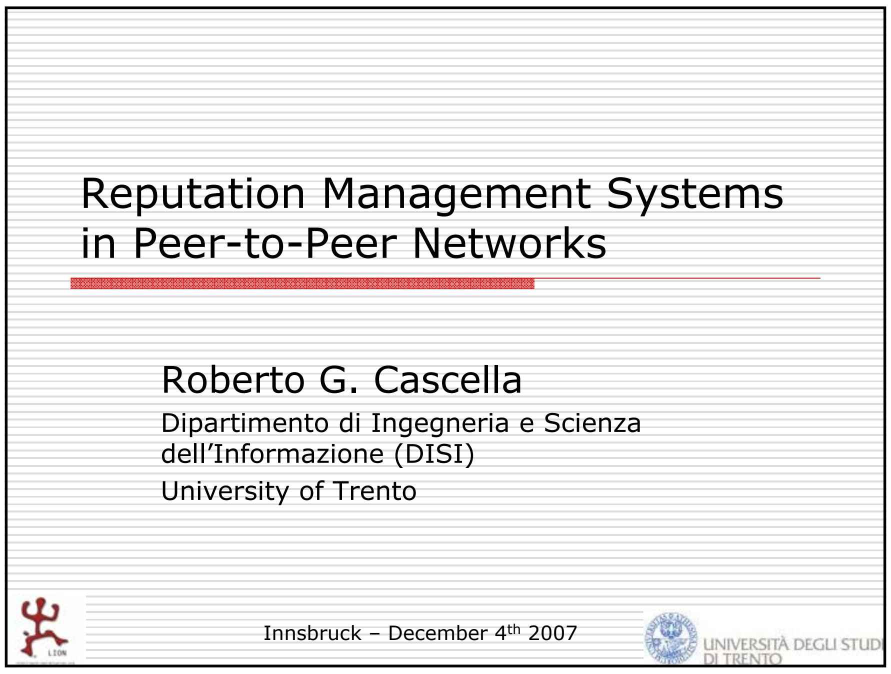

Reputation Management Systems in Peer-to-Peer Networks

Roberto G. Cascella

Dipartimento di Ingegneria e Scienza dell'Informazione (DISI)

University of Trento

Innsbruck – December 4th 2007

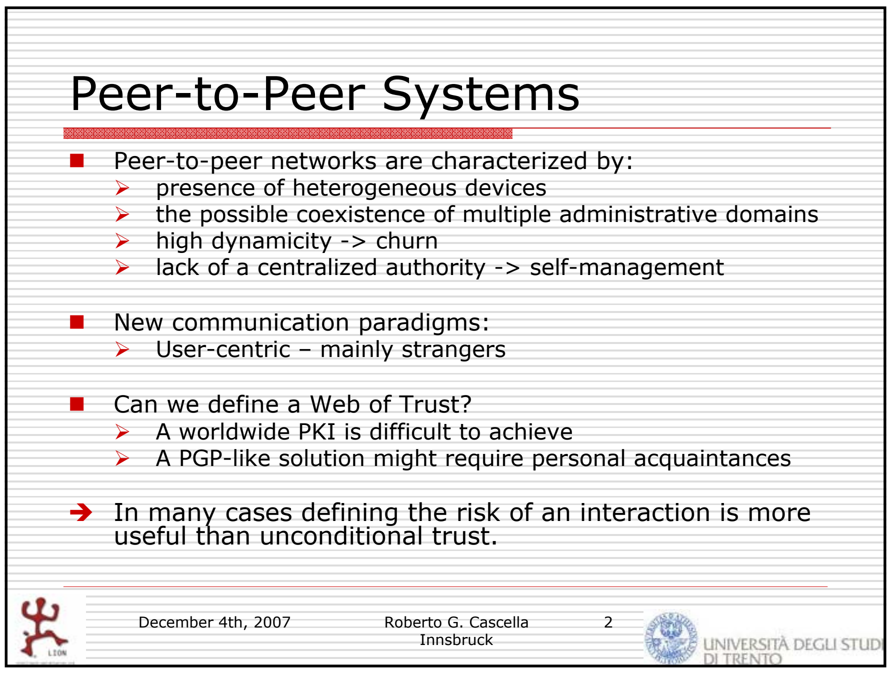

Peer-to-Peer Systems

- Peer-to-peer networks are characterized by:
  - presence of heterogeneous devices
  - the possible coexistence of multiple administrative domains
  - high dynamicity -> churn
  - lack of a centralized authority -> self-management
- New communication paradigms:
  - User-centric – mainly strangers
- Can we define a Web of Trust?
  - A worldwide PKI is difficult to achieve
  - A PGP-like solution might require personal acquaintances

➔ In many cases defining the risk of an interaction is more useful than unconditional trust.

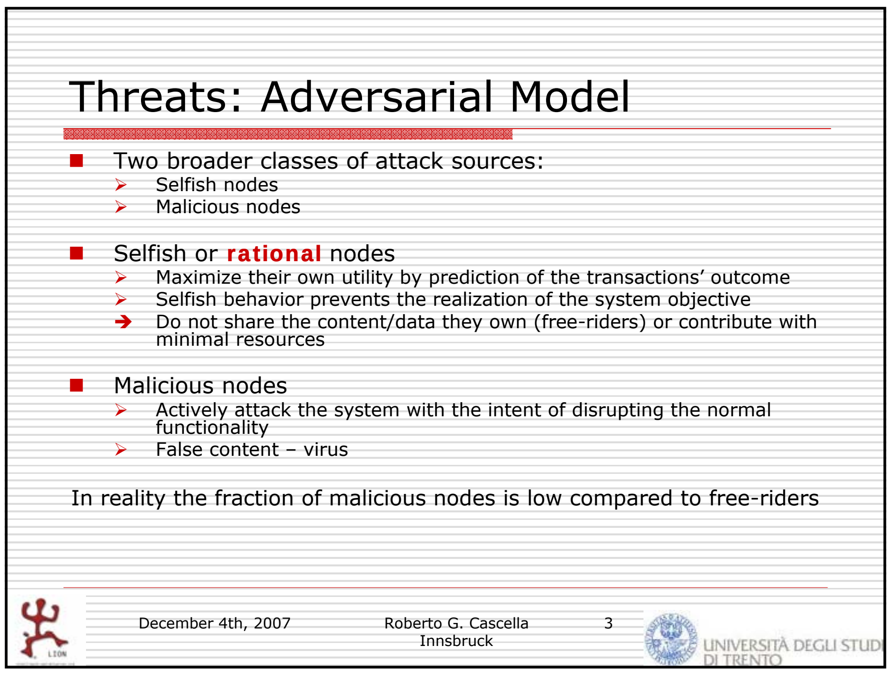

Threats: Adversarial Model

- Two broader classes of attack sources:
  - Selfish nodes
  - Malicious nodes
- Selfish or **rational** nodes
  - Maximize their own utility by prediction of the transactions' outcome
  - Selfish behavior prevents the realization of the system objective
  - ➔ Do not share the content/data they own (free-riders) or contribute with minimal resources
- Malicious nodes
  - Actively attack the system with the intent of disrupting the normal functionality
  - False content – virus

In reality the fraction of malicious nodes is low compared to free-riders

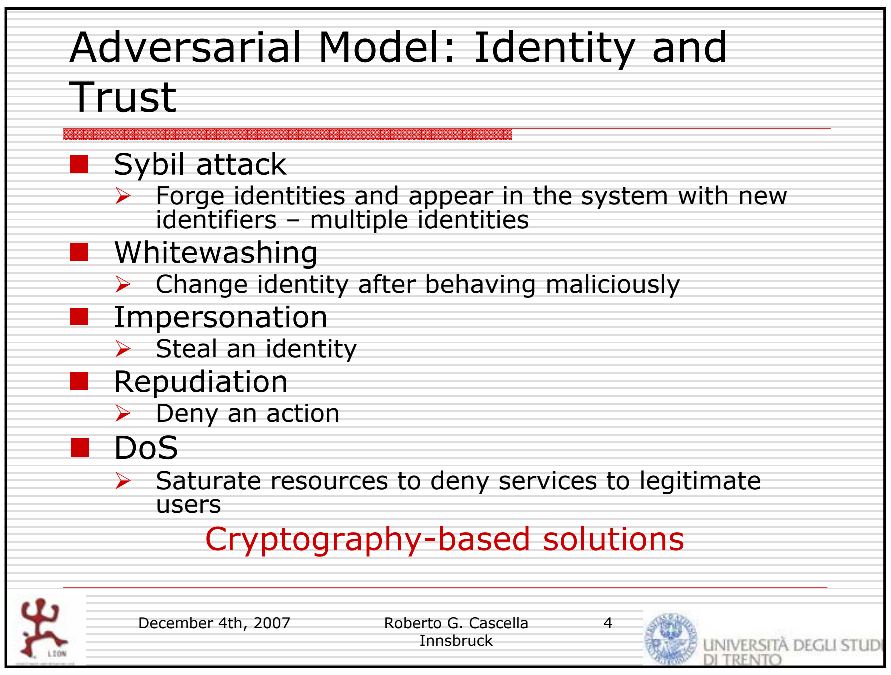

Adversarial Model: Identity and Trust

- Sybil attack
  - Forge identities and appear in the system with new identifiers – multiple identities
- Whitewashing
  - Change identity after behaving maliciously
- Impersonation
  - Steal an identity
- Repudiation
  - Deny an action
- DoS
  - Saturate resources to deny services to legitimate users

Cryptography-based solutions

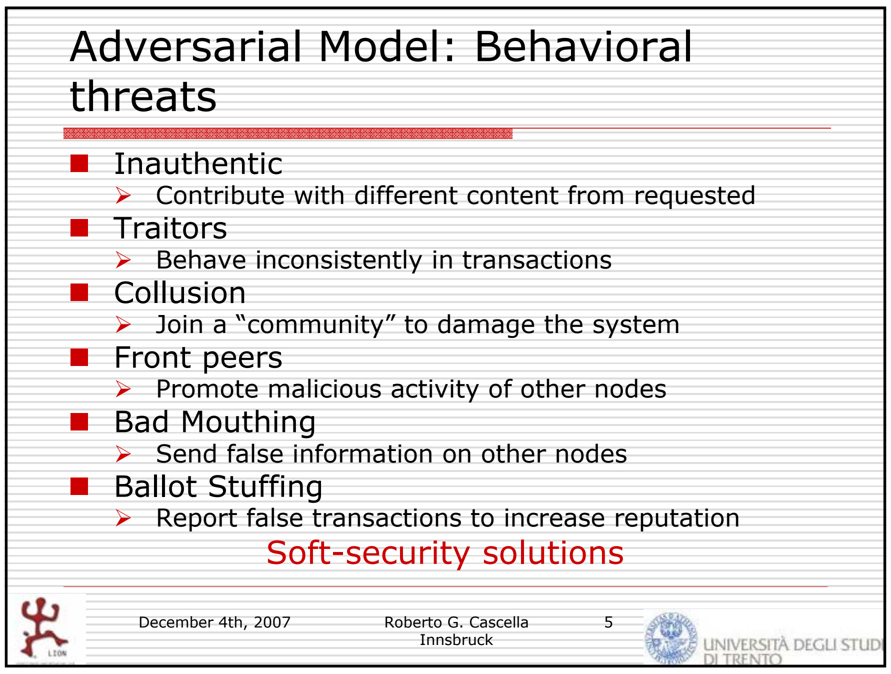

Adversarial Model: Behavioral threats

- Inauthentic
  - Contribute with different content from requested
- Traitors
  - Behave inconsistently in transactions
- Collusion
  - Join a "community" to damage the system
- Front peers
  - Promote malicious activity of other nodes
- Bad Mouthing
  - Send false information on other nodes
- Ballot Stuffing
  - Report false transactions to increase reputation

