An e-Market Framework for Informed Trading

John Debenham & Simeon Simoff

Faculty of IT
University of Technology, Sydney

{debenham, simeon}@it.uts.edu.au

http://e-markets.org.au/ WWW 2006, Edinburgh

Introduction

- Despite the potential size of the electronic business market, fully automated trading, such as e-procurement, using the Internet is virtually unheard of today.

- Trading involves the maintenance of effective business relationships, and is the complete process of: need identification, product brokering, supplier brokering, offer-exchange, contract negotiation, and contract execution.

- The e-Market Framework is available on the World Wide Web[1]. This project aims to make informed automated trading a reality.

---

[1] `http://e-markets.org.au`

Three Technologies

Three technologies are needed to fully automate the trading process:

- data mining — real-time data mining technology to tap information flows from the marketplace and the World Wide Web, and to deliver timely information at the right granularity.
- trading agents — intelligent agents that are designed to operate in tandem with the real-time information flows received from the data mining systems.
- virtual institutions — virtual places on the World Wide Web in which informed trading agents can trade securely both with each other and with human agents in a natural way — not to be confused with the term "virtual organisations" as used in Grid computing.

Data Mining

Real-time embedded data mining is an essential component of the proposed framework. In this framework the trading agents make their informed decisions, based on utilising two types of information

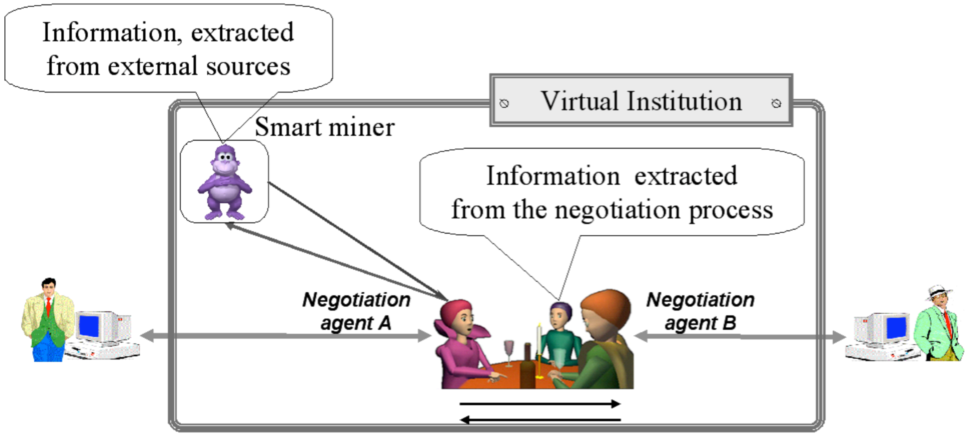

- information extracted from the negotiation process (i.e. from the exchange of offers), and;
- information from external sources, extracted and provided in condensed form.

Embedded Data Mining System

- Complements and services the information-based architecture.
- The information request and the information delivery format is defined by the interaction ontology.
- The interaction ontology evolves from the agent interaction in terms of negotiation parameters with a discrete set of feasible values. As agents proceed with negotiation they have a topic of negotiation and a shared ontology that describes that topic.
- The collection of parameter sets of the negotiation topic constitutes the input to the data mining system. Continuous numerical values are replaced by finite number of ranges of interest.

Focused Data Sets

The data mining system initially constructs data sets that are “focused” on requested information.

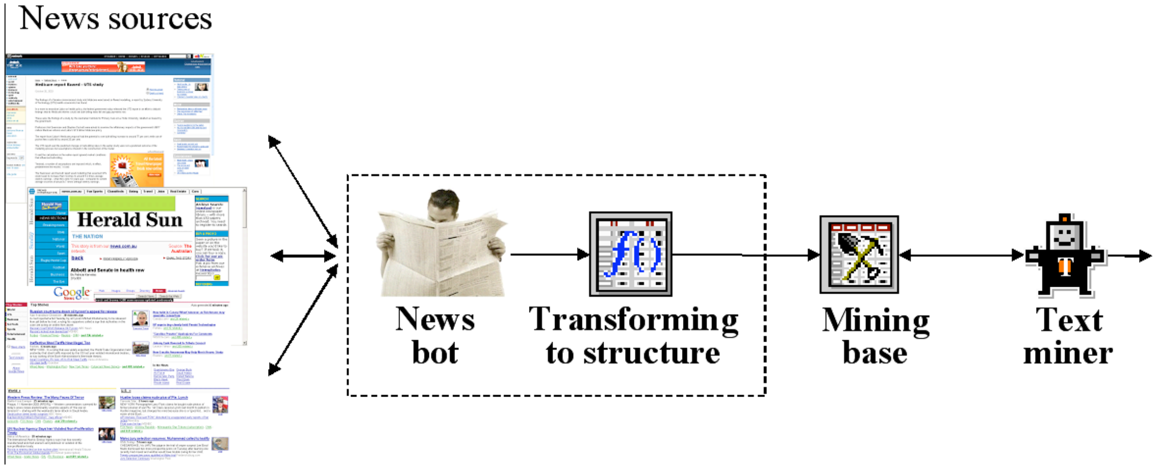

Technically, the automatic retrieval of the information pieces utilises the universal news bot architecture.

