Trust-oriented Utility-based Community Structure in Multiagent Systems

Georgia Kastidou and Robin Cohen

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada

{rcohen,gkastidou} @uwaterloo.ca

Background

- Artificial intelligence
  - Intelligent agents
  - Multiagent systems: distributed problem solving

  Example: P2P system
  - Agents modeling users

- Trust Modeling in Multiagent systems
  - Direct experience
  - Social networks

  Applications: e-marketplaces

Topic & Motivation

- Modeling trustworthiness of agents in communities
- Communities (e.g. P2P: limewire, freenet, gnunet etc):
  - Initially Open and Free
  - Currently… moving towards a more close and incentive-based direction. Reason?
    - The existence of malicious and selfish agents can have a significant negative impact on a community.
- Proposal:
  - Exploit the information that each community accumulates regarding its agents
    - Agents provide better contributions
    - Communities enjoy better agents

The Problem

- *"Increment of the social welfare of a set of communities based on the exchange of agents' trust and reputation information with concern for privacy protection"*
- **Solution**:

Central Idea: *Provide incentive for agents to be good contributors/participants* **and** *allow communities to exploit information of agents' past behavior.*

How? Apply Matching of Agents and Communities.

- Agents Reason About Communities
- Communities Reason About Agents

Our Goal

- **Discourage:**
  - malicious behaviour
  - free agents
  - white washers

  by letting them know that their behaviour will be spread in case they want to join another community.

- **Encourage** the agents to ask honestly by providing them with better chances to enter a new community.

Overview

- Agents decide which communities to join with respect to
  - their resources
  - the special characteristics of each community.

  Goal: Maximize Their Utility

- Communities decide which agents to accept/trust based on their reputation and trustworthiness.

  Goal: Maximize Their Agents' Overall "Utility"
  - Attract agents which are honest/good contributors