Soft-security solutions

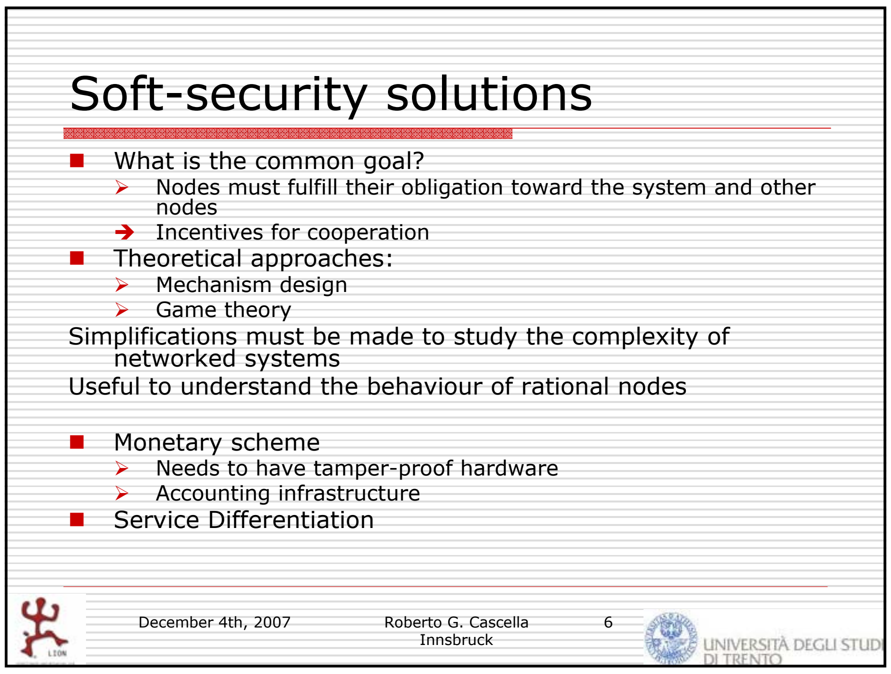

Soft-security solutions

- What is the common goal?
  - Nodes must fulfill their obligation toward the system and other nodes
  - ➔ Incentives for cooperation
- Theoretical approaches:
  - Mechanism design
  - Game theory

Simplifications must be made to study the complexity of networked systems

Useful to understand the behaviour of rational nodes

- Monetary scheme
  - Needs to have tamper-proof hardware
  - Accounting infrastructure
- Service Differentiation

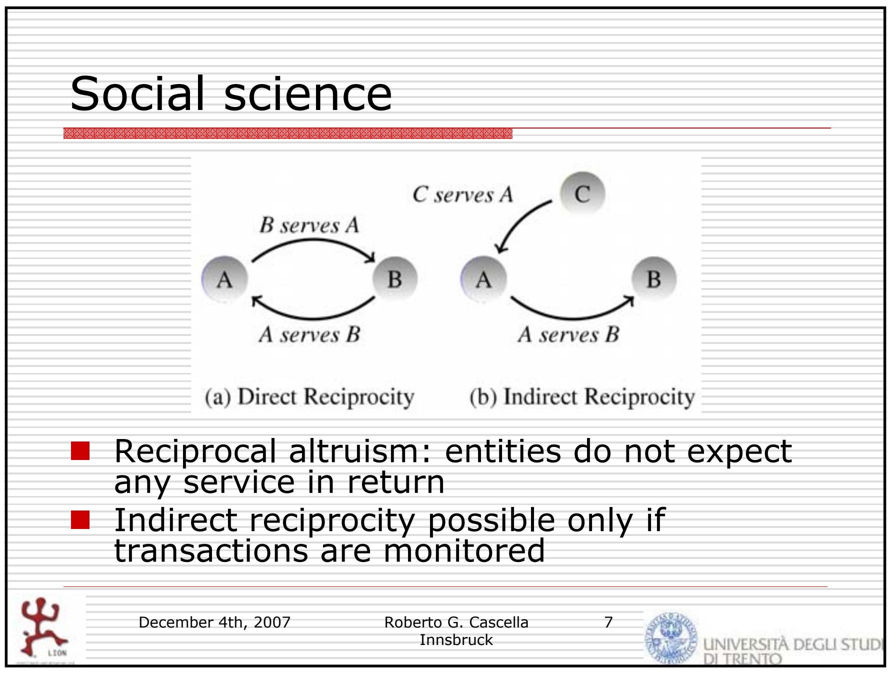

Social science

B serves A
A serves B
C serves A
C
A
B
A serves B

(a) Direct Reciprocity (b) Indirect Reciprocity

- Reciprocal altruism: entities do not expect any service in return
- Indirect reciprocity possible only if transactions are monitored

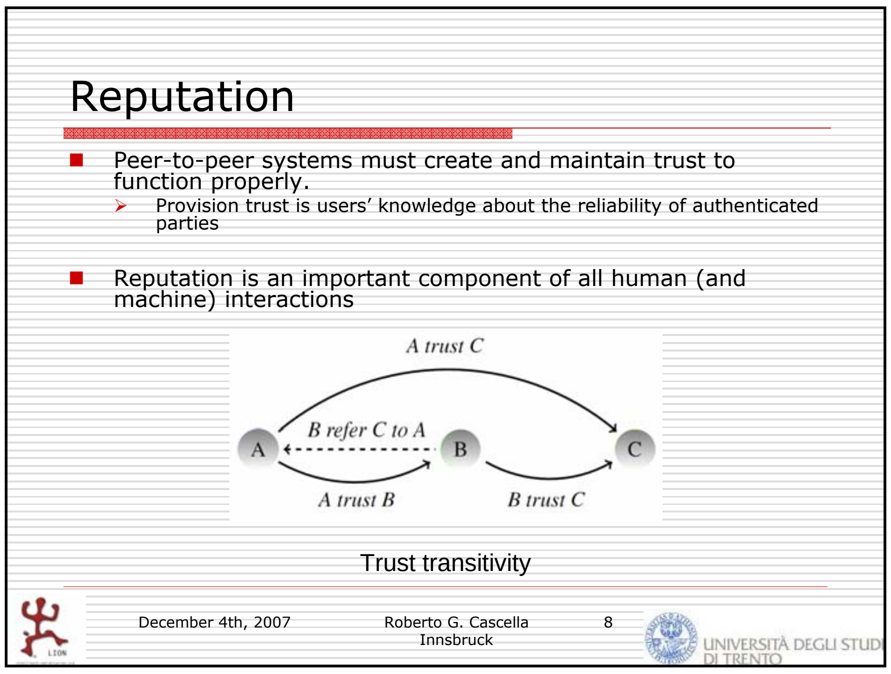

Reputation

- Peer-to-peer systems must create and maintain trust to function properly.
  - Provision trust is users' knowledge about the reliability of authenticated parties
- Reputation is an important component of all human (and machine) interactions

A trust C
B refer C to A
A
B
C
A trust B
B trust C

Trust transitivity

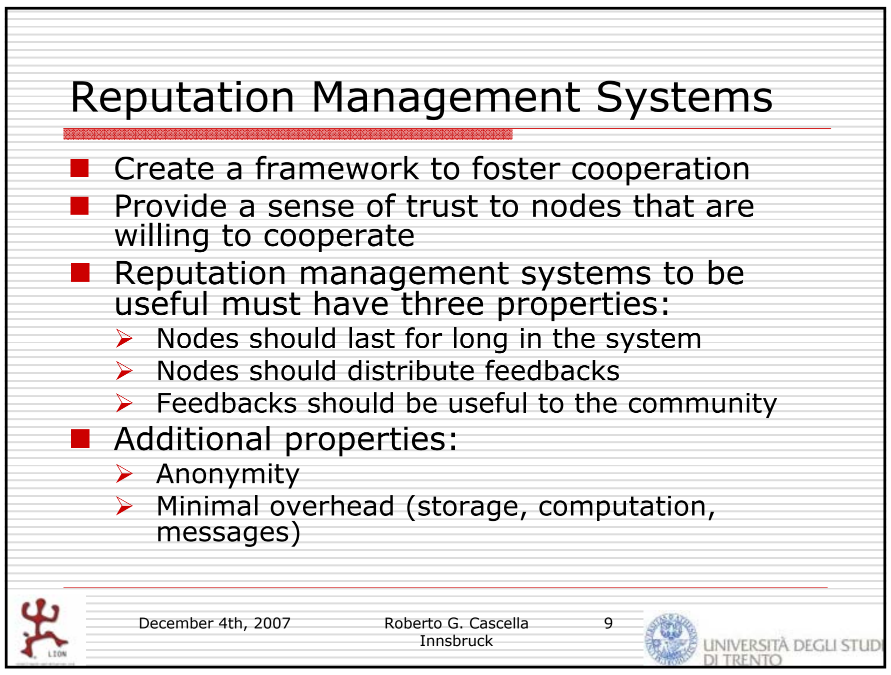

Reputation Management Systems

- Create a framework to foster cooperation
- Provide a sense of trust to nodes that are willing to cooperate
- Reputation management systems to be useful must have three properties:
  - Nodes should last for long in the system
  - Nodes should distribute feedbacks
  - Feedbacks should be useful to the community
- Additional properties:
  - Anonymity
  - Minimal overhead (storage, computation, messages)

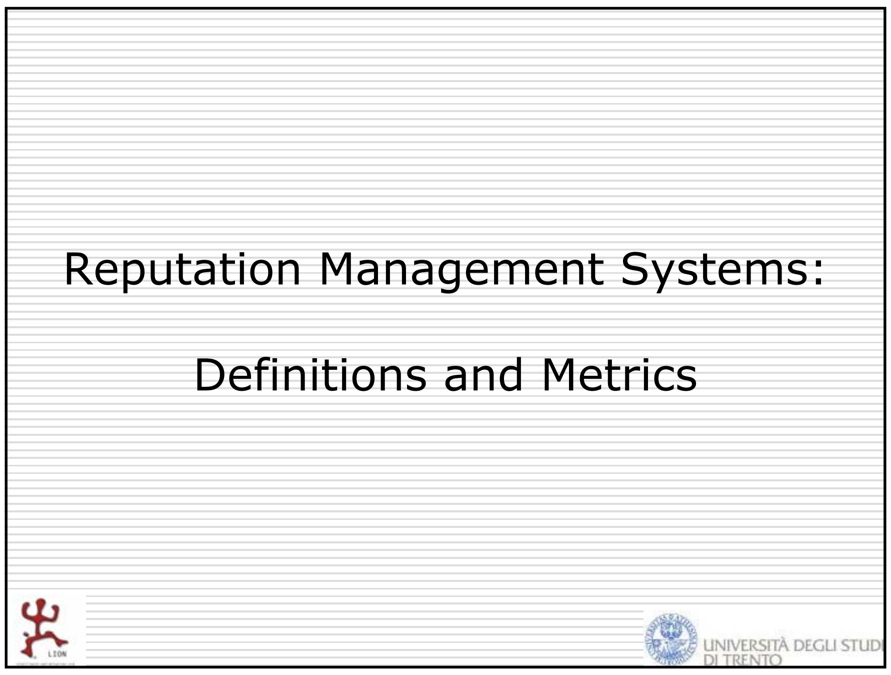

Reputation Management Systems:

Definitions and Metrics

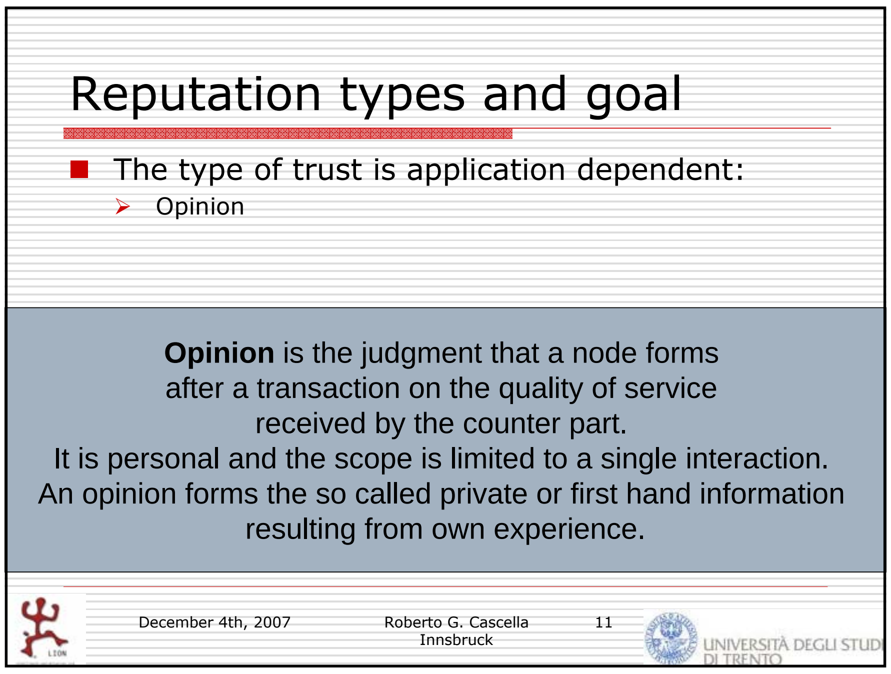

Reputation types and goal

- The type of trust is application dependent:
  - Opinion

Opinion is the judgment that a node forms after a transaction on the quality of service received by the counter part.
It is personal and the scope is limited to a single interaction.
An opinion forms the so called private or first hand information resulting from own experience.