The data mining agent constructs the news data set according to the concepts in the query. Each concept is represented as a cluster of key terms (a term can include one or more words), defined by the proximity position of the frequent key terms.

Advising Model

The data sets required further automatic preprocessing, related to possible redundancies in the information encoded in the set that can bias the analysis algorithms.

Once the set is constructed, building the “advising model” is reduced to a classification data mining problem. As the model is communicated back to the information-based agent architecture, the classifier output should include all the possible class labels with an attached probability estimates for each class. Hence, we use probabilistic classifiers (e.g. Naïve Bayes, Bayesian Network classifiers) without the min-max selection of the class output [e.g., in a classifier based on Naïve Bayes algorithm, we calculate the posterior probability $\mathbb{P}_p(i)$ of each class $c(i)$ with respect to combinations of key terms and then return the tuples $< c(i), \mathbb{P}_p(i) >$ for all classes, not just the one with maximum $\mathbb{P}_p(i)$].

Architecture of Agent-Based Data Mining System

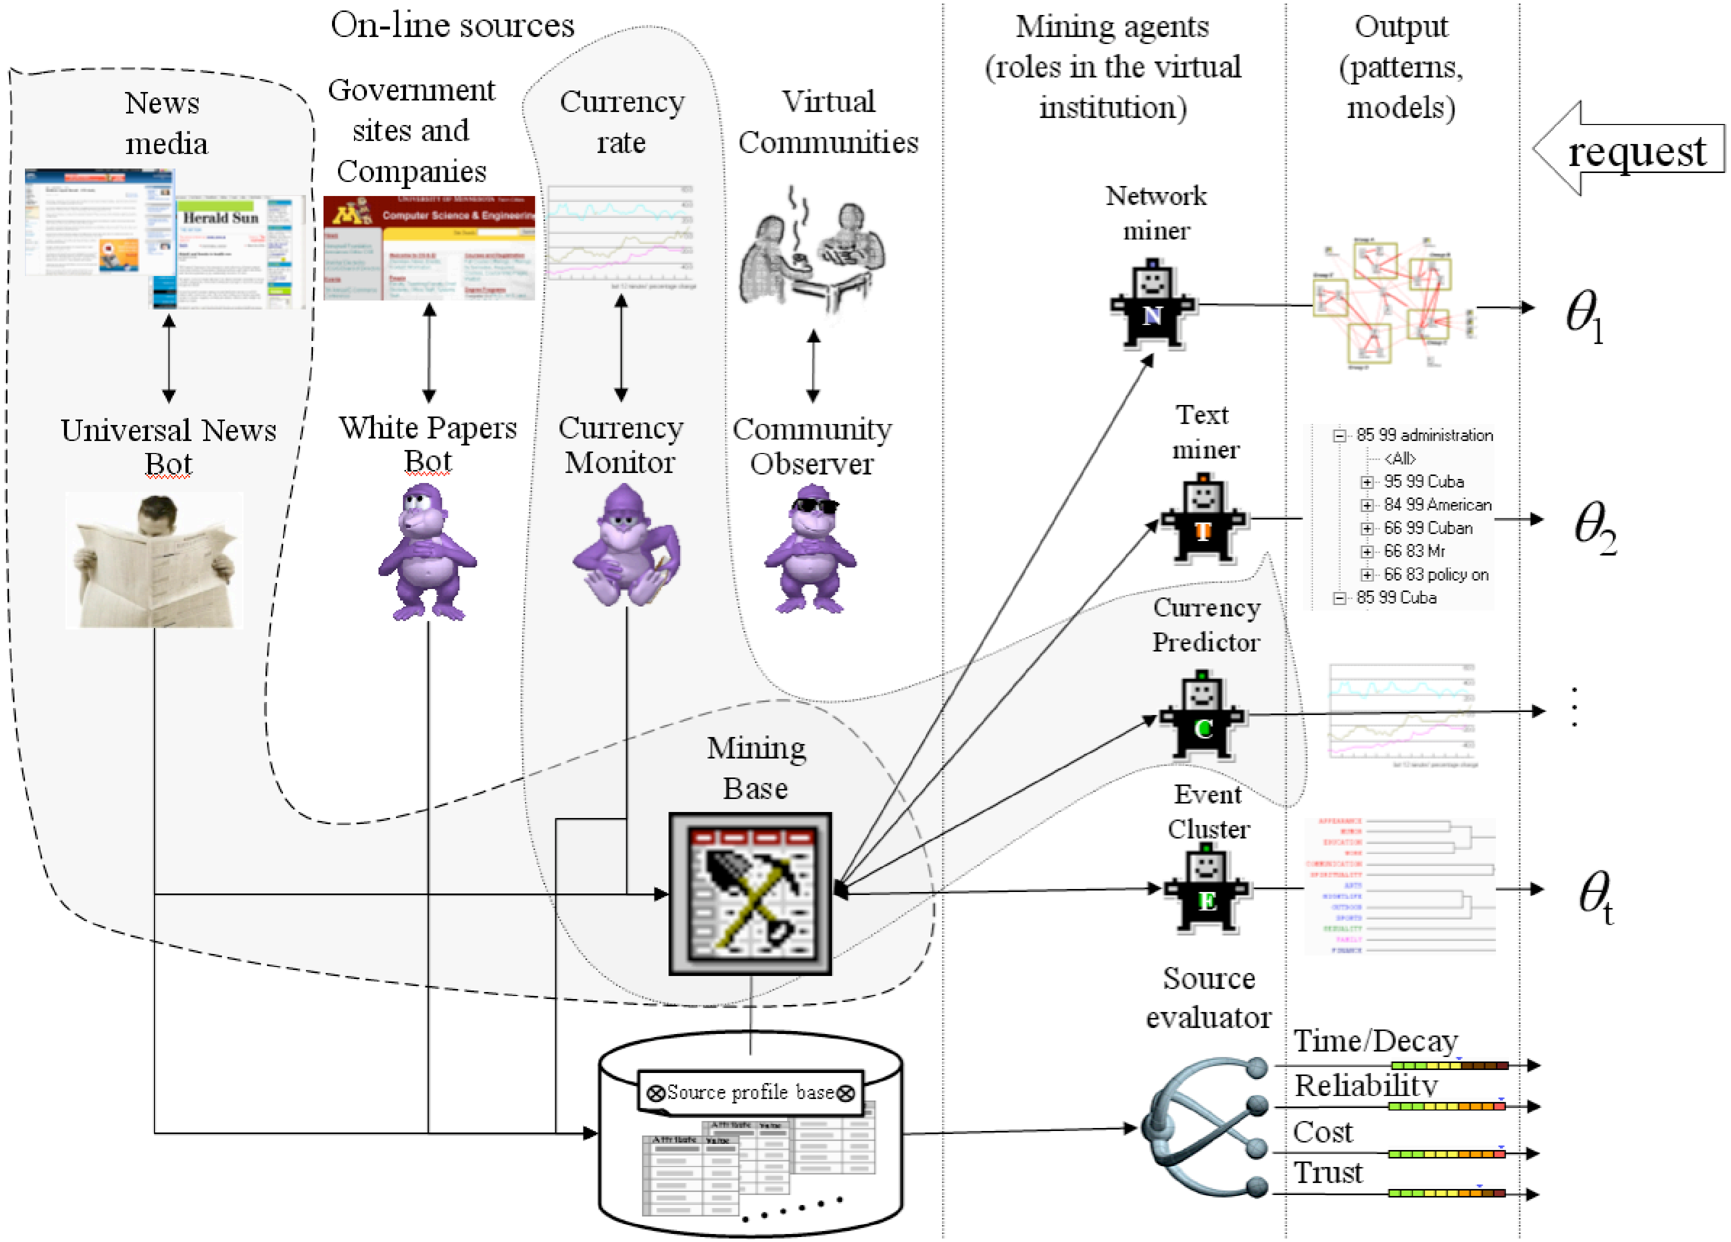

Trading Agents

Our “information-based” agents operate in real-time in response to market information flows. We have addressed the central issues of trust in the execution of contracts, and the reliability of information. Our agents understand the value of building business relationships as a foundation for reliable trade.

An inherent difficulty in automated trading — including e-procurement — is that it is generally multi-issue. Even a simple trade, such as a quantity of steel, may involve: delivery date, settlement terms, as well as price and the quality of the steel. The “information-based” agent’s reasoning is based on a first-order logic world model that manages multi-issue negotiation as easily as single-issue.

Information-Based Architecture

Agent Π engages in multi-issue negotiation with a set of other agents: $\{\Omega_1, \cdots, \Omega_o\}$. In addition to the information derived from its opponents, Π has access to a set of information sources $\{\Theta_1, \cdots, \Theta_t\}$ that may include the marketplace in which trading takes place, and general information sources such as news-feeds accessed via the Internet. Together, Π, $\{\Omega_1, \cdots, \Omega_o\}$ and $\{\Theta_1, \cdots, \Theta_t\}$ make up a multiagent system.