The Problem: Matching Communities & Agents

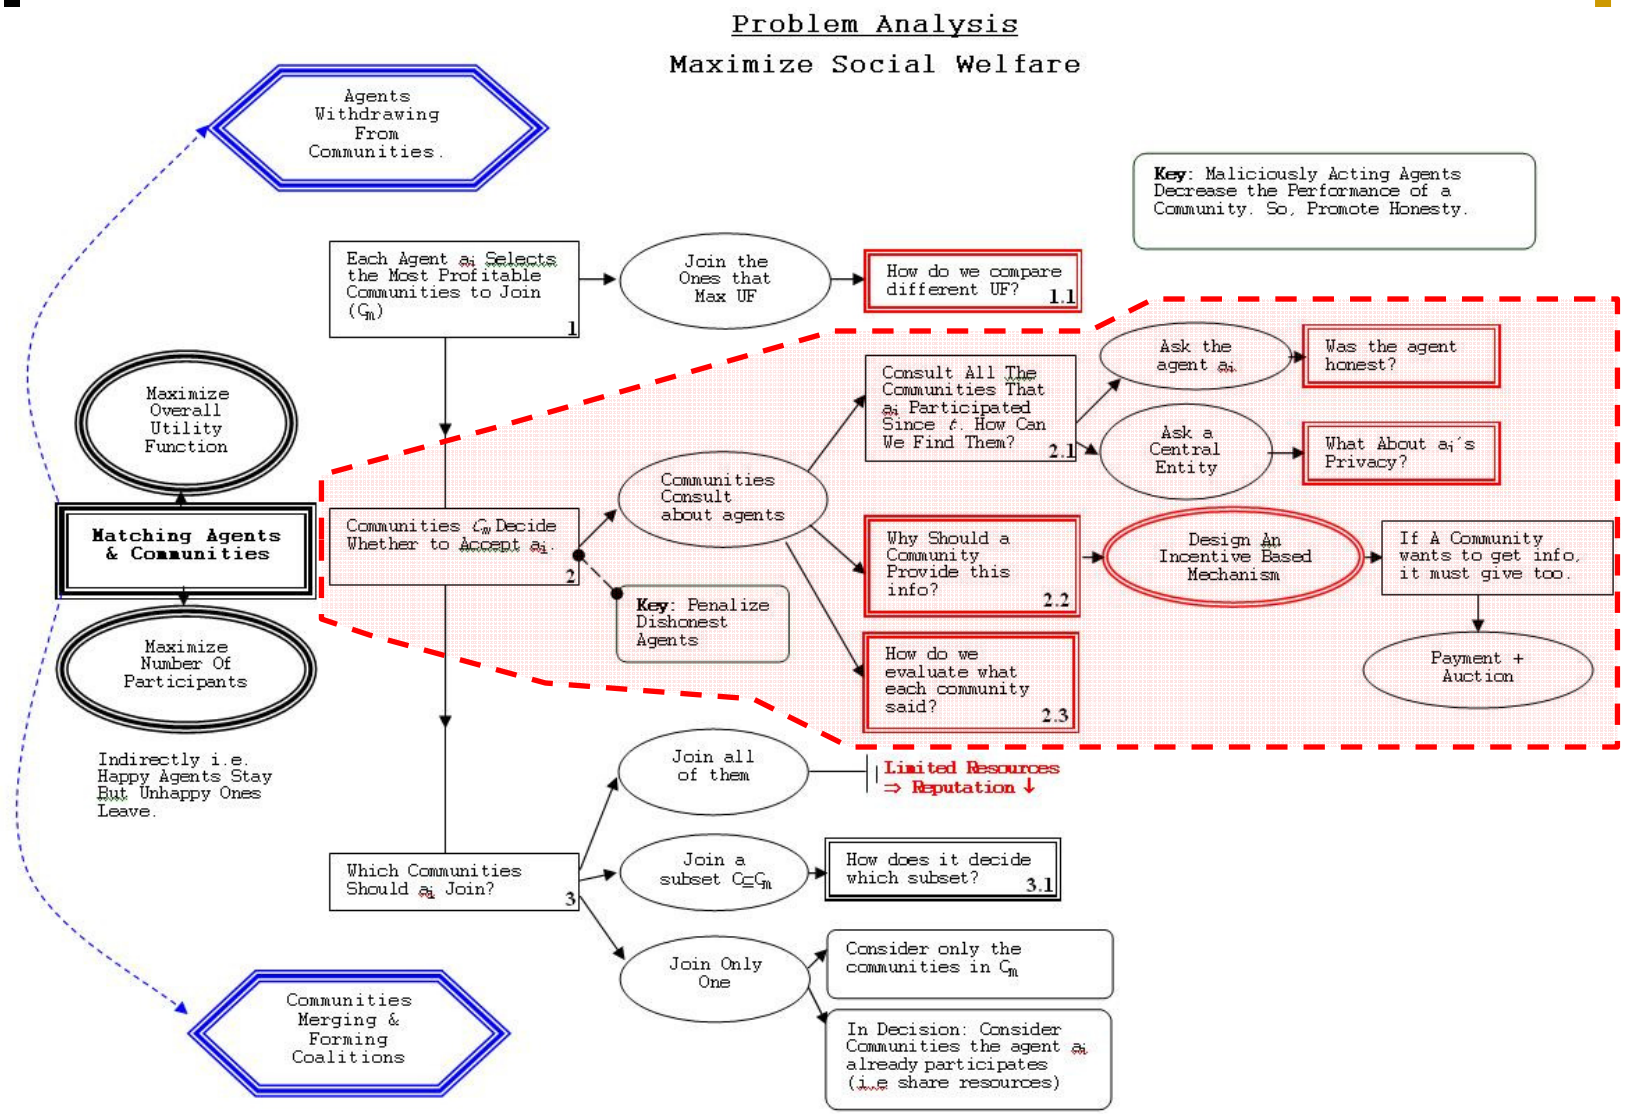

Overview

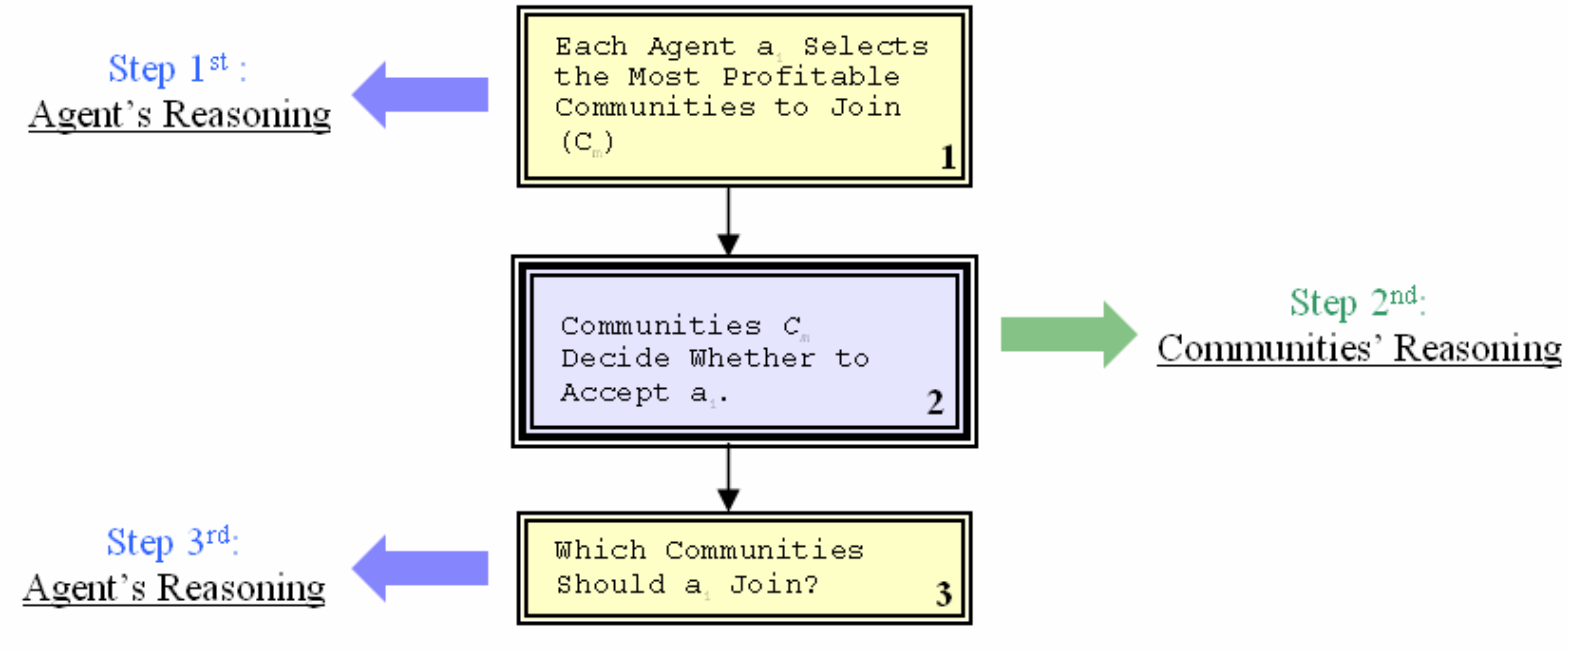
Step 1st :
Agent's Reasoning
Each Agent a_i Selects the Most Profitable Communities to Join (C_m)
1
Communities C_m Decide Whether to Accept a_i.
2
Step 2nd:
Communities' Reasoning
Step 3rd:
Agent's Reasoning
Which Communities Should a_i Join?
3

