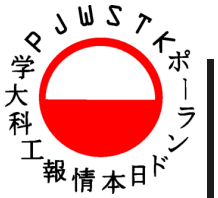

The uTrust project

- **uTrust: Universal Trust**
- Goal of the project:
  - Designing and implementing universal trust management algorithms and protocols
    - supported by formal analysis
    - computationally efficient
    - applied in a wide range of Open Distributed Systems
  - Universal adversary models and environments
  - A mathematical and computational model of trust management
    - Based on decision support under uncertainty
  - Output: a library of Trust Management services

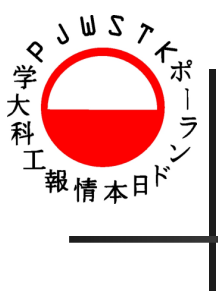

The uTrust project

- Project co-funded by Polish Ministry or Science
  - grant N N516 4307 33 ;]
  - the grant is open to participants from Poland
  - facilities: 3 quadcore servers, 2 in Planetlab
  - manpower: >=5 participants from PJIIT, >=5 outside participants
- Trust Management Library will be developed as an open-source project
  - … and distributed under public licence

Concepts, Methods, and Problems in Trust Management Research

Dr Adam Wierzbicki

PJIIT

Polish-Japanese Institute of Information Technology

Basic concepts of trust management

- Trust management (***TM***)
- Trust
- Reputation
- Uncertainty and risk
- Fairness and reciprocity

Trust management

- Most general definition: allow users (not necessarily human!) to make decisions in a situation of uncertainty
  - The uncertainty is created by actions of other users in an open distributed system
  - Usually, trust management just tells users what to expect about actions of others – the decision making aspect is not covered
- Can trust management work without trust?

Trust

- Two basic definitions:
  - An expectation (subjective, context dependent) about the behavior of another user
  - The tolerance (subjective, context dependent) to risk in an interaction with another user
- Difference: second definition more restrictive?
- Can trust be only created through knowledge about past actions (direct or indirect)?
  - Does escrow create trust? Or, does it reduce the need for trust?

Reputation

- Two definitions:
  - The perception, created by past actions, about a users intentions and norms
  - An estimate of the probability of a user's future behavior based on his past actions
- Second definition more computational
  - First definition usually requires an arbitrary scale of reputation
- Reputation is also context-dependent and subjective
- Reputation is influenced by evidence
  - Can reputation be a-priori?

Uncertainty and risk

- TM is only required under uncertainty; it aims to eliminate uncertainty
- Uncertainty: a user knows the possible outcomes, but cannot determine their probabilities
- Risk: a user knows the possible outcomes and can determine their (subjective, context-dependent) probabilities
  - Connection to reputation: reputation can be viewed as the probabilities

Fairness and reciprocity

- Agent fairness
  - Fulfilling a contract, obeying rules or social norms by individual agents
  - Often, reciprocity is considered as agent fairness
- System fairness
  - Providing strong incentives for fairness
  - In some contexts, fairness can be identified with distributional fairness

Basic problems of TM

- Establishing initial trust
- Trust propagation
  - Sometimes: formal reasoning about trust
- Reputation calculation
  - Centralized
  - Distributed
- Gathering evidence (observation)
  - Through direct interaction or as a third party
  - Expressing outcomes
  - Validating received evidence
- Expressing TM policies
  - Under what conditions can we decide to act?
  - How to express context?
- Resilience to attacks, correctness

Establishing initial trust

- When nothing is known about the agent, TM should assign him a default trust level
- Choosing too low default levels can result in excluding newcomers from the system
- Choosing too high default levels can result in abusing trust
- Is there another way?
  - Trust negotiation
  - Escrow

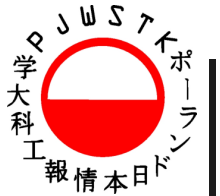

Trust propagation

- When agent A trusts agent B, and agent B trusts agent C, does A trust C?
  - Is trust transitive?
  - How to compute propagated trust?
- Trust delegation: when A trusts B and gives him his credentials/access rights, and C trusts A, can C trust B?
  - How to compute delegated trust?

