Maximum Likelihood Estimation of Peers' Trustworthiness in P2P Networks

Zoran Despotovic, Karl Aberer

EPFL-I&C

Distributed Information Systems Laboratory

{zoran.despotovic, karl.aberer}@epfl.ch

Overview

- Computational Models of Trust
  - Problem Definition
  - P2P Perspective
- Trust Models Classification
  - Classification Criteria
  - Classes of Approaches
- Maximum likelihood estimation
  - Model and Algorithm
  - Simulation Results
- Conclusions and Future Work

Computational Models of Trust - Problem

- Weighted trust (multi)graph
  made from interactions among the peers
- Example: weights in two contexts - trust and recommendation

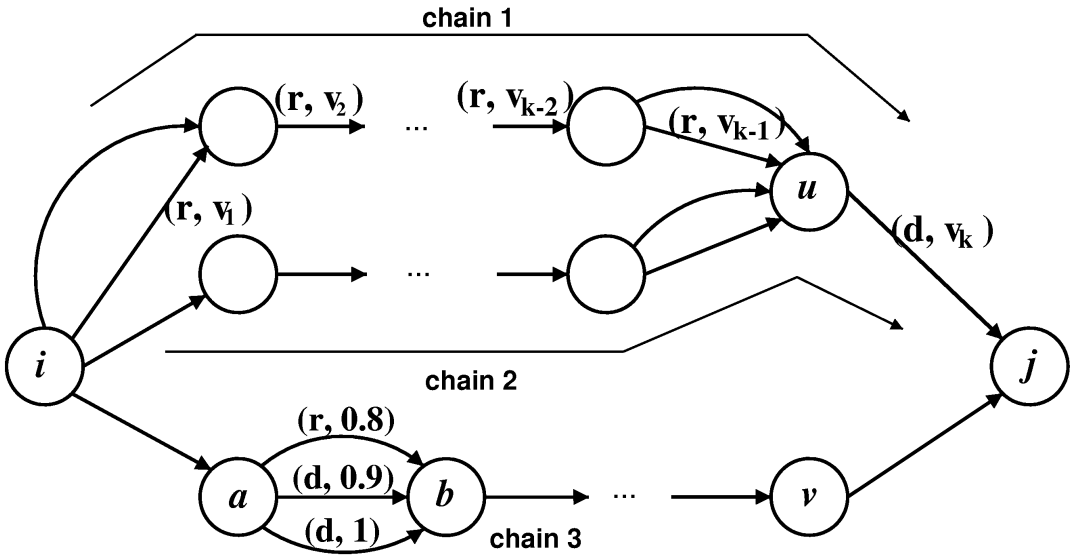

- What feedback to take and how to aggregate it
  to promote trust in the community?

P2P Perspective

- There is an underlying P2P overlay network
  - Structured
  - Unstructured
- Trust multigraph $\neq$ P2P overlay network
- Which network to use to get relevant data?
  - In principle the underlying P2P overlay

Trust Models Classification

- Classification dimensions:
  - Assumed behavior
  - Implementation overhead
  - Trust semantics
- Trust:
  - Peer A (trustor) trusts peer B (trustee) if
    1) some gain is to be shared with peer B
    2) peer A is exposed to a risk of loss
- Trust models:
  1) reducing the opportunism of the trustee
  2) reducing vulnerability of the trustor and then
  3) deciding if and when to enter an interaction

Three Classes of Approaches

- Social Networks:
  - Unclear trust semantics (meaning of the outputs)
  - Huge implementation overhead
  - Robust
- Probabilistic Estimation Techniques:
  - Clean interpretation of the values - probabilities
  - Low implementation overhead
  - Robustness?
- Game-Theoretic Models:
  - Rational behavior
  - Clean interpretation - trust is built into equilibria
  - P2P?

Maximum likelihood estimation

- Model:
  - Binary outcomes of the interactions
  - Innate probabilities of performing honestly and lying when reporting on others: (θ_k, l_k)
- Algorithm
  - Probability of observing report y_k from peer k on peer j:

$$P[Y_k = y_k] = \begin{cases} l_k(1-\theta_j) + (1-l_k)\theta_j & \text{if } y_k = 1 \\ l_k\theta_j + (1-l_k)(1-\theta_j) & \text{if } y_k = 0 \end{cases}$$

  - Collect all reports $y_1, y_2, \ldots, y_n$ on peer j and select θ_j that maximizes

$$L(\theta_j) = P[Y_1 = y_1]P[Y_2 = y_2] \cdots P[Y_n = y_n]$$

Properties of the Method

- Small implementation overhead:
  - Communication costs to retrieve reports of the witnesses
  - No storage costs or computation overhead
- Clean interpretation of the output values - probabilities
  - Unambiguous interpretation on the absolute $[0, 1]$ scale
  - Compute the utilities of various choices
- Collusions not considered

