

CmpE 593 Multiagent Systems

Pinar Yolum
pyolum@cmpe.boun.edu.tr

Department of
Computer Engineering
Boğaziçi University

Chapter 7 Trust

Based largely on
Service-Oriented Computing: Semantics, Processes, Agents
– Munindar P. Singh and Michael N. Huhns, Wiley, 2004

Service Selection

- Finding the right service provider from a set of providers
- Yellow-pages
 - Lookup based on service criteria
 - May not always exist
 - May return many results
- Economic service selection
- Semantic service selection

Economic Service Selection

- Market-oriented programming
 - Design an environment with mechanisms for buying and selling
 - Little interaction between agents; mostly for exchanging goods at different prices
 - Preferences or abilities of agents are not considered
- Consumer and producer agents
 - Self-interested
 - Maximize their utility

Prices

- The computational state is described completely by current prices for the various goods
- Communications are between each participant and the market, and only in terms of prices
- Participants reason about others and choose strategies entirely in terms of prices being bid

Functions of a Market

- Provides this information to participants
- Takes requests (buy, sell bids) from participants, enforcing rules such as bid increments and time limits
- Decides outcome based on messages from participants, considering rules such as reserve prices.

Equilibrium

- At equilibrium, the market has computed the allocation of resources and dictates the activities and consumptions of the agents.
- Under certain conditions, a simultaneous equilibrium of supply and demand across all goods exists
 - Reachable via distributed bidding
 - *Pareto optimal*: you cannot make the allocation better for one agent without making it worse for another

Auctions

- Market where prices are determined dynamically
- Online auctions where agents participate
 - Must be fair and secure
 - Preserve privacy
- Auction types vary
 - Ascending (English) vs. Descending (Dutch)
 - Silent (auctioneer names a price; bids are silent) vs. outcry (bids name prices; auctioneer listens)
 - Hidden identity or not.
 - Combinatorial: involve *bundles* or sets of goods

English and Dutch Auctions

- English
 - Prices start low and increase
 - Highest bidder gets the object at price bid
 - Variations:
 - Minimum bid increment
 - Reserve price (no sale if too low)
 - Limited time
- Dutch
 - Prices start high and decrease
 - First to interrupt wins

Sealed-Bid Auctions

Also known as *tenders*: bidding to buy.

- One-shot bidding without knowing what other bids are being placed.
- Used by governments and large companies to give out certain large contracts (lowest price quote for stated task or procurement).
 - All bids are gathered.
 - Auctioneer decides outcomes based on given rules (e.g., highest bidder wins and pays the price it bid).
- Vickrey Auction:
 - Second price sealed bid auction
 - Highest bidder wins, but pays the *second* highest price

Continuous Double Auction

As in stock markets.

- Multiple sellers and buyers, potentially with multiple sell and buy bids each.
- Buy bids are like upper bounds
- Sell bids are like lower bounds
- *Clears* continually:
 - The moment a buyer and seller agree on a price, the deal is done and the matching bids are taken out of the market
 - Possible, a moment later a better price may come along, but it will be too late then.

Auction Management (1)

- Bidding rules to govern, e.g.,
 - Whose turn it is
 - What the minimum acceptable bid is, e.g., increments
- Information disclosure
 - What information is revealed to participants?
 - Bid value
 - Bidder
 - Winning bid
 - Winner
 - How often

Auction Management (2)

- Bids are cleared when they are executed and taken out of the market.
 - How are bids matched
 - Who?
 - What prices?
 - How often?
 - Until when?

Problems

- More applicable for services that differ only on price
- Services differ on quality?
- Negotiation of service descriptions
- Suggestions of service providers
- Semantic service selection
 - Requires deciding on which service provider will do the job best for the user
 - Take into account provider's reputation or customer's trust in the provider

Reputation

- Consider a society of principals, potentially each having opinions about the others.
 - The opinions are applied implicitly in whether and how different parties do business with each other
- Someone's reputation is a *general opinion* about that party
 - Sometimes partially probed by asking others
 - Never explicitly fully aggregated, except in current computational approaches

Reputation

The agency (e.g., eBay) is the authority that

- Authenticates users
- Records, aggregates, and reveals ratings
- Provides the conceptual schema for
 - How to capture ratings (typically a number and text)
 - How to aggregate them
 - How to decay them over time

Why a Decentralized Approach?