The Problem: Matching Communities & Agents

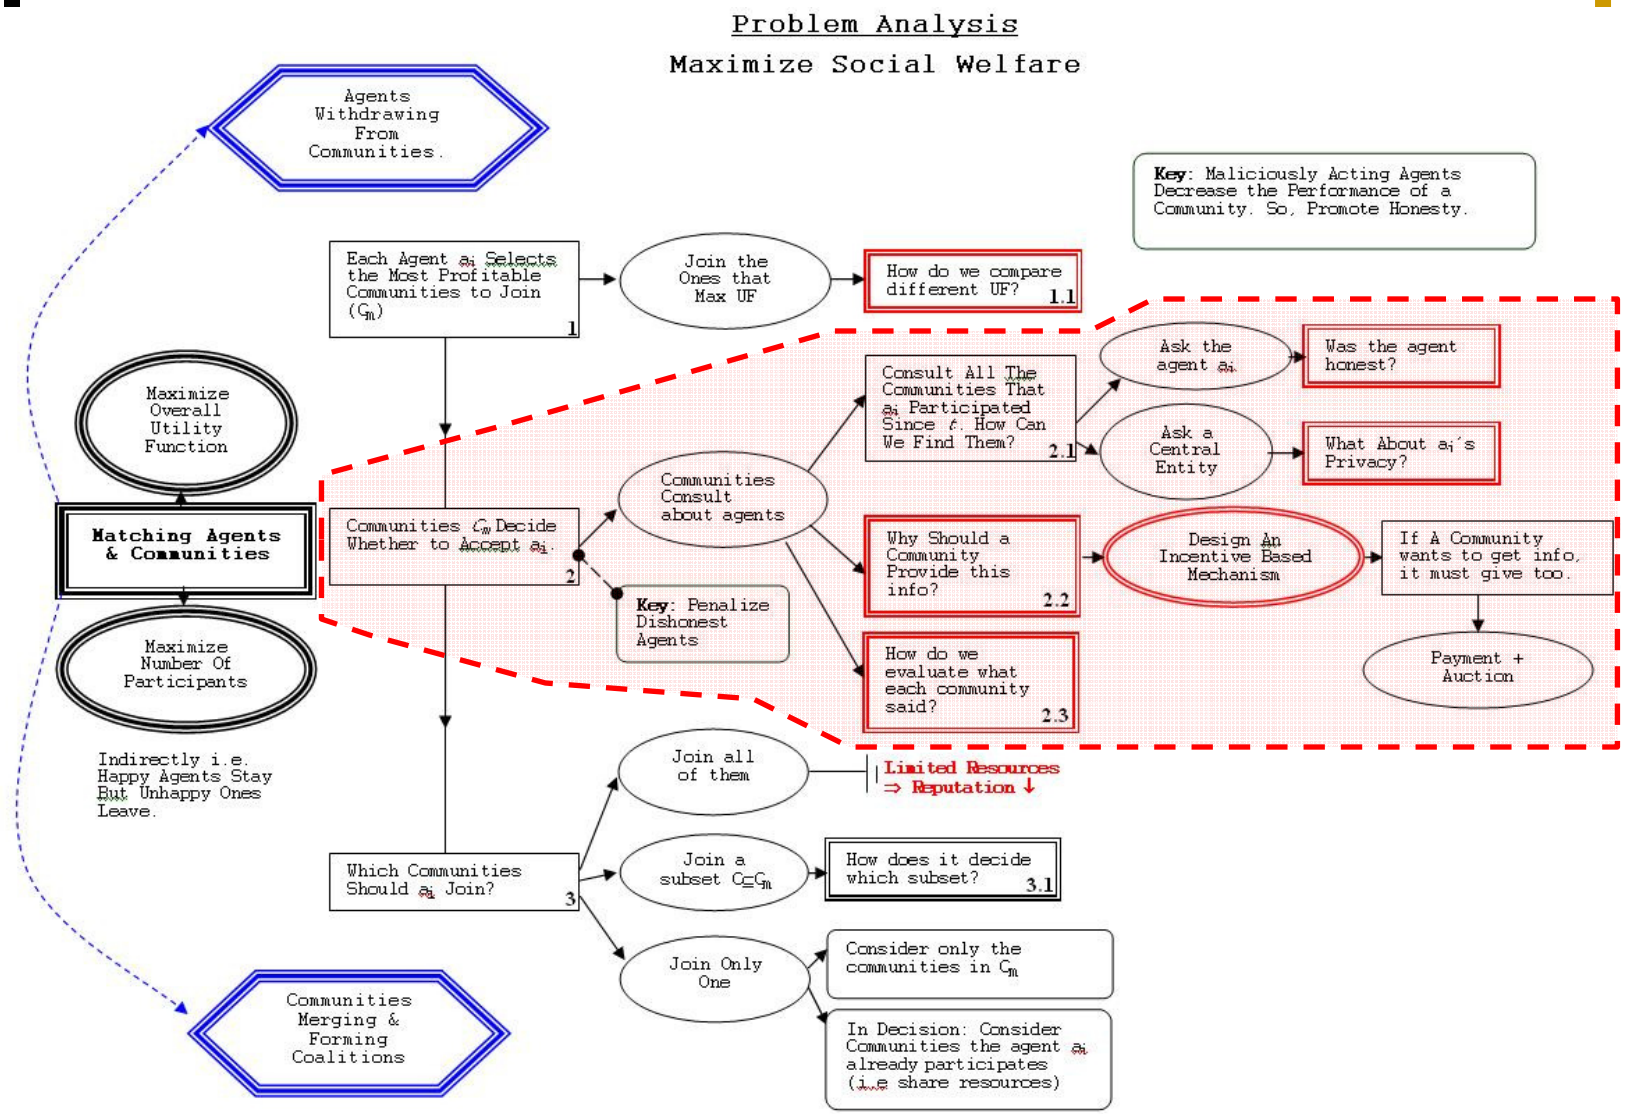

Communities Reasoning

- **Problem 2:** A community decides whether to accept an agent or not.

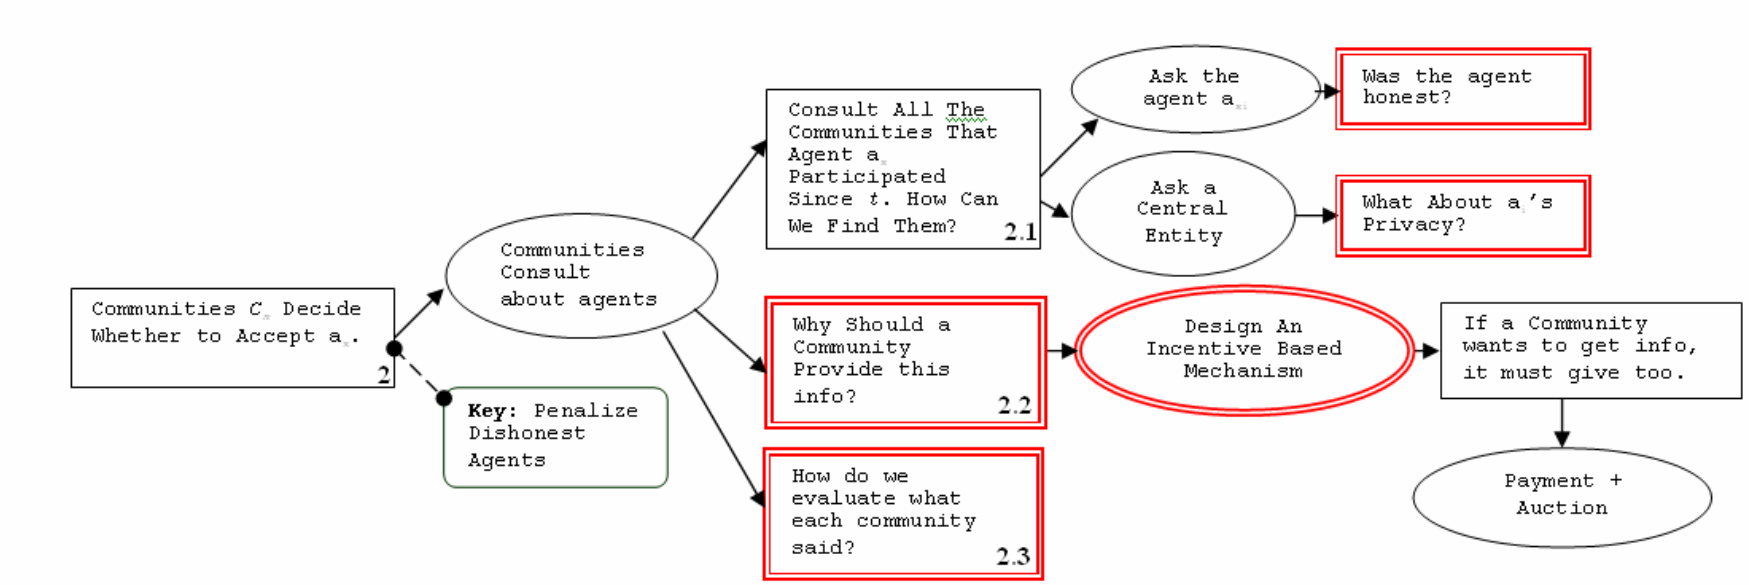

How Does a Community Decide Whether to Accept an Agent?

