A Multi-Agent System for e-barter including Transaction and Shipping Costs*

Natalia López, Manuel Núñez, Ismael Rodríguez, and Fernando Rubio
Dept. Sistemas Informáticos y Programación
Universidad Complutense de Madrid
28040 Madrid. Spain
{natalia,mn,isrodrig,fernando}@sip.ucm.es

ABSTRACT

An *e-barter* multi-agent system consists of a set of agents exchanging goods. In contrast to e-commerce systems, transactions do not necessarily involve the exchange of money. Agents are equipped with a *utility function* to simulate the preferences of the customers that they are representing. They are grouped into *local* markets, according to the localities of the corresponding customers. Once these markets are *saturated* (i.e. no more exchanges can be performed) new agents, representing those local markets, are generated and combined into new markets. By reiteratively applying this process we finally get a *global market*.

Even though a formalism to define e-barter architectures has been already introduced, that framework had a strong drawback: Neither transaction nor shipping costs were considered. In this paper we extend that framework to deal with systems where fees have to be paid to the *owner* of the system. These fees are computed according to the goods involved in the corresponding exchanges. In addition, shipping costs have also to be paid. These modifications complicate the setting because the utility that customers receive after exchanging goods is not directly given by the original utility function. That is, the returned utility after an exchange is performed has to be computed as a combination of the *former* utility and the derived costs. In particular, some exchanges may be disallowed because those costs exceed the increase of utility returned by the new basket of goods.

Keywords

e-commerce, process algebras.

1. INTRODUCTION

During the last years there has been a growing interest in the study and implementation of systems/communities where intelligent (electronic) agents replace their (human) owners. In particular, there is ongoing research in technologies able to model users by means of agents which autonomously perform electronic transactions (see [6] for a survey on the topic). In order to increase the power of these agents they should be able to know the preferences of the corresponding user. In this case, the concept of *utility function* is very useful. A *utility function* returns a real number for each possible basket of goods: The bigger this number is, the happier the owner is with this basket. Intuitively, agents should act as the customer by considering the utility function that the corresponding user has in mind (see e.g. [18, 5, 3, 9, 13, 8]). Besides, a formal definition of the preferences of the user provides the agent with some negotiation capacity when interacting with other agents [7, 19, 10]. Let us remark that, in most cases, utility functions take a very simple form. For instance, they may indicate that a customer C is willing to exchange the item a by the items b and c. Moreover, there exists several proposals showing how agents can be trained on the user preferences (see e.g. [1, 3, 19]). Finally, a customer has always the possibility of changing the utility function. Once the agent is notified of these changes, it will change its negotiation strategy.

An e-barter system [12] consists of a set of agents performing exchanges of products. In contrast with the usual understanding of e-commerce, e-barter does not necessarily reduce all the transactions to *money*: An exchange is made if the involved parts are *happy* with their new items. Besides, e-barter allows a richer structure of exchanges. Suppose a very simple circular situation where for each $1 \leq i \leq r$ we have that agent A_i owns the good a_i and desires the good $a_{(i \mod r)+1}$ (see Figure 1). This multi-agent transaction can be easily performed within this framework. On the contrary, it would not be so easy to perform it if these items must be first *converted* into money. In fact, in case items are to be converted into money, the agent who desires the most expensive item would be unable to obtain it. So, the whole exchange will be deadlocked, even though all of the agents would get happier performing it. Actually, it could be thought that agents would be able to exchange the items provided that the price of all of the items is the same, but in that case we are not really using money: If all the items have the same price, any item can be used as currency unit, and what we obtain is a barter environment where money is not needed.

*Research supported in part by the CICYT project TIC2000-0701-C02-01.

SAC 2003, Melbourne, Florida, USA
©2003 ACM 1-58113-624-2/03/03 ...$5.00.

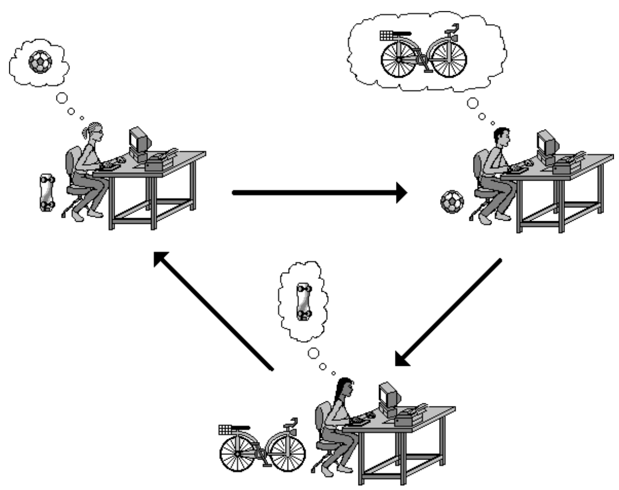

Figure 1: Exchange of items in the presence of circular dependencies.

However, in our environment money[1] can be considered just as another possible resource: An agent may be willing to exchange the good a by (m units of the good) money.

