fractional maximal trade set for $\mathcal{B}$. Construct the corresponding circulation f on (G, l, u, γ) as follows:

i. For each c_i and p_j, if $c_i \in S_j$, set $f(c_i, p_j) = \alpha_j T_{ij}$. If $c_i \in N_j$, set $f(p_j, c_i) = \alpha_j |T_{ij}|$.
ii. For each company c_i, if $b_i \ge 0$, let $f(s, c_i) = b_i$; otherwise, let $f(c_i, t) = |b_i|$.
iii. Let $f(t, s) = \sum_{b_i < 0} |b_i|$.
iv. Let all other edges have a flow of 0.

Lemma 4.2. *The flow f is a a feasible integer circulation on (G, l, u, γ), and its cost is ab_T.*

Proof: Step (i) ensures the that net flow out of p_j is $\sum_{i:T_{ij}<0} \alpha_j |T_{ij}| - \sum_{i:T_{ij}>0} \alpha_j T_{ij} = -\sum_{i=1}^{m} \alpha_j T_{ij} = 0$ because T is a valid trade set for $\mathcal{B}$. Hence, flow is conserved at each p_j. From the same step, the net flow out of each c_i is $\sum_{j:T_{ij}>0} \alpha_j T_{ij} - \sum_{j:T_{ij}<0} \alpha_j |T_{ij}| = b_i$. Step (ii) forces $f(c_i,t) - f(s,c_i) = -b_i$ so flow is conserved at c_i, for each i. And since $\sum_{i=1}^{n} f(s,c_i) = \sum_{b_i>0} b_i = \sum_{i=1}^{n} f(c_i,t)$ because of step (ii) and Lemma 4.1, step (iii) ensures flow conservation at nodes s and t. Finally, it is straightforward to verify that the flows on all the edges satisfy the upper and lower bound constraints so the circulation we have constructed is feasible. Its cost is $\sum_{i=1}^{m} f(s,c_i) + \sum_{i=1}^{m} f(c_i,t) = \sum_{i=1}^{m} |b_i| = ab_T$. □

• *Feasible integer circulations to fractional maximal trade sets.* Given a feasible integer circulation f on (G,l,u,γ), constructing the corresponding trade set T is straightforward. For each c_i and p_j, if $(c_i,p_j) \in E$, set $T_{ij} = f(c_i,p_j)/\alpha_j$. If $(p_j,c_i) \in E$, set $T_{ij} = -f(p_j,c_i)/\alpha_j$. Otherwise, set $T_{ij} = 0$.

Lemma 4.3. *The resulting fractional trade set T is maximal for $\mathcal{B}$.*

Proof: Because of the way the lower and upper bound capacities are assigned on the edges of the network, and because f is a feasible circulation, T is a valid trade set. When the supply of p_j exceeds its demand, for each $c_i \in N_j$, $T_{ij} = R_{ij}$ because $l(p_j,c_i) = u(p_j,c_i) = \alpha_j |R_{ij}|$. Similarly, when the supply of p_j does not exceed its demand then, for each $c_i \in S_j$, $T_{ij} = R_{ij}$. Thus, the volume of T is maximized so it is a fractional maximal trade set for $\mathcal{B}$. □

Theorem 4.4. *Let f^* be an integer minimum cost circulation for (G,l,u,γ). The fractional maximal trade set T^* resulting from the transformation of f^* is a most balanced fractional maximal trade set for $\mathcal{B}$. Furthermore, the cost of f^* equals ab_{T^*}.*

Proof: At each node c_i, $\sum_{e'=(c_i,p_j)\in E} f^*(e') - \sum_{e=(p_j,c_i)\in E} f^*(e) = b_i$, the balance of company c_i in T^*. Since f^* is a feasible circulation, $f^*(c_i,t) - f^*(s,c_i) = -b_i$. If $b_i \geq 0$ then $f^*(c_i,t) = 0$ and $f^*(s,c_i) = b_i$; otherwise, we can decrease the flow on the three edges (c_i,t), (s,c_i) and (t,s) by $f^*(c_i,t)$, and the new circulation will have a cost less than that of f^*. By the same reasoning, if $b_i < 0$ then $f^*(c_i,t) = -b_i$ and $f(s,c_i) = 0$. Consequently, the cost of f^* is $\sum_{b_i>0} b_i + \sum_{b_i<0} |b_i| = ab_{T^*}$.

Let T be any fractional maximal trade and f be its corresponding integer circulation on (G,l,u,γ) as defined above. Since f^* is an MCC for (G,l,u,γ), $\gamma(f^*) \leq \gamma(f)$ so $ab_{T^*} \leq ab_T$. Therefore, T^* is a most balanced fractional maximal trade set in (G,l,u,γ). □

4.1.1 An efficient algorithm for solving MCC in (G,l,u,γ).

Given a network (G,l,u,γ) and a feasible circulation f on the network, its *residual network*, which we denote as $R(G,f)$ has the same vertex set as G. The edges also have costs (specified by γ_R), and upper and lower bound capacities (specified by u_R and l_R). An edge e of G is in the residual network if and only if $f(e) < u(e)$. Furthermore, $u_R(e) = u(e) - f(e)$, $l_R(e) = 0$, and $\gamma_R(e) = \gamma(e)$. The reverse of e, which we shall denote as $\bar{e}$, is in the residual network if and only if $l(e) < f(e)$, and $u_R(\bar{e}) = f(e) - l(e)$, $l_R(\bar{e}) = 0$, and $\gamma_R(\bar{e}) = -\gamma(e)$.

To solve for the MCC in (G,l,u,γ), we make use of the minimum mean cycle-canceling algorithm. The algorithm starts with a feasible circulation f. At each iteration, it finds a minimum mean (directed) cycle[2] in the residual graph $R(G,f)$. If the cycle has negative cost, then the algorithm augments f using this cycle. (Step 2 of algorithm BALANCE specifies this step precisely.) Otherwise, $R(G,f)$ contains no negative-cost directed cycle, and the algorithm outputs f, which must be an MCC. Goldberg and Tarjan [7] showed that the number of augmentations is $O(|V||E|)$. We note that there are other algorithms for MCC whose theoretical guarantees are better than the minimum mean cycle-canceling algorithm (e.g., see reference in [1, 7]). We shall show, however, that this algorithm is very simple to implement for our network.

Given a circulation f in (G,l,u,γ), we show next that the negative-cost directed cycles in $R(G,f)$ have a nice structure. As a result, the search for the minimum mean cycle in $R(G,f)$ becomes straightforward.

Lemma 4.5. *All negative-cost directed cycles in $R(G,f)$ cost $-\$2$ and consist of a sequence of nodes of the form $t, c_{i_1}, p_{j_1}, c_{i_2}, p_{j_2}, \ldots, p_{j_k}, c_{i_{k+1}}, s, t$. (Proof: See appendix.)*

We are now ready to present our algorithm.

[2] If Y is a directed cycle, then its *mean cost* is $\sum_{e\in Y}$ cost$(e)/|Y|$. The *minimum mean (directed) cycle* is the cycle with the smallest mean cost.

