Posit Spaces: A Performative Model of e-Commerce

Peter McBurney
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF UK
p.j.mcburney@csc.liv.ac.uk

Simon Parsons
Department of Computer & Information Science
Brooklyn College, City University of New York
Brooklyn NY 11210 USA
parsons@sci.brooklyn.cuny.edu

ABSTRACT

What distinguishes e-commerce from ordinary commerce? What distinguishes it from distributed computation? In this paper we propose a performative theory of e-commerce, drawing on speech act theory, in which e-commerce exchanges are promises of future commercial actions, whose real-world meanings are constructed jointly and incrementally. We then define a computational model for this theory, called Posit Spaces, along with the syntax and semantics for an agent interaction protocol, the Posit Spaces Protocol or *PSP*. This protocol enables participants in a multi-agent commercial interaction to propose, accept, modify and revoke joint commitments. Our work integrates three strands of prior research: the theory of Tuple Spaces in distributed computation; formal dialogue games from argumentation theory; and the study of commitments in multi-agent systems.

Categories and Subject Descriptors

D.2.11 [**Software Architectures**]: Patterns; F.1.1 [**Models of computation**]; I.2.11 [**Distributed Artificial Intelligence**]: Coherence and co-ordination, Languages and Structures, Multiagent systems.

General Terms

Design, Languages, Standardization, Theory

Keywords

Agent Communications Languages, Commitments, e-Commerce, Javaspaces, Negotiation, Performatives, Speech Acts, Tuple Spaces

1. INTRODUCTION

Despite the success of the Internet, we know of no formal theory of e-commerce. In what way, if at all, does e-commerce differ from other forms of commerce? In what way, if at all, does it differ from parallel or distributed computing? In what way, if at all, do e-commerce systems differ from multi-agent systems? A theory of e-commerce should provide answers to these questions, in order to distinguish, if this is possible, e-commerce from these other activities. In other words a theory of e-commerce should describe activities we would recognize as e-commerce and should describe *all* such activities, from online auctions to complex multi-party commercial negotiations. In this paper, we present such a theory, drawing on speech act theory, the philosophy of argumentation, distributed computation and the study of commitments in multi-agent systems.

The paper is structured as follows: In Section 2 we present our theory of e-commerce, and distinguish it from other forms of commerce. In Section 3, we present a list of requirements for a computational model for our theory, and review the three main antecedents of our work: Tuple Space theory as a model of distributed computation; dialogic commitment stores in formal dialogue games; and Singh's treatment of commitments in multi-agent systems. In Section 4 we present the syntax and semantics of our Posit Spaces Protocol (PSP) and give a brief example of its use. In Section 5 we compare our protocol with Tuple Space theory and with Singh's framework. Section 6 concludes the paper with a discussion of related and future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'03, July 14-18, 2003, Melbourne, Australia.
Copyright 2003 ACM 1-58113-683-8/03/0007 ...$5.00.

2. WHAT IS E-COMMERCE?

Can e-commerce be distinguished from other forms of commerce? To answer this we begin by discussing commerce. In accordance with standard approaches in economics, we define a **good** as a product or service, either of which may be material or intangible. We define a **legal person** as a human person, a legally-constituted company, society or charity, or a Government agency. We define a **commercial transaction** as the exchange, between two or more legal persons, of two or more goods for one another. In the simplest such transactions, both goods are material products, as when two subsistence farmers barter their agricultural outputs, e.g., maize for cotton. In modern societies, however, most commercial transactions involve the exchange of one good for money or a money-equivalent. Initially, money comprised coins made of rare heavy metals, such as gold and silver; paper money, when it was introduced, expressed a promise by the issuer to exchange the paper for a designated amount of some rare metal upon demand. Even today, monetary equivalents such as cheques and credit card vouchers express promises to exchange them for money upon demand. In other words, we may view financial instruments from paper money onwards as encoding commitments by a legal person to undertake some future action or to bring about some future state.

Of course, all the goods in an exchange may encode such commitments, as when one currency is exchanged for another, or when a customer uses a cheque to purchase a sofa from a furniture store for later delivery; Here, the store is committing to undertake a future action — delivery of a specific sofa to the customer's home — in exchange for the customer also committing to instruct his

bank, via the message on his signed cheque, to transfer money to the store's account. So what are the *goods* in this transaction? They are both promises of future actions, and one can readily see that a given customer may view particular promises more favourably than others: a sofa delivered the next day may be preferred to the very same sofa delivered in a year's time.

So what is distinctive about e-commerce? Commercial transactions executed over the Internet do not, with only a few exceptions permit the actual exchange there and then of the goods in question.[1] Accordingly, what is exchanged electronically are normally commitments of this sort, i.e., promises of future action by one or more persons involved in the transaction. In executing a transaction involving the exchange of such commitments, the persons involved are making utterances of the following sort:

- Buyer: *I agree to give you payment of monetary amount p, to be paid by means w, in exchange for the good g, under conditions $a, b, c, \ldots$*
- Seller: *I agree to give you the good g in exchange for payment of monetary amount p, to be paid by means w, under conditions $a, b, c, \ldots$*

These expressions change the world external to the electronic domain where they are uttered, and thus become true by virtue of their utterance: they are *performatives* in the terminology of speech act theory [1]. Their utterance implies, as with the exchange of cheques or promises of sofa delivery, commitments to future action or achievement of some future state, namely the actual exchange between the parties of the goods to which they refer.