In order to formalize e-barter systems, the notions of *utility function*, *fair exchange*, and *equilibrium* are borrowed from Microeconomics (see [14] for a very formal and rigorous presentation of microeconomic theory). Let us suppose a system with k agents where n different products can be exchanged. Each agent has as information from the customer a pair $(\overline{x}, u)$, with $\overline{x} \in \mathbb{R}_+^n$ and $u : \mathbb{R}_+^n \to \mathbb{R}_+$. The first component of the pair denotes the amount of resources that the customer owns of each product. The second component is the utility function indicating the preferences of the customer with respect to the different products. That is, $u(\overline{x}) < u(\overline{y})$ denotes that the basket $\overline{y}$ is preferred to the basket $\overline{x}$. For example, $u(2,3) < u(3,1)$ means that it is preferred to own 3 units of the first product and 1 unit of the second one that to own 2 and 3 units, respectively, of the corresponding goods. A subset of agents will be willing to exchange resources if none of them decreases its utility and at least one of them improves. These exchanges are called *fair*. Formally, let us consider $\mathcal{A} = \{i_1, \ldots, i_m\} \subseteq \{1, \ldots, m\}$ and the pairs $(\overline{x_i}, u_i)$ for any $i \in \mathcal{A}$. Let us suppose that after the exchange we have that the information associated with the agents belonging to $\mathcal{A}$ is given by the pairs $(\overline{y_i}, u_i)$. The exchange is fair if for any $i \in \mathcal{A}$ we have $u_i(\overline{x_i}) \leq u_i(\overline{y_i})$ and there exists $j \in \mathcal{A}$ such that $u_j(\overline{x_j}) < u_j(\overline{y_j})$. Let us remark that a necessary condition for a exchange is that no products are created/destroyed, that is, $\sum_{i\in\mathcal{A}} \overline{x_i} = \sum_{i\in\mathcal{A}} \overline{y_i}$. Eventually, the system will reach a situation where no more exchanges can be performed. In other words, it is not possible to improve the situation of one agent without deteriorating another one. Such a situation is called *equilibrium* (also called *Pareto optimum*). In order to determine the equilibria of a system, techniques inherited from game theory can be used (see e.g. [21, 17, 20]).

In e-barter systems, agents are grouped according to the localities of the corresponding customers (see Figure 2. First, agents are combined into *local* markets (e.g. customers living in the same city). Once this market is *saturated*, that is, when no more exchanges can be performed, an agent representing the interests of all the agents in the market is created. The new agents will be again grouped into markets (e.g. agents are grouped by counties). This situation is repeated until a *global* market is created. This hierarchical structure presents at least two advantages. First, shipping costs are diminished because agents will exchange resources as close to the location of the customer as possible. Second, by creating new (representative) agents once a market is saturated and by combining them into higher order markets, we keep a small number of agents belonging to a certain market. This is very relevant if we take into account that a big number of agents would make very difficult to find the products that they are looking for. This is so because the number of messages that agents send to communicate with each other dramatically increases with the number of agents in the market. Finally, let us note that if an agent does not find the product that it is looking for in a local market, there will be a new agent looking for the same product (and taking into account the preferences of the original agent) in a wider market.

In order to formally specify e-barter systems we will use a process algebraic notation based on the language `PAMR` [16]. This language is very suitable for our purposes because it was specially developed to deal with the specification and analysis of concurrent and distributed systems where resources play a fundamental role. Unfortunately, `PAMR` does not provide a *higher order* constructor as the one needed in e-barter systems, so the language has to be extended. In addition to a syntax, we will provide an operational semantics for the new language. By doing so, every stage of the creation of an e-barter system may be formally specified, avoiding ambiguities and providing a clear structure of the system. Moreover, a designer of e-barter systems does not need to go through all the semantic machinery. It is enough to understand how the syntax of our language works.

In this paper we provide a formalism to specify e-barter systems where two of the main limitations of [12] are overcome. First, we consider that customers have to pay a fee depending on the products that they exchange. Second, shipping costs will be also collected. In order to compute shipping costs we have to take into account not only the products that the customer receives, but also the distance between the sender and the receiver. These two changes modify the behavior of systems. In particular, the set of final distributions of products will be reduced because some fair exchanges will not be performed due to the additional costs. Moreover, we have to adapt the notion of Pareto optimum to our framework. In microeconomics terms, the problem is that we partially lose the notion of *contract curve* because the induced generalized Edgeworth box shrinks after an exchange. Specifically, money is taken out from the system due to the extra costs. Let us illustrate these changes with a simple example.

EXAMPLE 1.1. *Let us consider a system with two users and two products. In addition, we consider money as the third product. Let us suppose that the initial distributions*

[1] Let us remark that this is not the usual treatment of money in general equilibrium theory. It is not considered as a resource and it is simply used to set *relative* prices.

Figure 2: Example of a hierarchical market structure.

are $(0, 3, 5)$ *and* $(2, 1, 10)$*, respectively, while the corresponding utility functions are defined as* $u_1(x_1, x_2, x_3) = 30 \cdot x_1 + 10 \cdot x_2 + x_3$ *and* $u_2(x_1, x_2, x_3) = 10 \cdot x_1 + 10 \cdot x_2 + x_3$*, respectively. Intuitively, the first user is indifferent between one unit of the first product and three units of the second product. If the first user gives two units of the second product in exchange for one unit of the first product, both users improve. That is,* $u_1(0, 3, 5) < u_1(1, 1, 5)$ *and* $u_2(2, 1, 10) < u_2(1, 3, 10)$*. However, this exchange could be disallowed if we consider transaction and shipping costs. In this case, we would have to decide whether we have both* $u_1(0, 3, 5) \leq u_1(1, 1, 5 - t(1, 0) - c_{1,2}(1, 0))$ *and* $u_2(2, 1, 10) \leq u_2(1, 3, 10 - t(0, 2) - c_{1,2}(0, 2))$*, where* t *is a function computing transaction costs and* $c_{1,2}$ *is a function computing shipping costs according to the distance between the users. Besides, in order to have a fair exchange, one of the previous inequalities must be strict. If the exchange is performed, the system will increase its amount of money by* $t(1, 0) + t(0, 2)$ *units. Thus, the total amount of money owned by the users is reduced in* $t(1, 0) + t(0, 2) + c_{1,2}(1, 0) + c_{1,2}(0, 2)$ *units.* □

The rest of the paper is structured as follows. In Section 2 we introduce some auxiliary notation. Section 3 represents the bulk of the paper. First, we give an informal description of the behavior of e-barter systems. Next, we present a formalization of all the necessary concepts to specify e-barter systems. Finally, in Section 4 we present our conclusions and some lines for future work.