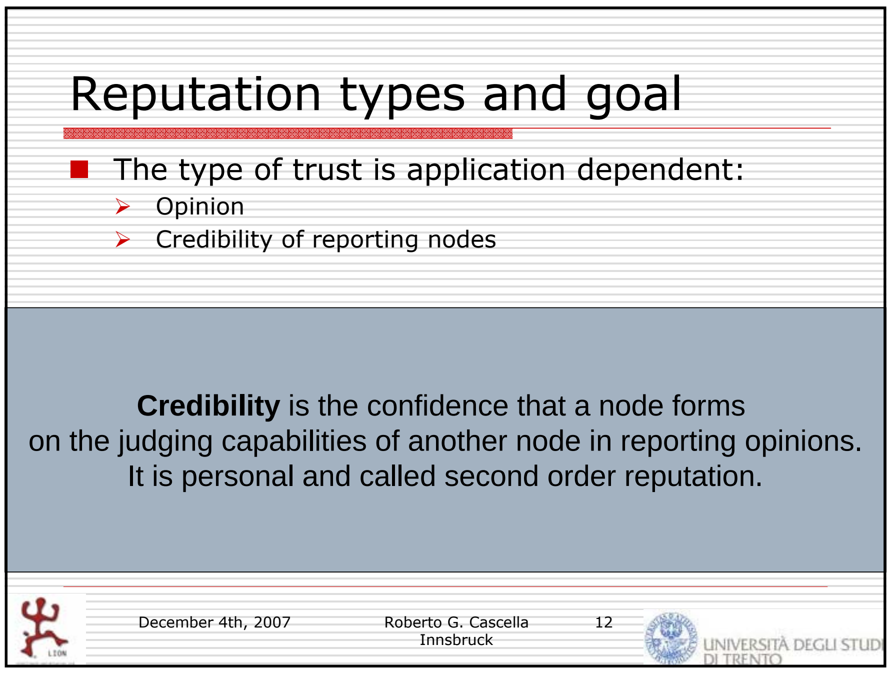

Reputation types and goal

- The type of trust is application dependent:
  - Opinion
  - Credibility of reporting nodes

Credibility is the confidence that a node forms on the judging capabilities of another node in reporting opinions.
It is personal and called second order reputation.

Reputation types and goal

- The type of trust is application dependent:
  - Opinion
  - Credibility of reporting nodes
  - Reputation (community judgment)

Reputation measures the trustworthiness of a peer in a system.
It is the global system-wide view of a node
or what is believed about this node.
In short, reputation is the collective measure of trustworthiness
based on the judgement of a community. It is quantified and it is calculated
by considering the action of a node in the view of a community of users.

Reputation types and goal

- The type of trust is application dependent:
  - Opinion
  - Credibility of reporting nodes
  - Reputation (community judgment)
- Reputation to be useful must be objective
  - Algorithms for aggregation of reported values
- The goal of the reputation might be context and application dependent:
  - A node can be trustworthy for providing service of type 1 or/and untrustworthy for providing service of type 2

Trust

Trust is a relationship of reliance and decision in social science.
A trusted party proves to benefit the belief of other peers to fulfill its obligation.
The definition of trust might include also the concept of risk,
when the value of the outcome of a transaction is high and
there exists the probability of failure.
The concept of trust is stronger than reputation as a node risks in person.

- The trustworthiness of the node is subjective
  - Function of reputation and opinion
  - Quantification of the risk

$$T_{xj} = (1 - w_p)O^{avg} + w_p \frac{\sum_d R_{dj}^{avg} \cdot C_{xd}}{\sum_d C_{xd}}$$

Peer-to-peer system: layered structure

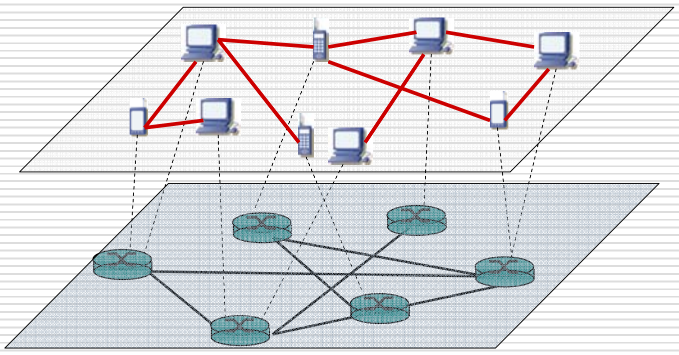

System definition

- In a reputation management system the reputation information needs to be
  1) collected from the feedback providers (how a node behaved in the past) – reactive, proactive or hybrid approach
  2) aggregated to form a useful measure of trustworthiness (where?)
  3) disseminated to members requesting the reputation value of a particular node
- A reputation management system needs to implement three distinct functions.

Reputation aggregation: where?

- Transacting Node
- All Nodes
- Central Database
- One-hop Neighbours
- Multi-hop Neighbours
- Designated agents: algorithm dependent -> Hash function

Sergio Marti and Hector Garcia-Molina. Taxonomy of trust: Categorizing P2P reputation systems. Computer Netowkrs, 50(4):472–484, 2006.

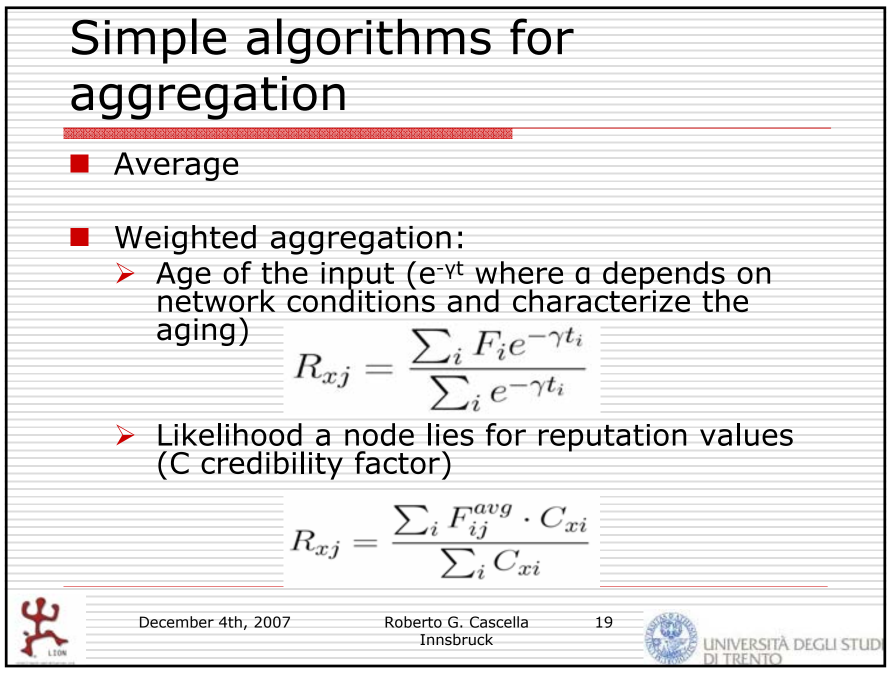

Simple algorithms for aggregation

- Average
- Weighted aggregation:
  - Age of the input ($e^{-\gamma t}$ where α depends on network conditions and characterize the aging)

$$R_{xj} = \frac{\sum_i F_i e^{-\gamma t_i}}{\sum_i e^{-\gamma t_i}}$$

  - Likelihood a node lies for reputation values (C credibility factor)

$$R_{xj} = \frac{\sum_i F_{ij}^{avg} \cdot C_{xi}}{\sum_i C_{xi}}$$

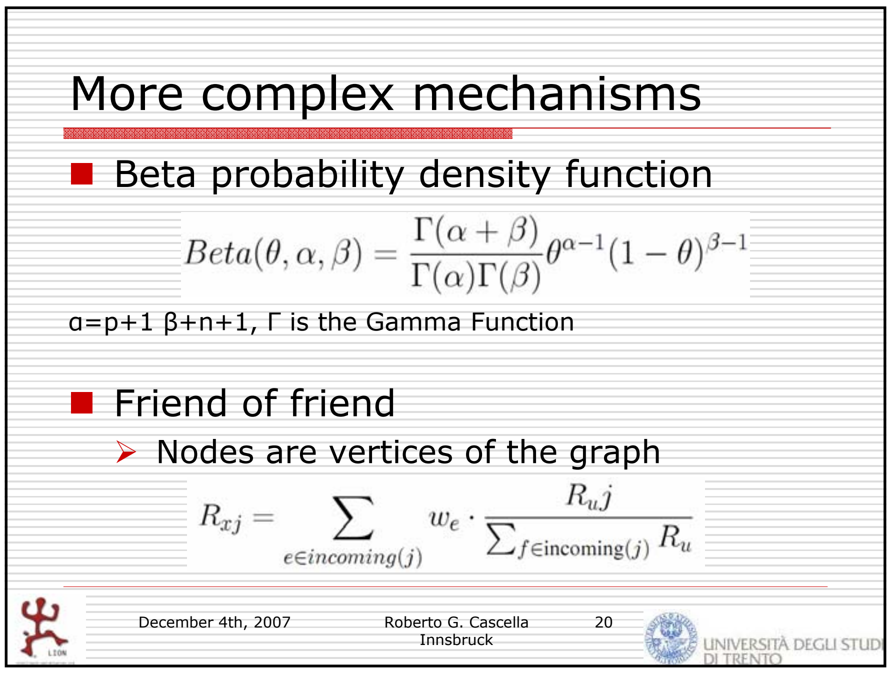

More complex mechanisms

- Beta probability density function

$$Beta(\theta, \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}$$

α=p+1 β+n+1, Γ is the Gamma Function

- Friend of friend
  - Nodes are vertices of the graph

$$R_{xj} = \sum_{e \in incoming(j)} w_e \cdot \frac{R_u j}{\sum_{f \in \text{incoming}(j)} R_u}$$