BALANCE(R)

1. Find an initial maximal trade set T. Construct (G, l, u, γ) and the corresponding integer circulation f of T.
2. Construct the residual graph $R(G, f)$. Determine if $R(G, f)$ contains a negative-cost cycle as follows. In $R(G, f)$, remove (t, s) and all the edges with a cost of \$1, and find the shortest path from t to s using breadth first search.

   If no such path exists, go to Step 3.

   Else, append the edge (s, t) to the path to form a negative-cost cycle Y. Augment f using Y. (That is, let $\epsilon \leftarrow \min\{u(e) - f(e),$ for all forward edges e on $Y\} \cup \{f(e) - l(e)$ for all backward edges e on $Y\}$. For all edges e in Y, if e is a forward edge, $f(e) \leftarrow f(e) + \epsilon$; if e is a backward edge, $f(e) \leftarrow f(e) - \epsilon$.) Repeat Step 2.
3. Construct the trade set T equivalent to the circulation f. Return(T).

Theorem 4.6. *BALANCE outputs a most balanced fractional maximal trade set in* $O((|V|+|E|)|V||E|)$ *time where* $|V| = O(|P| + |C|)$ *and* $|E| = O(\sum_{j=1}^{n}(|S_j| + |N_j|))$.

Proof: If we can show that Step 2 of BALANCE correctly finds the minimum mean cycle in $R(G, f)$ then BALANCE is essentially an implementation of the minimum mean cycle-canceling algorithm for (G, l, u, γ). Since it outputs the fractional maximal trade set that corresponds to the MCC in (G, l, u, γ), according to Theorem 4.4, the trade set must be most balanced.

Consider $R'(G, f)$, constructed by deleting edge (t, s) and all edges with cost \$1 in $R(G, f)$. From Lemma 4.5, $R(G, f)$ contains a negative-cost directed cycle if and only if $R'(G, f)$ has a directed path from t to s. Step 2 finds the shortest $t - s$ path in $R'(G, f)$ so the negative-cost directed cycle Y has the fewest number of edges. Since all negative-cost directed cycles in $R(G, f)$ have the same cost, Y is a minimum mean cycle in $R(G, f)$.

Doing a breadth-first search in $R(G, f)$ takes $O(|V| + |E|)$ time. Augmenting the flow on Y takes $O(|E|)$ time. According to Goldberg and Tarjan, the number of augmentations of BALANCE is $O(|V||E|)$ so step 2 takes $O((|V| + |E|)|V||E|)$ time where $|V| = O(|P| + |C|)$ and $|E| = O(\sum_{j=1}^{n}(|S_j| + |N_j|))$. Steps 1 and 3 take $O(|P||C|) = O(|V|^2)$ time. Therefore, the overall running time of BALANCE is $O((|V| + |E|)|V||E|)$ time. □

An extended example of the algorithm is contained in the appendix. We note that, in practice, $|E|$ is likely going to be small compared to $|P| \times |C|$. BizXchange, for example, currently has 600 member companies and 125 product categories. Companies typically only sell between one to three goods and buy between five to 25 goods. Furthermore, as the trade exchange grows, the number of companies grows but the number of goods remains relatively constant.

4.2 When the maximal trade sets have to be integral

Since the flows on the edges between the company and good nodes need not be multiples of the good costs, the trade set T^* in Theorem 4.4 may contain non-integer values. There is, however, an exception – when all the goods have $\alpha_j = 1$. And since the maximal trade set that minimizes ab_T when $\alpha_j = 1$ for all j also minimizes ab_T when $\alpha_j = \alpha$, α a constant, the statement below is true.

Theorem 4.7. *Assuming all the goods in* $\mathcal{B}$ *have the same cost* α*, a most balanced (integer) maximal trade set can be found in* $O((|V| + |E|)|V||E|)$ *time where* $|V| = O(|P| + |C|)$ *and* $|E| = O(\sum_{j=1}^{n}(|S_j| + |N_j|))$.

If the goods in $\mathcal{B}$ have different costs, then we need a method that will transform the most balanced fractional maximal trade set T^* into an integer maximal trade set T^{**} in such a way that $ab_{T^{**}}$ is approximately equal to ab_{T^*}.

4.2.1 Rounding Fractional Maximal Trade Sets

Given a fractional maximal trade set T, we wish to transform it into an integer maximal trade set T' so that $ab_{T'}$ is approximately equal to ab_T. In this section, we shall try to achieve this goal by employing a randomized rounding procedure that sets each T'_{ij} equal to $\lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$ so that (i) for each j, $\sum_{i=1}^{m} T'_{ij} = 0$, and (ii) for each i, $E[b'_i] = E[\sum_{j=1}^{n} \alpha_j T'_{ij}] = b_i$. In the process of rounding T, condition (i) guarantees that the number of units of p_j sold remains equal to the number of units of p_j bought, while condition (ii) states that the *expected balance* of each company in T' is equal to its balance in T. It is, of course, possible for $|b'_i|$ to differ significantly from $|b_i|$ so that $ab_{T'} = \sum_{i=1}^{m} |b'_i|$ is much larger than $\sum_{i=1}^{m} |b_i| = ab_T$. We hope, however, that if we generate several integer maximal trade sets, one of them will have an absolute balance close to ab_T. First, we make the following observation:

Lemma 4.8. *Let T be a fractional maximal trade set for the barter universe $\mathcal{B} = (C, P, \alpha, R)$. Suppose T' is an integer matrix obtained by setting each T'_{ij} equal to $\lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$ so that $\sum_{i=1}^{m} T'_{ij} = \sum_{i=1}^{m} T_{ij} = 0$ for each j. Then T' is an integer maximal trade set for $\mathcal{B}$ and $ab_{T'} \leq ab_T + m \sum_{j=1}^{n} \alpha_j$. (Proof: See appendix.)*

Without loss of generality, assume that, in the jth column of the fractional maximal trade set T, exactly k entries are not integers. The fractional entries have to be all positive if the supply of p_j exceeds its demand, and all negative otherwise. For $i = 1, \ldots, m$, let $x_i = T_{ij} - \lfloor T_{ij} \rfloor$ if $T_{ij} \geq 0$, and $\lceil T_{ij} \rceil - T_{ij}$ if $T_{ij} < 0$, so that $0 \leq x_i \leq 1$. Since $\sum_{i=1}^{m} T_{ij} = 0$, the fractional parts of the T_{ij}'s must also sum up to an integer; i.e., $\sum_{i=1}^{m} x_i = z$ for some positive integer $z < k$. Setting $T'_{ij} = \lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$ so that $\sum_{i=1}^{m} T'_{ij} = 0$ is equivalent to setting $x'_i = \lceil x_i \rceil$ or $\lfloor x_i \rfloor$ so that $\sum_{i=1}^{m} x'_i = z$. The latter can be done by choosing exactly z out of the k fractional x_i's, and rounding them all to 1, and rounding the remaining $k - z$ x_i's to 0. Thus, creating an integer maximal trade set T' with the properties described in the previous lemma can be done in $O(|P||C|)$ time. We would like to choose a rounding of the x_i's, however, in a randomized manner so that $Prob(x'_i = 1) = x_i$ for each i. To do so, we make use of the following lemma.