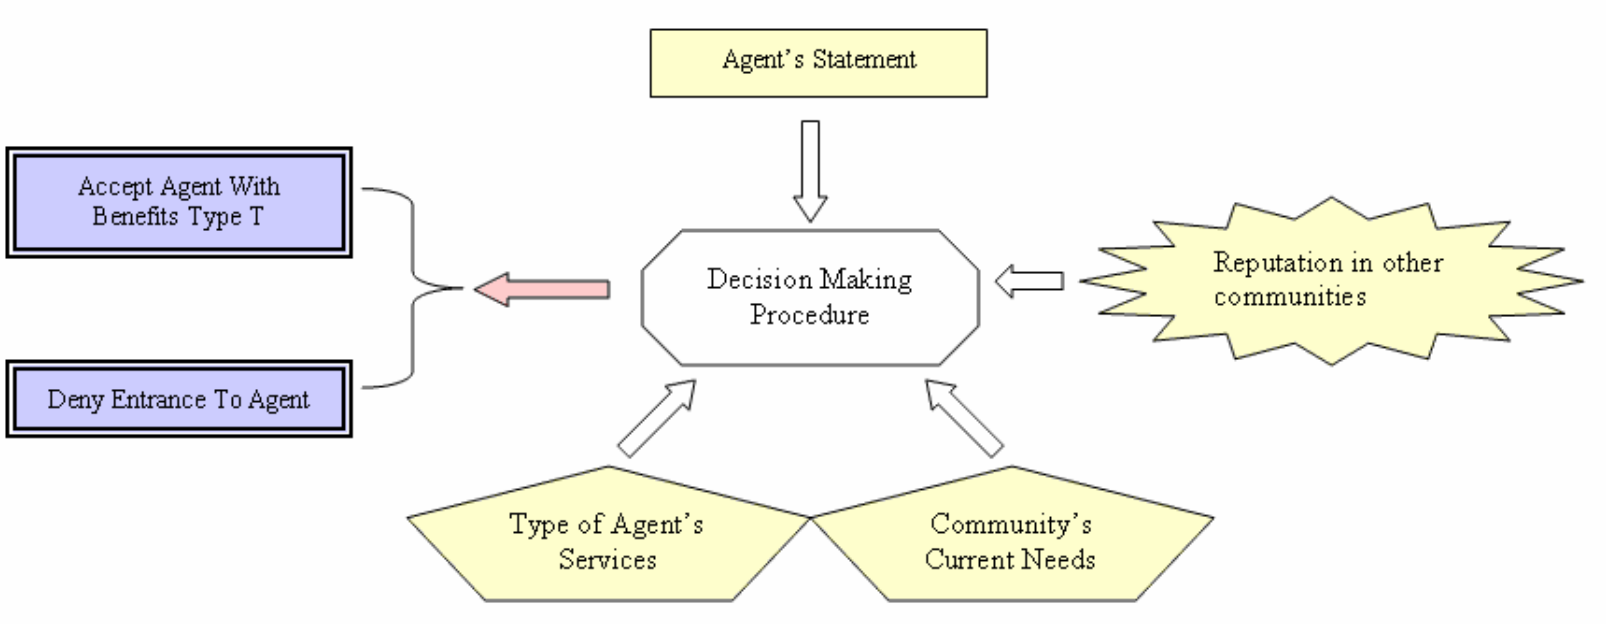

Modeling The Trustworthiness of an Agent

- Quality of contribution
  - Assessed by other agents in community
  - Local reputation: aggregation
    - Ratings received
    - How often rated
    - Importance of services
- Level of participation
  - Agent declares anticipated level: honesty
  - Reward honest, even if low level
  - Target agents to be monitored, even if high level

Communities Reasoning

- **Subproblem 2.1:** *How does the interested community know which communities to consult?*

- ***Solution:***

  agent reports its history

  privacy: control of agent

- **Challenges:**

  prevent lying

  detect lying

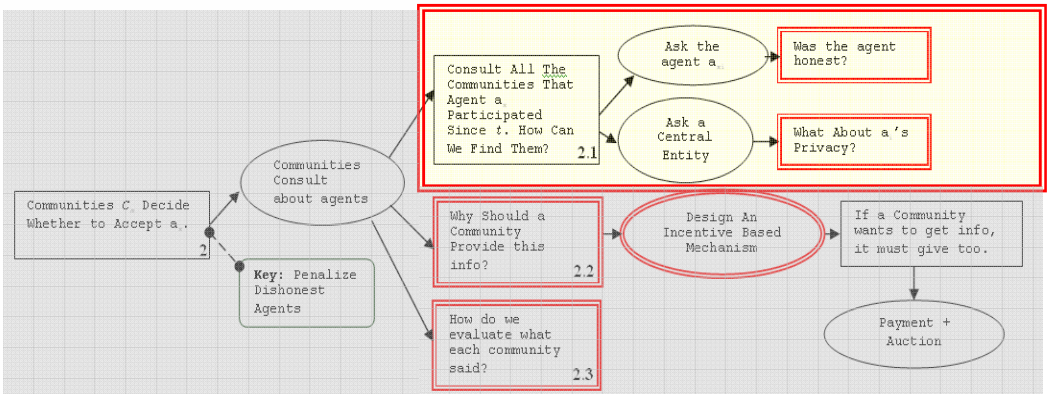

Communities Reasoning

- **Subproblem 2.2:** *Why should a community provide any information about its agents to any other community?*
- ***Solution:***
  - Provide incentive e.g. “want to receive info? then you must provide too”
- **Challenges:**
  - Design a Mechanism