The implied commitments of such utterances share a number of characteristics. Firstly, they are made by persons (or their electronic agents) who are, at least for the purposes of the transaction concerned, autonomous. The persons making the commitments cannot be ordered to make them or not to make them by the other parties involved in the transaction; at best, those others can attempt to *persuade* an agent to adopt a particular commitment. There is a subtle consequence of agent autonomy involved here. Since each agent is autonomous, no commitment binds an agent unless that agent first agrees to it. But, once so bound, an agent cannot normally unilaterally modify or revoke the commitment without the prior agreement of all other agents party to the original commitment; *their* autonomy requires this. Thus, the commitments are made jointly by the parties to the transaction.

Secondly, the commitments made in an electronic transaction are promises to establish or maintain a specified real-world state, as a result of executing, incurring or maintaining an action or course of action.[2] They are not commitments in the sense of the persistence of an agent's beliefs or intentions [29, p. 205], although they may reflect the existence of such internal commitments.[3] Nor are they merely an expression of a willingness to defend a particular statement in a dialogue [10, p. 257], although again such a defense may be required in the interaction. In other words, these commitments have a semantics in the real world external to the electronic interaction in which the performative statements expressing them are uttered.

Thirdly, the external meaning is determined by the context of the interaction in which the performatives are uttered, and this context includes both the agreed rules of the interaction, such as the procedural rules of an auction, and the prior statements of the parties to the interaction. Thus, the real-world meaning of any deal achieved is potentially constructed incrementally, in the course of the negotiation, rather than existing — whether in the real-world or in the minds of the agents — prior to commencement. As the external meaning of the dialogue evolves along with it, so too may the beliefs, desires and intentions of the participants; a consumer's preferences between products, for example, may alter radically depending on what alternatives the consumer believes are available [13].

In asserting that e-commerce commitments have meanings which are constructed incrementally, we are saying that the meaning of utterances may depend on the sequence of interactions which lead to them. We can view the communication and negotiation process prior to a transaction as a joint search by the participants through a space of possible deals [11], in which each party may only know at the outset its own subspace of acceptable deals. As proposed deals and responses are communicated by the parties to one another, each gains a better understanding of the overlap between its subspace and those of others. Each may also gain a better understanding of its own subspace, since the other party may propose possible deals of which it previously was not aware, and of its own preferences across the elements of this sub-space. The process of negotiation incrementally creates the space of possible deals and the subspaces of acceptable and feasible deals. This semantics is thus a semantics for the interaction itself, not merely of the statements expressed in it. It is thus analogous to the possible-worlds semantics for human language dialogues studied by linguists under the name of Discourse Representation Theory [12], in which participants to a conversation jointly and incrementally construct the meaning of the utterances and of the dialogue itself.

Following these comments we now define e-commerce transactions:

*An **e-commerce transaction** is an exchange via electronic media of performative statements by two or more persons, or software agents acting on their behalf, in which commitments to achievement of future states involving the exchange of goods with one another are expressed. The real-world meanings of these commitments are constructed jointly and incrementally by the participants in the course of the electronic interaction leading to the utterance of the performatives. Their meaning will depend on the context of the exchange, including any rules governing the interaction and the prior dialogue itself.*

Thus, according to this definition, the two features distinguishing e-commerce from ordinary commerce is the use of electronic media for communications and the explicit presence of performative statements with a jointly- and incrementally-created semantics. Of course, any commercial transaction may involve the utterance of performatives prior to, or coincident with, the exchange of goods, but in e-commerce, under our formulation, this exchange of performatives is always present explicitly.

3. E-COMMERCE AND COMPUTATION

We desire a computational model for our notion of e-commerce. We may define, in broad outline, the requirments of such a model as follows. The model must:

- Support all forms of electronic commercial transactions, including auctions, listing boards, structured negotiations, and unstructured argumentation-based interactions.
- Support multi-agent negotiations and transactions.[4]

[1] The exceptions are intangible products and services, such as music, medical advice or access to networked resources.

[2] In viewing commitments as specifying future world-states rather than actions, we follow the treatment of [25, 27].

[3] In this paper, we limit attention only to publicly-observable behavior, and so do not discuss such internal commitments.

[4] Much of the agent e-commerce literature has focused on two-party commercial trans-

- Support a joint and incremental search for deals, and the joint making and changing of commitments.
- Support a notion of commitments as performatives.
- Allow for agent autonomy in the making, modifying and revoking of commitments.
- Permit spatial and temporal de-coupling of the communications between the agents involved, so as to allow for trade listing boards, market aggregators, etc.

Some of these requirements may be met by existing models and approaches, and so we begin by considering three broad strands of prior research on which we have drawn: Tuple Spaces; formal dialogue games; and models of multi-agent commitments.

3.1 Tuple spaces

David Gelernter's theory of tuple spaces [3, 9] was proposed as a model of communication between distributed computational entities.[5] This theory, and the associated programming language *Linda*, have formed the basis of SUN's popular *Javaspaces* technology [8].[6] The essential idea is that computational agents connected together may create named object stores, called *tuples*, which persist, even beyond the lifetimes of their creators, until explicitly deleted. In their Javaspaces manifestation, tuples may contain data, data structures, programs, objects or devices. They are stored in tuple-spaces, which are blackboard-like shared data stores, and are normally accessed by other agents by associative pattern matching. The use of shared stores means that communication between multiple agents can be spatially and temporally decoupled. There are three basic operations on tuple spaces:

- *out*, with which an agent creates a tuple with the specified contents and name in a shared space accessible to all agents in the system.
- *read*, with which an agent makes a copy of the contents of the specified tuple from the shared space to some private store.
- *in*, with which an agent makes a copy of the contents of the specified tuple from the shared space to some private store, and then deletes it from the shared space.