Π has two languages: $\mathcal{C}$ and $\mathcal{L}$. $\mathcal{C}$ is an illocutionary-based language for communication. $\mathcal{L}$ is a first-order language for internal representation. Messages expressed in $\mathcal{C}$ from $\{\Theta_i\}$ and $\{\Omega_i\}$ are received, time-stamped, source-stamped and placed in an *in-box* $\mathcal{X}$. The messages in $\mathcal{X}$ are then translated using an *import function* I into sentences expressed in $\mathcal{L}$ that have integrity decay functions (usually of time) attached to each sentence, they are stored in a *repository* $\mathcal{Y}^t$

The Architecture

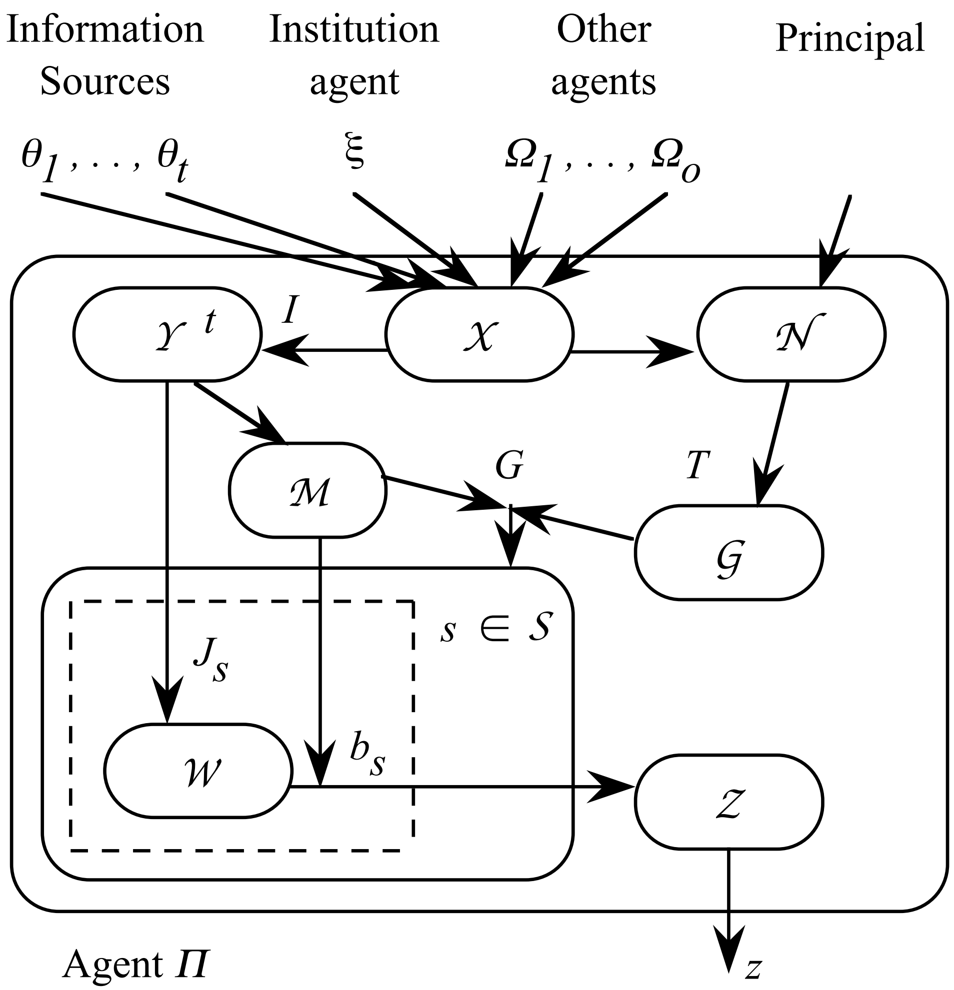

Π triggers a goal, $g \in \mathcal{G}$, in two ways: first in response to a message received from an opponent $\{\Omega_i\}$ "I offer you €1 in exchange for an apple", and second in response to some internal need, $\nu \in \mathcal{N}$.