Lemma 4.9. *Let m and z be positive integers with $z \leq m$. Let $\mathcal{V}(m, z)$ be the set of all vectors in $\mathbf{R}^m$ that have z components equal to 1 and $m - z$ components equal to 0. (Note that $\mathcal{V}(m, z)$ contains $\binom{m}{z}$ vectors.) Let $\vec{x} = (x_1, \ldots, x_m) \in [0, 1]^m$ so that $\sum_{i=1}^{m} x_i = z$. Then $\vec{x}$ can be expressed as a convex combination of the elements in $\mathcal{V}(m, z)$. That is, $\vec{x} = \sum_{\vec{v}_r \in \mathcal{V}(m,z)} \lambda_r \vec{v}_r$ so that $0 \leq \lambda_r \leq 1$ for each r, and $\sum_r \lambda_r = 1$. (Proof: See appendix.)*

In the constructive proof of Lemma 4.9 given by Volkmer [18], the convex combination of $\vec{x}$ contains at most k vectors from $\mathcal{V}(m, z)$, where k is the number of fractional entries in $\vec{x}$. The construction takes $O(km)$ time. We note that $\mathcal{V}(m, z)$ contains all the possible roundings of $(x_1, \ldots, x_m)$. Let us use the coefficients in the convex combination of $\vec{x}$ as probabilities over the vectors in $\mathcal{V}(m, z)$. That is, for each $\vec{v}_r \in \mathcal{V}(m, z)$, define $Prob(\vec{v}_r) = \lambda_r$. Choose one vector $\vec{u} = (u_1, u_2, \ldots, u_m)$ in $\mathcal{V}(m, z)$ at random using this probability distribution. For each i, $Prob(u_i = 1) = \sum_{v_{r,i}=1} Prob(\vec{v}_r) = \sum_{r:v_{r,i}=1} \lambda_r = x_i$, where the last equality is true because $x_i = \sum_{r:v_{r,i}=1} \lambda_r v_{r,i}$. We are now ready to present our randomized rounding procedure.

RANDOM_ROUND(T)
For $j = 1$ to n, do:

1. For $i = 1$ to m, if $T_{ij} \geq 0$, set $x_i = T_{ij} - \lfloor T_{ij} \rfloor$. Else, set $x_i = \lceil T_{ij} \rceil - T_{ij}$.
   Set $z = \sum_{i=1}^{m} x_i$.
2. Express $\vec{x} = (x_1, \ldots, x_m)$ as a convex combination of the vectors in $\mathcal{V}(m, z)$ using CONVEX_COMB [3].
   For each $\vec{v}_r \in \mathcal{V}(m, z)$, set $Prob(\vec{v}_r)$ equal to the coefficient of $\vec{v}_r$ in the convex combination of $\vec{x}$.
3. Pick one vector $\vec{u}$ randomly from $\mathcal{V}(m, z)$ using the probability distribution in step 2.
   For $i = 1$ to m, if $T_{ij} \geq 0$, set $T'_{ij} = \lfloor T_{ij} \rfloor + u_i$. Else, set $T'_{ij} = \lceil T_{ij} \rceil - u_i$.

Return(T').

Theorem 4.10. *RANDOM_ROUND(T) outputs an integer maximal trade set T' such that for each i and j, $E[T'_{ij}] = T_{ij}$. Thus, for each i, $E[b'_i] = b_i$ and $ab_{T'} \leq ab_T + m \sum_{j=1}^{n} \alpha_j$. The procedure runs in $O(K|C| + |P||C|)$ time where K is the number of fractional entries in T. (Proof: See appendix.)*

5 Matching Sellers and Buyers

Once a desirable maximal trade set T has been found, sellers and buyers for each good have to be matched, and the number of units traded between them specified. We make use of a greedy heuristic that attempts to minimize the

[3] The pseudocode for CONVEX_COMB can be found in the appendix.

average number of sellers matched to a buyer per good. At each iteration j, the company with currently the largest supply of p_j is matched with the company with currently the largest demand of p_j. Below is our algorithm.

GREEDY_MATCHING (T)
For $j = 1$ to n, do:

1. Initialize $S_j \leftarrow \{c_i : T_{ij} > 0\}$, $N_j \leftarrow \{c_i : T_{ij} < 0\}$, and $M \leftarrow \emptyset$.
2. While $(N_j \neq \emptyset)$

    Find a and b so that $T_{aj} = \max\{T_{ij}, c_i \in S_j\}$ and $T_{bj} = \min\{T_{ij}, c_i \in N_j\}$.
    $\epsilon \leftarrow \min\{T_{aj}, |T_{bj}|\}$
    $M \leftarrow M \cup \{(a, b, \epsilon)\}$
    If $T_{aj} - \epsilon = 0$, set $S_j \leftarrow S_j - \{c_a\}$. Else, set $T_{aj} \leftarrow T_{aj} - \epsilon$.
    If $|T_{bj}| - \epsilon = 0$, set $N_j \leftarrow N_j - \{c_b\}$ Else, set $T_{bj} \leftarrow T_{bj} + \epsilon$.

3. Return(M).

It is straightforward to verify the following lemma:

Lemma 5.1. *GREEDY_MATCHING outputs a valid matching for each good p_j. It runs in $O(\sum_{j=1}^{n}(|S_j| + |N_j|)^2)$ time.*

6 Empirical Evaluation

We evaluated the quality of the solutions from our stochastic rounding algorithm, as well as its running time by comparing its results to those from a commercial mixed integer programming (MIP) package on 100 randomly generated matrices. We varied two parameters: product price range and current balance range, taking uniformly distributed random values over the ranges and producing ten matrices for each combination of parameter settings. Since a barter pool represents a closed economy, the sum of the starting balances was always zero. Each matrix represented 100 companies and 50 products. Supply and demand were randomly generated in such a way that their distributions mirrored those in the BizXchange barter pool. The stochastic rounding algorithm was run 10 times for each matrix and the best value was selected. Increasing the number of runs to 50 made little difference in the results. Results are shown in Table 1. For each matrix, the degree of sub-optimality of the stochastic rounding algorithm solution was computed as the difference between the value produced by the stochastic rounding algorithm and the value produced by the MIP package, and expressed as a percentage of the range of possible values of absolute balance for that matrix. It is necessary to express the degree of sub-optimality relative to the range of possible values since a small absolute difference in solutions is more signinicant if the range of solutions is small than if it is large. Since we have no algorithm for exactly solving the integer problem, we had to use our best estimate of the range of possible values of the absolute balance. For the minimum value, we took the lesser of the MIP and stochastic rounding solutions. For the maximum value, we used the MIP package, setting the objective to maximize absolute balance. Note that this approach may underestimate the width of the range but will not overestimate it since all solutions are valid but not guaranteed to be optimal. Thus we are never overestimating the quality of stochastic rounding algorithm's solution. Negative values for degree of sub-optimality indicate that the stochastic rounding algorithm produced a better solution than MIP. The solution from the stochastic rounding algorithm is always within 0.7% of the MIP solution, with little variance. Running times for the stochastic rounding algorithm and for the MIP package, as well as their standard deviations are shown in the last four columns of the table. Running times for the stochastic rounding algorithm are the sum of the running times for BALANCE and ten runs of RANDOM_ROUND. In every case but one, the stochastic rounding algorithm has better running time and smaller variance in running time than MIP.