2. BASIC CONCEPTS

In this section we introduce some concepts that we will use during the rest of this paper. Specifically, we present the notions of utility function and we explain how operational rules for a process algebra are defined. First we present some mathematical notation.

DEFINITION 2.1. *We consider* $\mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$*. We will usually denote vectors in* $\mathbb{R}^n$ *(for* $n \geq 2$*) by* $\overline{x}, \overline{y}, \ldots$*. Given* $\overline{x} \in \mathbb{R}^n$*,* x_i *denotes its* i*-th component. We extend to vectors some usual arithmetic operations. Let* $\overline{x}, \overline{y} \in \mathbb{R}^n$*. We define* $\overline{x} + \overline{y} = (x_1 + y_1, \ldots, x_n + y_n)$*, and* $\overline{x} \leq \overline{y}$*, if for any* $1 \leq i \leq n$ *we have* $x_i \leq y_i$*.*

We will usually denote matrices in $A^{n \times m}$ *(for* $n, m \geq 2$*, and a set* A*) by calligraphic letters* $\mathcal{E}, \mathcal{E}_1, \ldots$ □

The relevant characteristics of the customers of an e-barter system are their *baskets of resources* (indicating the items that they own) and their *utility functions* (indicating preference among different baskets of resources).

DEFINITION 2.2. *Let us suppose* $m > 0$ *different kinds of resources. Baskets of resources are defined as vectors* $\overline{x} \in \mathbb{R}_+^m$*. A utility function is a function* $u : \mathbb{R}_+^m \longrightarrow \mathbb{R}$*.* □

In microeconomic theory there are some restrictions that are usually imposed on utility functions (mainly strict monotonicity, convexity, and continuity). Intuitively, given a utility function u we have that $u(\overline{x}) < u(\overline{y})$ means that the basket $\overline{y}$ is preferred to $\overline{x}$.

Process Algebras (see [2] for a good overview on the topic) are formal languages used for the specification and verification of distributed and concurrent systems. As we pointed out in the introduction of this paper, we will use such a language to formalize e-barter systems. The syntax of these languages is given as an EBNF expression. In order to assign meaning to syntactic terms, an operational semantics is usually defined. Operational rules will be defined as usual deduction rules. That is, a rule

$$\frac{\text{Premise}_1 \ \wedge \ \text{Premise}_2 \ \wedge \ \ldots \ \wedge \ \text{Premise}_n}{\text{Conclusion}}$$

indicates that if all of the premises hold, then we can deduce the conclusion. Premises indicate individual behavior of components of a system, while conclusions indicate how the system behaves according to individual performances. Let us remark that when a rule has not got any premise, the conclusion trivially holds.

3. FORMALIZING E-BARTER SYSTEMS

In this section we present how e-barter systems are organized. First, we show the basic algorithm underlying the definition of e-barter systems. Next we introduce the formal framework to specify e-barter systems.

3.1 Basic Behavior of e-barter Systems

Customers willing to participate in an e-barter system are *represented* by (electronic) agents.[2] These agents are provided with two data: The *basket of resources* that the customer is willing to exchange and a *utility function*. The utility function relates the preference that a customer has for the owned items with respect to the desired items. Besides, once an agent has reached a (possibly multilateral) deal, it must be notified to the customer. If all the customers give their approval, the deal will be effectively performed, transaction fees will be added and shipping costs will be computed according both to the amount of received items and to the distance between the involved customers.

From now on we concentrate on the behavior of the different electronic entities. The behavior of an e-barter system works according to the following algorithm:

1. Each agent generates the barters that its customer would be willing to perform (according to the corresponding basket of resources and utility functions).

2. Agents exchange goods inside their local market. A multilateral exchange will be made if (at least) one of the involved agents improves its utility and none of them decreases its utility. This is repeated until no more exchanges are possible. In this case, we say that the local market is *saturated*.

3. Once a market is saturated, their agents are combined to create a new agent. The new agent will have as basket of resources the addition of the corresponding to each agent. Its utility function will encode the utilities of the combined agents. Let us remark that this new agent behaves as a representative of the combined agents. *First order* agents will be combined again into markets, according to *proximity* reasons.

4. Higher order agents trade between them until their market is saturated.

5. Once a (higher order) market is saturated, the agents start to allocate the resources in a top-down way through the tree of markets until the resources arrive to the leaves of the tree (i.e. the *original* agents). Then, they create a new agent (as indicated in step 3).

6. Once their markets are saturated, new markets are created by combining agents until there exists a unique market. Once this market is saturated, and the resources are conveniently allocated, the whole tree of agents is reset, and we start again at the first step.

As we already pointed out in the introduction, the previous behavior ensures some good properties. In particular, exchanges are made between agents located as near as possible, that is, shipping costs are minimized.

3.2 e-barter Systems: Syntax and Semantics

In this section we provide a formal syntax and semantics for the definition of e-barter systems. Even though we use a process algebraic notation (mainly when defining the operational rules) we do not need most of the usual operators for this kind of languages (choice, restriction, etc). In fact, our constructions remind a parallel operator as the one presented, for example, in CCS [15].

[2]In terms of [22], our agents present as information attitude belief (vs. knowledge), while as pro-attitudes we may consider commitment and choice (vs. intention and obligation).