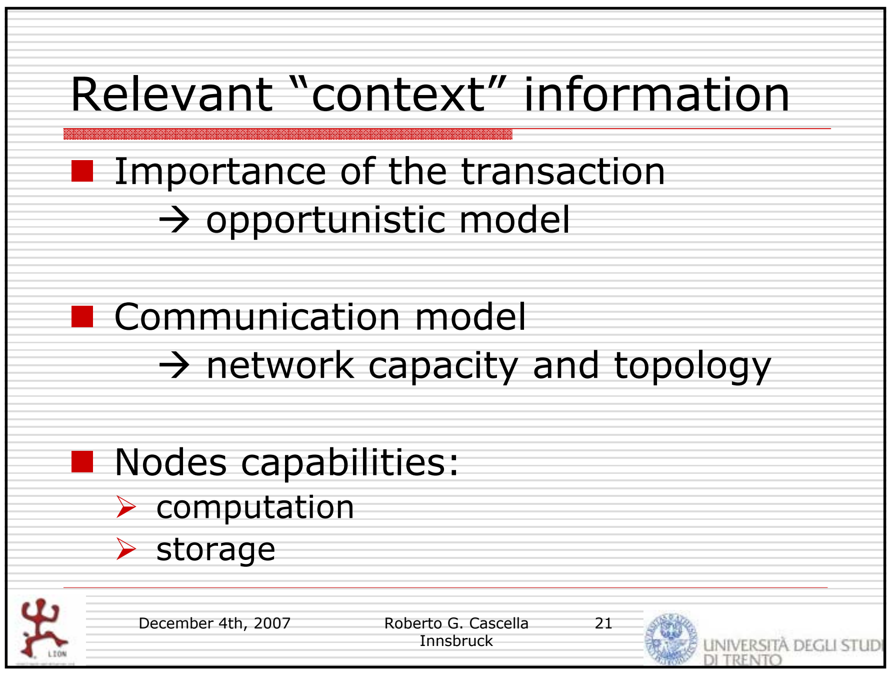

Relevant “context” information

- Importance of the transaction
  → opportunistic model
- Communication model
  → network capacity and topology
- Nodes capabilities:
  - computation
  - storage

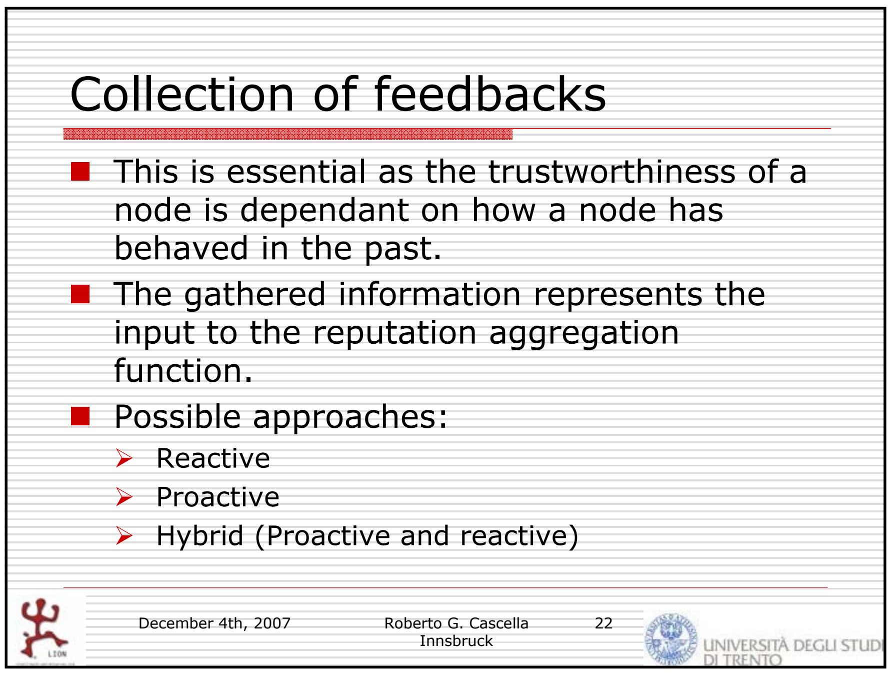

Collection of feedbacks

- This is essential as the trustworthiness of a node is dependant on how a node has behaved in the past.
- The gathered information represents the input to the reputation aggregation function.
- Possible approaches:
  - Reactive
  - Proactive
  - Hybrid (Proactive and reactive)

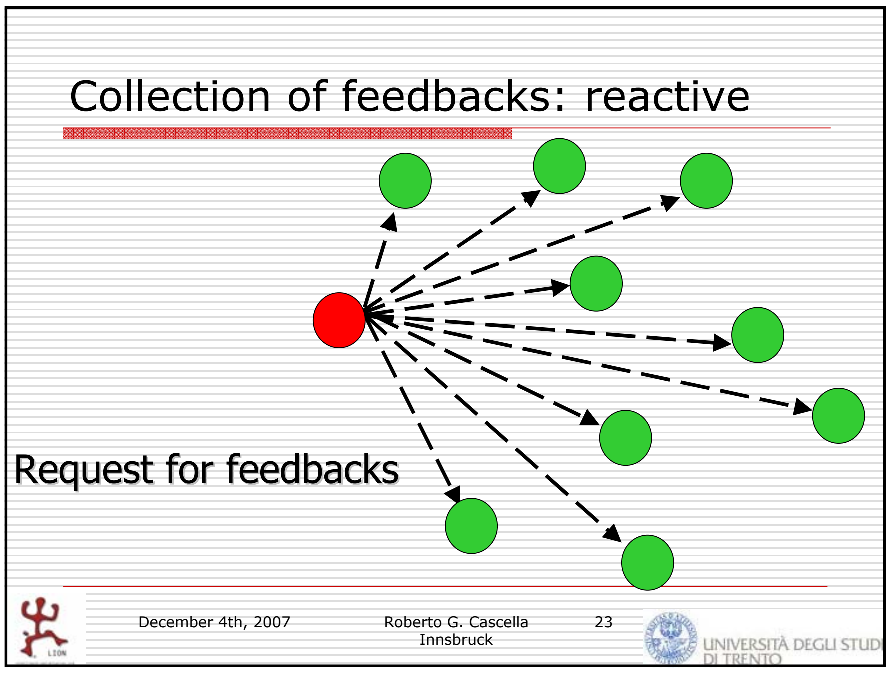

Collection of feedbacks: reactive

Request for feedbacks

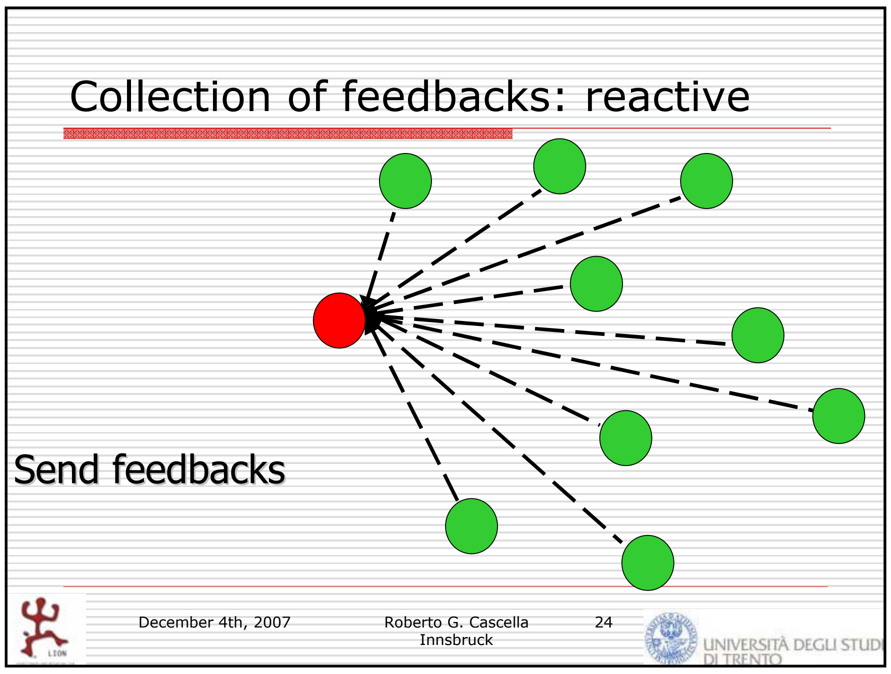

Collection of feedbacks: reactive

Send feedbacks

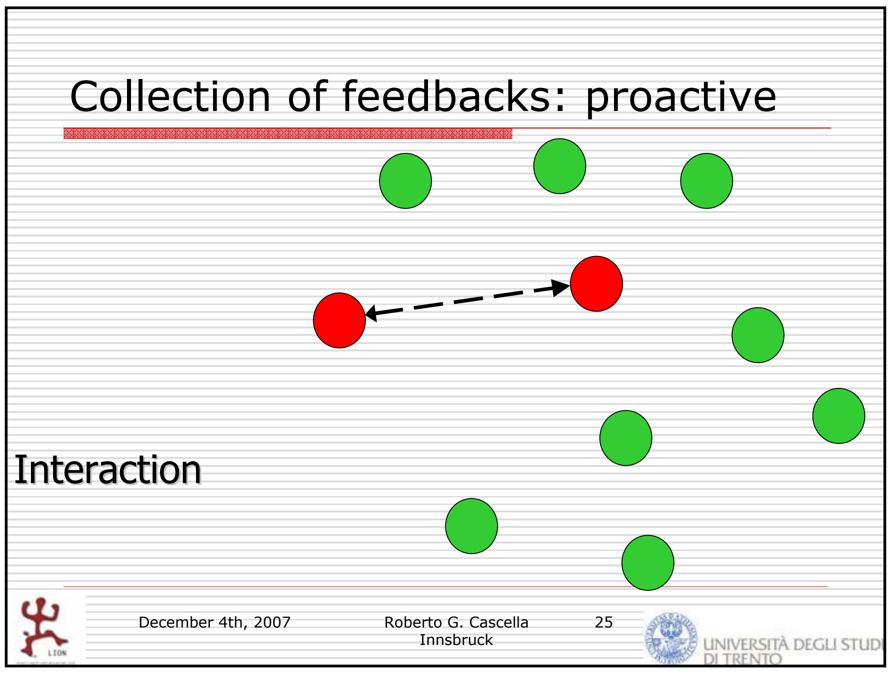

Collection of feedbacks: proactive

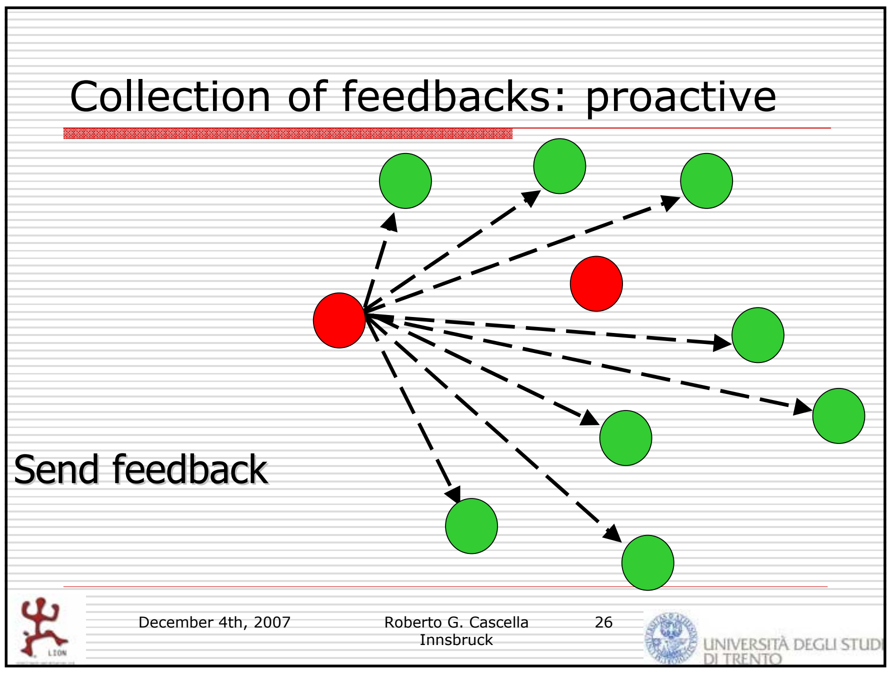

Collection of feedbacks: proactive

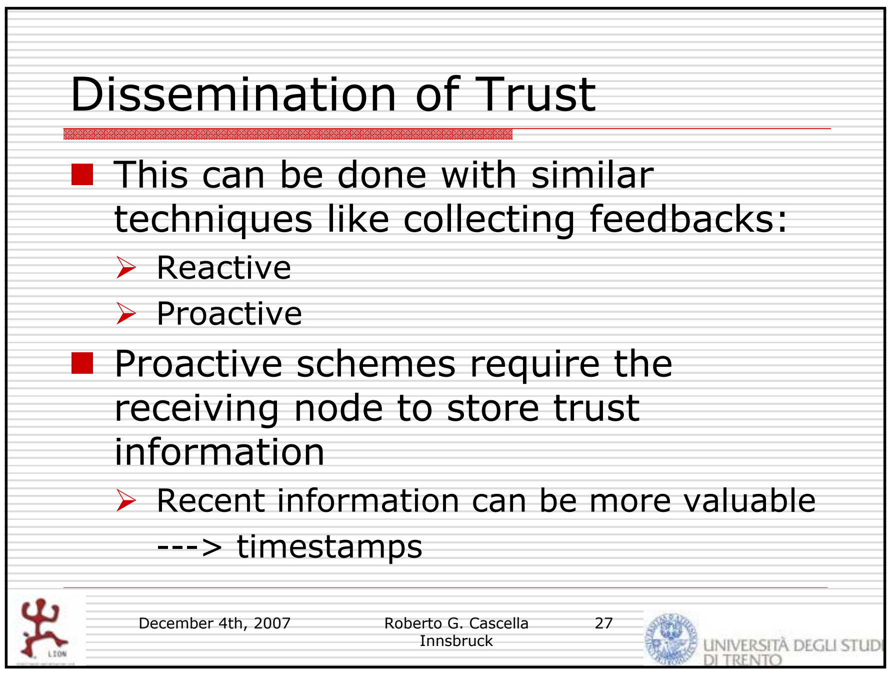

Dissemination of Trust

- This can be done with similar techniques like collecting feedbacks:
  - Reactive
  - Proactive
- Proactive schemes require the receiving node to store trust information
  - Recent information can be more valuable ---> timestamps