Using one year of transaction history data from BizXchange, we built a barter pool trade simulator by learning Bayesian network models to predict company product demands. The BizXchange data also specified the product supplied by each company (one per company) and the company's credit line, which ranged from \$500 to \$20,000. Company names were purged in order to protect confidentiality. We ran a number of simulations to compare the absolute balance and the trade volume with and without balance optimization. In the simulations without balance

Description		Ave.	Std.	MIP	Stochastic	Std.	Std.
Price Range ($)	Current Balance ($)	Degree of sub-opt. Stochastic (%)	deviation of deg. of sub-opt. Stochastic (10)(%)	Ave. Runtime (msec)	Ave. Runtime (msec)	deviation of MIP Runtime	deviation of Stochastic Runtime
[10, 10]	0	0.00	0.00	400.00	1,600.00	0.52	0.84
[10, 10]	[-4000, 10,000]	0.01	0.01	45,250.00	1,750.00	120.00	0.50
[10, 100]	0	0.14	0.10	41,100.00	1,500.00	54.59	0.53
[10, 100]	[-4000, 10,000]	0.07	0.09	29,600.00	1,100.00	39.32	0.32
[10, 200]	0	0.21	0.26	12,700.00	900.00	4.24	0.32
[10, 200]	[-4000, 10,000]	0.21	1.15	7,400.00	400.00	7.34	0.70
[10, 500]	0	0.41	0.58	11,400.00	500.00	8.68	0.53
[10, 500]	[-4000, 10,000]	0.39	0.59	5,600.00	600.00	5.82	0.52
[10, 1000]	0	-0.21	1.73	12,500.00	600.00	24.74	0.52
[10, 1000]	[-4000, 10,000]	0.63	0.02	1,000.00	500.00	1.25	0.53

Table 1: Comparison of stochastic rounding algorithm with mixed integer programming.

optimization, maximal trade sets were determined using a simple greedy algorithm. Each simulation used the same set of 180 companies and 36 product categories, with 3 - 7 suppliers per product category. Company balances all started at zero trade dollars. Simulations were run for 100 trade cycles, each cycle being one week. To observe the effectiveness of the optimization for various financial operating ranges, we varied the credit limit and upper limit. The credit limit ranged from that given in the data (CL) to four times that value (4CL). The upper bound ranged from twice the credit limit (2CL) to eight times the credit limit (8CL). Product demands were determined from the Bayes net predictive model. Product supply was determined by taking the difference between each company's upper limit and their current balance. Companies were not allowed to exceed their credit limit or upper limit. As the results of the simulations displayed in Table 2 show, using the optimizer results in a reduction in the average absolute balance as well as in the standard deviation of absolute balances over companies. This reduction is accompanied by an increase in the total barter pool trade volume per week. Somewhat surprisingly, these differences in results between the optimized and non-optimized runs of the simulator tend to increase as the financial operating ranges of the companies increase. Although our simulation is far from incorporating all the complexities of trade dynamics, it does suggest that maintaining balance of trade tends to increase trade volume over the long run.

Description	Average ab_T per week ($)	Decrease in ab_T O vs NoO (%)	Std deviation of ab_T over companies averaged over 100 weeks ($)	Decrease in std deviation O vs NoO (%)	Average Total Trade Volume per week ($)	Increase in Trade Volume per week O vs NoO (%)
[CL, 2CL] NoO	800,918.86		4,511.52		44,788.68	
[CL, 2CL] O	636,285.06	20.56	4,281.34	5.10	51,097.98	14.09
[CL, 4CL] NoO	951,813.20		6,480.54		37,965.80	
[CL, 4CL] O	673,938.34	29.19	5,517.07	14.87	50,767.88	33.72
[CL, 8CL] NoO	1,035,472.18		7,896.79		32,935.22	
[CL, 8CL] O	651,756.68	37.06	6,638.58	15.93	49,751.32	51.06
[2CL, 4CL] NoO	1,587,551.24		9,566.14		50,476.72	
[2CL, 4CL] O	1,125,356.30	29.11	9,437.32	1.35	68,914.98	36.53
[4CL, 4CL] NoO	2,535,481.44		17,050.14		66,734.66	
[4CL, 4CL] O	1,574,940.20	37.88	14,819.81	13.08	89,485.88	34.09

Table 2: Results of simulation runs with (O) and without (NoO) balance optimization.

7 Related Work

Segev and Beam [14] point out that while there exists a large body of literature on auction theory, there is no equivalent theory for brokered marketplaces. They examine the effect of search costs on brokerage costs where negotiation is only on price. Their model assumes a single commodity where each seller has one unit and each buyer desires one unit and where supply equals demand. They assume a single broker who receives bids from buyers and price quotes from sellers, each consisting of a single price point. The broker matches buyers and sellers by finding a seller whose asking price is below the bid of a buyer. Through simulation, they examine combinations of search and brokerage cost that make the broker a more attractive option than direct searching.

Dailianas et al [6] present algorithms for matching buyers and sellers in e-marketplaces for trading soft composable commodities such as bandwidth. Such quantities admit fractional solutions. Each seller submits one or more offer curves indicating the price per unit for various quantities. Each buyer submits bids specifying a quantity and unit price or a bid curve. It is assumed that the price per unit of good traded drops as the quantity increases. So the marketplace can make a profit by aggregating demand and keeping some of the difference between the selling price and the amount each buyer is willing to pay. Dailianas et al explore three optimization strategies: maximize profit for the marketplace; find the allocation of resources from the sellers that will maximally satisfy the demands of the buyers; and satisfy at least a given percentage of the buyers and then maximize profit. They present an exact algorithm for the first objective (maximize profit) but point out that it is too computationally complex to use in practice. They go on to present heuristic algorithms for each of the three objectives and show that the algorithms produce solutions that are very close to optimal in most cases.

Tewari and Maes [16] describe the MARI (Multi-Attribute Resource Intermediary) project, which aims at developing an agent-mediated e-marketplace infrastructure for matching buyers and sellers. Each buyer and each seller is represented by an agent, which represents that buyer's or seller's preferences via a multi-attribute utility function. The approach is illustrated with a market for translation services, where each buyer is looking for a certain number of words to be translated and each seller has a certain translation capacity to sell. They present an algorithm for matching buyers and sellers so that the welfare, measured by the aggregate surplus of all transaction parties, is maximized. The amount a buyer is willing to pay for a given translation service is determined by his utility function. The algorithm works by representing the problem as a bipartite graph in which each buyer and each seller is represented by a node and there is a link between potential buyer-seller pairs, based on compatibility between hard buyer constraints and seller characteristics. Each link is assigned a reward equal to the bid-ask spread between the buyer and seller at the ends of the link. The objective is to satisfy all the buyers' needs by matching with sellers such that the sum of the rewards is maximized. They solve the weighted bipartite graph matching problem by transforming it to a linear program. Tewari and Maes are working with a single good characterized by multiple attributes and are performing matching to optimize aggregate surplus, while we are working with multiple goods and are matching based on more complex optimization criteria, resulting in more complex minimum cost flow problems. The work of Tewari and Maes focuses less on the combinatorial optimization aspects and more on providing a general framework for agent mediated buying and selling of non-tangible goods and services.