Π's Reasoning

- Maximum entropy inference, J_s^+, first constructs an *information base* $\mathcal{I}_s^t$ as a set of sentences expressed in $\mathcal{L}$ derived from $\mathcal{Y}^t$, and then from $\mathcal{I}_s^t$ constructs the world model, W_s^t, as a set of complete probability distributions.
  - Of all models that are consistent with $\mathcal{Y}^t$, W_s^t is the model that has maximum entropy — it is "maximally non-commital" with respect to that which is unknown.
- Given a prior world model, W_s^u, where $u < t$, minimum relative entropy inference, J_s^-, first constructs the incremental information base $\mathcal{I}_s^{(u,t)}$ of sentences derived from those in $\mathcal{Y}^t$ that were received between time u and time t, and then from W_s^u and $\mathcal{I}_s^{(u,t)}$ constructs a new world model, W_s^t.
  - W_s^t is model with minimum relative entropy with respect to W_s^u that is consistent with $\mathcal{I}_s^{(u,t)}$.

Π Manages Information

A chunk of $[info]$ may not be directly related to one of Π's chosen distributions or may not be expressed naturally as constraints, and so some inference machinery is required to derive these constraints — this inference is performed by model building functions, J_s, that have been activated by a plan s chosen by Π.

$J_s^D([info])$ denotes the set of constraints on distribution D derived by J_s from $[info]$.

In the absence of new $[info]$ the integrity of distribution $D = (q_i)_{i=1}^n$ decays:

$$q_i^{t+1} = (1 - \rho^D) \times d_i^D + \rho^D \times q_i^t, \text{ for } i = 1, \ldots, n$$

$\rho^D \in (0, 1)$ is the decay rate, and $(d_i^D)_{i=1}^n$ is the *decay limit distribution*.

Updating the world model with [*info*]

A single chunk of $[\mathit{info}]$ may effect a number of distributions. Suppose that a chunk of $[\mathit{info}]$ is received from Ω and that Π attaches the epistemic belief probability $R^t(\Pi, \Omega, O([\mathit{info}]))$ to it. Each distribution models a facet of the world. Given a distribution $D^t = (q_i^t)_{i=1}^n$, q_i^t is the probability that the possible world ω_i for D is the true world for D. The effect that a chunk $[\mathit{info}]$ has on distribution D is to enforce the set of linear constraints on D, $J_s^D([\mathit{info}])$. If the constraints $J_s^D([\mathit{info}])$ are taken by Π as valid then Π could update D to the posterior distribution $(p_i^{[\mathit{info}]})_{i=1}^n$ that is the distribution with least relative entropy with respect to $(q_i^t)_{i=1}^n$ satisfying the constraint:

$$\sum_i \{p_i^{[\mathit{info}]} \;:\; J_s^D([\mathit{info}]) \text{ are all } \top \text{ in } \omega_i\} = 1.$$

Updating the world model with [*info*] — contd.

But $R^t(\Pi, \Omega, O([info])) = r \in [0,1]$ and Π should only treat the $J_s^D([info])$ as valid if $r = 1$. r determines the extent to which the effect of $[info]$ on D is closer to $(p_i^{[info]})_{i=1}^n$ or to the prior $(q_i^t)_{i=1}^n$:

$$p_i^t = r \times p_i^{[info]} + (1 - r) \times q_i^t$$

But, we should only permit a new chunk of $[info]$ to influence D if doing so gives us new information. Ie: the resulting distribution should have more information *relative to* the decay limit distribution than the existing distribution has. Precisely, this is measured using the Kullback-Leibler distance measure, and $[info]$ is only used if:

$$\sum_{i=1}^n p_i^t \log \frac{p_i^t}{d_i^D} > \sum_{i=1}^n q_i^t \log \frac{q_i^t}{d_i^D}$$

In addition, the integrity of each distribution D will decay in time.

Updating the world model — Summary

Distribution D is revised to:

$$q_i^{t+1} = \begin{cases} (1-\rho^D) \times d_i^D + \rho^D \times p_i^t & \textit{if} \text{ usable } [\textit{info}] \text{ is} \\ & \text{received at time } t \\ (1-\rho^D) \times d_i^D + \rho^D \times q_i^t & \textit{otherwise} \end{cases}$$

for $i = 1, \cdots, n$, decay rate ρ^D, and decay limit distribution $(d_i^D)_{i=1}^n$.

We have yet to estimate $R^t(\Pi, \Omega, O([\mathit{info}]))$. We do this by measuring the error in information. Π's plans will have constructed a set of distributions. We measure the 'error' in information as the error in the effect that information has on each of Π's distributions.