Problems with explicit aggregation

- *Context and understanding*: The contexts of usage may not be in agreement.
- *Privacy*: The parties providing their ratings are stating publicly (or to the reputation agency) what they may only wish to reveal in private.
- *Trust*: The parties using the ratings don't necessarily know where the ratings come from.

Views of Trust

- Property of individuals
 - lacks context
 - confused with cooperative mentality, honesty, loyalty, sincerity, altruism
- Property of social relationships
 - social resource used to achieve collective goals
- Property of the social system (based on individual actions)
 - emerges from the social system
 - affected by social roles
 - building trust in micro level results in trust in macro level (your trust in a doctor, increases your trust in the medical system).

Trust

- To trust is to believe that the results of somebody's intended actions will be appropriate from our point of view. Misztal
- Trust is a bet about the future contingent actions of others. Szompka
- A nation's well-being, as well as its ability to compete, is conditioned by a single pervasive cultural characteristic: the level of trust inherent in a society. Fukuyama
(Gambetta)

Confidence vs. trust

Luhmann's distinction:

- Hope: Wish it will come true (no basis).
- Confidence: Think it will come true (based on evidence).
- Trust: Commit to action with partly uncertain consequences.

- Risk (vs. Danger): Unexpected results may be a consequence of our decisions (not just fate).

Varieties of Commitment

Szompka relates trust to commitments:

- *Anticipatory trust*: Expecting the other party to do its work normally (choosing an airline based on past reliability).
- *Responsive trust*: Giving up control on an object and giving it to someone else (leaving a child with a babysitter).
- *Evocative trust*: Expecting to initiate reciprocal trust.

Circumstances of Commitment

- Range of consequences (employment vs. trip)
- Expected duration (marriage vs. date)
- Possibility of withdrawal (pre-commitment)
- Amount of risk relative to its probability (flight vs. surgery)
- Presence of insurance (bank with state guarantees)
- Value of object to be trusted (lending a car vs. a book)

Primary trustworthiness

- Reputation
 - Contagious (buying a bestseller, Bloomingdale's effect)
 - Despite its value, may become a burden.
- Performance
- Appearance
 - Dress, bodily discipline, civility
 - Easy to fake

“People tend to trust others who are similar to them and to distrust those who are dissimilar from them”
(Earle & Cvetkovich)

Estimation of trust

- Anonymity limits trust.
- Clarity of criteria
 - Results of doing a job (athlete vs. professor)
- Competence to perceive
 - How can you decide on the performance of a scientist?

Secondary trustworthiness (1)

- **Accountability**
 - Presence of monitoring agencies (court, police)
 - Incentive to be trustworthy
- **Precommitment**
 - Covenant marriages in Louisiana
 - More serious, higher degrees of love
 - Lowers the trustworthiness of other marriages

Secondary trustworthiness (2)

- **Situational facilitation (effect of setting)**
 - Symphony vs. soccer game
 - Visibility & closeness (eliminating anonymity)
 - Taping conversations in White House
 - Name badges for soccer fans
 - Sacred quality of the setting
 - Self-enforcing setting (traffic, skiing)

Social Capital

- Figure out one's reputation from background information (Dasgupta).
- Reputation is acquired gradually, but destroyed quickly.
- Meet few honest people and no dishonest ones => revise your opinion about people (micro-macro link) (Hirschman).
- Supply of trust may well increase than decrease through use.
- Do not remain intact if they stay unused (like speaking a foreign language, playing the piano).

Functions of Trust

- Endowing trust
 - Trust creates trust.
 - May be reciprocated by mutual trust.
- Receiving trust
 - Creates non-conformity, freedom of action
 - Invoke trust from others (word travels)

Functions of Trust

Trust


- encourages sociability
 - enriches interpersonal ties, moral density (Durkheim)
 - social capital (Putnam)
- encourages tolerance (vs. rumors, prejudice)
- strengthens the tie of an individual with community (vs. mafia) (Gambetta)
- lowers transaction costs, increases chances of cooperation (Fukuyama).