Tuple spaces are public-write, public-read spaces: any entity in the system may create a new tuple, and any entity may delete an existing one. A refinement of Linda, Law-Governed Linda [18], established an administrative layer which authorizes all attempts to execute *out, in* and *read* commands according to pre-defined security and privacy policies. Although this adds some security features, tuples are still entities created or modified individually, not jointly.

3.2 Dialogue Games

Because commercial deals are typically reached after a process of interaction between the agents concerned, it is appropriate to consider the nature interaction between them as a form of dialogue. Thus, the second strand of research we will draw on is the study of formal dialogue games from the philosophy of argumentation. Although originally due to Aristotle, this subject was revived by philosopher Charles Hamblin's use of dialogue games to study non-deductive reasoning [10]. These games have recently become important for the design of protocols for agent interactions, e.g. [15, 17]. A key concept, formalized initially by Hamblin, is that of an agent's Commitment Store, associated to each participant in an interaction [10, p. 257]. These stores keep track, through the course of a dialogue, of the dialogical commitments incurred by each agent, i.e., the claims which each agent is willing to defend if challenged by others. Commitment stores are different to tuple spaces: firstly, there is not a central store, but one for each participant; secondly, entries are made to the store as a result of specific utterances made by the associated agent [16]. All agents may see the contents of each others' stores, but only the associated agent may delete its contents, and only then if the rules of the dialogue game provide a locution effecting this. Thus, commitment stores are private-write, public-read spaces.

3.3 Commitments

An influential treatment of agent commitments has been presented by Munindar Singh and his colleagues [25, 27, 32]. In this account, commitments are promises made by a *debtor* to a *creditor* to establish or maintain a certain world-state. Formally, a commitment c is denoted by $c = C(P_1, P_2, \mathcal{I}, p)$ where debtor agent P_1 promises creditor agent P_2 in the context of multi-agent system $\mathcal{I}$ to achieve the world-state described by proposition p (which may include temporal references). Conditional commitments [32], where an agent promises to achieve the world-state identified in one proposition provided another proposition is true, can be notated similarly.

Singh's framework also permits meta-commitments, which are commitments about commitments, and rules or norms in the context of the interaction between the agents. Both meta-commitments and contextual rules may govern the invoking, modifying or revoking of commitments. One norm may be, for example, that a commitment is not delegated without prior agreement of the delegatee. Singh defines six primary operations on commitments, which we summarize here:

Create: This action creates a commitment, and is typically undertaken by the debtor.

Discharge: This action satisfies a commitment, and is performed by the debtor of the commitment when the final state condition of the commitment is satisfied.

Cancel: This action revokes a commitment, and may be performed by the debtor. Depending on the meta-commitments obtaining, a cancellation of one commitment may create another.

Release: This action eliminates a commitment, and may be undertaken by the creditor or arise from the context.

Delegate: This actions shifts the role of debtor to another agent within the same multi-agent system and may be performed by the old debtor or by the context. The creditor is informed of an act of delegation.

Assign: This action transfers a commitment to another creditor within the same multi-agent system, and can be performed by the old creditor or the institution. The debtor is informed of an act of assignation.

Although commitments are understood as joint promises, the formalisation of Singh does not make this explicit. In particular, the rules governing modification or revocation of commitments allow the debtor to discharge, cancel or delegate a commitment, or the creditor to assign a commitment, without the prior approval of the

actions. Yet many transactions, from everyday house-purchases to Exotic Options contracts, involve multiple parties.

[5] See [20] for a review of tuple-space models.

[6] See: `http://java.sun.com/products/javaspaces/`.

other parties concerned, unless these are required by the context or by any meta-commitments. Thus, the default position is unilateral amendment, discharge or revocation, with meta-commitments and the context covering special cases. In e-commerce, by contrast, the default position should be, we believe, that all parties to a commitment need to give their approval for its amendment, discharge or revocation. It is easy to imagine a debtor, for example, claiming to have discharged a prior commitment on the basis of the achievement of some world-state, and a creditor to that commitment contesting that the world-state has in fact been realized.

A second comment about Singh's framework is important, since it reveals a lack of generality. The framework is applied in [27] to an auction marketplace (the fish market of [24]), a domain where the rules of the auction (the context) treat bids as proposed commitments which are then accepted or rejected by the auctioneer at each round. Thus calls-for-bids and bids-in-response-to-a-call are made using the *Create* locution, with conditional commitments as the contents: the auctioneer promises to provide fish at a certain price if he receives money; the bidders promise to buy fish, if the price is a certain level. These statements are conditional promises, and the rules of the interaction context (the auction) turn them into commitments once certain other locutions are uttered. So, effectively the *Create* locution is being used here for proposing commitments, not for their creation, because the framework has no locution for proposing a commitment.[7]

That this is possible in the fish-market domain is because the rules of the auction institution are sufficiently constraining, and because the commitments concerned can be expressed as conditional commitments, each to be undertaken upon achievement of some defined world state. However, if the commitments were to involve simultaneous achievement of different world states, or simultaneous execution of actions, by different parties, then they could not be expressed as conditional commitments. They could only be expressed as conjunctions, e.g., $c = C(P_1, P_2, \mathcal{I}, p)$ & $d = D(P_2, P_1, \mathcal{I}, q)$. If such a joint commitment were proposed inside a *Create* locution, which agent, P_1 or P_2, would utter it? Agent autonomy means neither can make commitments on behalf of the other, so neither could create it. How could it then be proposed by one agent, and how accepted or refused by the other? The problem here lies in the absence in Singh's framework of a locution with the sole effect of proposing a commitment, an action distinct from creating it.