DEFINITION 3.1. *A market system is given by the following EBNF:*

$$\begin{array}{lcl} MS & ::= & ms(M) \\ M & ::= & A \mid \mathbf{unsat}((M,\ldots,M), sh, pr) \mid \sigma(M) \\ A & ::= & (S, u, \overline{x}, sh, pr) \\ S & ::= & [\,] \mid [A,\ldots,A] \end{array} \qquad \square$$

First, in order to avoid ambiguity of the grammar we annotate market systems with the non-terminal symbol ms. Intuitively, $M = (S, u, \overline{x}, sh, pr)$ (that is, $M = A$) represents a *saturated* market, that is, a market where no more exchanges can be performed among its agents. u denotes the utility function of M and $\overline{x}$ represents the basket of resources owned by M. We consider that there are m different products,[3] that is $\overline{x} \in \mathbb{R}^m_+$, and that the amount of *money* is placed in the last component of the tuple. Besides, sh is the *shipping function* indicating the shipping cost of each possible transaction in this market. In turn, pr is the profit collected by the market due to transaction costs. We will assume the existence of another function which will be common to all of the markets, the *transaction function*, denoted by tr. The function tr computes the transaction costs for each of the agents involved in an exchange by taking into account the goods that this agent receives. Let us remark that while the shipping costs will depend on the market in which the transaction is performed, the transaction costs will not. By doing so we can formally specify that shipping costs increase with the distance between customers.

Regarding the first argument of M, there are two possible situations. Either S is an empty list or not. In the first case, we have that M represents an *original* agent, that is, a direct representative of a customer (note that a single agent is trivially saturated). In the second case, if $S = [A_1, \ldots, A_n]$ we have that M represents an agent associated with the (possible higher order) agents $A_1, \ldots, A_n$ belonging to a saturated market. $M = \mathtt{unsat}((M_1, \ldots, M_n), sh, pr)$ represents an *unsaturated* market consisting of the markets $M_1, \ldots, M_n$, the shipping function sh and the profit value pr. Let us remark that in this case some of the submarkets may be saturated. Once all the markets of the system are saturated, the whole system is turned again into unsaturated. The term $\sigma(M)$ will represent that such operation must be performed on M.

Next, we present an example showing how an e-barter system may be constructed. In this example we will also (informally) introduce operational transitions of the language.

EXAMPLE 3.2. *Let us consider a total of six agents* $A_i = ([\,], u_i, \overline{x_i}, sh_0, 0)$*, for* $1 \leq i \leq 6$*, where* sh_0 *denotes a dummy shipping function. We suppose that these agents are grouped*

[3]We are assuming that all the items are *goods*. Nevertheless, agents could also trade *bads*. For example, a customer would be willing to give an apple pie if he *receives* minus s brown leaves in his garden. However, bads are usually not considered in microeconomic theory, as they can be easily turned into goods: Instead of considering the amount of leaves, one may consider the absence of them.

into three different markets which are initially unsaturated, so we make the following definitions:

$$\begin{array}{rcl} M_1 & = & \textbf{\textit{unsat}}((A_1, A_2), sh_1, 0) \\ M_2 & = & \textbf{\textit{unsat}}((A_3, A_4), sh_2, 0) \\ M_3 & = & \textbf{\textit{unsat}}((A_5, A_6), sh_3, 0) \end{array}$$

Suppose that the first two markets are linked, and the resulting market is also linked with the remaining M_3. We should add the following definitions:

$$\begin{array}{rcl} M_4 & = & \textbf{\textit{unsat}}((M_1, M_2), sh_4, 0) \\ M_5 & = & \textbf{\textit{unsat}}((M_4, M_3), sh_5, 0) \end{array}$$

Finally, the global market is defined as $M = ms(M_5)$.

Following the philosophy explained in the previous section, transactions will be made within a market only between saturated submarkets. So, only M_1, M_2, and M_3 are allowed to perform transactions (note that original agents are trivially saturated).

We will denote exchange of resources by $\leadsto$. Suppose that after some exchanges, M_1 becomes saturated. That is, there exists a sequence of exchanges $M_1 \leadsto M_1^1 \leadsto M_1^2 \cdots \leadsto M_1^n = M_1'$ such that $M_1' \not\leadsto$. In this case, the market grouping the first two agents should be labeled as saturated. So, the agents effectively perform all the achieved transactions becoming A_1' and A_2'. Then, the first market will be turned into $([A_1', A_2'], f(u_1, u_2), \overline{x_1} + \overline{x_2} - (0, 0, \ldots, costs_1), sh_1, pr_1)$, where f is a function combining utility functions (such a function will be formally defined), pr_1 is the amount of money that the system has obtained due to the fees applied to the exchange of goods, and $costs_1$ denotes the transaction and shipping costs associated with the exchanged products in the market M_1. In parallel, M_2 will have a similar behavior.

Once both M_1 and M_2 are saturated, the transactions between them will be allowed. Note that these transactions (inside the market M_4) will be performed according to the new utility functions, $f(u_1, u_2)$ and $f(u_3, u_4)$ respectively, and to the new baskets of resources, $\overline{x_1} + \overline{x_2} - (0, 0, \ldots, costs_1)$ and $\overline{x_3} + \overline{x_4} - (0, 0, \ldots, costs_2)$, respectively.

The process will iterate after M_5 gets saturated. Finally, we will have a market as $\sigma(M_5')$. Then, the global market is structurally reset. □

In order to simplify forthcoming operational rules we introduce the following notation to deal with utility functions. Utility functions associated with original agents (that is, $A = ([\ ], u, \overline{x}, sh, 0)$, where sh is useless) will behave as explained in Definition 2.2. That is $u(\overline{z})$ indicates the relative preference shown by A towards the basket of resources $\overline{z}$. Nevertheless, if $A = ([A_1, \ldots, A_n], u, \overline{x}, sh, pr)$ then we will consider that in addition to its usual meaning, the utility function also keeps track of how a basket of resources is distributed among the (possible higher order) agents $A_1, \ldots, A_n$. That is, $u(\overline{z}) = (r, \overline{z_1}, \ldots, \overline{z_n})$, where r still represents the utility, while $\sum \overline{z_i} = \overline{z}$ and $\overline{z_i}$ denotes the portion of the basket $\overline{z}$ assigned to A_i. Overloading the notation, if we simply write $u(\overline{z})$ we are referring to the first component of the tuple, while $u(\overline{z}).i$ denotes the $(i+1)$-*th* component of the tuple.