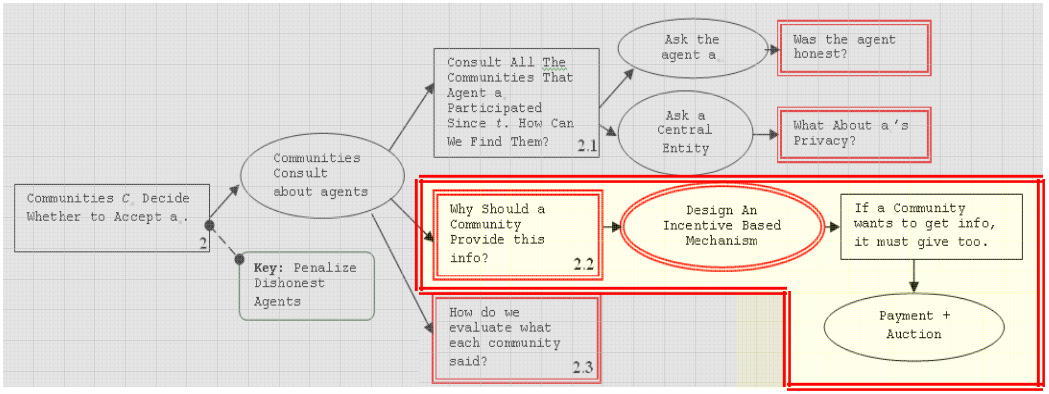

Communities Reasoning

- **Subproblem 2.3:** *How can a community evaluate what another community said regarding the reputation/trustworthiness of an agent?*
- **Challenges:**
  - Detect Untrustworthy Communities
  - Consider the different styles of reputation/trust modeling
  - Design a Metric

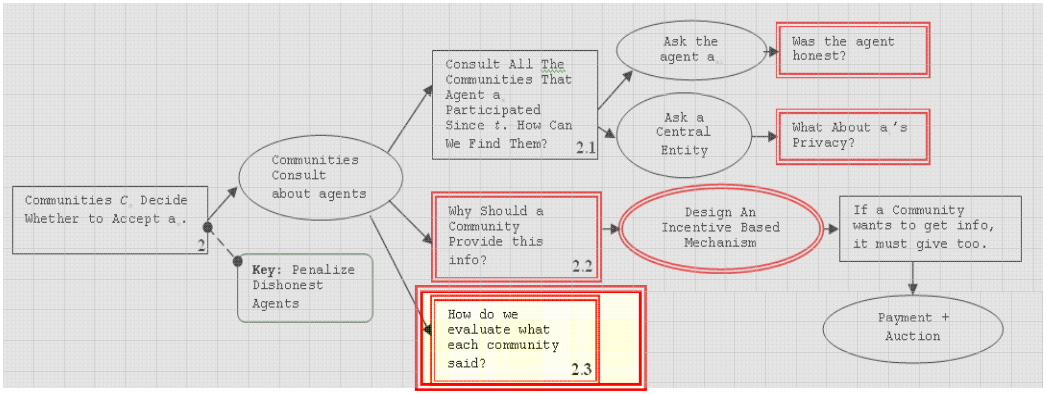

First Steps: Communities Reasoning Subproblem 2: Exchange of Information

- General Idea:
  - Lose “money” when you acquire information ⇒Decrease Utility.
  - Gain “money” when you provide information ⇒ Increase Utility.
  - Detecting/penalizing lying communities
  - The price should be strongly related to the reputation/trustworthiness of the agent.

First Steps: Communities Reasoning Subproblem 2: Exchange of Information

- First Approach: Each community has a utility function:

$$U_{c_j}^{ExRp} = U_{c_j}^{Gain} + U_{c_j}^{Comp} - U_{c_j}^{Cost} - U_{c_j}^{Pnlt}$$

Where:

$U_{c_j}^{Gain}$: represents the rewards the community c_j gained from providing advice about its agents. Currently:

$$U_{c_j}^{Gain} = \sum_{(c_k, a_i) \in PAd_j} rel(c_j, c_k) * rep_{c_j}^{a_i}$$

$U_{c_j}^{Cost}$: represents the cost the community c_j has to pay for requesting information regarding agents. Currently:

$$U_{c_j}^{Cost} = \sum_{(c_k, a_i) \in PAd_j} rel(c_j, c_k) * rep_{c_j}^{a_i}$$

First Steps: Communities Reasoning Subproblem 2: Exchange of Information

$$U_{c_j}^{ExRp} = U_{c_j}^{Gain} + U_{c_j}^{Comp} - U_{c_j}^{Cost} - U_{c_j}^{Pnlt}$$

$U_{c_j}^{Pnlt}$: represents the penalty the community cj has to pay for providing false information about the reputation of its agents. Currently:

$$U_{c_j}^{Pnlt} = \omega_{c_j} * \sum_{(c_k, a_i) \in PNAd_j} rel(c_j, c_k) * rep_{c_j}^{a_i}$$

$$\omega_{c_j} = \begin{cases} 0 & n_f(c_j) = 0 \\ \vartheta * \exp^{\frac{n_f(c_j)}{n_{total}(c_j)}} & n_f(c_j) \geq 1 \end{cases}$$

$U_{c_j}^{Comp}$: represents the compensation the community c_j gains when it proved that it received wrong information from another community.

This compensation is equal to the penalty that the community will have to pay for providing inaccurate information

First Steps: Communities Reasoning Subproblem 2: Exchange of Information

- Scenario: Agent a_i who participates in c_A is interested in the communities c_B, and c_C. It's reputation in c_A is 0.95