4. POSIT SPACES

We now propose a computational formalism to represent our performative theory of e-commerce, drawing on these three strands of prior research. Our formalism is intended to be general and to achieve the requirements specified at the start of Section 3.

4.1 Syntax of PSP

We suppose we have multi-agent system comprising n autonomous computational agents, denoted $P_1, \ldots, P_n$, each with the goal of exploring the possibility of executing an e-commerce transaction. Following [19], we call the electronic space in which they interact an *institution*; we assume that it has explicit rules of interaction, and that these become known to each participant upon entry to it. We will denote the institution under which a particular interaction is conducted by $\mathcal{I}$, and call it the *governing* institution for that interaction. An auction-space is an example of such an institution, and the rules of the auction define the rules of interaction.

We assume further there are $n + 1$ stores which the n agents in the interaction may have read-access, depending on the rules of $\mathcal{I}$. Associated to each agent P_i is a **Proposal Store**, denoted $PS(P_i)$, to which only that agent has write-access. Thus, each Proposal Store is private-write and public-read. In addition, there is a **Deal Store** to which no agent has direct write-access, but to which all may have read-access. The Deal Store will hold the commitments entered into by the participants in their interaction. We denote the Deal Store by $DS(P_1, \ldots, P_n, \mathcal{I})$, or simply DS when this causes no confusion. Depending on the rules of the institution, the Deal Store may also be partitioned, so that only particular agents have read-access to its partition elements; for example, in a multi-party negotiation, the Deal Store may be partitioned into sub-spaces to each of which only 2 parties have read-access; this would facilitate private bilateral side deals within the public space. Likewise, each agent may partition its own Proposal Store, so as to allow private read-access to sub-spaces for other particular agents.[8] For simplicity, from here on we assume that no Store is partitioned in this way.

The contents of the Proposal Stores and Deal Store are persistent entities called **posits**, which are essentially one or more proposed commitments. We denote posits by lower-case greek letters, $\alpha, \beta, \ldots$ Once such proposals are accepted, in accordance to the rules of the governing institution, they enter the Deal Store. We assume that commitments are represented in a suitable formal language, such as that of Singh mentioned above. For e-commerce domains, we assume all proposed deals involve exchanges, that is two or more joint commitments. A buyer commits to transfer money to the seller if and only if the seller agrees to transfer a certain good to the buyer. Thus, a posit α could consist of two (or more) such commitments, of the form $c = C(P_1, P_2, \mathcal{I}, p)$ & $d = D(P_2, P_1, \mathcal{I}, q)$. If the seller only agrees to supply the good upon first receiving the money, the proposition q in the second commitment may refer to a world-state where the first commitment — to achievement of proposition p — has already been fulfilled. Because we allow any notation for commitments, we do not define a specific notation for posits.

We next define a set of locutions which enable participants to create and delete posits from their own Proposal Stores.

- **PROPOSE**(P_i, α), which creates a new posit α, with specified name and contents, in the Proposal Store $PS(P_i)$ of the speaker P_i.
- **ACCEPT**($P_j, PS(P_i), \alpha$), which copies an existing posit α and its contents from the Proposal Store $PS(P_i)$ in which it is currently held to the Proposal Store $PS(P_j)$ of the speaker P_j.
- **DELETE**(P_i, α), which deletes an existing posit α from the Proposal Store $PS(P_i)$ of the speaker P_i.

As mentioned, we assume the rules of the electronic institution $\mathcal{I}$ define when a proposed commitment becomes binding on the participants. In an auction with many potential buyers and one potential seller, for example, a commitment may only come into force when accepted by the seller and by one of the buyers. In a multi-party negotiation, such as the aircraft supply-chain domain of [5, Section 8.5], a commitment may require acceptance from all parties to the interaction before coming into force. We also assume the rules of $\mathcal{I}$ also specify which agents are required to agree before a commitment can be cancelled or modified in any way. Our model is general across any such set of defined rules, by means of the following two conditions:

[7] This conflation of the act of proposing a commitment with its creation, it should be noted, is contrary to the stated definition of the locution *Create*.

[8] One could imagine such access privileges being policed by designated support agents not themselves engaged in the e-commerce interaction, similar to the space-administration objects in [4].

- A posit α enters the Deal Store $DS(P_1, \ldots, P_n, \mathcal{I})$ when and only when the rules of the governing institution $\mathcal{I}$ have been satisfied to turn a proposed exchange of commitments into a firm agreement between the parties. Such a posit is said to have become binding, and a deal is said to have been struck.

- A posit α is deleted from the Deal Store $DS(P_1, \ldots, P_n, \mathcal{I})$ when and only when the rules of the governing institution $\mathcal{I}$ have been satisfied to permit its revocation. Such a posit is said to be revoked.

We also define two locutions allowing agents to express their desires regarding posits in the Deal Store:

- **SUGGEST_REVOKE**(P_i, α), an utterance which expresses that participant P_i desires the deletion of posit α from the Deal Store DS, and that P_i desires that other participants agree to its deletion from the Deal Store.

- **RATIFY_REVOKE**(P_j, α), an utterance which expresses that participant P_j desires the deletion of posit α from the Deal Store DS.