A fair amount of work has recently been done on matchmaking between buyers and sellers based on semantically matching the descriptions of the goods. Approaches include the use of description logics [11] and use of techniques from information retrieval [17]. This work is complementary to ours in the sense that the degree of match between descriptions can be used as yet another measure to factor into the decision of which suppliers to refer a business to.

There has recently been a large amount of work on the use of auctions to efficiently match buyers with sellers. Our balance problem is closely related to the combinatorial auction problem (CAP), in which buyers may bid on combinations of goods, and the value of a good to a buyer may be a function of the other goods that he wins. In the typical formulation of the CAP, the auction house is faced with a set of price offers for various bundles of goods and its objective is to allocate the goods so as to maximize its own revenue. This optimization problem is intractable in general, although researchers have identified numerous tractable special cases, typically expressed in terms of constraints on the bidding language [10, 15]. Other work has addressed developing heuristic techniques for general combinatorial auctions [13]. While early work on the CAP dealt only with single-unit CAPs, more recent work has dealt with multi-unit CAPs, in which there are multiple units of some goods available [9, 8]. Our balance problem can be transformed into a multi-unit CAP:

Theorem 7.1. *The balance problem can be reduced to a sealed-bid multi-unit CAP in time linear in the size of the requirements matrix. (Proof: See Appendix)*

8 Future Research

This research has raised a number of interesting questions which remain open for future research. We would like to make the simulator more realistic by adding things such as the probability that a business will follow a recommendation and the tendency of businesses to stick with suppliers. Our algorithm for matching buyers and suppliers may recommend that some companies go to many suppliers to satisfy a single product need. We need to be able to incorporate a bound on the number of suppliers. While we have separate solutions for the balance and fairness problems,

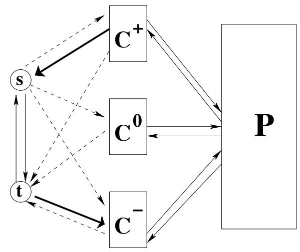

Figure 2: The edges with cost −\$1 are with thicker lines while the edges with cost \$1 are drawn with dotted lines.

we have yet to find a method of merging the solutions. Finally, we hope to combine the optimization techinques with prediction and recommendation modules to produce a more complete barter trade management system.

Acknowledgements
We thank Bob Bagga and Chris Haddawy at BizXchange for their time and patience in helping us to understand the operations of trade exchanges and for generously providing the transaction data. We thank Sumanta Guha for help in formalizing an earlier qualitative version of the balance problem. The convex hull lemma is due to Hans Volkmer. Thanks to Yongyos Kaewpitakkun and Thai Bui for help in implementing the matrix rounding algorithm. Thanks to Khaimook Dhananaiyapergse who built the simulator and ran the simulation experiments.

Appendix: Proofs and Examples

Lemma 4.5. *All negative-cost directed cycles in $R(G, f)$ cost $-\$2$ and consist of a sequence of nodes of the form* $t, c_{i_1}, p_{j_1}, c_{i_2}, p_{j_2}, \ldots, p_{j_k}, c_{i_{k+1}}, s, t$.

Proof: Let T be a maximal trade set, and f be its corresponding integer circulation on (G, l, u, γ). Recall that if (i) $b_i < 0$, then $f(c_i, t) = b_i$ and $f(s, c_i) = 0$, (ii) if $b_i > 0$, then $f(c_i, t) = 0$ and $f(s, c_i) = b_i$, and (iii) if $b_i = 0$, then $f(c_i, t) = 0 = f(s, c_i)$.

Let C^+, C^-, C^0 denote, respectively, the set of companies with positive, negative and zero balances. It is easy to verify that, in $R(G, f)$, only the edges from C^+ to s and the edges from t to C^- cost $-\$1$. All other edges between $\{s, t\}$ and C cost \$1 and all edges between C and P cost \$0. (See Figure 2.) Thus, every negative-cost directed cycle in $R(G, f)$ must contain exactly one edge from C^+ to s and/or exactly one edge from t to C^-; otherwise, the cycle must contain at least two edges from C^+ to s or at least two edges from t to C^-, violating the definition of a directed cycle. Such a cycle costs at least $-\$2$.

If the cycle contains an edge from t to C^-, but not an edge from C^+ to s, then the cycle can only return to t by using an edge from C to t. This leads to a contradiction because the cost of such a cycle is at least \$0. By a similar argument, the cycle cannot contain an edge from C^+ to s only, and not an edge from t to C^-. Let (t, c_{i_1}) be the directed edge out of t in the cycle and let $(c_{i_{k+1}}, s)$ be the directed edge into s in the cycle. To complete the cycle, there must be (i) a directed path from c_{i_1} to $c_{i_{k+1}}$ that does not make use of nodes s and t and (ii) a directed path from s to t. The directed path from c_{i_1} to $c_{i_{k+1}}$ has the form $c_{i_1}, p_{j_1}, c_{i_2}, p_{j_2}, \ldots, p_{j_k}, c_{i_{k+1}}$ and has cost \$0. The directed path from s to t either consists of the edge (s, t), which costs \$0, or contains edges from s to C and from C to t, which costs at least \$2. The former must be true; otherwise, the cycle would cost at least \$0, and it will not be a negative-cost directed cycle. $\square$

Lemma 4.8. *Let T be a fractional maximal trade set for the barter universe $\mathcal{B} = (C, P, \alpha, R)$. Suppose T' is an integer matrix obtained by setting each T'_{ij} equal to $\lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$ so that $\sum_{i=1}^{m} T'_{ij} = \sum_{i=1}^{m} T_{ij} = 0$ for each j. Then T' is an integer maximal trade set for $\mathcal{B}$ and $ab_{T'} \leq ab_T + m \sum_{j=1}^{n} \alpha_j$.*

Proof: We note that $0 \leq |T'_{ij}| \leq |R_{ij}|$ for each i and j because $0 \leq |T_{ij}| \leq |R_{ij}|$, $T'_{ij} = \lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$, and R is an integer matrix.