Computational Trust

- Institutional Trust
 - Organizations monitor members' actions
 - Ensure a quality of service
 - Realized by digital certificates
- Local Trust
 - Based on personal evidence
 - Prior interactions
- Social Trust
 - Based on evidence from others
 - Information sources should be trustworthy

Social Trust: Referrals

- An agent represents a principal offering or searching for services.



- An agent generates a query for a service and sends it to its *neighbors* (a small subset of its acquaintances). Each neighbor may provide the service or refer to other agents (based on its *referral policies*).

- Each agent models the *expertise* (quality of a service) and *sociability* (quality of the referrals) of its acquaintances.
- Based on these models, each agent can change its set of neighbors (using its *neighbor selection policy*).
- Referral network*: as induced by the neighborhood relation.



Spring 2005—Pinar Yolum

31

Application Domains

Commerce:

- Distinct service producers and consumers.
- Producers have expertise, consumers have sociability.
- Answers are easy to evaluate.
- Expertise of consumers does not increase.



Knowledge Management:

- All agents can be producers and consumers.
- Answers are harder to evaluate.
- Expertise of consumers may increase (expertise of the producers can be cached by others).

Spring 2005—Pinar Yolum

32

Basic Experimental Setup

- Interests used to generate queries
- Query, answer, interest, and expertise are vectors from Vector Space Model where each dimension corresponds to a domain
- Dimension of the vectors is 4
- Sociability is scalar
- 400 agents, with 10 to 25% service providers
- 8 neighbors per consumer
- Initial neighbors picked randomly
- Reselect neighbors after every 2 queries
- 4 to 20 neighbor changes

Spring 2005— Pinar Yolum

33


Referral Policies

- **Refer all neighbors:**
 - Does not consider which neighbors would be more likely to answer.
- **Refer all matching neighbors:**
 - Calculate how capable each neighbor is in answering a particular query.
$$Q \otimes E = \frac{\sum_{t=1}^n (q_t e_t)}{\sqrt{n \sum_{t=1}^n q_t^2}}$$
 - Refer those neighbors with “sufficient” expertise (those greater than a given threshold \mathbf{T}).
- **Refer best neighbor:**
 - Refer the most capable neighbor.
 - Guarantees that at least one neighbor is referred.

Spring 2005— Pinar Yolum

34

Efficiency of Referral Policies



Policies:

- Refer All
- Refer All Matching
- Refer Best


$$\text{Efficiency} = \frac{\text{\# of good answers}}{\text{\# of contacted agents}}$$

Too many agents are contacted.

Not enough good answers are found.

Spring 2005— Pinar Yolum

Effectiveness of Referral Policies



For each agent A

- Rank all others (e.g., B) using $\frac{I_A \otimes E_B}{\text{path}(A, B)}$
- Take the n th highest
- n is twice the neighbor set size (16)


Network quality = Average over all agents

Low quality even though answers are found

Low efficiency but high quality

Spring 2005— Pinar Yolum

Authoritativeness



PageRank: Pages pointed to by authorities are also authoritative.

$$P(i) = d \sum_{j \in \kappa_i} \frac{P(j)}{|\mathcal{N}_j|} + (1-d)$$

- The extent of authoritativeness increases as the agents exchange more referrals.

Spring 2005— Pinar Yolum

Neighbor Selection Policies

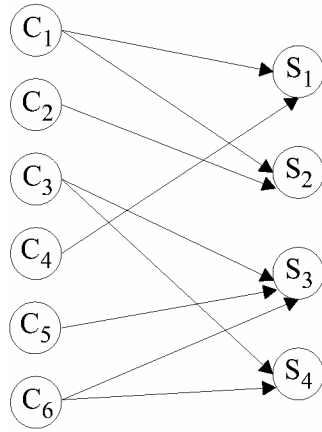
Consider weighted sociability (W) and expertise.

- **Weighted Average ($W=0.25, 0.50, 0.75$)**
 - Choose the best m based on weighing both the expertise and the sociability of the acquaintances.
- **Providers ($W=0$)**
 - Choose the best m agents whose expertise matches the agent's interests.
- **Sociables ($W=0.90, 1$)**
 - Choose the most sociable m acquaintances.