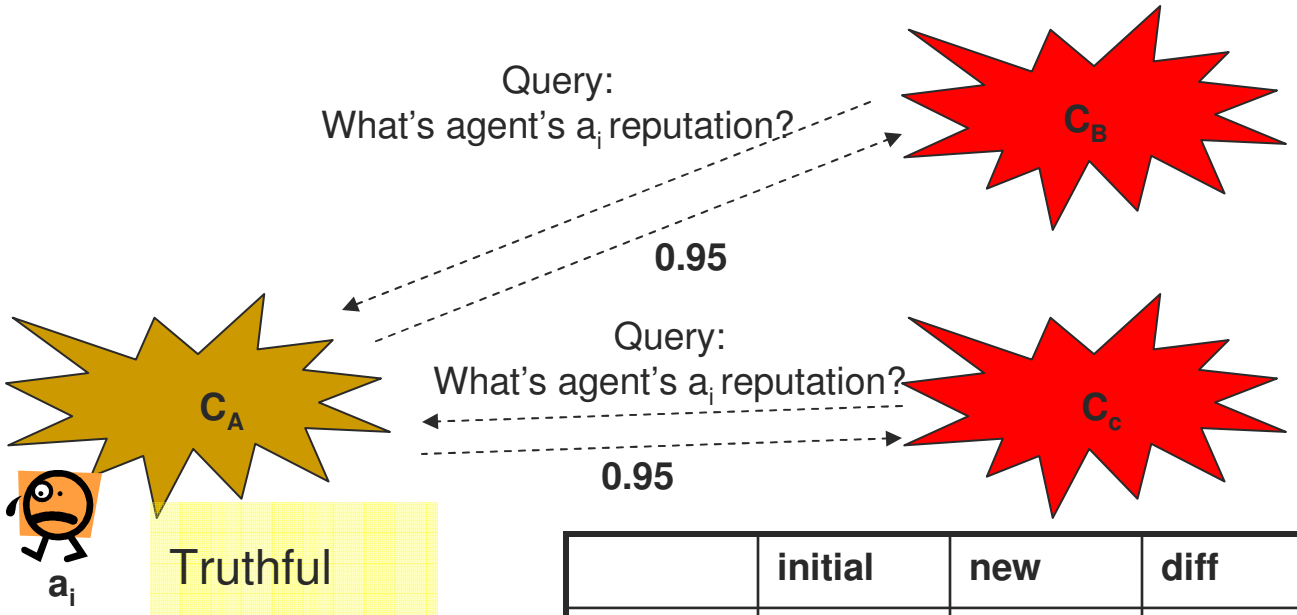

	initial	new	diff
$U_{c_A}^{ExRp}$	8.05	9.95	**+2 x 0.95**
$U_{c_B}^{ExRp}$	10.95	10.00	**-0.95**
$U_{c_c}^{ExRp}$	8.5	7.55	**-0.95**

First Steps: Communities Reasoning Subproblem 2: Exchange of Information

- Second Approach (*Prevent Communities From Lying*):
  - Auction The Reputation Of Each Agent.
  - Consider a Vickrey Auction with reserve prices
    - Winning bidder highest bid, price second highest bid
    - Incentive-compatible

Vickrey Auction Approach

- c_y: as both auctioneer and bidder
- Good=information about reputation of agent
- Price=reputation of agent
- Rules
  - Do not hide reputable agents
  - Avoid coalitions to acquire goods at low prices

Vickrey Auction Approach

- Second Approach: Each community has a utility function

$$U_{c_j}^{ExRp} = U_{c_j}^{Gain} - U_{c_j}^{Cost}$$

Where:

$U_{c_j}^{Gain}$: represents the rewards the community c_j gained from providing advice about its agents.

$U_{c_j}^{Cost}$: represents the cost the community c_j has to pay for requesting information regarding agents.

Vickrey Auction Approach

- Trusted Entity E
- c_j provides $[a_i, rep_{a_i}^{c_j}, cop_{a_i}^{c_j}]$ to E

The Auction:

- c_k submits b_{c_k}
- Highest bid? $b_{c_k} >= rep_{a_i}^{c_j}$
- Winner pays $payment = \max\{b_{c_p}, rep_{a_i}^{c_j}\}$

  where b_{c_p} is the 2nd highest bid
- c_k decreases utility by *payment*
- c_j increases utility by *payment*

First Steps: Communities Reasoning
Subproblem 3: Evaluating Communities

- Ask agent its communities
- Ask communities reputation/contribution of agent
- Penalize the untruthful party
- Subjective differences: Bayesian reasoning.

Agent's Reasoning

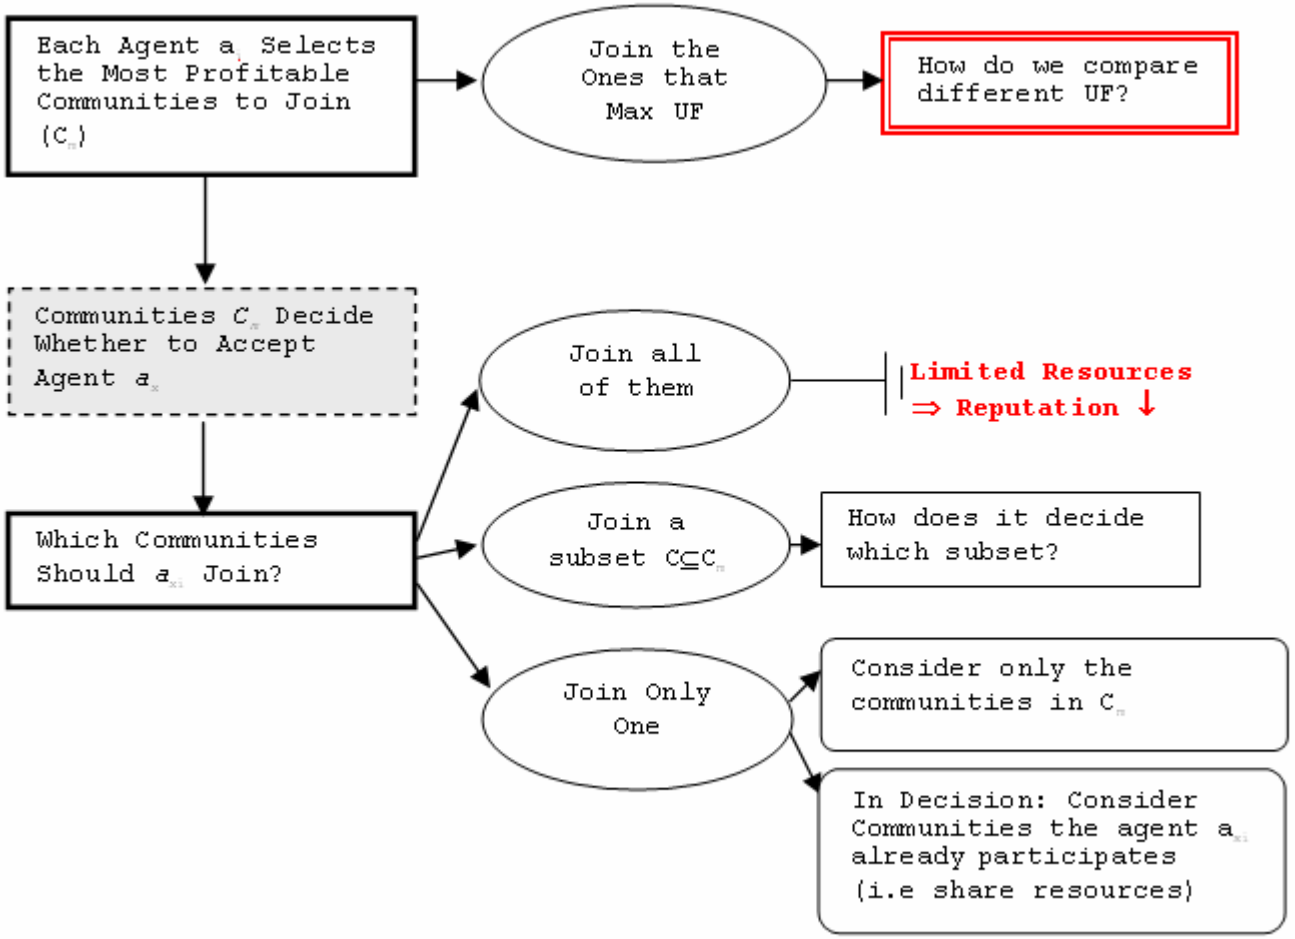

First Steps: Agents Reasoning

Evaluating Expected Utility through the average utility that the agents of the community have acquire:

$$gcop_{c_p}^{a_i} = \frac{E[U_{c_p}^{a_i}]}{AVG(U_{c_p})}$$

Example of a utility function in a P2P (Golle et al. 2001):

$$U_i = [f_i^{AD}(AD) + f_i^{NV}(NV) + f_i^{AL}(AL)] - [f_i^{DS}(DS) + f_i^{BW}(BW)] - FT$$

Privacy Consideration

- Allow agents to choose
  - privacy sensitive
  - privacy concerned
  - Privacy disinterested
- Avoid unintentional leak of information
  - Policies to determining access to information
- Communities sensitive to privacy

Discussion

- Effective community structure in multiagent systems promoted by sharing reputation ratings of agents.

- P2P networks: more than just utility-based, monetary incentive mechanisms

Discussion

We have two key aims:

- i) to exploit the information that each community accumulates by exchanging it with other communities
  - … indirectly, putting pressure on the agents by informing them that any misbehavior will be spread to the communities they might wish to join in the future, and
- ii) to prevent agents from overstretching themselves among many communities.

Future Work

- Optimize: communities with best agents and agents in best communities
  - High utility
- Re-examine assumptions
  - Agent reporting expected contribution
  - Average utility modeling
- Address challenges
  - Auction-overprice
  - Learning evaluation function of communities

Questions?

First Steps
Communities Reasoning

- Scenario: Agent a_i who participates in C_A is interested in the communities C_B, and C_C. It's reputation in C_A is 0.95.

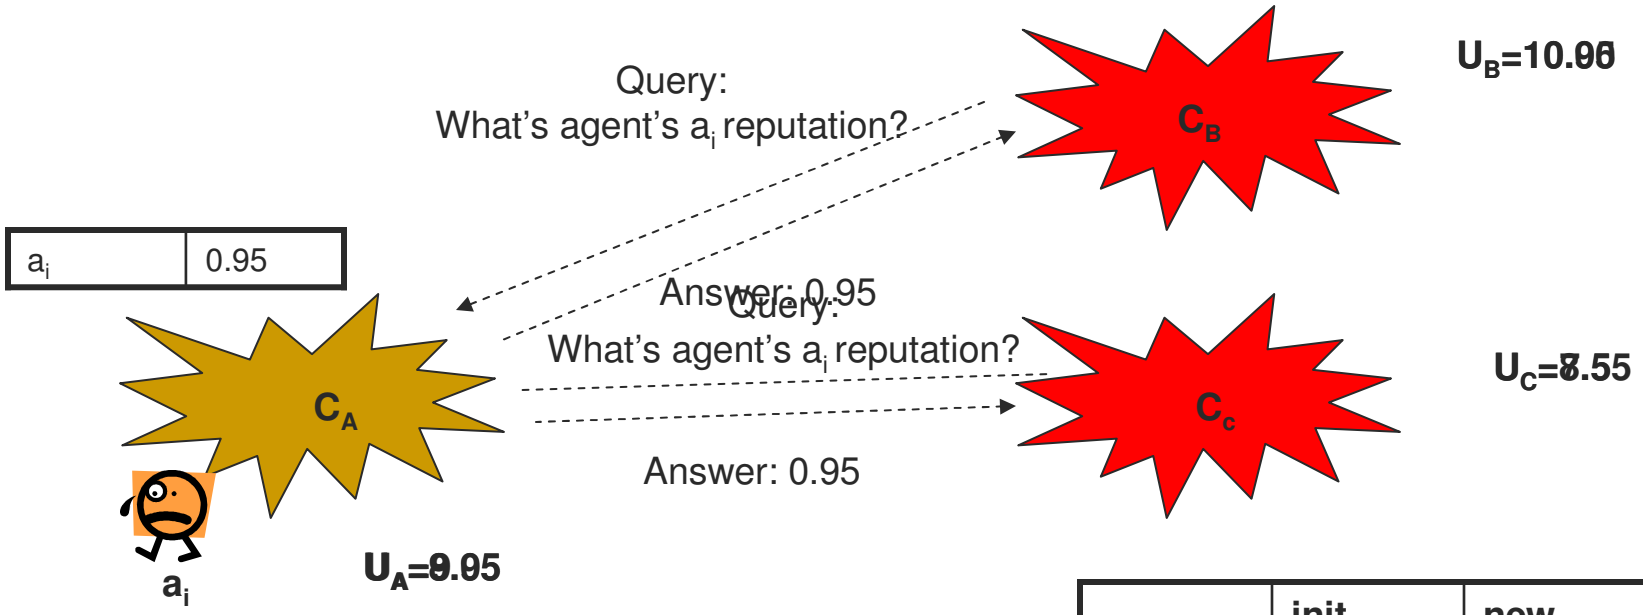

The agent a_i is a very reputable agent

	init	new	diff
U_A	8.05	9.95	+2 x 0.95
U_B	10.95	10.00	-0.95
U_C	8.5	7.55	-0.95

Background: Trust & Reputation

- Sophisticated approaches are required to prevent malicious users from manipulating a reputation/trust system.
- Consider the number of contributions the agent made to a community but also the value/quality of the items that he/she contributed.
- The beta probability function can be used as the basis for developing sophisticated trust and reputation systems.
- Reputation and trustworthiness of an agent in a system may depend on the special characteristics of the system.
- The uncertainty of the adequacy of the information an agent has regarding the reputation of the other agents can significantly influence the "correctness" of its decisions, while the accuracy of an agent's ratings is dependent on his/her personal style of evaluation.
- The degree of the suitability of a reputation/trust mechanism is related to the size of the environment in which it is applied.

Background Comparison

- Ebay: subjectivity of ratings, ballot box stuffing, doesn't consider the value of the product
- Amazon: ballot box stuffing, doesn't consider the type of the reading
- Thomas: Purchase of same product
- Beta/other: Do not consider
  - the frequency of the contribution
  - quality of the contribution
  - The degree the contribution benefits the community
    (e.g. provide new services)
  - they general style of the "rater".