In the next definition we present the *anchor case* of our operational semantics. In order to perform complex exchanges, agents should first indicate the barters they are willing to accept.

Definition 3.3. *Let $A = (S, u, \overline{x}, sh, pr)$ be a saturated market. The exchanges the agent A would perform are given by the following operational rules:*

$$\frac{u(\overline{x}+\overline{y}) \geq u(\overline{x}) \ \wedge \ (\overline{x}+\overline{y}) \geq \overline{0}}{(S,u,\overline{x},sh,pr) \xrightarrow{\overline{y}} (S,u,\overline{x}+\overline{y},sh,pr)}$$

$$\frac{u(\overline{x}+\overline{y}) > u(\overline{x}) \ \wedge \ (\overline{x}+\overline{y}) \geq \overline{0}}{(S,u,\overline{x},sh,pr) \stackrel{\overline{y}}{\longmapsto} (S,u,\overline{x}+\overline{y},sh,pr)}$$

where $\overline{y} \in \mathbb{R}^n$, being n the different kinds of available resources. □

Let us remark that $\overline{y}$ may have negative components. Actually, these tuples will contain the barters offered by the agent. For example, if $\overline{y} = (1, -1, 0, -3)$ fulfills the premise, then the agent would accept a barter where it is offered one unit of the first product in exchange of one unit of the second good and three units of money. Regarding the rules, the first premise simply indicates that the agent would not decrease (resp. would increase) its utility. The second premise indicates that the agent does not run into *red numbers*, that is, an agent cannot offer a quantity of an item if it does not own enough. Thus, a transition as $\longrightarrow$ denotes that the market does not worsen, meanwhile a transition $\longmapsto$ denotes that the market does improve. Finally, let us comment that even though transaction and shipping costs do not explicitly appear in the previous rules, they are implicitly reflected in the last component of the corresponding tuples $\overline{y}$. We will later formalize how these costs are assigned to the *owner* of the system and to the shipping companies. Next we will show how the offers made by the agents are combined.

Definition 3.4. *Let $M = \textbf{\textit{unsat}}((M_1, \ldots, M_n), sh, pr)$. Let $I = \{s_1, \ldots, s_r\} \subseteq \{1, \ldots, n\}$ be a set of indexes denoting the saturated markets belonging to M (that is, for any $i \in I$ we have $M_i = (S_i, u_i, \overline{x_i}, sh_i, pr_i)$). We say that the matrix $\mathcal{E} \in (\mathbb{R}_+^m)^{n \times n}$ is a valid exchange matrix for M under the cost tuple $\overline{c}$, denoted by $\textbf{\textit{valid}}(M, \mathcal{E}, \overline{c})$, if for any $1 \leq i \leq n$ we have $\sum_j \mathcal{E}_{ij} \leq \overline{x_i} - (0, \ldots, 0, c_i)$, $\mathcal{E}_{ii} = \overline{0}$, and $\forall\ 1 \leq k \leq n$ such that $k \notin I$, $\mathcal{E}_{ki} = \overline{0}$ and $\mathcal{E}_{ik} = \overline{0}$.* □

First, let us remark that the notion of *valid* matrix is considered only in the context of unsaturated markets: If a market is saturated then no more exchanges can be performed. Second, only saturated markets belonging to an unsaturated market may perform exchanges among them. This restriction is imposed in order to give priority to transactions performed by *closer* agents belonging to unsaturated submarkets. Regarding the definition of *valid* matrix, let us note that matrixes $\mathcal{E}$ have as components baskets of resources (that is, elements belonging to $\mathbb{R}_+^m$). $\mathcal{E}_{ij}$ represents the basket of resources that the market M_i would give to M_j. In the tuple $\overline{c}$, c_i denotes the transaction and shipping cost that agent M_i will have to afford. So, the condition $\sum_j \mathcal{E}_{ij} \leq \overline{x_i} - (0, \ldots, 0, c_i)$ indicates that the total amount of resources given by market M_i must be less than or equal to the basket of resources owned by that market minus the money paid by the transaction. Finally, let us comment that an exchange does not need to include all of the saturated markets. For example, if we have an exchange where only r' markets participate, then the rows and columns corresponding to the remaining $r - r'$ saturated markets will be filled with $\overline{0}$, as they are for the unsaturated markets.

Next we introduce the rules defining the exchange of resources. Intuitively, if we have a valid exchange matrix, where (at least) one of the involved agents improves and no one worsens, then the corresponding exchange will be performed.

Definition 3.5. *Let* $M = \boldsymbol{unsat}((M_1, \dots, M_n), sh, pr)$. *Let* $I = \{s_1, \dots, s_r\} \subseteq \{1, \dots, n\}$ *be a set of indexes denoting the saturated markets belonging to* M *(that is, for any* $i \in I$ *we have* $M_i = (S_i, u_i, \overline{x_i}, sh_i, pr_i)$*). The operational transitions denoting exchange of resources that* M *may perform are given by the rule shown in Figure 3. We say that* M *is a local equilibrium, denoted by* $M \not\rightsquigarrow$*, if there do not exist* $M', \mathcal{E}$ *such that* $M \overset{\mathcal{E}}{\rightsquigarrow} M'$. □

The operational rule shown in Figure 3 is applied under the same conditions appearing in the definition of a valid exchange matrix: It is applied to unsaturated markets and the exchange is made among a subset of the saturated submarkets. The premises indicate that (at least) an unsaturated market will improve after the exchange and that no one deteriorates. Let us remind that, in general, a market may generate both $M_i \xrightarrow{\overline{y}} M_i'$ and $M_i \overset{\overline{y}}{\longmapsto} M_i'$. So, the previous rule also considers situations where more than a market improves (we only require that at least one improves). Besides, let us remark that $M_i \xrightarrow{\overline{0}} M_i'$ always holds. So, a market not involved in the current exchange does not disallow the exchange. The costs required to have a valid exchange matrix will be computed both from the transaction and shipping costs. Regarding the conclusion, submarkets belonging to M are modified according to both the corresponding exchange matrix and the costs of the exchange, while unsaturated submarkets do not change. Let us remark that the costs of each exchange will be paid by the receiver. Besides, only the transaction costs will be added to the cumulated profit of the market. The following result follows from the previous definition. It indicates that exchanges allowed by the previous rule are fair.