These five locutions and the two rules together comprise an interaction protocol, called the **Posit Spaces Protocol (PSP)**, which is a parsimonious formalization of the multi-agent performative theory of e-commerce of Section 2. The protocol draws from Tuple Space theory the notion of shared, persistent stores, decoupled spatially and temporally from the agents in the institution. So, for example, posits may be entered by an agent into its Proposal Space and then read at a later time by other agents, just as potential buyers may read a listing of goods for sale on a bulletin board [21]. PSP adds to Tuple Space theory the notion of private-write, public-read spaces, a concept taken from the Commitment Stores of formal dialogue games. The rules governing insertion and deletion of posits from the Deal Store are also motivated by the rules of formal dialogue games, where particular combinations of utterances may have effects on the Commitment Stores of dialogue participants. In addition, the requirements of joint creation and joint revocation imposed by the protocol rules make as default the conditions which appear as exceptions in Singh's framework.

PSP ensures the requirements arising from agent autonomy are met: No proposed commitment becomes binding until all those required to agree to it do so; and once it becomes binding, it can only be revoked with the agreement of every relevant party. In each case the specification of which agents are required is given by the rules of the governing institution. Moreover, in another reflection of agent autonomy, there is nothing in the PSP which precludes posits being entered in the Deal Store which express conflicting commitments. If agents wish to propose or accept mutually-incompatible commitments, then they are free to do so. Of course, other agents recognizing a conflict in the commitments made by a first agent may refuse to accept such posits from it, thus preventing these posit becoming binding.

4.2 Semantics of PSP

Having defined an abstract language in our theory of posit spaces, we now consider its semantics. An **axiomatic semantics** for a programming language [26] defines the pre-conditions and post-conditions for each locution, as in the Semantic Language SL of the FIPA ACL [6]. We give an axiomatic semantics for the locutions of PSP in terms of the beliefs and intentions of the participating agents, and this is presented in the Appendix. Because these conditions specify internal states of each agent concerned, they will generally not be verifiable by other agents involved [28].

An **operational semantics** treats the entire multi-agent system as a single virtual machine. Here, the locutions of the protocol are viewed as commands in a programming language which executes on the machine, each locution acting to alter the machine's state. Given specific internal architectures for the agents in the system (such as a BDI architecture) it would be straightforward to define an operational semantics for the locutions of PSP, in the same way as has been done for social commitments in [31] and for a multi-agent purchase negotiation protocol in [15]. Because we desire PSP to be applicable regardless of the internal architecture of the participating agents, we have not developed an operational semantics for the protocol.

4.3 Example

We present a simple example of a dialogue between agents using PSP to propose, accept and modify joint commitments. Assume three agents, P_1, P_2 and P_3 are engaged in a negotiation, and agent P_1 proposes a set of commitments described in a posit α. Suppose this posit is accepted by the other two agents, and thus becomes binding. Assume further that agent P_2 subsequently wishes to modify one of the commitments contained in α, with the modified posit denoted as β. Agent P_2 therefore suggests the revocation of α and the adoption of β. The interaction under PSP between these agents could proceed as follows. Here we have numbered the utterances in order, and indicated in italics any effects on the contents of spaces.

1: PROPOSE(P_1, α)

Posit α enters $PS(P_1)$.

2: ACCEPT($P_2, PS(P_1), \alpha$)

Posit α enters $PS(P_2)$.

3: ACCEPT($P_3, PS(P_1), \alpha$)

Posit α enters $PS(P_3)$, and then also the Deal Space.

4: SUGGEST_REVOKE(P_2, α)

5: PROPOSE(P_2, β)

Posit β enters $PS(P_2)$.

6: RATIFY_REVOKE(P_1, α)

7: ACCEPT($P_1, PS(P_2), \beta$)

Posit β enters $PS(P_1)$.

8: PROPOSE(P_3, γ)

Posit γ enters $PS(P_3)$.

...

In interactions involving proposed amendment or cancellation of binding posits, the generality of the Posit Spaces Protocol permits great representational flexibility. For instance, in this example, once the deal regarding α is struck, each agent has complete freedom regarding the order in which they utter subsequent locutions. At utterance 6, we see agent P_1 signal its acceptance for the revocation of α. However, this posit will not be revoked until agent P_3 also signals acceptance. But in utterance 8, agent P_3 has decided to propose an alternative posit, γ, instead; perhaps P_3 does

not wish to revoke α until agreement has been reached on an alternative. Agents may have tactical or strategic reasons to prefer some sequences of locutions over others when modifying or revoking posits, and the flexibility of the protocol permits agents to select the most appropriate sequence according to circumstances.

5. COMPARISON OF PSP

We now consider the relationships between both PSP and Tuple Spaces, and between PSP and Singh's commitments framework.

5.1 PSP and Tuple Spaces

One could ask whether Tuple Space theory could provide a computational model for e-commerce. Indeed, it has previously been proposed for this domain [21]. How could this be achieved? Agents could use the *out* locution to propose commitments, the *read* locution to consider them, and the *in* locution to accept them. The only drawback is that, for any tuple, the tuple space has only two states: the tuple is either present in the tuple space or it is absent. Thus, for commitments involving actions by only two agents, this is fine: one agent proposes a commitment and one other agent from a group of agents either accepts it or does not. Thus, in the example of [21], sales are proposed by sellers and individual buyers either accept them or do not.