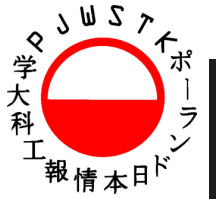

Reputation calculation

- A general problem: given a matrix of all available evidence about past behavior of agents, how to calculate the reputation of an agent?
  - ...in the eyes of another agent: reputation of A in the view of B?
  - Are reports by other agents the only source of information?
  - Can we have reports from objective (trusted) third parties?
- What if the calculation has to be distributed?
  - each agent has his own matrix
  - does the computation converge?
  - how does the TM system work with incomplete information?

Gathering evidence

- Different types of evidence:
  - All in context
  - reports: A reports to C about the behavior of B in a previous encounter with A
  - observations: A observes B in an encounter with C and reports this to D (can also be first-hand observation; C=D=A)
  - recommendations: A recommends B to C with a certain level of trust
- Different methods of gathering evidence:
  - all based on encounters of agents
  - yet, encounters can also model a centralized gathering of evidence
  - encounters can be associated with a transaction. If the transaction has no/little risk, then it can be made just for the purpose of gathering information. There can be different types of tansactions during 1 encounter

Expressing TM policies

- TM policies are used to evaluate trust
- Example: express the credentials needed to obtain a service
  - in other words: „to be considered sufficiently trusted to obtain a service”
- Complexity:
  - expressing context
  - expressing business rules
  - reasoning about policies – needed in trust negotiation
    - ex: are the presented credentials sufficient to satisfy a policy that required a different set of credentials? Perhaps by trust delegation?

Resilience to attacks

- What is the environment of TM operation?
  - What security services are available: authentication/privacy/integrity/...
  - What kind of attacks are possible: man-in-the-middle? DoS?
- What are the characteristics of adversaries?
  - Do we assume the worst case of adversary intelligence/malice/cooperation?
    - do adversaries trust each other?
  - Are the adversaries resource-constrained?
  - Proportion of computing power of adversaries/computing power of normal agent
- All these questions have no general answer
- Result: little is known about TM resilience

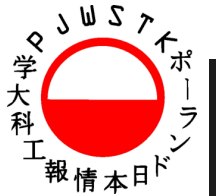

New research problems of TM

- Effectiveness of Trust Management
  - Does TM reduce risk?
  - Does TM increase the total number of transactions, or the total revenue?
- Resistence to adversaries
  - No benchmarks
- Scalability and efficiency
  - Nothing is known

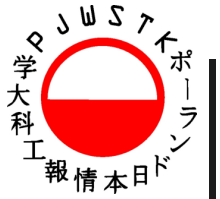

Universal Trust Management

- TM can be viewed as a *service*
  - Like other security services: authentication, privacy, integrality, authorization
- A service needs to be universal (general)
  - On the other hand, it cannot be abstract
  - And should be configurable (and scalable)
- Purpose of uTrust: build a library that offers a Universal TM service

The need for new models

- Contemporary TM models:
  - Do not model all relevant aspects of TM
    - Few models explicitly consider risk and uncertainty
    - Few models consider the decision making process, and model utilities and incentives of users
  - Do not allow for a universal evaluation of TM methods and systems
    - No generally agreed criteria
    - No general adversary models and environments
  - Do not explain all different TM methods
    - Various kinds of evidence
    - TM that does not use reputation

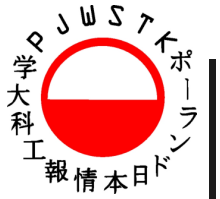

Components of a TM model

- User: trustor, trustee, observer(?)
- Context (social network?)
- Encounter: action, outcome
  - Outcomes:
    - price, profit, compensation;
    - own reputation;
    - own security;
    - quality of service, performance
- Evidence: security? (authentication, integrity; verifiability?)

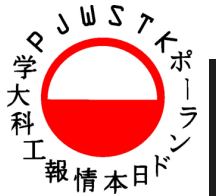

An example of new criteria

- Consider trust management in auction systems
  - Use simplistic reputation mechanisms
  - No risk calculation
    - Would a user trust the risk calculation by the auction provider?
  - Auction provider criteria: *increase revenue*!
    - Equivalent to increasing number of auctions
    - Contrary to criteria of reducing user risk!
- Consider a new criterion: *reducing risks* of users
  - Result: a trust management system for e-auctions ***must be independent*** of the auction provider
    - P2P?