Proposition 3.6. *Let* $M = \boldsymbol{unsat}((M_1, \dots, M_n), sh, pr)$ *be an unsaturated market. Let* $I = \{s_1, \dots, s_r\} \subseteq \{1, \dots, n\}$ *be a set of indexes denoting the saturated markets belonging to* M *(that is, for any* $i \in I$ *we have* $M_i = (S_i, u_i, \overline{x_i}, sh_i, pr_i)$*). Let us suppose that it is possible to perform the transition* $\boldsymbol{unsat}((M_1, \dots, M_n), sh, pr) \overset{\mathcal{E}}{\rightsquigarrow} \boldsymbol{unsat}((M_1', \dots, M_n'), sh, pr')$*, where for all* $i \in I$ *we have* $M_i' = (S_i, u_i, \overline{x_i'}, sh_i, pr_i)$*). Then for any* $i \in I$ *we have* $u_i(\overline{x_i}) \le u_i(\overline{x_i'})$ *and there exists* $j \in I$ *such that* $u_j(\overline{x_j}) < u_j(\overline{x_j'})$. □

We need to consider two more exchanging rules.

$$\frac{M_k \overset{\mathcal{E}}{\rightsquigarrow} M_k'}{\mathtt{unsat}((M_1,\dots,M_k,\dots,M_n),sh,pr) \overset{\mathcal{E}}{\rightsquigarrow} \mathtt{unsat}((M_1,\dots,M_k',\dots,M_n),sh,pr)}$$

$$\frac{M \overset{\mathcal{E}}{\rightsquigarrow} M'}{ms(M) \overset{\mathcal{E}}{\rightsquigarrow} ms(M')}$$

The first one indicates that if an unsaturated submarket produces an exchange, then the market must take that situation into account. The second rule reflects modifications in the environment of the constructor ms.

If a market reaches an equilibrium then we need to modify the attribute of the market, replacing a term of the form $\mathtt{unsat}((M_1, \dots, M_n), sh, pr)$ by a term as $(S, u, \overline{x}, sh, pr')$. Once a market is saturated, the money collected in the different submarkets as transaction costs will be transferred to it. In addition, resources are recursively moved from the corresponding agents to the leaves of the tree (indicating customers). Let us remark that a necessary condition for a market to be saturated is that all of its submarkets are also saturated. The following rule uses two auxiliary notions that will be formally presented in the forthcoming Definition 3.8.

Definition 3.7. *Let* $M = \boldsymbol{unsat}((M_1, \dots, M_n), sh, pr)$ *be a market, where we have that* $M_i = (S_i, u_i, \overline{x}_i, sh_i, pr_i)$ *for any* $1 \le i \le n$*. The following rule modifies the market from unsaturated to saturated:*

$$\frac{M \not\rightsquigarrow}{M \rightsquigarrow ([M_1', \dots, M_n'], u, \sum \overline{x_i}, sh, pr + \sum_i pr_i)}$$

where $u = \boldsymbol{CreateUtility}(u_1, \dots, u_n, \overline{x_1}, \dots, \overline{x_n})$ *and for any* $1 \le i \le n$ *we have that* $M_i' = (S_i', u_i, \overline{x_i}, sh_i, 0)$ *and* $S_i' = \boldsymbol{Deliver}(S_i, u_i, \overline{x_i})$. □

Let us remark that in this rule we do not label the transition $\rightsquigarrow$. These transitions play a role similar to internal transitions in classical process algebras. We need to add two more rules, as in the previous case, to record this transformation in the context of different constructors:

$$\frac{M_k \rightsquigarrow M_k'}{\mathtt{unsat}((M_1,\dots,M_k,\dots,M_n),sh,pr) \rightsquigarrow \mathtt{unsat}((M_1,\dots,M_k',\dots,M_n),sh,pr)}$$

$$\frac{M \rightsquigarrow M'}{ms(M) \rightsquigarrow ms(M')}$$

Next we present the pending functions. Intuitively, the function $\mathtt{Deliver}(S, u, \overline{x})$ distributes the basket of resources $\overline{x}$ among the original agents appearing in the leaves of the tree S. This distribution considers both the utility functions of the agents and the quantities of resources contributed by each of the agents. In addition to that, the function $\mathtt{CreateUtility}(u_1, \dots, u_n, \overline{x_1}, \dots, \overline{x_n})$ computes a combined utility function from the ones provided as arguments, so that it is possible to negotiate for maximizing the overall profit of the represented agents. Let us remind that, in this section, utility functions associated with higher order agents do not only reveal preference. In addition, they also take into account how resources will be distributed among agents. Thus, if we are considering an agent representing n agents, a new utility function returning a tuple with $n + 1$ components will be created. The first component (the value of the utility function) will return the worst utility (0) if any of the represented agents worsens. In this way, it is guaranteed that the market does not make any exchange which deteriorates any of its clients. Otherwise, the value of the utility will be the addition of individual utilities in the distribution of resources which maximizes this value.