Withdrawal for a Community

- Reasons an agent might wish to leave from a community:
  - its contribution to the community exceeds its gains.
  - it wishes to join another community instead with higher anticipated profit.
  - its reputation in the community *cj* results in decreasing its global reputation.
  - it has no more interest in the services that the community c_j offers.
  - it has decided that the community is malicious, thus its reputation might be harmed.

- An agent can:
  - reason about communities
  - join multiple communities simultaneously
  - participate in multiple communities at the same time
  - decide to withdraw from a community
- A community can:
  - reason about agents
  - communicate with other communities for exchanging information
  - merge with other communities

The Problem: Matching Communities & Agents

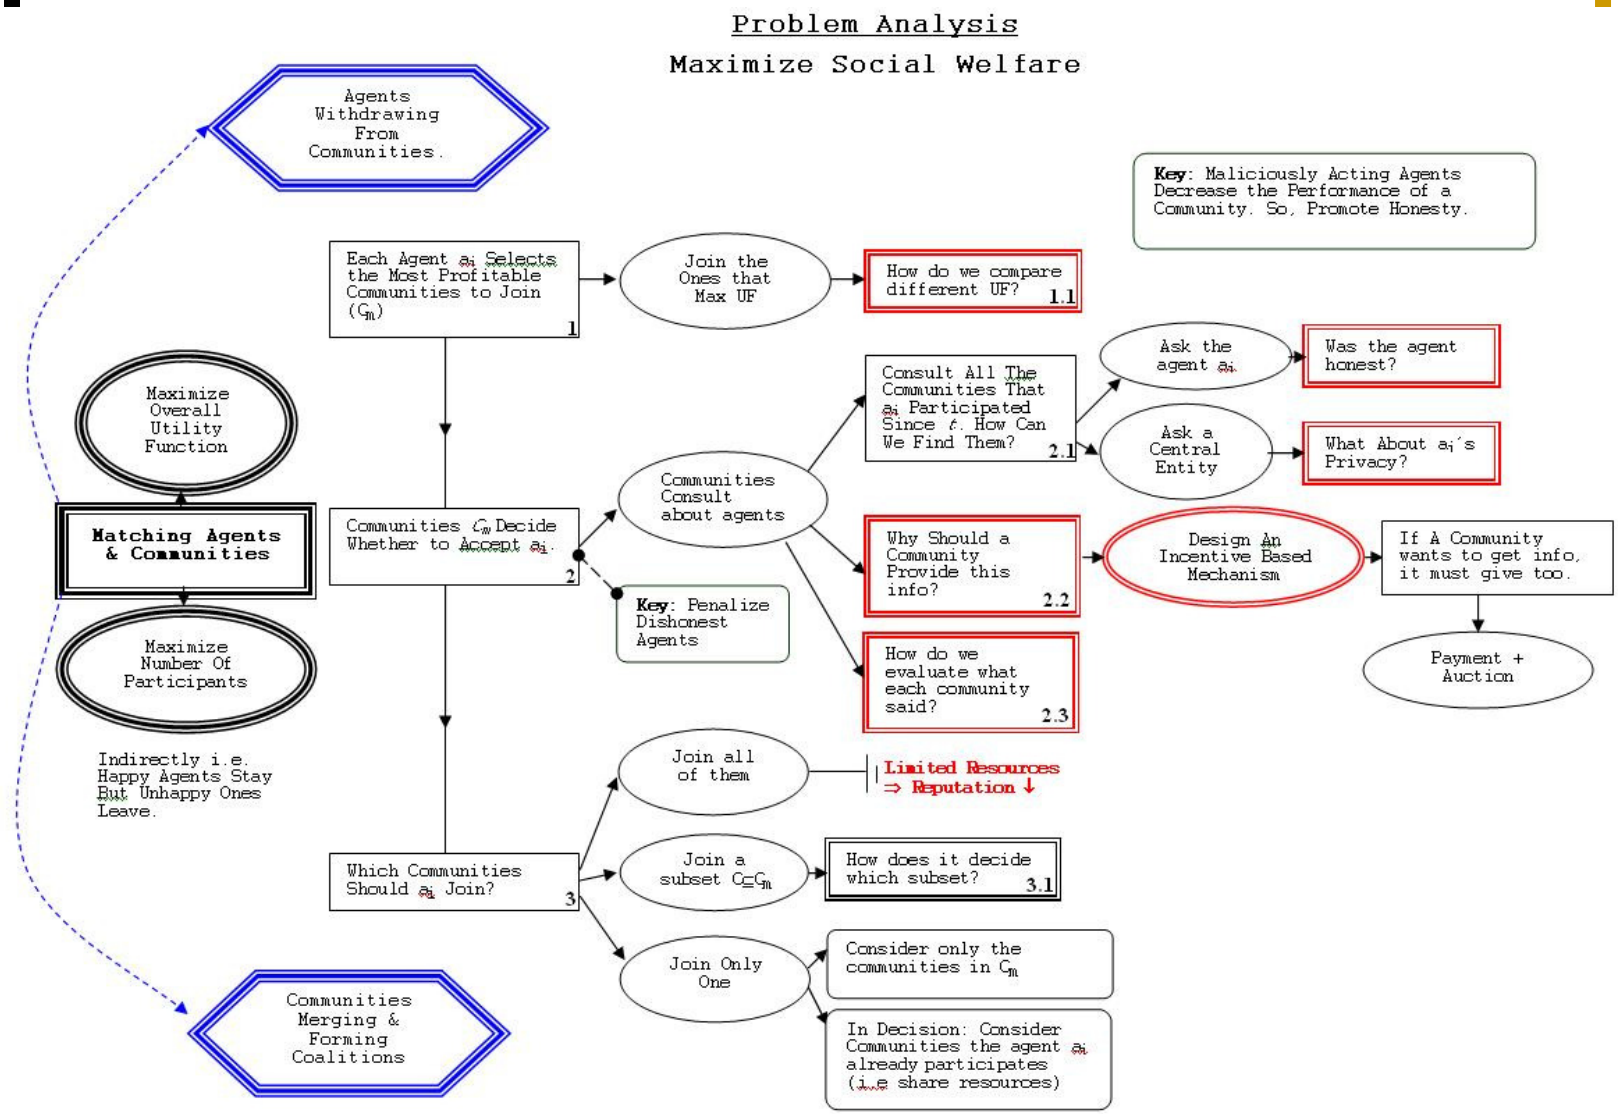

First Steps

- Assumptions:
  - Computation of Global Reputation
  - Computation of Average Agents' Utility
  - Each agent is able to report its expected contribution to a community
  - Privacy attitudes of agents
    - Rewards for sacrificing privacy