Spring 2005— Pinar Yolum

38

Graph Structures




Bipartite Graphs

Find how close a graph is to being bipartite by removing k edges

Providers policy creates a bipartite graph.

Spring 2005—Pinar Yolum

Clustering



Measures how similar the neighbors of an agent are as well as how similar the agent is to its neighbors.

$$\frac{e^{-\|I_A - I_B\|^2} - e^{-n}}{1 - e^{-n}}$$


Sociability increases interest clustering.

Agents with similar interests

- May be looking for similar providers.
- May give useful referrals.
- Thus, will be considered sociable, and kept as neighbors.

Spring 2005—Pinar Yolum

Clustering



Quality decreases when interest clustering increases.

- Agents with similar interests point to same providers.
- Other providers are not reachable.

Spring 2005—Pinar Yolum

Discussion

- Study emerging properties
 - How the referrals are exchanged.
 - How the neighbors are chosen.
- Identify undesirable structures

Use these properties to design mechanisms to build applications of referral systems

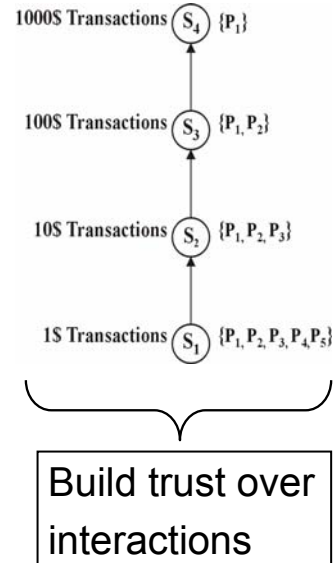
Referral framework with

- autonomous
- heterogeneous
- adaptive agents.

Spring 2005—Pinar Yolum

Local Trust

- Each node denotes transactions with different value.
- Most providers can be trusted for low value transactions.
- Fewer providers are trusted for high value transactions.
- Gradually *promote* providers to try for transactions with greater values.

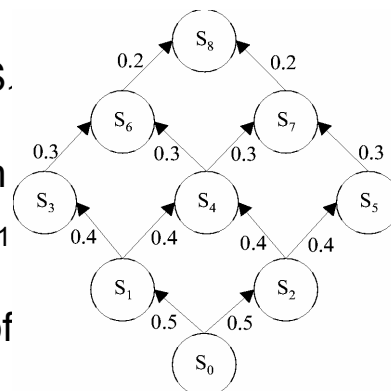


Spring 2005— Pinar Yolum

43

Service Graphs

- Each node denotes a service type.
- An edge between S_0 and S_1 denotes that an agent that can perform service S_0 can *possibly* perform service S_1 .
- The weight of an edge determines the likelihood of performing the upper service.



Spring 2005— Pinar Yolum

44

Example

- Ali's agent is looking for a restaurant to buy baklava.
- Ali's agent bought cheesecake from TGI Fridays twice before and Ali rated TGI Fridays as 0.8 and 0.9 for the overall service.
- Ali's agent has to decide if it can trust TGI Fridays to buy baklava considering the following:
 - How related serving cheesecake and serving baklava is?
 - How satisfied Ali was with TGI Fridays' cheesecake?
- If Ali's agent promotes TGI Fridays
 - Evaluates TGI Fridays for serving baklava.
 - Updates the weight between serving cheesecake and serving baklava based on the evaluation.

Spring 2005—Pinar Yolum

45

Representations

- Query: A service type (e.g., serving cheesecake).
- Answer: Identifier of a service provider (e.g., TGI Fridays).
- Expertise:
 1. A service graph where each node denotes a distinct service.
 2. A vector where each dimension corresponds to a domain.
 - Vector representation does not capture relations between service types.