Information Reliability

Suppose that a chunk of $[info]$ is received from agent Ω at time s, and at some later time t it is found that $[info]$ should have been $[fact]$. Let $(p_i^{[info]})_{i=1}^n$ be the minimum relative entropy distribution given that $[info]$ has been received, and let $(p_i^{[fact]})_{i=1}^n$ be that distribution if $[fact]$ had been received instead. The *observed reliability* for distribution D, $R_D^{([info]|[fact])}$, is the value of r that minimises the Kullback-Leibler distance between $(p_i^s)_{i=1}^n$ and $(p_i^{[fact]})_{i=1}^n$:

$$\arg\min_r \sum_{i=1}^n (r \cdot p_i^{[info]} + (1-r) \cdot q_i^s) \log \frac{r \cdot p_i^{[info]} + (1-r) \cdot q_i^s}{p_i^{[fact]}}$$

If $E^{[info]}$ is the set of distributions that $[info]$ effects, then the overall *observed reliability* on the basis of the verification of $[info]$ with $[fact]$ is: $R^{([info]|[fact])} = 1 - (\max_{D \in E^{[info]}} |1 - R_D^{([info]|[fact])}|)$.

Estimating $R^t(\Pi, \Omega, O([\textbf{info}]))$

For each ontological context o_j, at time t when, perhaps, a chunk of $[\mathit{info}]$, with $O([\mathit{info}]) = o_k$, may have been verified with $[\mathit{fact}]$:

$$R^{t+1}(\Pi, \Omega, o_j) =$$
$$(1 - \rho) \times R^t(\Pi, \Omega, o_j) + \rho \times R^{([\mathit{info}]|[\mathit{fact}])} \times \mathrm{Sem}(o_j, o_k)$$

where $\mathrm{Sem}(\cdot, \cdot) : O \times O \rightarrow [0, 1]$ measures the semantic distance between two sections of the ontology, and ρ is the learning rate.

This leads to a overall expectation of the *reliability* that agent Π has for agent Ω:

$$R^t(\Pi, \Omega) = \sum_j P^t(o_j) \times R^t(\Pi, \Omega, o_j)$$

Valuing Information

Suppose that a set of stamped messages $X = \{x_i\}$ is received in $\mathcal{X}$. The *information* in X at time t with respect to a particular distribution $D^t_{s,i} \in W^t_s$, strategy s, goal g and environment e is:

$$\mathbb{I}(X \mid D^t_{s,i}, s, g, e) \triangleq \mathbb{H}(D^t_{s,i}(\mathcal{Y}^t)) - \mathbb{H}(D^t_{s,i}(\mathcal{Y}^t \cup I(X)))$$

for $i = 1, \cdots, n$, where the argument of the $D^t_{s,i}(\cdot)$ is the state of Π's repository from which $D^t_{s,i}$ was derived. Then:

$$\mathbb{I}(X \mid s, g, e) \triangleq \sum_i \mathbb{I}(X \mid D^t_{s,i}, s, g, e)$$

and:

$$\mathbb{I}(X \mid g, e) \triangleq \sum_{s \in \mathcal{S}(g)} \mathbb{P}(s) \cdot \mathbb{I}(X \mid s, g, e)$$

Virtual Institutions

In collaboration with "Institut d'Investigacio en Intel.ligencia Artificial[2]", Spanish Scientific Research Council, UAB, Barcelona.

Electronic Institutions are composed of autonomous agents, that interact according to predefined conventions and that guarantee that certain norms of behaviour are enforced.

Virtual Institutions enable rich interaction, based on natural language and embodiment of humans and software agents in a "liveable" vibrant environment. This view permits agents to behave autonomously and take their decisions freely up to the limits imposed by the set of *norms* of the institution. An important consequence of embedding agents in a virtual institution is that the predefined conventions on language and protocol greatly simplify the design of the agents.

---

[2] `http://www.iiia.csic.es/`

Virtual Institutions — Requirements

A Virtual Institution is in a sense a natural extension of the social concept of institutions as regulatory systems that shape human interactions.

Virtual Institutions are electronic environments designed to meet the following requirements towards their inhabitants:

1. enable institutional commitments including structured language and norms of behaviour which enable reliable interaction between autonomous agents and between human and autonomous agents;
2. enable rich interaction, based on natural language and embodiment of humans and software agents in a “liveable” vibrant environment.

VI's — Requirement 1

Has been addressed to some extent by the Electronic Institutions (EI) methodology and technology for multi-agent systems, developed in the Spanish Government's IIIA Laboratory in Barcelona.

The EI environment is oriented towards the engineering of multiagent systems. The Electronic Institution is an environment populated by autonomous software agents that interact according to predefined conventions on language and protocol. Following the metaphor of social institutions, Electronic Institutions guarantee that certain norms of behaviour are enforced. This view permits that agents behave autonomously and make their decisions freely up to the limits imposed by the set of norms of the institution. The interaction in such environment is regulated for software agents. The human, however, is "excluded" from the electronic institution.

VI's — Requirement 2

The second requirement is supported to some extent by the distributed 3D Virtual Worlds technology.

Emulating and extending the physical world in which we live, Virtual Worlds offer rich environment for a variety of human activities and multi-mode interaction. Both humans and software agents are embedded and visualised in such 3D environments as avatars, through which they communicate.