And since we assumed that $\sum_{i=1}^{m} T'_{ij} = 0$ for each j, T' is a valid trade set for $\mathcal{B}$. Furthermore, if the supply of p_j exceeds its demand, then each $c_i \in N_j$ has to buy all of its demand for p_j, so $T_{ij} = R_{ij}$, which is an integer. Hence,

$T'_{ij} = T_{ij} = R_{ij}$. The same equation holds true for all $c_i \in S_j$ when the supply of p_j is no more than its demand. Therefore, T' is a maximal trade set for $\mathcal{B}$. Finally, $ab_{T'} = \sum_{i=1}^m |b'_i| \le \sum_{i=1}^m |b_i| + \sum_{i=1}^m |b'_i - b_i|$. So

$$\begin{aligned} ab_{T'} &\le ab_T + \sum_{i=1}^m \left| \sum_{j=1}^n \alpha_j (T'_{ij} - T_{ij}) \right| \\ &\le ab_T + \sum_{i=1}^m \sum_{j=1}^n \alpha_j \left| (T'_{ij} - T_{ij}) \right| \\ &= ab_T + \sum_{i=1}^m \sum_{j=1}^n \alpha_j \\ &= ab_T + m \sum_{j=1}^n \alpha_j \end{aligned}$$

where the first equality makes use of the fact that T'_{ij} equals $\lceil T_{ij} \rceil$ or $\lfloor T_{ij} \rfloor$. □

Lemma 4.9. *Let m and z be positive integers with $z \le m$. Let $\mathcal{V}(m,z)$ be the set of all vectors in $\mathbf{R}^m$ that have z components equal to 1 and $m - z$ components equal to 0. (Note that $\mathcal{V}(m,z)$ contains $\binom{m}{z}$ vectors.) Let $\vec{x} = (x_1, \ldots, x_m) \in [0,1]^m$ so that $\sum_{i=1}^m x_i = z$. Then $\vec{x}$ can be expressed as a convex combination of the elements in $\mathcal{V}(m,z)$. That is, $\vec{x} = \sum_{\vec{v}_r \in \mathcal{V}(m,z)} \lambda_r \vec{v}_r$ so that $0 \le \lambda_r \le 1$ for each r, and $\sum_r \lambda_r = 1$.*

Proof: We shall show by induction that the following statement is true for every $m \in \mathbf{Z}^+$.

A_m: For all $\vec{x} = (x_1, x_2, \ldots, x_m) \in [0,1]^m$ such that $\sum_{i=1}^m x_i = z$, $z \in \{1, 2, \ldots, m\}$, $\vec{x}$ can be expressed as a convex combination of the elements in $\mathcal{V}(m,z)$.

When $m = 1, 2, 3$, it is easy to verify that A_m is true. Let us now prove the induction step from A_{m-1} to A_m. When $z = 1$, $\mathcal{V}(m,1)$ contains exactly m vectors: $\vec{e}_1, \vec{e}_2, \ldots, \vec{e}_m$, where $\vec{e}_i$ is the vector whose ith component equals 1, and whose other components equal 0. Hence, $\vec{x} = \sum_{i=1}^m x_i \vec{e}_i$. So let us assume that $1 < z \le m$.

Case 1. Suppose one of the x_i's equal 0, say x_m. Then consider the vector $\vec{x}\,' = (x_1, x_2, \ldots, x_{m-1})$. We know that $\sum_{i=1}^{m-1} x_i = z$. Since A_{m-1} is true, then there exists constants λ_r so that $\vec{x}\,' = \sum_{\vec{u}_r \in \mathcal{V}(m-1,z)} \lambda_r \vec{u}_r$ where $0 \le \lambda_r \le 1$ for each r, and $\sum_r \lambda_r = 1$.

For each $\vec{u}_r$, let $\vec{v}_r$ be the m-vector obtained by appending a 0 to $\vec{u}_r$. That is, if $\vec{u}_r = (u_{r,1}, \ldots, u_{r,m-1})$ then $\vec{v}_r = (u_{r,1}, \ldots, u_{r,m-1}, 0)$. Since $\sum_{i=1}^{m-1} u_{r,i} = z$, $\sum_{i=1}^m v_{r,i} = z$ as well. Therefore each $\vec{v}_r \in \mathcal{V}(m,z)$, and $\vec{x} = \sum_{\vec{v}_r \in \mathcal{V}(m,z)} \lambda_r \vec{v}_r$ because the mth component of $\vec{x}$ is zero.

Case 2. Suppose one of the x_i's equal 1, say x_m. Again, consider the vector $\vec{x}\,' = (x_1, x_2, \ldots, x_{m-1})$. This time, $\sum_{i=1}^{m-1} x_i = z - 1$. Since $1 < z \le m$, $1 \le z - 1 \le m - 1$. Because A_{m-1} is true, then there exists constants λ_r so that $\vec{x}\,' = \sum_{\vec{u}_r \in \mathcal{V}(m-1,z-1)} \lambda_r \vec{u}_r$ where $0 \le \lambda_r \le 1$ for each r, and $\sum_r \lambda_r = 1$.

For each $\vec{u}_r$, let $\vec{v}_r$ be the m-vector obtained by appending a one to $\vec{u}_r$. Using the same reasoning in case 1, $\vec{x} = \sum_{\vec{v}_r \in \mathcal{V}(m,z)} \lambda_r \vec{v}_r$.

Case 3. Suppose $0 < x_i < 1$ for $i = 1, \ldots, m$. Define $t_i = \frac{x_i}{1-x_i}$ if $i = 1, 2, \ldots, z$ and $\frac{1-x_i}{x_i}$ if $i = z+1, \ldots, m$. Let $t = \min\{t_i, 1 \le i \le m\}$, which is a positive number. Rewrite $\vec{x}$ as follows:

$$\vec{x} = \frac{t}{1+t}\vec{e} + \frac{1}{1+t}\vec{y}$$

where $\vec{e}$ is the m-vector whose first z components are 1 and whose remaining $m - z$ components are 0. Let $\vec{y} = (y_1, y_2, \ldots, y_m)$. It is straightforward to verify that (i) $0 \le y_i \le 1$ for each $i = 1, \ldots, m$, (ii) $\sum_{i=1}^m y_i = z$, and (iii) at least one y_i is 0 or 1. Hence, we have expressed $\vec{x}$ as the convex combination of a vector in $\mathcal{V}(m,z)$ and another vector $\vec{y}$, whose components sum to z, and one of which is equal to 0 or 1. According to case 1 or 2, $\vec{y}$ can be expressed as a convex combination of the vectors in $\mathcal{V}(m,z)$. By induction, A_m is true for all $m \in \mathbf{Z}^+$. □

CONVEX_COMB($\mathbf{x}[\mathbf{1}\ldots\mathbf{m}]$)

$k \leftarrow$ number of non-zero entries in x.
Initialize all entries of $A[1\ldots k][1\ldots m]$, $t[1\ldots k]$,$E[1\ldots m]$, and $\lambda[1\ldots k]$ to 0.
$iter \leftarrow 0$.
$I \leftarrow \emptyset$. /* I will keep track of all entries in x set to 1 */
While x is not a zero-vector
    $iter \leftarrow iter + 1$
    $I \leftarrow I \cup \{i : x[i] = 1\}$
    For each $i \in I$
        $A[iter][i] \leftarrow 1$
        $x[i] \leftarrow 0$
    endfor
    If x is a zero-vector
        $\lambda[iter] \leftarrow 1/\prod_{j=1}^{iter-1}(1 + t[j])$
        exit while loop
    Else do
        $(t[iter], E, x) \leftarrow$ **SPLIT-VECTOR($\mathbf{x}[\mathbf{1}\ldots\mathbf{m}]$)**
        For $i = 1$ to m
            $A[iter][i] = A[iter][i] + E[i]$
        endfor
        $\lambda[iter] \leftarrow t[iter]/\prod_{j=1}^{iter}(1 + t[j])$
    endif
endwhile
Return($A[1\ldots k][1\ldots m], \lambda[1\ldots k]$).