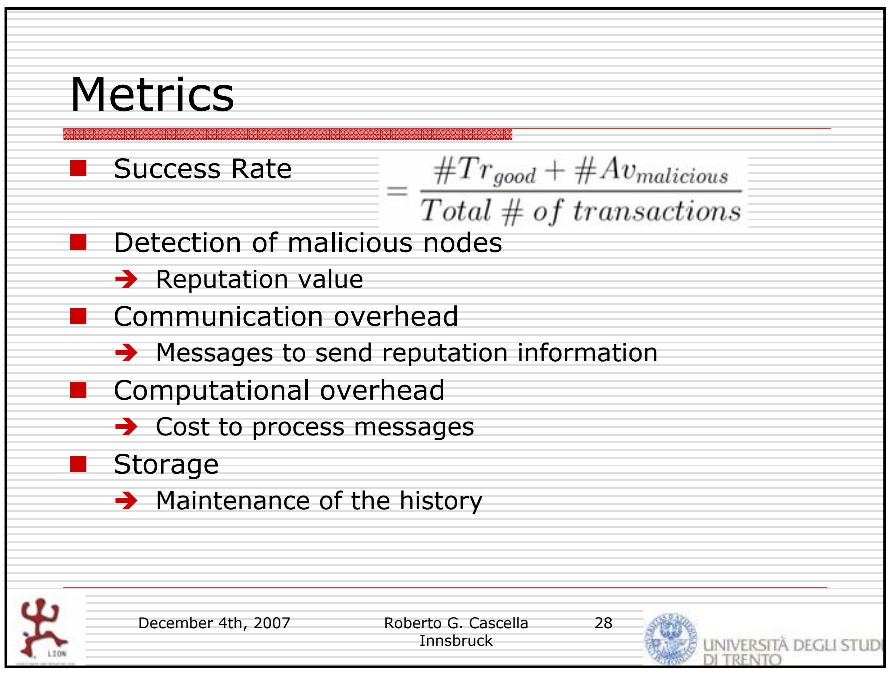

Metrics

- Success Rate $= \frac{\#Tr_{good} + \#Av_{malicious}}{Total\ \#\ of\ transactions}$
- Detection of malicious nodes
  - Reputation value
- Communication overhead
  - Messages to send reputation information
- Computational overhead
  - Cost to process messages
- Storage
  - Maintenance of the history

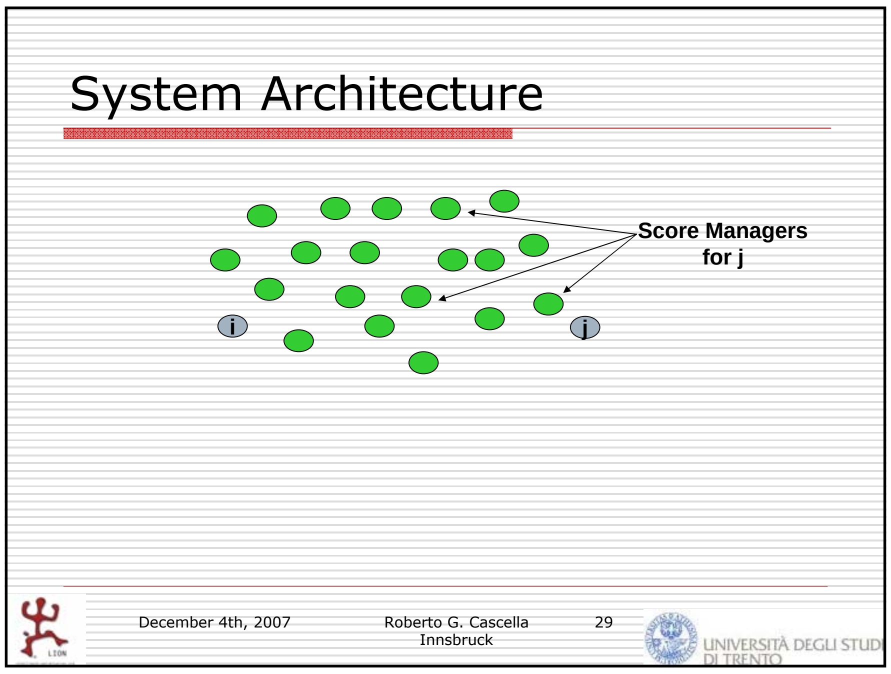

System Architecture

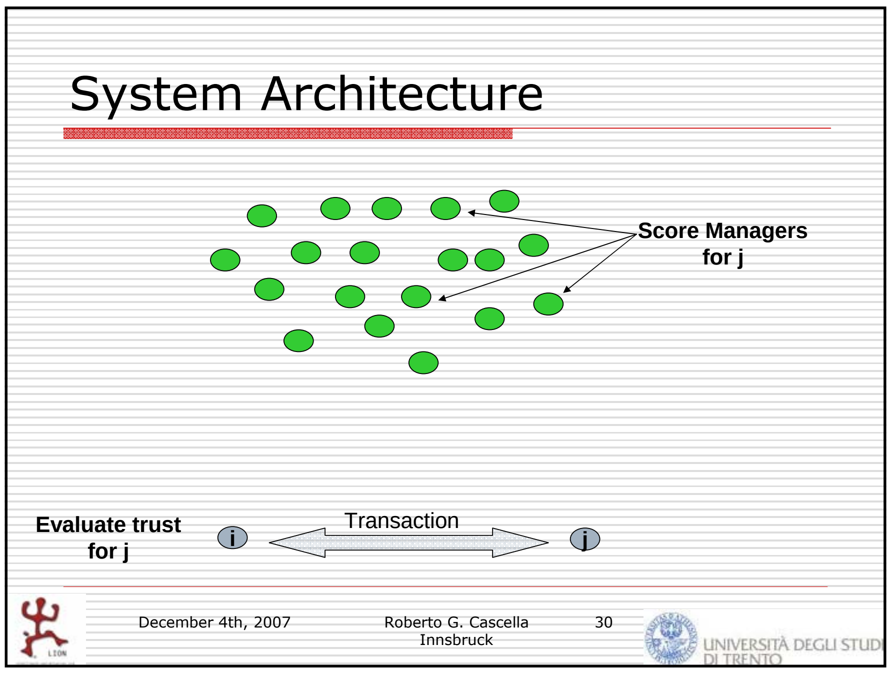

System Architecture

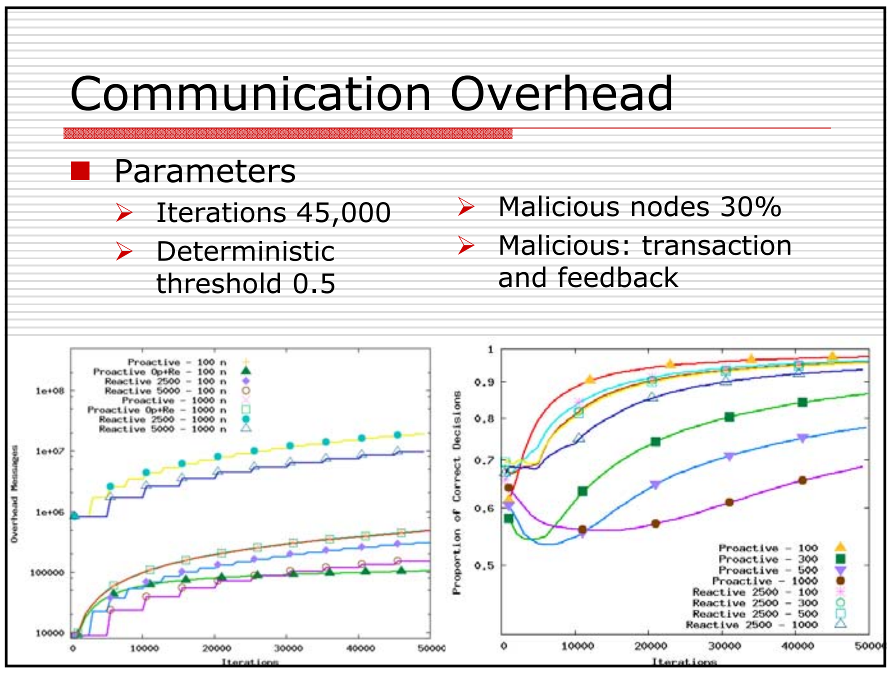

Communication Overhead

- Parameters
  - Iterations 45,000
  - Deterministic threshold 0.5
  - Malicious nodes 30%
  - Malicious: transaction and feedback

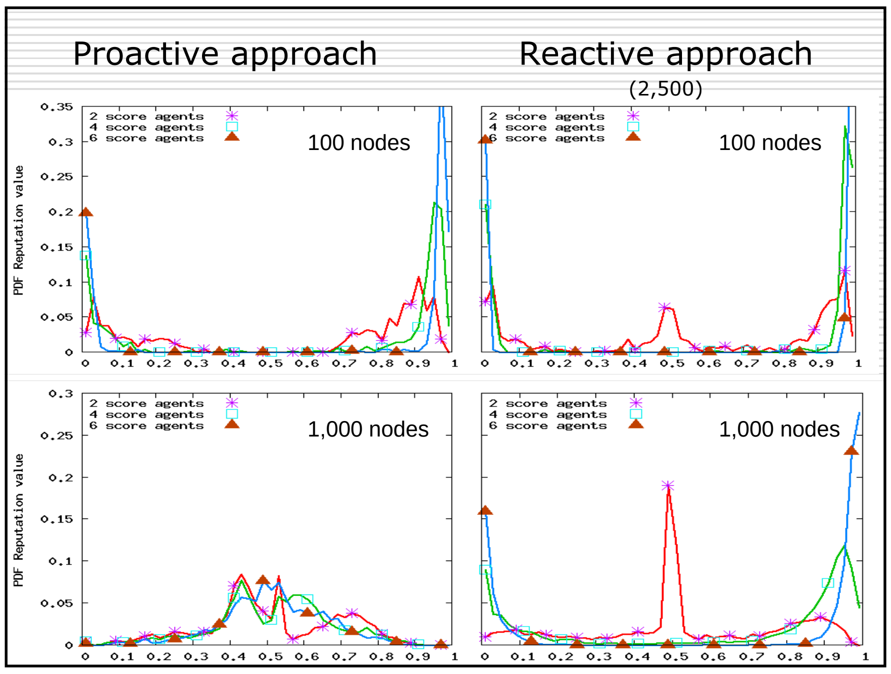

Proactive approach

Reactive approach (2,500)