Following the metaphor of the physical world, these environments do not impose any regulations (in terms of language) on the interactions and any restrictions (in terms of norms of behaviour). These environments do not provide means for enabling some behavioural norms, for example, fulfilling commitments, penalisation for misbehaviour and others.

Virtual Institutions address both Requirements

interaction richness


Agent Name Information Assistant

Description Find out information about the e-marketplace

Virtual Worlds
e-market solutions

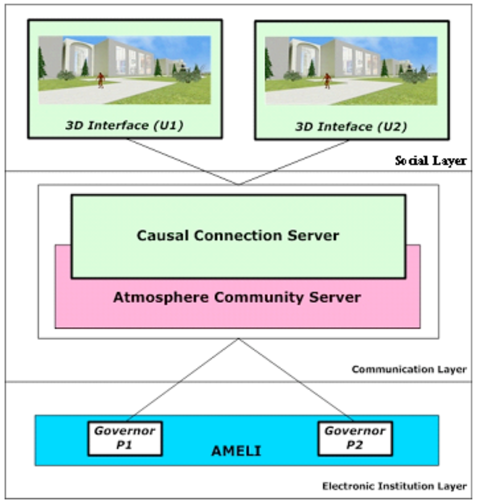

Virtual Institutions
e-market solutions

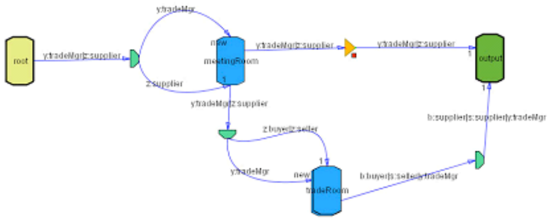

Electronic Institutions
e-market solutions

level of regulation

Virtual Institution Implementation Framework

1. The Electronic Institution Layer hosts the environments that support the Electronic Institutions technological component. At runtime, the Electronic Institution layer loads the institution specification and mediates agents interactions while enforcing institutional rules and norms.

2. Communication Layer connects causally the Electronic Institutions layer with the 3D representation of the institution, which resides in the Social layer. The causal connection is the integrator.

3. In the Social Layer a Virtual Institution representation is a graph and its topology can structure the space of the virtual environment in different ways. This layer is represented in terms of a 3D Virtual World technology, structured around rooms, avatars, doors (for transitions) and other graphical elements.