SPLIT-VECTOR($\mathbf{x}[\mathbf{1}\ldots\mathbf{m}]$)
Initialize $t[1\ldots m]$, $e[1\ldots m]$ and $y[1\ldots m]$ to 0.
$COUNT \leftarrow 0$.
$SUM \leftarrow \sum_{i=1}^{m} x[i]$.
For $i = 1$ to m
    If $x[i] = 0$
        $t[i] \leftarrow \infty$ (i.e., a very large number)
        $e[i] \leftarrow 0$
    Else
        $COUNT \leftarrow COUNT + 1$
        if $COUNT \leq SUM$
            $t[i] \leftarrow x[i]/(1 - x[i])$
            $e[i] \leftarrow 1$
        else
            $t[i] \leftarrow (1 - x[i])/x[i]$
            $e[i] \leftarrow 0$
        endif
    endif
endfor
$t^* \leftarrow \min\{t[i], 1 \leq i \leq m\}$
$y[1\ldots m] \leftarrow (1 + t^*)x[1\ldots m] - t^* e[1\ldots m]$
Return $(t^*, e[1\ldots m], y[1\ldots m])$.

Figure 3: Let $x = (x_1, x_2, \ldots, x_m) \in [0,1]^m$ such that $\sum_{i=1}^{m} x_i = z$, $z \in \{1, 2, \ldots, m\}$. CONVEX_COMB expresses x as a convex combination of the elements in $\mathcal{V}(m, z)$.

The proof above leads us to the algorithm in Figure 3.

An example. Let $x = [0.3, 0.5, 0.8, 0.5, 0.9]$ so that $SUM = 3$.

Iteration 1: SPLIT-VECTOR(x) will produce
$t^* = 1/9$, $e = [1, 1, 1, 0, 0]$ and $y = [2/9, 4/9, 7/9, 5/9, 9/9]$.
So $A[1] = [1, 1, 1, 0, 0]$ and $\lambda[1] = 1/10$.

Iteration 2: $I = \{5\}$, $A[2]$ is set to $[0, 0, 0, 0, 1]$ and $x = [2/9, 4/9, 7/9, 5/9, 0]$ is passed to SPLIT-VECTOR. It will produce
$t^* = 2/7$,
$e = [1, 1, 0, 0, 0]$ and
$y = [0, 2/7, 1, 5/7, 0]$.
So $A[2] = [1, 1, 0, 0, 1]$ and $\lambda[2] = 2/10$.

Iteration 3: $I = \{3, 5\}$, $A[3]$ is set to $[0, 0, 1, 0, 1]$ and $x = [0, 2/7, 0, 5/7, 0]$ is passed to SPLIT-VECTOR. It produces
$t^* = 2/5$
$e = [0, 1, 0, 0, 0]$
$y = [0, 0, 0, 1, 0]$.
$A[3] = [0, 1, 1, 0, 1]$ and $\lambda[3] = 2/10$.

Iteration 4: $I = \{3, 4, 5\}$ and $A[4] = [0, 0, 1, 1, 1]$. Since all the components of x are now zero, $\lambda[4] = 5/10$.

Hence, $x = \sum_{i=1}^{4} \lambda[i]A[i]$.

Theorem 4.10. *RANDOM_ROUND(T) outputs an integer maximal trade set T' such that for each i and j, $E[T'_{ij}] = T_{ij}$. Thus, for each i, $E[b'_i] = b_i$ and $ab_{T'} \le ab_T + m\sum_{j=1}^{n} \alpha_j$. The procedure runs in $O(K|C| + |P||C|)$ time where K is the number of fractional entries in T.*
Proof: When $T_{ij} \ge 0$, $T'_{ij} = \lfloor T_{ij} \rfloor + v^i$. Hence,

$$
\begin{aligned}
E[T'_{ij}] &= \lfloor T_{ij} \rfloor + E[v^i] \\
&= \lfloor T_{ij} \rfloor + (1 \times x_i + 0 \times (1 - x_i)) \\
&= \lfloor T_{ij} \rfloor + T_{ij} - \lfloor T_{ij} \rfloor \\
&= T_{ij}.
\end{aligned}
$$

The same result holds when $T_{ij} < 0$. Therefore, for each i, $E[b'_i] = E[\sum_{j=1}^{n} \alpha_j T'_{ij}] = \sum_{j=1}^{n} \alpha_j E[T'_{ij}] = \sum_{j=1}^{n} \alpha_j T_{ij} = b_i$. The bound for $ab_{T'}$ follows from Lemma 4.8. Rounding column j of T takes $O(k_j m)$ time, where k_j is the number of fractional entries in the column. Thus, the total running time of step 2 is $\sum_{j=1}^{n} O(k_j m) = O(Km)$, where K is the number of fractional entries in T. Steps 1 and 3 takes $O(nm)$ time in the worst case so RANDOM_ROUND takes $O(Km + nm)$ time. □

Theorem 7.1. *The balance problem can be reduced to a sealed-bid multi-unit CAP in time linear in the size of the requirements matrix.*
Proof: For the formal framework, we generalize the single-unit model of Nisan [10] to cover multi-unit CAPs. We assume a sealed-bid auction in which each bidder has his own private valuation function v_i, which specifies his valuation for each possible set of items he may receive.

We have m bidders who wish to bid on a set of goods $P = \{p_1, p_2, \ldots, p_n\}$. Let $q(j)$ be the number of available units of good j. An *allocation* is an $m \times n$ matrix G where G_{ij} is the number of units of good j allocated to buyer i such that $\sum_i G_{ij} \le q(j)$. A *full allocation* [8] is an allocation in which $\sum_i G_{ij} = q(j)$. Each bidder i has a valuation function v_i associated with it, where $v_i(a_1, a_2, \ldots, a_n)$ denotes the value of the bundle of goods in which a_j units of good j are allocated to bidder i. The allocation problem is then to find a full allocation G so that

$$\sum_{i=1}^{m} v_i(G_{i1}, G_{i2}, \ldots, G_{in})$$

is maximized. We shall consider the *constrained allocation problem*, which has the same objective, except that some of the entries in G have upper and lower bounds.

We now show how to represent the balance problem in the form of the above multi-unit CAP. Consider an $m \times n$ requirements matrix R. For each good p_j, set $q(j) = 2 \times \min(\sum_{i:c_i \in N_j} |R_{ij}|, \sum_{i:c_i \in S_j} R_{ij})$. If the supply of p_j exceeds its demand, then twice the total demand is auctioned (half of which will be assigned to all the buyers of p_j and the other half to all the sellers of p_j) and vice versa. Recall that α_j is the cost of good p_j in our barter universe $\mathcal{B}$. For each i and j, define the function $f_{ij}(a)$ as $-\alpha_j a$ if $c_i \in N_j$, and $\alpha_j a$ if $c_i \in S_j$, and 0 otherwise.