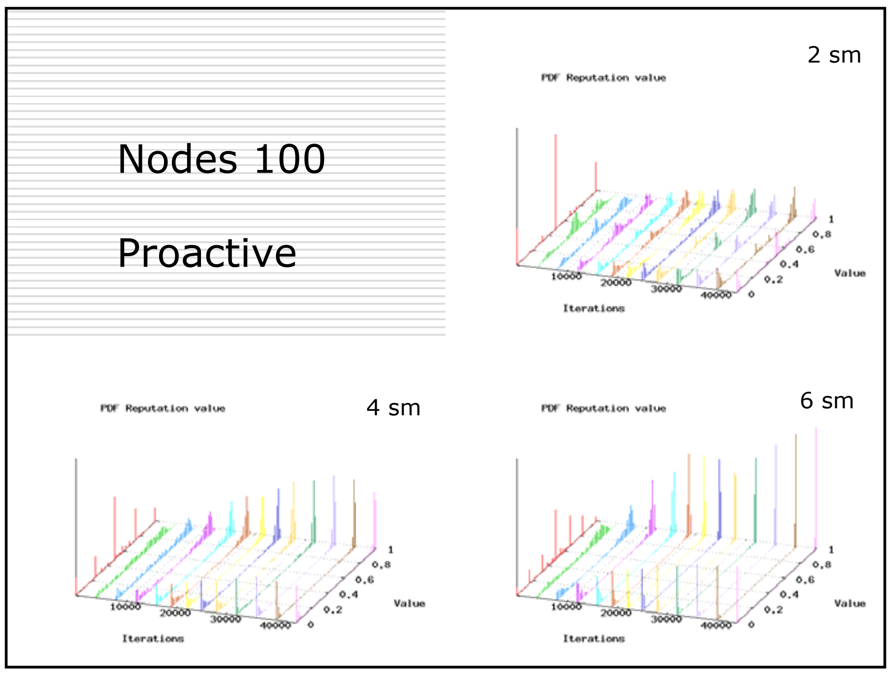

Nodes 100

Proactive

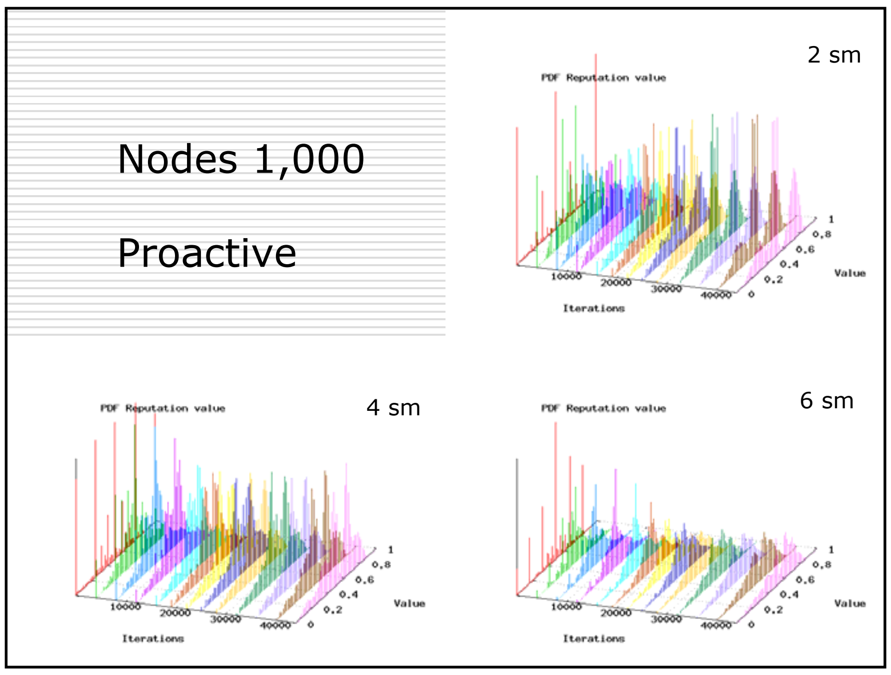

Nodes 1,000

Proactive

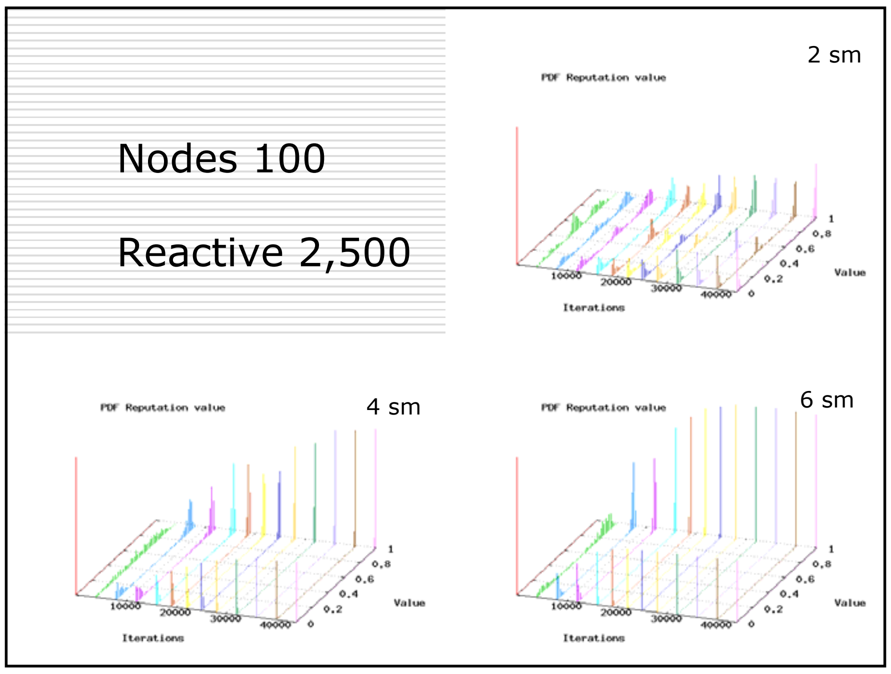

Nodes 100

Reactive 2,500

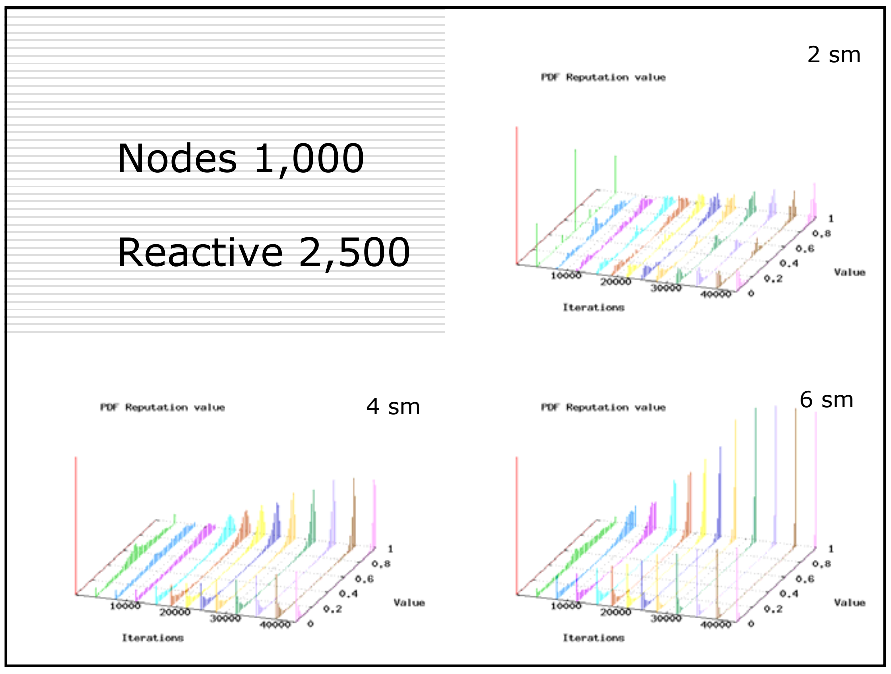

Nodes 1,000

Reactive 2,500

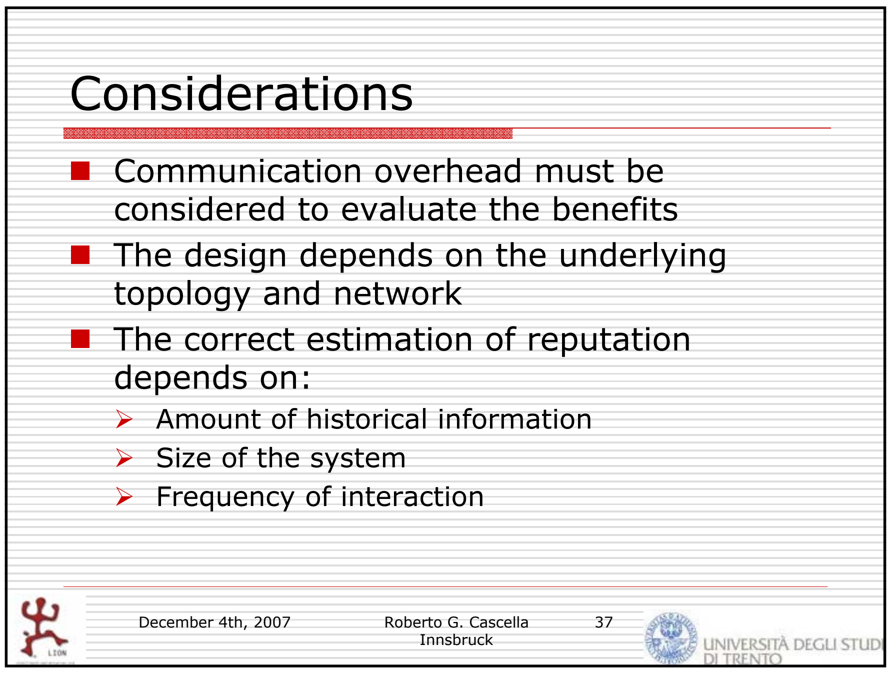

Considerations

- Communication overhead must be considered to evaluate the benefits
- The design depends on the underlying topology and network
- The correct estimation of reputation depends on:
  - Amount of historical information
  - Size of the system
  - Frequency of interaction

ROCQ

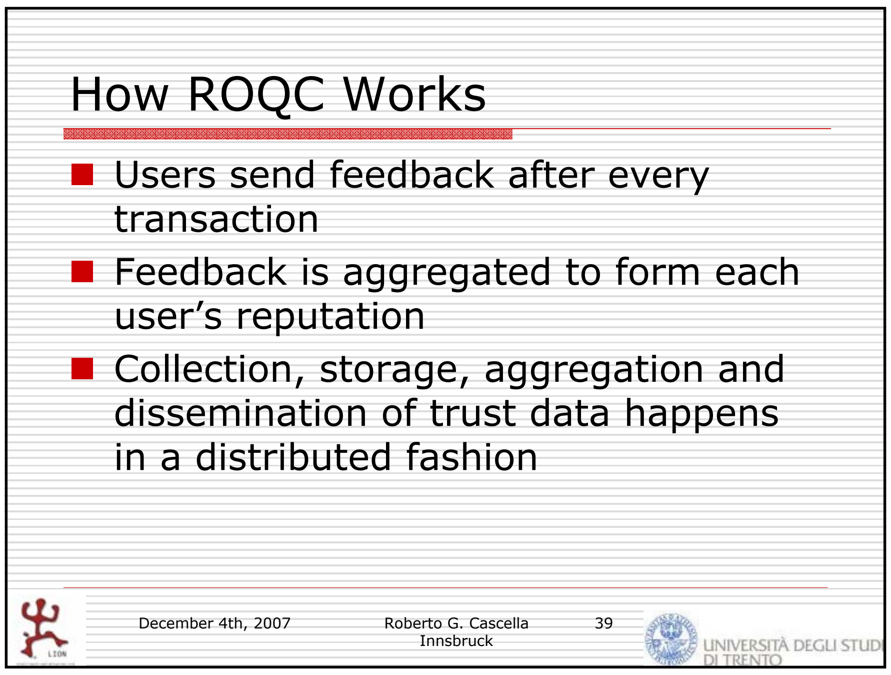

How ROQC Works

- Users send feedback after every transaction
- Feedback is aggregated to form each user's reputation
- Collection, storage, aggregation and dissemination of trust data happens in a distributed fashion