Definition 3.8. *Let* $A = (S, u, \overline{x}, sh, pr)$ *be an agent. The allocation of the basket of resources* $\overline{x}$ *among the agents belonging to* S *with respect to the utility function* u*, denoted by* $\boldsymbol{Deliver}(S, u, \overline{x})$*, is recursively defined as:*

$$\boldsymbol{Deliver}(S, u, \overline{x}) = \begin{cases} [\,] & \text{if } S = [\,] \\ [M_1', \dots, M_n'] & \text{if } S = [M_1, \dots, M_n] \end{cases}$$

where for any $1 \le i \le n$ *we have* $M_i = (S_i, u_i, \overline{x_i}, sh_i, pr_i)$*,* $M_i' = (S_i', u_i, u(\overline{x}).i, sh_i, pr_i)$*, and* $S_i' = \boldsymbol{Deliver}(S_i, u_i, u(\overline{x}).i)$.

$$\frac{\exists\, k \in I : M_k \overset{\overline{y_k}}{\longmapsto} M'_k \ \wedge\ \forall\, i \in I, M_i \overset{\overline{y_i}}{\longrightarrow} M'_i \ \wedge\ \texttt{valid}(M, \mathcal{E}, (cost_1, \ldots, cost_n))}{M \overset{\mathcal{E}}{\rightsquigarrow} \texttt{unsat}\,((M'_1, \ldots, M'_n), sh, pr + \sum_i tr(\sum_j \mathcal{E}_{ji})) \qquad \left[\ M'_i = \begin{cases} M_i & i \notin I \\ (S_i, u_i, \overline{x_i} + \overline{y_i}, sh_i, pr_i) & \text{otherwise} \end{cases}\ \right]}$$

$$\begin{array}{lrcl} \text{where} & \mathcal{E} \in (\mathbb{R}_+^m)^{n\times n} & & \\ & \overline{y_i} & = & \sum_j \mathcal{E}_{ji} - \sum_j \mathcal{E}_{ij} - (0, \ldots, 0, cost_i) \\ & cost_i & = & tr(\sum_j \mathcal{E}_{ji}) + sh(\sum_j \mathcal{E}_{ji}) \end{array}$$

Figure 3: Operational rule for the exchange of resources in an unsaturated market.

Given n pairs $(u_i, \overline{x_i})$ the utility function is defined from the utility functions $u_1, \ldots, u_n$ with respect to the initial baskets of resources $\overline{x_1}, \ldots, \overline{x_n}$, and we denote it by using $\mathtt{CreateUtility}(u_1, \ldots, u_n, \overline{x_1}, \ldots, \overline{x_n})$, *as:*

$$\mathtt{CreateUtility}(u_1, \ldots, u_n, \overline{x_1}, \ldots, \overline{x_n}) = u_{\mathrm{market}}$$

where $u_{market}(\overline{x}) = max(\{(r, \overline{x'_1}, \ldots, \overline{x'_n}) | r = \sum_{1\leq i\leq n} u_i(\overline{x'_i}) \wedge \sum_{1\leq i\leq n} \overline{x'_i} = \overline{x} \wedge \bigwedge_{1\leq i\leq n} u(\overline{x'_i}) \geq u(\overline{x_i})\})$, *maximizing over the first argument (representing the utility), and assumming* $\max(\emptyset) = (0, \overline{0}, \ldots, \overline{0})$. □

Next, in order to define how a market evolves, we compose sequences of transitions.

Definition 3.9. *We say that a market M evolves into a market M', and we write $M \rightsquigarrow^* M'$, if there exist markets $M_1, \ldots, M_{n-1}$ such that*

$$M \overset{a_1}{\rightsquigarrow} M_1 \overset{a_2}{\rightsquigarrow} M_2 \overset{a_3}{\rightsquigarrow} \cdots M_{n-1} \overset{a_n}{\rightsquigarrow} M'$$

where for any $1 \leq i \leq n$ we have that a_i is an empty label or an exchange matrix. □

Finally, we provide a mechanism to reset a global market. If the root of the tree becomes a saturated market, then the whole tree of markets is created again. This is done by considering the five rules shown in Figure 4. The first one defines how it is launched the process to turn all the markets back to unsaturated mode, provided that the global market has become saturated. In this case, the transition is labelled by the global amount of money collected by the whole market system as transaction fees. The other rules define how to recursively reset the tree from the root to the leaves (original customer).

As we have already commented, the addition of transaction and shipping costs does not allow to properly speak about Pareto optimums. First, as the following result states, if we set these costs to zero then we obtain that the last saturated market represents one of the possible Pareto optimums for the whole set of original agents (regardless their locality in the market structure).

Theorem 3.10. *Let M, M' be markets and A be an agent such that*

$$M \rightsquigarrow^* M' \hookrightarrow ms(\sigma(A))$$

If transition and shipping costs are set to zero then we have that the distribution of resources provided by $\sigma(A)$ represents a Pareto optimum with respect to the original agents belonging to M.

Proof Sketch: If the final situation were not a Pareto equilibrium then there would exist at least one more fair exchange. Nevertheless, according to the operational semantics, this exchange would have been performed in the lowest market which embraces all of the agents involved in this hypothetical exchange, which makes a contradiction. □

Regarding the general situation of our system where the transaction and shipping costs are greater than zero, the classical Pareto equilibrium concept does not apply. For instance, a desirable exchange in a free or low transaction fee environment could not be a desirable exchange in an expensive one. Then, the equilibriums are different because of the effect of the fees. This framework cannot be defined in terms of one more agent representing the effect of the market fees, because this agent would have the special ability to enable and disable some exchanges. Therefore, we need a more general concept of equilibrium.

Definition 3.11. *A (t, s)-Pareto equilibrium is a distribution of resources in which no more fair exchanges are possible according to the transaction costs function t and the shipping costs function s.* □

The following result relates (t, s)-Pareto equilibria and sequences of transitions. The proof is similar to the one for Theorem 3.10.