TM system model

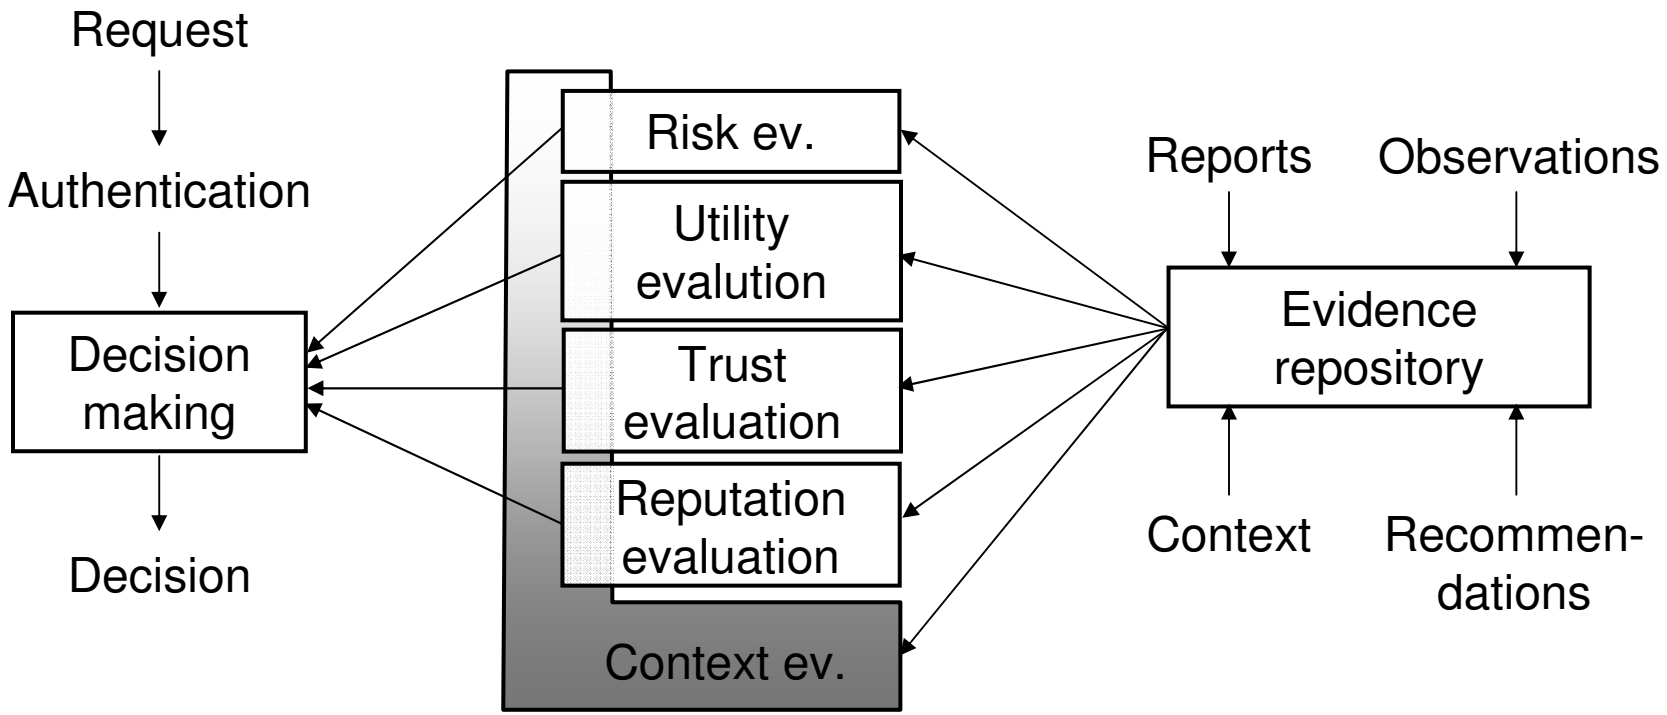

- General (?) model of TM system
- Context can be a very complex thing
  - social network, semantic information, information about encounter or action context, ...

Building a universal TM library

- The library should include all functions needed to build a TM system
- The library should be capable to work in a distributed or centralized manner
  - ...or hybrid
- The library should be used and tested in various application domains
  - auctions
  - P2P systems
  - ad-hoc networks
  - Web Services workflow systems, grids