We specify the valuation functions as follows: for each i, let

$$v_i(a_1, a_2, \ldots, a_n) = -\left|\sum_{j=1}^{n} f_{ij}(a_j)\right|.$$

(This is an asymmetric valuation function in the terminology of Nisan.) Our goal is to find the full allocation G that maximizes $\sum_{i=1}^{m} v_i(G_{i1}, G_{i2}, \ldots, G_{in})$ subject to the condition that for each j, (i) if supply of p_j exceeds its demand then $G_{ij} = |R_{ij}|$ for each $c_i \in N_j$ and $0 \le G_{ij} \le R_{ij}$ for each $c_i \notin N_j$; (ii) if the supply of p_j is no more than its demand then $0 \le G_{ij} \le |R_{ij}|$ for each $c_i \notin S_j$ and $G_{ij} = R_{ij}$ for each $c_i \in S_j$. It is straightforward to verify that the optimal full allocation G^* is also the most balanced maximal trade set. Finally, computing $q(j)$ and the valuation functions from the requirements matrix takes at most $O(|P||C|)$ time so our theorem follows. $\square$

Example of Algorithm BALANCE

Recall the requirements matrix in Section 4. Let the initial trade T be the matrix shown below.

	p_1 (\$ 2)	p_2(\$3)
c_1	20	20
c_2	10	0
c_3	−20	15
c_4	−10	−35

The circulation that corresponds to the trade set is shown on the left. The residual graph $R(G, f)$ is shown on the right. Edge (t, s) and the other edges that have a cost \$1 are drawn with dotted lines. We do not indicate their costs, upper and lower bound capacities because they are not used for augmenting the circulation. In $R(G, f)$, the shortest negative-cost cycle from t to s is t, c_4, p_2, c_2, s, t and $\epsilon = 20$.

Augmenting flow along this cycle yields the residual graph on the left below. Its shortest negative-cost cycle is $t, c_4, p_2, c_2, p_1, c_1, s, t$ with $\epsilon = 40$. Augmenting the flow along this cycle, we end up the the residual graph on the right.

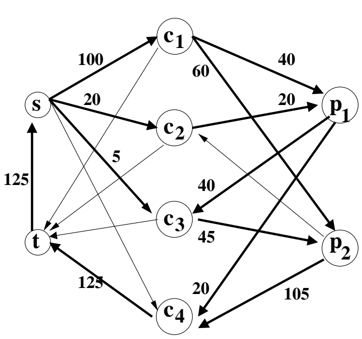

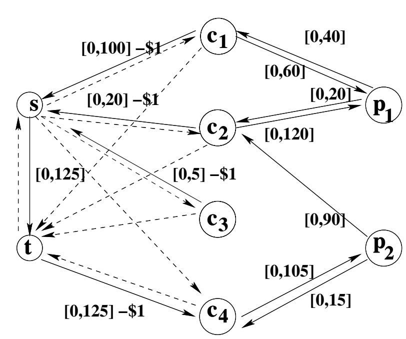

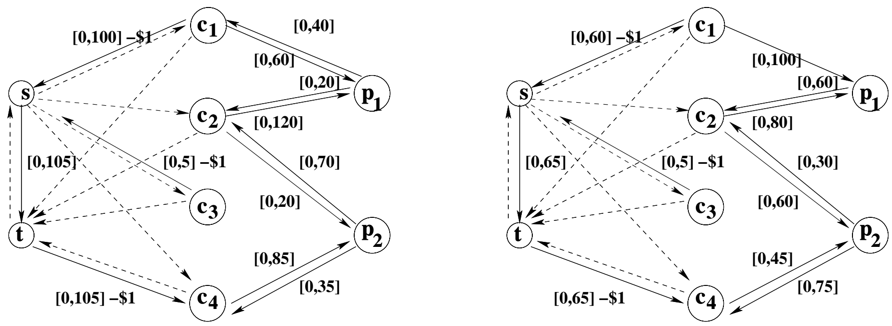

$R(G, f)$ no longer has a directed $s - t$ path so the trade set that corresponds to this circulation, which is shown below, is the most balanced fractional maximal trade set. Its absolute balance is $60 + 5 + 65 = 130$.

	p_1 (\$ 2)	p_2(\$3)
c_1	0	20
c_2	30	-20
c_3	-20	15
c_4	-10	-15

References

[1] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows*. Prentice Hall, 1993.

[2] International Reciprocal Trade Association, 1999. http://www.irta.net/barterstatistics.html.

[3] International Reciprocal Trade Association, 2001. www.irta.net.

[4] B. Cresti. Welfare improving barter in imperfect competition. Technical Report N. 2002/71, Center for Operations Research and Econometrics (CORE), 2002.

[5] B. Cresti. U.S. domestic barter: An empirical investigation. Technical Report N. 2003-5, Institut de Recherches Economiques et Sociales (IRES), 2003. http://www.ires.ucl.ac.be/DP/IRES_DP/2003-5.pdf.

[6] Apostolos Dailianas, Jakka Sairamesh, Vibby Gottemukkala, and Anant Jhingran. Profit-driven matching in e-marketplaces: Trading composable commodities. In *Agent Mediated Electronic Commerce (IJCAI Workshop)*, pages 153–179, 1999.

[7] A. Goldberg and R. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM*, 36(4):873–886, 1989.

[8] A. Kothari, D. Parkes, and S. Suri. Approximately-strategyproof and tractable multi-unit auctions. In *ACM Conference on Electronic Commerce (EC'03)*, pages 166–175, 2003.

[9] Kevin Leyton-Brown, Yoav Shoham, and Moshe Tennenholtz. An algorithm for multi-unit combinatorial auctions. In *AAAI/IAAI*, pages 56–61, 2000.

[10] Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce (EC'00)*, pages 1–12, 2000.

[11] T. Di Noia, E. Di Sciascio, F.M. Donini, and M. Mongiello. A system for principled matchmaking in an electronic marketplace. In *Proc. of the Twelfth Int'l Conf. on World Wide Web (WWW'03)*, pages 321–330, Budapest, Hungary, May 2003.

[12] M. Rossi, T. Saarinen, and V.K. Tuunainen. Introduction to special issue on new business models for electronic commerce. *ACM SIGMIS Database*, 34(2):7, 2003.

[13] Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *IJCAI*, pages 1102–1108, 2001.

[14] Arie Segev and Carrie Beam. Brokering strategies in electronic commerce marketplaces. In *ACM Conference on Electronic Commerce (EC'99)*, pages 167–176, 1999.

[15] Moshe Tennenholtz. Some tractable combinatorial auctions. In *AAAI/IAAI*, pages 98–103, 2000.

[16] Gaurav Tewari and Pattie Maes. Design and implementation of an agent-based intermediary infrastructure for electronic marketplaces. In *ACM Conference on Electronic Commerce (EC'00)*, pages 86–94, 2000.

[17] D. Veit, J. Muller, M. Schneider, and B. Fiehn. Matchmaking for autonomous agents in electronic marketplaces. In *Proc. Fifth Int'l Conf. on Autonomous Agents (AGENTS '01)*, pages 65–66, Montreal, Quebec, Canada, 2001.

[18] H. Volkmer. Department of Mathematical Sciences, University of Wisconsin-Milwaukee, 2003.