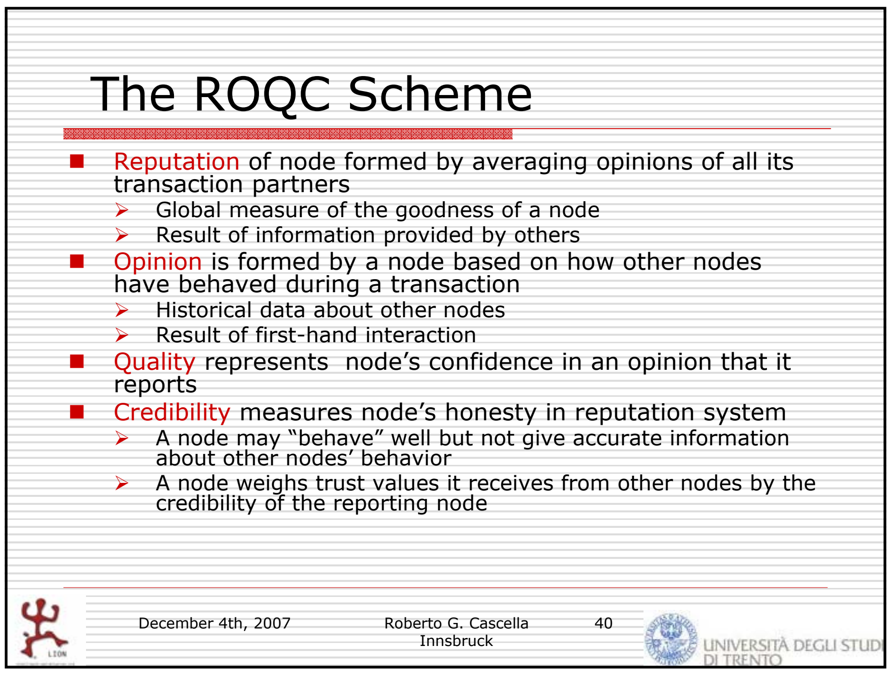

The ROQC Scheme

- Reputation of node formed by averaging opinions of all its transaction partners
  - Global measure of the goodness of a node
  - Result of information provided by others
- Opinion is formed by a node based on how other nodes have behaved during a transaction
  - Historical data about other nodes
  - Result of first-hand interaction
- Quality represents node's confidence in an opinion that it reports
- Credibility measures node's honesty in reputation system
  - A node may "behave" well but not give accurate information about other nodes' behavior
  - A node weighs trust values it receives from other nodes by the credibility of the reporting node

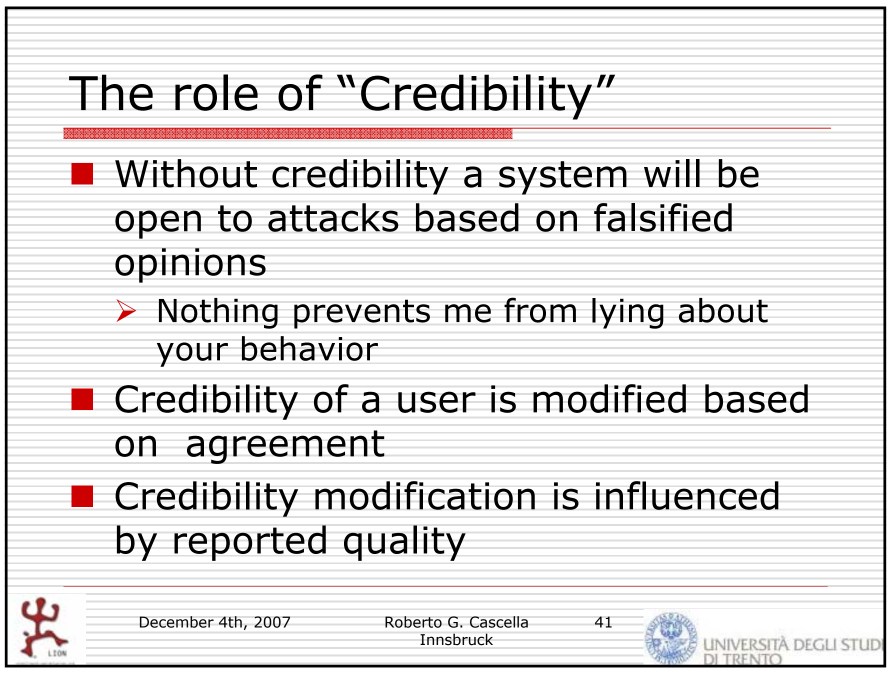

The role of "Credibility"

- Without credibility a system will be open to attacks based on falsified opinions
  - Nothing prevents me from lying about your behavior
- Credibility of a user is modified based on agreement
- Credibility modification is influenced by reported quality

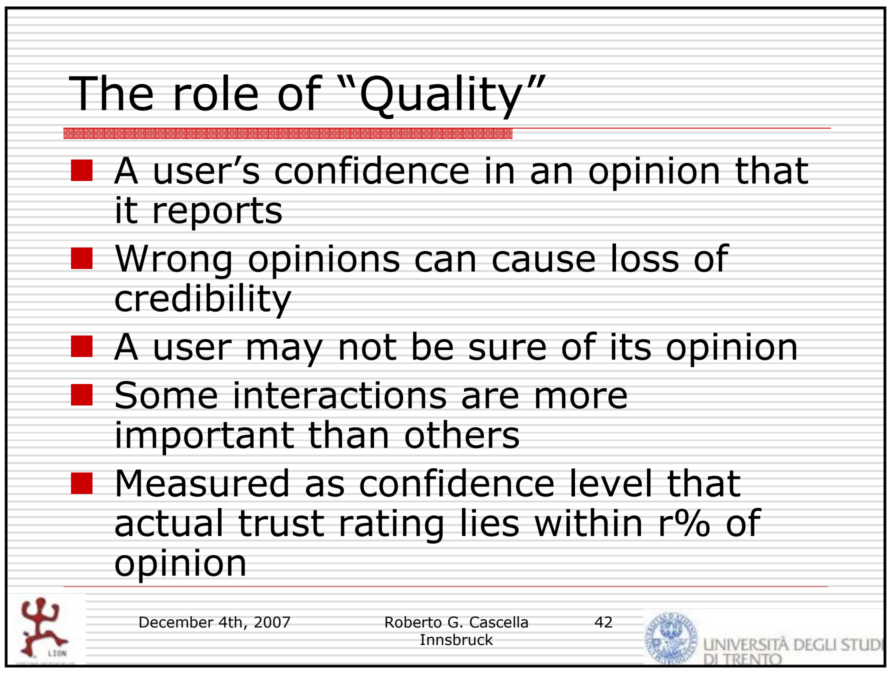

The role of "Quality"

- A user's confidence in an opinion that it reports
- Wrong opinions can cause loss of credibility
- A user may not be sure of its opinion
- Some interactions are more important than others
- Measured as confidence level that actual trust rating lies within r% of opinion

ROCQ: Equations

$$R_{mj} = \frac{\sum_i O_{ij}^{avg} \cdot C_{mi} \cdot Q_{ij}}{\sum_i C_{mi} \cdot Q_{ij}}$$

$$C_{mi}^{k+1} = \begin{cases} C_{mi}^k + \frac{(1-C_{mi}^k Q_{ij})}{2}\left(1 - \frac{|R_{mj}-O_{ij}^{avg}|}{s_{mj}}\right), & \text{if } |R_{mj} - O_{ij}^{avg}| < s_{mj} \\ C_{mi}^k - \frac{C_{mi}^k Q_{ij}}{2}\left(1 - \frac{s_{mj}}{|R_{mj}-O_{ij}^{avg}|}\right), & \text{if } |R_{mj} - O_{ij}^{avg}| \geq s_{mj} \end{cases}$$

ROCQ: Equations

Quality is the likelihood that actual trust value lies within this range →

$$O_{ij}^{avg} \cdot (1 \pm \frac{r}{100})$$

$$Q_{ij} = 1 - B\left(\frac{(N_{ij}-1)}{(N_{ij}-1)+t^2}; \frac{1}{2} \cdot (N_{ij}-1), \frac{1}{2}\right)$$

The t-value for the *Student's t-distribution* is given by the following equation:

$$t = \frac{r}{100} \cdot \frac{O_{ij}^{avg} \cdot \sqrt{N_{ij}}}{s_{ij}}$$

System Architecture

- Assume a structured overlay network that uses Distributed Hash Tables
- DHT is used to assign Score Managers (SM)
- Multiple SMs to ensure reliability and guard against malicious SMs
- SM for a peer stores all trust information related to that peer
- Opinions about a peer are reported to all of its SMs

6 score managers
Deterministic threshold of 0.5
Proactive dissemination
Nodes 200
Transactions 50,000

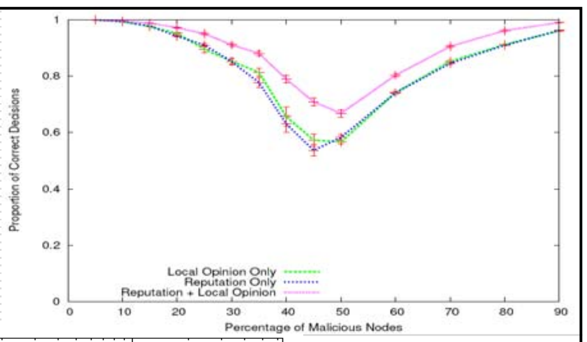

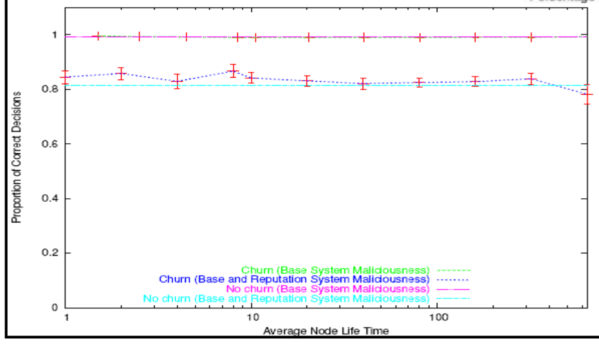

EigenTrust

- The Eigentrust algorithm is based on the notion of transitive trust
- Local Rating: $s_{ij} = \mathrm{sat}(i,j) - \mathrm{unsat}(i,j)$
- Normalized rating: $c_{ij} = \frac{\max(s_{ij}, 0)}{\sum_j \max(s_{ij}, 0)}$
- Local trust values: $t_{ik} = \sum_j c_{ij} c_{jk}$
- Friend of friend: $\vec{t} = (C^T)^n \vec{c_i}$

Note that for large values of n
➔ t will converge to the same vector **Left principal eigenvector of C**

EigenTrust: refinements

- The algorithm has faster converge with a set of pre-trusted peers
  - Malicious peers lies

Definitions:

- A_i: set of peers which have downloaded files from peer i
- B_i: set of peers from which peer i has downloaded files

Algorithm:

```
Each peer $i$ do {
Query all peers $j \in A_i$ for $t_j^{(0)} = p_j$;
repeat
    Compute $t_i^{(k+1)} = (1-a)(c_{1i}t_1^{(k)} + c_{2i}t_2^{(k)} + \ldots + c_{ni}t_n^{(k)}) + ap_i$;
    Send $c_{ij}t_i^{(k+1)}$ to all peers $j \in B_i$;
    Compute $\delta = |t_i^{(k+1)} - t_i^{(k)}|$;
    Wait for all peers $j \in A_i$ to return $c_{ji}t_j^{(k+1)}$;
until $\delta < \epsilon$;
}
```