However, if the commitments involve promises by more than two agents, say by n agents, then there are 2^{n-1} possible outcomes to any one proposal, since every agent involved other than the proposer may accept or reject it. Since, for any given tuple, a tuple space has only 2 states, Tuple Space theory cannot express all 2^{n-1} possible outcomes, whenever $n > 2$. A clumsy alternative would be for the proposing agent to use the *out* locution $n-1$ times, with each utterance containing almost the same tuple, the only difference being some field indicating which of the $n-1$ other agents was to consider the tuple. In this approach, the space has 2^{n-1} possible states, and so could express the 2^{n-1} possible possible responses to the proposal. But what is to stop an agent — malevolent or badly-coded — using the *in* locution to remove from the space one or more of the tuples which are intended by the proposer for consideration by other agents. Here, the 2^{n-1} possible states of the space have more than 2^{n-1} meanings, since the absence of a tuple following its insertion into the space could mean that the designated agent accepts it, or it could just mean that another agent has deleted it. Rules permitting only the designated agent to delete a tuple, as in Law-Governed Linda, would thus be required.

Suppose that a commitment involving n agents is agreed in this way. Thus, all $n-1$ versions of the tuple have been legitimately removed from the tuple space, thereby indicating the agreement of the designated agents to the proposed commitment. Suppose, in the fullness of time, that one agent wishes to amend or revoke the commitment. How is this to be done? There is no publicly-accessible record of the commitment, since the tuple is no longer present in the shared space; so the agent would need to utter $n-1$ *out* locutions, each with a version of a tuple which contained a proposal to amend or revoke the original commitment. Thus, for a proposed commitment which was accepted by all agents, and then revoked by all, Tuple Space theory would require $4(n-1)$ utterances. To achieve the same effect, PSP, by contrast, would require only one *PROPOSE* locution, and $n-1$ *ACCEPT* locutions, one *SUGGEST_REVOKE* and $n-1$ *RATIFY_REVOKE* locutions, i.e., $2n$ utterances in all. If there are three or more agents, i.e., if $n > 2$, then PSP is more efficient than tuple spaces used in the way described here. Of course, a more efficient use of Tuple Spaces would be for the proposer to insert just one tuple, and for every other agent to indicate its acceptance by inserting another tuple, marked uniquely with that agent's identifier. This would require only n utterances to create a joint commitment. However, if agents are precluded from impersonating others, this would be equivalent to partitioning the Tuple Space into n components, one for each participant, working in the same way as Proposal Spaces.

The application of Tuple Spaces to the design of an e-commerce listing board described in [21] requires agents to share their private message ports with one another to finalize a negotiation that begins in the public interaction space. PSP, through the use of partitioned Proposal and Deal Spaces and explicit rules regarding the making, modification and revocation of commitments can ensure both privacy and compliance with the rules of the governing institution. This may be important if other agents, such as regulatory authorities, require oversight of any transactions completed.

5.2 PSP and Commitments

We also consider the relationship between PSP and the framework of Singh *et. al* for commitments, which we summarized in Section 3.3. Commitments are created when they enter the Deal Store, which occurs when they are placed inside the Proposal Stores of each agent required to indicate acceptance. This corresponds to the creation of a commitment, equivalent to the definition of Singh's *Create* locution. For the other five operations: *Discharge, Cancel, Release, Delegate* and *Assign*, PSP enables each of these via the removal and/or amendment of commitments in the Deal Store. Such actions need the agreement of all parties required under the rules of the governing institution; in e-commerce applications this would include all parties to the original commitment, and all new parties in any amended version (such as a new assignee or delegatee). In Singh's framework, by contrast, as mentioned in Section 3.3, the default position is unilateral amendment, discharge or revocation, with meta-commitments and the institutional context covering special cases.

Another difference is that PSP makes no explicit distinction between delegating, assigning or amending a commitment. These involve syntactical differences between commitments (and hence posits) with semantic consequences, and each is achievable through a succession of utterances proposing a new posit and revoking a prior one, as shown in Section 4.3. PSP also makes no distinction between discharging, cancelling and releasing a commitment, actions which involve semantic but not syntactical differences between posits in PSP. Each is achievable through revocation. The importance of such semantic distinctions will differ by Institution, by occasion (such as the dialogical context of the utterances) and by posit, and so we believe a model aimed at generality needs to abstract away from these distinctions. In any particular case, the agents engaged in an interaction will bring to bear whatever considerations they deem relevant to their decision to accept particular posits or particular revocations.

6. CONCLUSIONS

In this paper, we have defined e-commerce in a manner which distinguishes it from traditional commercial activities not mediated electronically. Our definition emphasizes the performative nature of the utterances between the participants in an electronic marketplace: these utterances express statements about future action-commitments by the participants, and become true by virtue of being uttered. Their external meaning depends on the institutional and dialogical context in which they are uttered, and is created jointly and incrementally by the participants in the course of their interaction. We then proposed a novel conceptual model for e-commerce activities, which we call Posit Spaces theory. Defining the locutions and the syntactical rules for their use in this model gave us a

multi-agent interaction protocol, the Posit Spaces Protocol (PSP). In this paper we also articulated an axiomatic semantics for PSP.

PSP draws on three strands of prior research: (a) Gelernter's Tuple Spaces theory of distributed computation, from which we took the concept of spatially- and temporally-decoupled persistent data stores as a model for distributed computation; (b) the use of Commitment Stores in the formal dialogue games of philosophy. This notion provided us with private-write, public-read stores, thereby enabling commitments to be proposed and accepted, and with the motivation for rules which lead to commitments becoming binding once certain utterances are made; and (c) Singh's treatment of commitments in multi-agent systems as promises to maintain or achieve specified world states, and which provided a formal framework for their representation and a defined set of operations over them.