Theorem 3.12. *Let M, M' be markets and A be an agent such that*

$$M \rightsquigarrow^* M' \hookrightarrow ms(\sigma(A))$$

We have that $\sigma(A)$ represents a (t, s)-Pareto equilibrium with respect to the original agents belonging to M. □

4. CONCLUSION AND FUTURE WORK

In this paper we have presented a formal framework to specify e-barter systems where transaction and shipping costs are considered. These costs modify the usual understanding of Pareto optimum because exchanges are now constrained by these additional costs. We are currently working in two opposite lines regarding our framework for e-barter. First, we are implementing a real e-barter system. By now, we restrict utility functions to be linear. Second, we are studying theoretical extensions of our model to allow a more thorough understanding of e-barter systems. In this line, we are studying properties of e-barter systems by using the testing semantics defined in [4]. This semantics was designed to

$$\overline{ms((S,u,\overline{x},sh,pr)) \Rightarrow_{pr} ms(\sigma((S,u,\overline{x},sh,0)))} \qquad \overline{\sigma(([\,],u,\overline{x},sh,pr)) \hookrightarrow ([\,],u,\overline{x},sh,pr)}$$

$$\frac{S=[A_1,\ldots,A_n]}{\sigma((S,u,\overline{x},sh,pr)) \hookrightarrow \texttt{unsat}((\sigma(A_1),\ldots,\sigma(A_n)),sh,pr)}$$

$$\frac{M_k \hookrightarrow M'_k}{\texttt{unsat}((M_1,\ldots,M_k,\ldots,M_n),sh,pr) \hookrightarrow \texttt{unsat}((M_1,\ldots,M'_k,\ldots,M_n),sh,pr)} \qquad \frac{M \hookrightarrow M'}{ms(M) \hookrightarrow ms(M')}$$

Figure 4: Reseting a global market.

find deadlocks in concurrent systems. Thus, it is very appropriate to determine the possible equilibria in an e-barter system. Moreover, we are considering the stochastic framework presented in [11] in order to (probabilistically) quantify the non-deterministic decisions taken by the agents.

5. ACKNOWLEDGMENTS

The authors thank the anonymous referees for valuable comments on the previous version of the paper.

6. REFERENCES

[1] F. Bacchus and A. Grove. Graphical models for preference and utility. In *Uncertainty in Artificial Intelligence (UAI-95)*, pages 3–10. Morgan Kaufmann, 1995.

[2] J.A. Bergstra, A. Ponse, and S.A. Smolka, editors. *Handbook of Process Algebra*. North Holland, 2001.

[3] M. Dastani, N. Jacobs, C.M. Jonker, and J. Treur. Modelling user preferences and mediating agents in electronic commerce. In *Agent Mediated Electronic Commerce, LNAI 1991*, pages 163–193. Springer, 2001.

[4] D. de Frutos, M. Núñez, and J. Quemada. Characterizing termination in LOTOS via testing. In *Protocol Specification, Testing and Verification XV*, pages 237–250. Chapman & Hall, 1995.

[5] T. Eymann. Markets without makers - a framework for decentralized economic coordination in multiagent systems. In *WELCOM 2001, LNCS 2232*, pages 63–74. Springer, 2001.

[6] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *The Knowledge Engineering Review*, 13(2):147–159, 1998.

[7] S. Kraus. Negotiation and cooperation in multi-agent systems. *Artificial Intelligence*, 94(1-2):79–98, 1997.

[8] J. Keppens and Q. Shen. A Calculus of Partially Ordered Preferences for Compositional Modelling and Configuration. In *AAAI Workshop on Preferences in AI and CP: Symbolic Approaches*, pages 39–46. AAAI Press, 2002.

[9] J. Lang, L. van der Torre, and E. Weydert. Utilitarian desires. *Autonomous Agents and Multi-Agent Systems*, 5(3):329–363, 2002.

[10] A.R. Lomuscio, M. Wooldridge, and N.R. Jennings. A classification scheme for negotiation in electronic commerce. In *Agent Mediated Electronic Commerce, LNAI 1991*, pages 19–33. Springer, 2001.

[11] N. López and M. Núñez. A testing theory for generally distributed stochastic processes. In *CONCUR 2001, LNCS 2154*, pages 321–335. Springer, 2001.

[12] N. López, M. Núñez, I. Rodríguez, and F. Rubio. A formal framework for e-barter based on microeconomic theory and process algebras. In *Innovative Internet Computer Systems, LNCS 2346*, pages 217–228. Springer, 2002.

[13] M. McGeachie and J. Doyle. Utility Functions for Ceteris Paribus Preferences. In *AAAI Workshop on Preferences in AI and CP: Symbolic Approaches*, pages 33–38. AAAI Press, 2002.

[14] A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[15] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.

[16] M. Núñez and I. Rodríguez. PAMR: A process algebra for the management of resources in concurrent systems. In *FORTE 2001*, pages 169–185. Kluwer Academic Publishers, 2001.

[17] S. Parsons and M. Wooldridge. Game theory and decision theory in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 5(3):243–254, 2002.

[18] L. Rasmusson and S. Janson. Agents, self-interest and electronic markets. *The Knowledge Engineering Review*, 14(2):143–150, 1999.

[19] T. Sandholm. Agents in electronic commerce: Component technologies for automated negotiation and coalition formation. In *CIA'98, LNCS 1435*, pages 113–134. Springer, 1998.

[20] W.C. Stirling, M.A. Goodrich, and D.J. Packard. Satisficing equilibria: A non-classical theory of games and decisions. *Autonomous Agents and Multi-Agent Systems*, 5(3):305–328, 2002.

[21] M. Tennenholtz. Game theory and artificial intelligence. In *Foundations and Applications of Multi-Agent Systems, LNAI 2403*, pages 49–58. Springer, 2002.

[22] M. Wooldridge and N.R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.