Three-layer Virtual Institution Framework

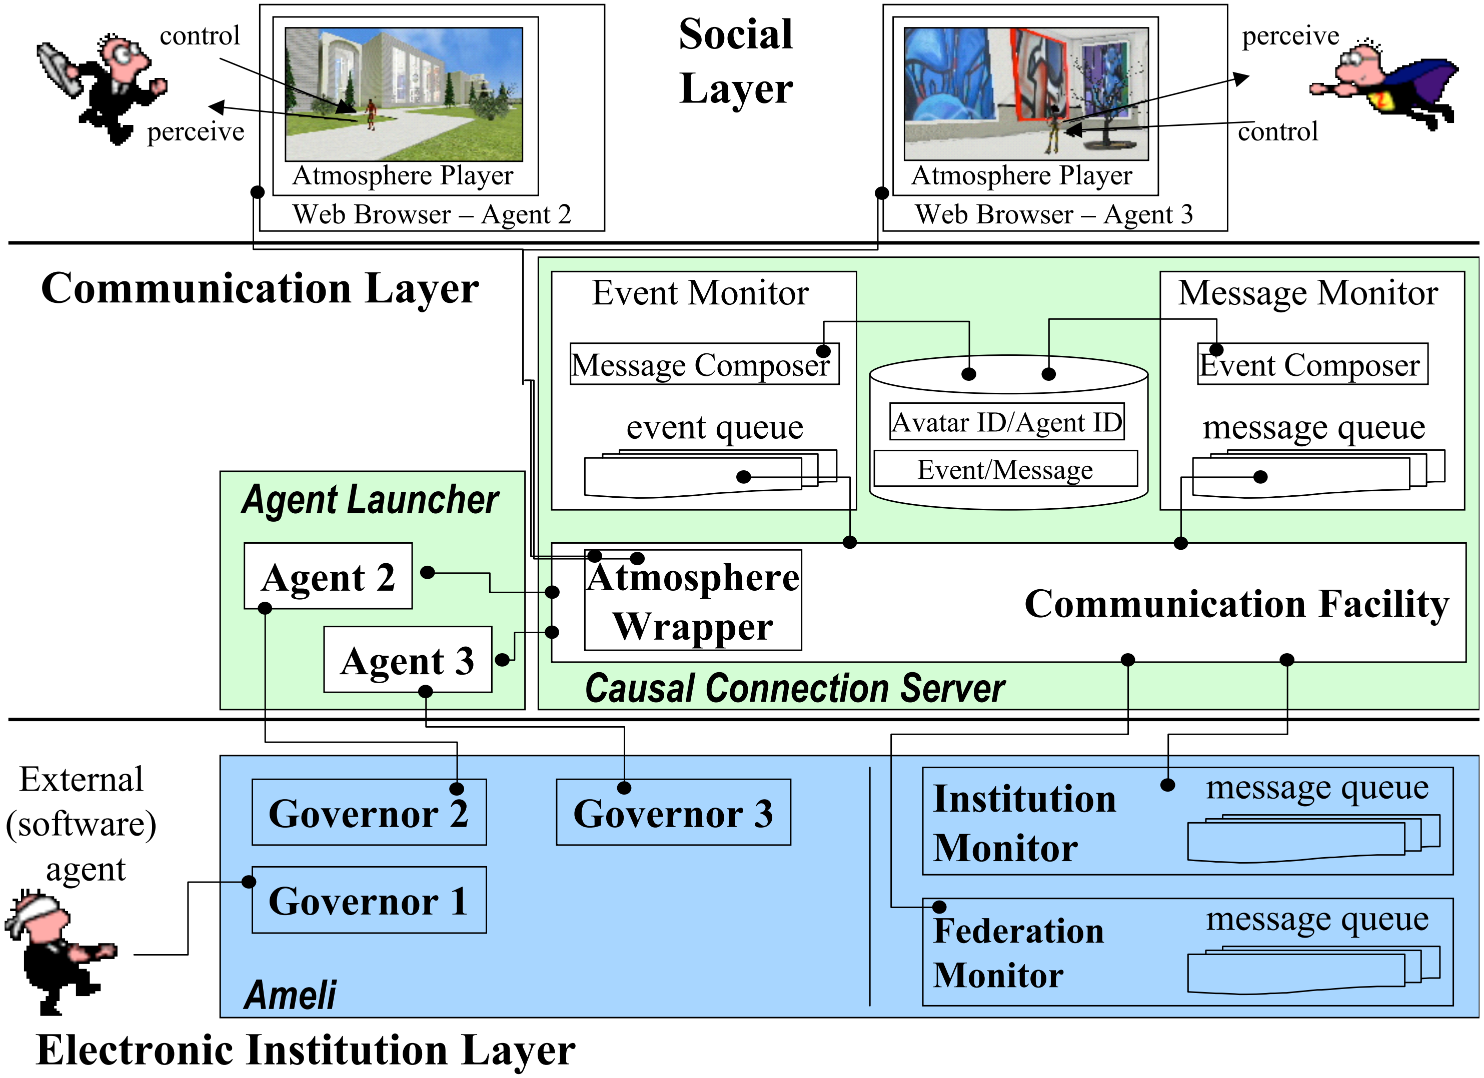

control

perceive

Atmosphere Player

Web Browser – Agent 2

Social Layer

Atmosphere Player

Web Browser – Agent 3

perceive

control

Communication Layer

Event Monitor

Message Composer

event queue

Avatar ID/Agent ID

Event/Message

Message Monitor

Event Composer

message queue

Agent Launcher

Agent 2

Agent 3

Atmosphere Wrapper

Communication Facility

Causal Connection Server

External (software) agent

Governor 2

Governor 3

Governor 1

Ameli

Institution Monitor

message queue

Federation Monitor

message queue

Electronic Institution Layer

Virtual Institutions — Core Technology

The core technology — the Causal Connection Server — enables the Communication Layer to act in two directions.

From the Electronic Institution layer, messages uttered by an agent have immediate impact on the Social layer. Transitions of agents between scenes in the EI layer must permit the corresponding avatar move within the Virtual World space accordingly.

From the Social layer, events caused by the actions of the human avatar in the Virtual World are transferred to the Electronic Institution layer and passed to an agent. This means that actions forbidden to the agent by the norms of the institution (encoded in the Electronic Institution layer), cannot be performed by the human. For example, if a human needs to register first before leaving for the auction space, the corresponding agent is not allowed to leave the registration scene. Consequently, the avatar is not permitted to open the corresponding door to the auction.

Conclusions

A demonstrable prototype e-Market system permits both human and software agents to trade with each other on the World Wide Web. The main contributions described are: the broadly-based and "focussed" data mining systems, the intelligent agent architecture founded on information theory, and the abstract synthesis of the virtual worlds and the electronic institutions paradigms to form "virtual institutions". These three technologies combine to present our vision of the World Wide Web marketplaces of tomorrow.