The PSP framework meets the criteria we presented in Section 3 for a computational model of e-commerce. It is clearly general enough to support all forms of electronic commercial transaction, from auctions to unstructured argumentation-based negotiations. By using shared spaces in a manner which extends Tuple Space theory, PSP permits the spatial and temporal decoupling of communications between the agents involved. It supports multi-agent interactions and, for commitments involving three or more agents, does at least as efficiently as Tuple Space theory. Moreover, the inclusion of a specific locution for proposing commitments enables the incremental search for deals. Similarly, the incorporation of specific protocol procedures for the creation and revocation of commitments tied to the rules of the electronic institution governing the agent interaction makes explicit the permissions required for commitments to be made and unmade. Thus, the protocol supports a notion of e-commerce commitments as performatives, complementing Singh's model of multi-agent commitments. It clearly supports an incremental and joint search for deals.

In e-commerce research considerable attention has focused on design of electronic marketplaces and institutions, and some of this is formal, e.g. [14, 19, 22, 30]. However, such research typically aims to model a particular type of interaction, such as auctions or argumentation-based dialogues, and not to capture all types of electronic commercial transactions under the one formalism. The Posit Spaces model is sufficiently general to capture any type of commercial interaction. Moreover, some work, such as [14], appears to conflate locutions uttered in an electronic interaction with the actions in the real-world which are promised to follow subsequently. Our approach, treating statements in the electronic interaction as speech-act performatives, does not do this.

A key feature of the Posit Space approach is the view that a multi-agent interaction which aims to achieve a commercial transaction involves the joint search through a deal space whose dimensions and contents are constructed incrementally. Related work in agent communications theory includes [23], in which agents in an open agent system jointly agree an axiomatic semantics for the agent communications language utterances they will use to communicate. However, [23] assumes the agents involved all start with a common semantic space, and then together assign particular locutions to specific points in this space. Such a structure would not appear to permit an incremental construction of the semantic space itself. In contrast, an incremental construction is possible with the *Protocol for Proposals* of [7]. But this protocol, since it arises from the conversational policies rather than the dialogue game tradition in agent communications, governs only dialogue segments, and not entire dialogues; moreover, it is also unclear how the *Protocol for Proposals* operates for commitments involving more than two parties.

Future work on PSP will proceed along several directions. Firstly, we aim to use the Posit Space theory to represent existing auction and argumentation protocols for electronic negotiation. This would provide a test of the practical usefulness of the model. The negotiation middleware developed in [2], for example, would be readily represented in PSP. Secondly, we aim to develop a denotational semantics for PSP, perhaps using graph theory or category theory, so as to gain a better understanding of the formal properties of the protocol. It should also be straightforward to express the axiomatic semantics we have given in the Appendix using modal operators for beliefs and intentions, similarly to the modal semantics SL of the FIPA ACL [6]. Thirdly, we aim to study the possible strategies for agents using the Posit Spaces Protocol, so as to provide guidance in negotiation contexts.

Appendix

We now present an axiomatic semantics for PSP:

- **PROPOSE**(P_i, α)

  Pre-conditions: Speaker P_i intends that each participant P_j, $j \neq i$, believe that P_i desires to transact the deal described by the commitments contained in α.

  Post-conditions: Each participant P_j, $j \neq i$, believes that participant P_i intends that each participant $P_j, j \neq i$, believe that P_i desires to transact the deal described by the commitments contained in α.

  Posit Stores: Posit α is added to $PS(P_i)$, the Proposal Store of speaker P_i.

- **ACCEPT**($P_j, PS(P_i), \alpha$)

  Pre-conditions: Speaker P_j intends that each listener $P_i, i \neq j$, believe that P_j desires to transact the deal described by the commitments contained in α.

  Post-conditions: Each participant P_i, $i \neq j$, believes that participant P_j intends that each participant $P_i, i \neq j$, believe that P_j desires to transact the deal described by the commitments contained in α.

  Posit Stores: Posit α is added to $PS(P_j)$, the Proposal Store of speaker P_j. In addition, when the particular participants or the requisite number of participants specified by the rules of the governing institution $\mathcal{I}$ have uttered this locution, the posit α is added to the Deal Store, $DS(P_1, \ldots, P_n, \mathcal{I})$.

- **DELETE**(P_i, α)

  Pre-conditions: Speaker P_i intends that each participant P_j, $j \neq i$, believe that P_i no longer desires to transact the deal described by the commitments contained in α.

  Post-conditions: Each participant P_j, $j \neq i$, believes that participant P_i intends that each participant $P_j, j \neq i$, believe that P_i no longer desires to transact the deal described by the commitments contained in α.

  Posit Stores: Posit α is deleted from $PS(P_i)$, the Proposal Store of speaker P_i.

- **SUGGEST_REVOKE**(P_i, α)

  Pre-conditions: Speaker P_i intends that each participant P_j, $j \neq i$, believe that P_i no longer desires that the commitments contained in α be fulfilled.

  Post-conditions: Each participant P_j, $j \neq i$, believes that participant P_i intends that each participant $P_j, j \neq i$, believe

that P_i no longer desires that the commitments contained in α be fulfilled.

Posit Stores: No effect.

- **RATIFY_REVOKE**(P_j, α)

  Pre-conditions: Speaker P_j intends that each participant P_i, $i \neq j$, believe that P_j no longer desires that the commitments contained in α be fulfilled.

  Post-conditions: Each participant P_i, $i \neq j$, believes that participant P_j intends that each participant $P_i, i \neq j$, believe that P_j no longer desires that the commitments contained in α be fulfilled.