Estimation Quality – [0,1] case

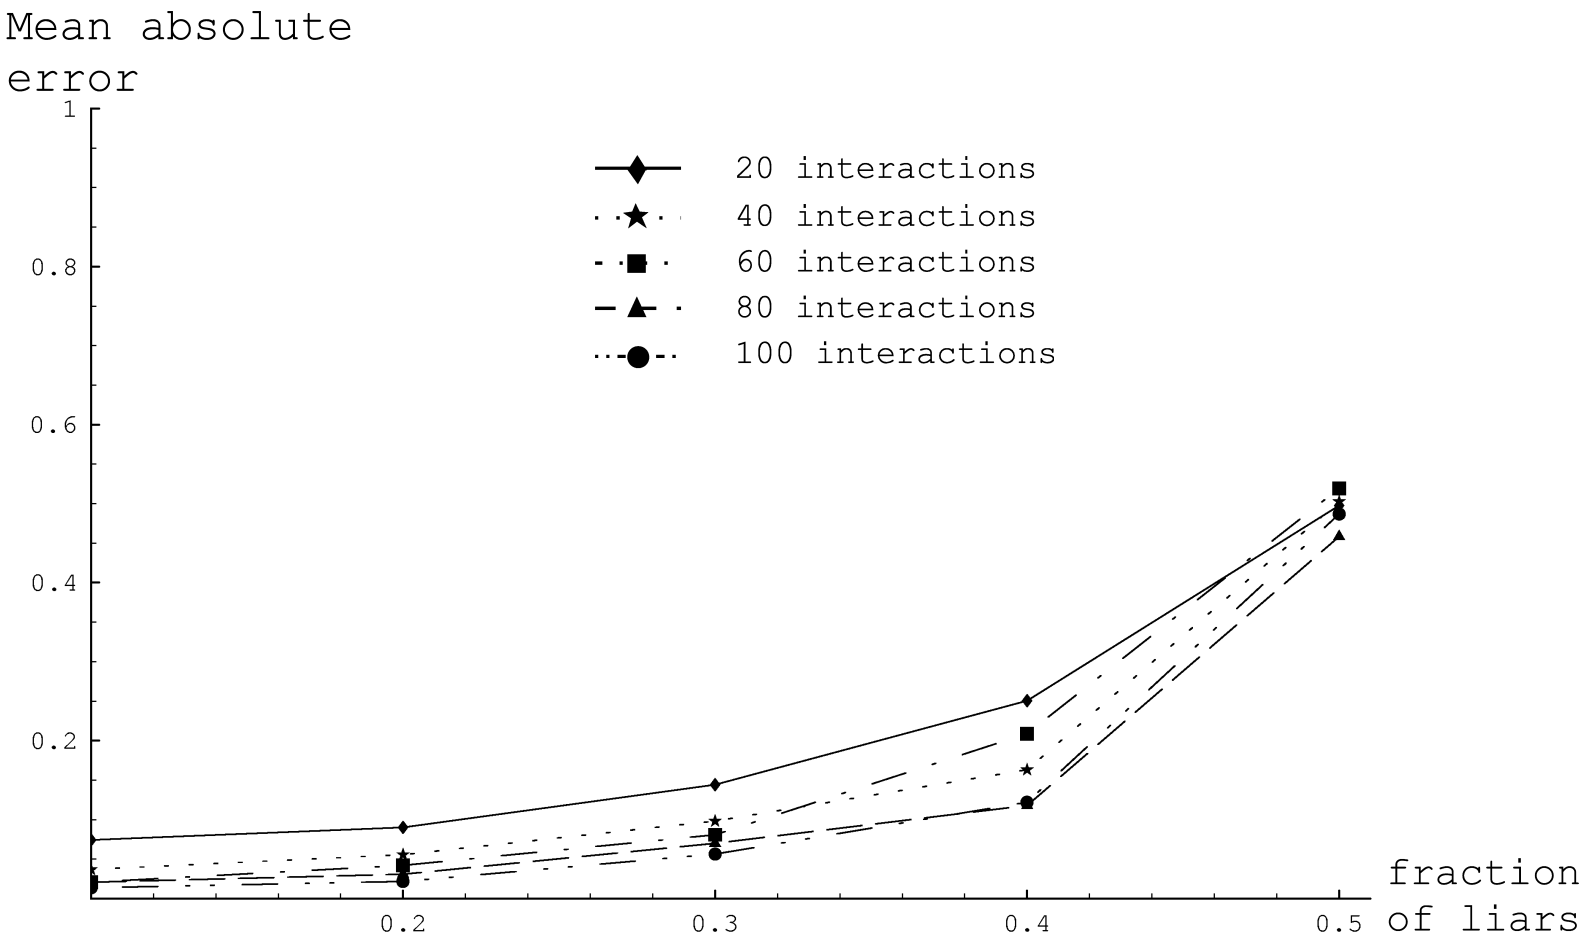

Estimation Quality – {0,1} case

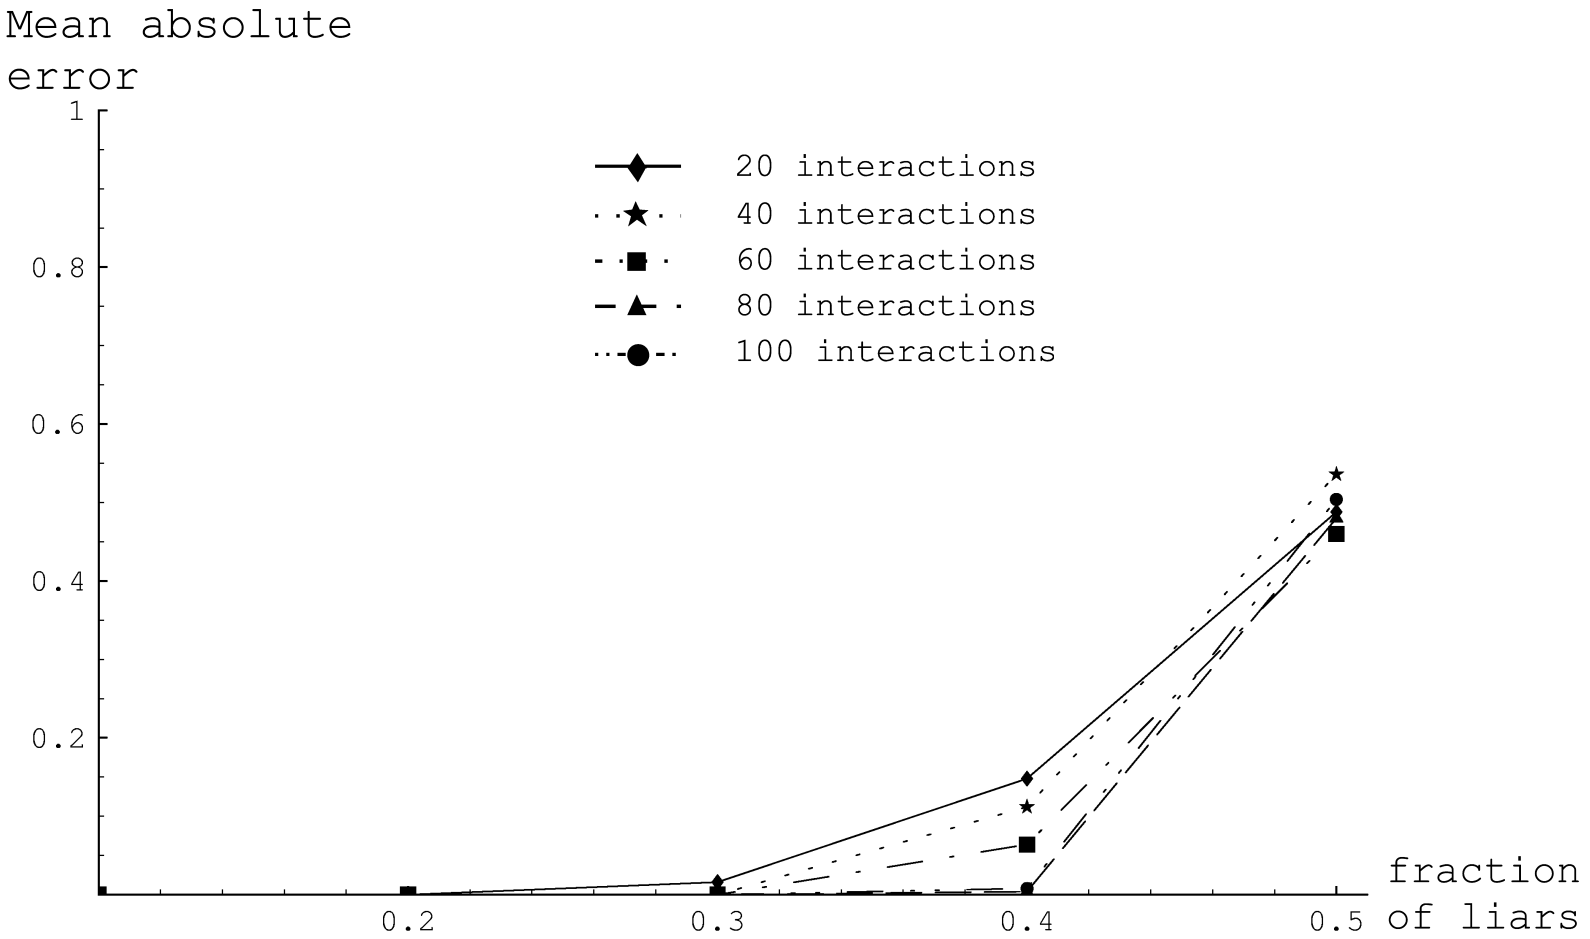

Conclusions

- Computational Models of Trust Classification
  - Two problems identified:
    - ▷ Implementation overhead
    - ▷ Trust semantics
- Solution based on Maximum Likelihood Estimation
- Future Work
  - Learn l_k for every peer separately
  - Modeling Collusions