Spring 2005—Pinar Yolum

46

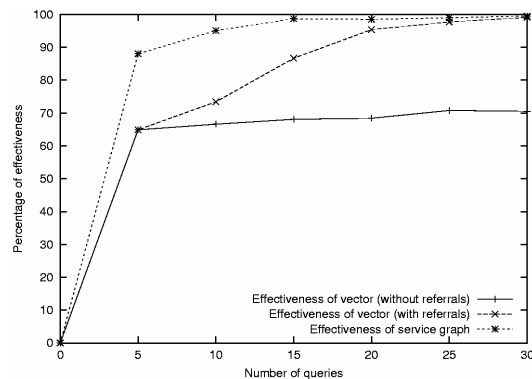
Experimental Setup

- Simulation runs
 - 100 agents, each with 3 neighbors
 - 30 queries per agent
 - Change neighbors every 3 queries
- Two cases of vector representation
 - With referrals
 - Without referrals

Spring 2005— Pinar Yolum

47

Effectiveness




How often can the desired service providers be found? (how many good answers does each agent get?)

- Service graphs achieve high effectiveness earlier than vector with referrals.
- Vector without referrals improves slightly.

Spring 2005— Pinar Yolum

Efficiency




How easily can the service providers be found (how many messages are sent?)

- Using service graphs, agents contact fewer agents than both vector approaches.
- The number of contacted agents drop with more queries.

Spring 2005—Pinar Yolum

Initial Setting


- *Homogenous:* Each agent has the same service graph.
- *Trusting:* Each agent has the same service graph, but the weights are higher.
- *Heterogeneous:* Each agent picks edges and weights randomly.



Spring 2005—Pinar Yolum

50

Promotion Error

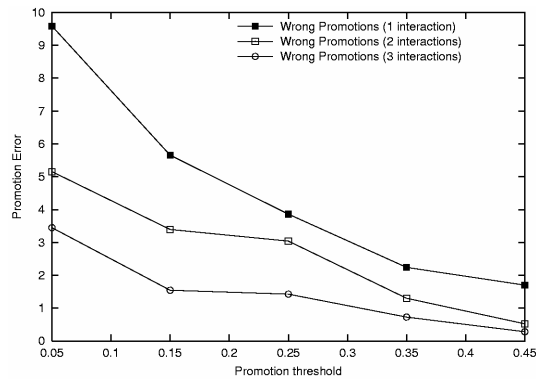


- Promotion error drops with more selective promotions.
- Initially having low trust (homogenous case) causes fewest errors.

Spring 2005—Pinar Yolum

51

Previous Interactions




- Increasing number of interactions decreases the number of errors.

Spring 2005—Pinar Yolum

52

Effectiveness (2)




- If agents are too cautious in their promotions, they locate fewer service providers (e.g., threshold of 0.45).

Spring 2005— Pinar Yolum

53

Effectiveness (3)

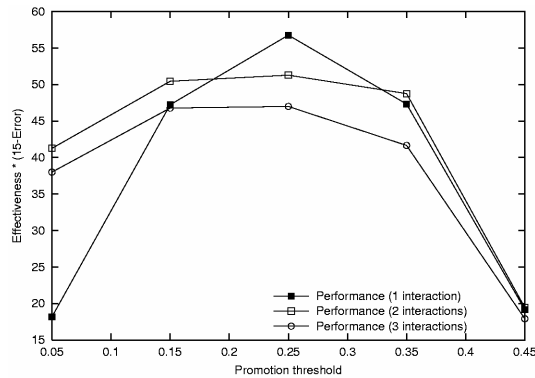


- It is easier to find service providers if they can be promoted from other services easily (i.e., after 1 interaction).

Spring 2005— Pinar Yolum

54

Effectiveness (4)



- With low threshold, there is high promotion error.
- With high threshold, providers cannot be found.
- Middle values keep error low, allow locating providers.
- Third interaction adds little value.

Spring 2005—Pinar Yolum

55

Small World Phenomenon

Milgram (1967): two individuals chosen at random in the U.S.A. are linked by a chain of 6 or fewer first-name acquaintances (empirical observation)

- Six degrees of separation
- Erdős numbers
- Diameter of the connected Web: 13

Spring 2005—Pinar Yolum

56

Small-World Network

- Generated by perturbing a regular ring
 - A highly structured (clustered) network with just a few random edges
 - Random edges correspond to shortcuts
- Yields high clustering and short paths
 - Direct relationships between agents who primarily participate in different subcommunities
 - Shortcuts: weak ties