  Posit Stores: When the particular participants or the requisite number of participants specified by the rules of the governing institution $\mathcal{I}$ have uttered this locution, the posit α is deleted from the Deal Store, $DS(P_1, \ldots, P_n, \mathcal{I})$.

7. REFERENCES

[1] J. L. Austin. *How To Do Things with Words*. Oxford University Press, Oxford, UK, 1962.

[2] M. Bratu, J. M. Andreoli, O. Boissier, and S. Castellani. A software infrastructure for negotiation within inter-organisational alliances. In J. Padget *et al.*, editor, *Agent-Mediated Electronic Commerce IV*, LNAI 2531, pages 161–179. Springer, Berlin, 2002.

[3] N. Carriero and D. Gelernter. Linda in context. *Communications of the ACM*, 32(4):444–458, 1989.

[4] S. Ducasse, T. Hofmann, and O. Nierstrasz. OpenSpaces: an object-oriented framework for reconfigurable coordination spaces. In A. Porto and G-C. Roman, editors, *Coordination Languages and Models*, LNCS 1906, pages 1–19. Springer, Berlin, 2000.

[5] R. van Eijk. *Programming Languages for Agent Communications*. PhD thesis, Department of Computer Science, Utrecht University, The Netherlands, 2000.

[6] FIPA. Communicative Act Library Specification. Technical Report XC00037H, Foundation for Intelligent Physical Agents, 10 August 2001.

[7] R. A. Flores and R. C. Kremer. To commit or not to commit: modelling agent conversations for action. *Computational Intelligence*, 18(2):120–173, 2002.

[8] E. Freeman, S. Hupfer, and K. Arnold. *JavaSpaces: Principles, Patterns and Practice*. Addison-Wesley, USA, 1999.

[9] D. Gelernter. Generative communication in Linda. *ACM Trans. Prog. Languages and Systems*, 7(1):80–112, 1985.

[10] C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.

[11] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.

[12] H. Kamp and U. Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht, The Netherlands, 1993.

[13] G. L. Lilien, P. Kotler, and K. S. Moorthy. *Marketing Models*. Prentice-Hall, Englewood Cliffs, NJ, 1992.

[14] N. López, M. Núñez, I. Rodríguez, and F. Rubio. A formal framework for e-barter based on microeconomic theory and process algebras. In H. Unger *et al.*, editor, *Innovative Internet Computing Systems (IICS-2002)*, LNCS 2346, pages 217–228. Springer, Berlin, Germany, 2002.

[15] P. McBurney, R. M. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *J. Autonomous Agents and Multi-Agent Systems*, 2003. *In press*.

[16] P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *J. Logic, Language and Information*, 11(3):315–334, 2002.

[17] P. McBurney, S. Parsons, and M. Wooldridge. Desiderata for agent argumentation protocols. In C. Castelfranchi and W. L. Johnson, editors, *Proc. AAMAS 2002*, pages 402–409, New York, 2002. ACM Press.

[18] N. H. Minsky and J. Leichter. Law-governed Linda communication model. In P. Ciancarini et al., editors, *Object-based Models and Languages for Concurrent Systems*, LNCS 924, pages 125–146. Springer, Berlin, 1995.

[19] P. Noriega and C. Sierra. Towards layered dialogical agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III: Proc. 3rd ATAL Workshop*, LNAI 1193, pages 173–188, Berlin, 1997. Springer.

[20] G. A. Papadopoulous and F. Arbab. Coordination models and languages. In M. V. Zelkowitz, editor, *Advances in Computers: The Engineering of Large Systems*, volume 46. Academic Press, The Netherlands, 1998.

[21] G. A. Papadopoulous and F. Arbab. Coordinating electronic commerce activities in MANIFOLD. *Netnomics*, 2:101–116, 2000.

[22] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *J. Logic and Computation*, 8(3):261–292, 1998.

[23] C. Reed, T. J. Norman, and N. R. Jennings. Negotiating the semantics of agent communications languages. *Computational Intelligence*, 18(2):229–252, 2002.

[24] J. A. Rodríguez, F. J. Martin, P. Noriega, P. Garcia, and C. Sierra. Towards a test-bed for trading agents in electronic auction markets. *AI Communications*, 11(1):5–19, 1998.

[25] M. P. Singh. An ontology for commitments in multiagent systems: toward a unification of normative concepts. *Artificial Intelligence and Law*, 7:97–113, 1999.

[26] R. D. Tennent. *Semantics of Programming Languages*. Prentice-Hall, Hemel Hempstead, UK, 1991.

[27] M. Venkatraman and M. P. Singh. Verifying compliance with commitment protocols. *J. Autonomous Agents and Multi-Agent Systems*, 2(3):217–236, 1999.

[28] M. J. Wooldridge. Semantic issues in the verification of agent communication languages. *J. Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.

[29] M. J. Wooldridge. *Introduction to Multiagent Systems*. John Wiley and Sons, New York, NY, 2002.

[30] P. R. Wurman, M. P. Wellman, and W. E. Walsh. A parametrization of the auction design space. *Games and Economic Behavior*, 35(1–2):304–338, 2001.

[31] P. Yolum and M. P. Singh. Commitment machines. In J-J. Meyer and M. Tambe, editors, *Proc. ATAL 2001*, pages 245–257, Seattle, 2001.

[32] P. Yolum and M. P. Singh. Flexible protocol specification and execution: applying event calculus planning using commitments. In C. Castelfranchi and W. L. Johnson, editors, *Proc. AAMAS 2002*, pages 527–534, New York, 2002. ACM Press.