- Secure trust storage
  - Nodes might report false trust values for themselves

Distributed version

Source: Sepandar D. Kamvar, Mario T. Schlosser, Hector Garcia-Molina. "The Eigentrust algorithm for reputation management in P2P networks". In Proceeding of WWW 2003: 640-651

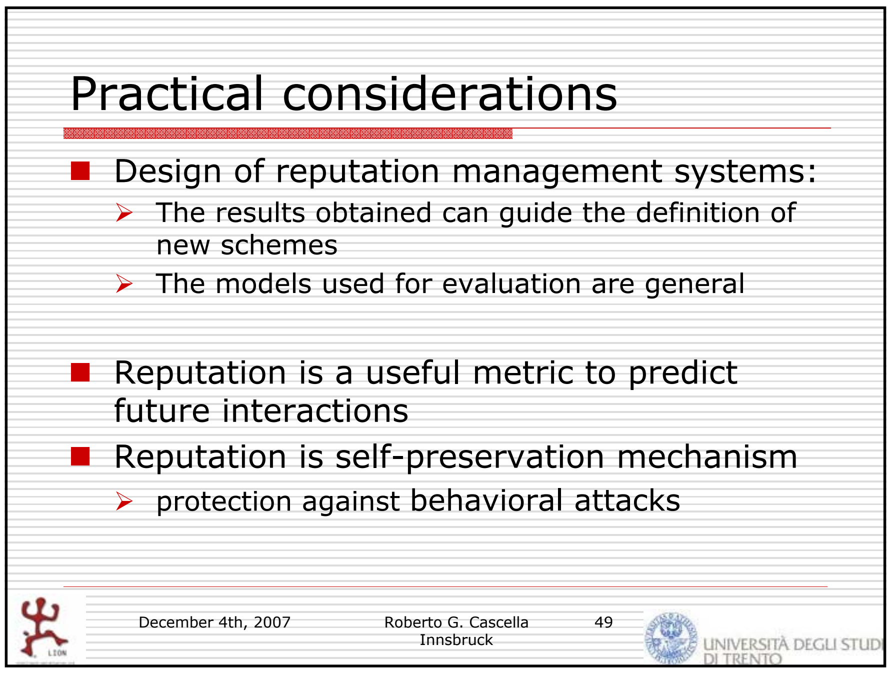

Practical considerations

- Design of reputation management systems:
  - The results obtained can guide the definition of new schemes
  - The models used for evaluation are general
- Reputation is a useful metric to predict future interactions
- Reputation is self-preservation mechanism
  - protection against behavioral attacks

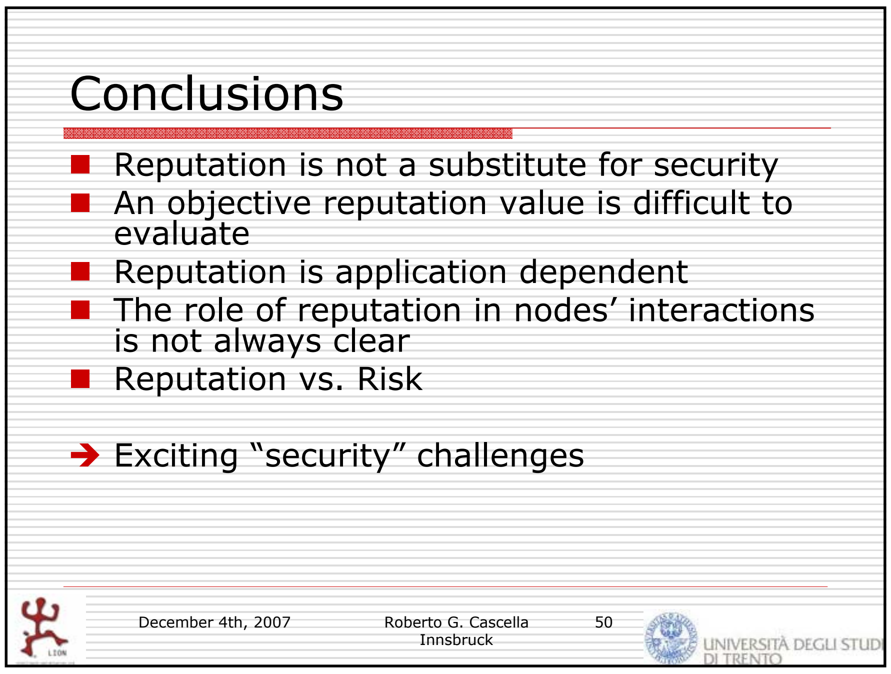

Conclusions

- Reputation is not a substitute for security
- An objective reputation value is difficult to evaluate
- Reputation is application dependent
- The role of reputation in nodes' interactions is not always clear
- Reputation vs. Risk

➔ Exciting "security" challenges

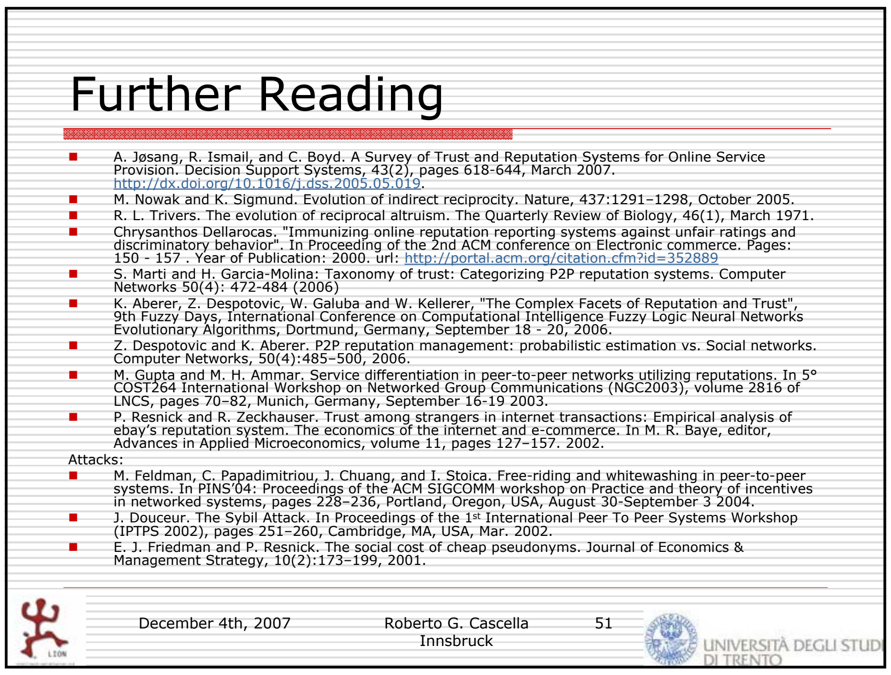

Further Reading

- A. Jøsang, R. Ismail, and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. Decision Support Systems, 43(2), pages 618-644, March 2007. http://dx.doi.org/10.1016/j.dss.2005.05.019.
- M. Nowak and K. Sigmund. Evolution of indirect reciprocity. Nature, 437:1291–1298, October 2005.
- R. L. Trivers. The evolution of reciprocal altruism. The Quarterly Review of Biology, 46(1), March 1971.
- Chrysanthos Dellarocas. "Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior". In Proceeding of the 2nd ACM conference on Electronic commerce. Pages: 150 - 157 . Year of Publication: 2000. url: http://portal.acm.org/citation.cfm?id=352889
- S. Marti and H. Garcia-Molina: Taxonomy of trust: Categorizing P2P reputation systems. Computer Networks 50(4): 472-484 (2006)
- K. Aberer, Z. Despotovic, W. Galuba and W. Kellerer, "The Complex Facets of Reputation and Trust", 9th Fuzzy Days, International Conference on Computational Intelligence Fuzzy Logic Neural Networks Evolutionary Algorithms, Dortmund, Germany, September 18 - 20, 2006.
- Z. Despotovic and K. Aberer. P2P reputation management: probabilistic estimation vs. Social networks. Computer Networks, 50(4):485–500, 2006.
- M. Gupta and M. H. Ammar. Service differentiation in peer-to-peer networks utilizing reputations. In 5° COST264 International Workshop on Networked Group Communications (NGC2003), volume 2816 of LNCS, pages 70–82, Munich, Germany, September 16-19 2003.
- P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. The economics of the internet and e-commerce. In M. R. Baye, editor, Advances in Applied Microeconomics, volume 11, pages 127–157. 2002.

Attacks:

- M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. In PINS'04: Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems, pages 228–236, Portland, Oregon, USA, August 30-September 3 2004.
- J. Douceur. The Sybil Attack. In Proceedings of the 1st International Peer To Peer Systems Workshop (IPTPS 2002), pages 251–260, Cambridge, MA, USA, Mar. 2002.
- E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. Journal of Economics & Management Strategy, 10(2):173–199, 2001.

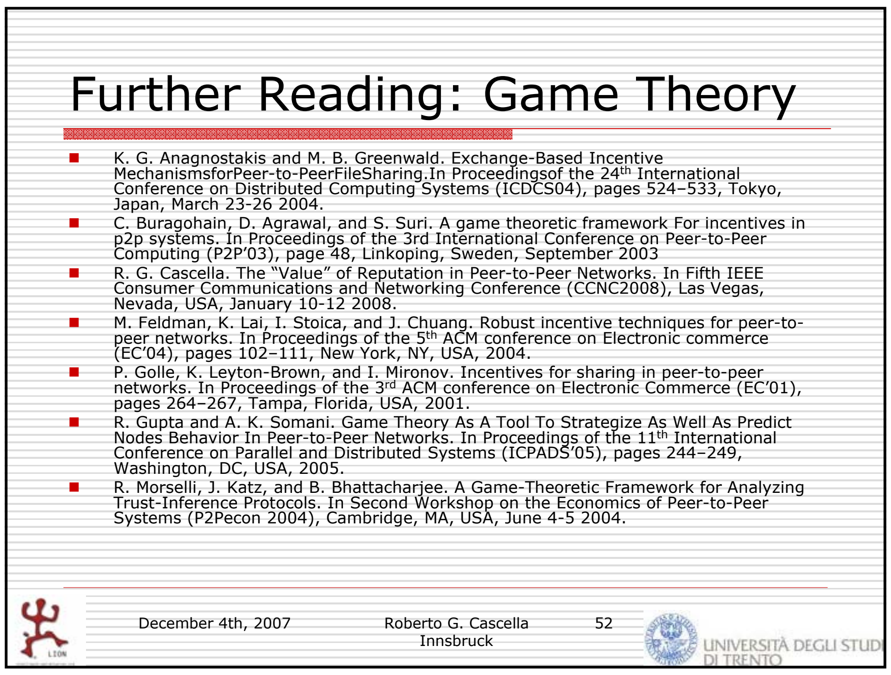

Further Reading: Game Theory

- K. G. Anagnostakis and M. B. Greenwald. Exchange-Based Incentive MechanismsforPeer-to-PeerFileSharing.In Proceedingsof the 24th International Conference on Distributed Computing Systems (ICDCS04), pages 524–533, Tokyo, Japan, March 23-26 2004.
- C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework For incentives in p2p systems. In Proceedings of the 3rd International Conference on Peer-to-Peer Computing (P2P'03), page 48, Linkoping, Sweden, September 2003
- R. G. Cascella. The "Value" of Reputation in Peer-to-Peer Networks. In Fifth IEEE Consumer Communications and Networking Conference (CCNC2008), Las Vegas, Nevada, USA, January 10-12 2008.
- M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In Proceedings of the 5th ACM conference on Electronic commerce (EC'04), pages 102–111, New York, NY, USA, 2004.
- P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In Proceedings of the 3rd ACM conference on Electronic Commerce (EC'01), pages 264–267, Tampa, Florida, USA, 2001.
- R. Gupta and A. K. Somani. Game Theory As A Tool To Strategize As Well As Predict Nodes Behavior In Peer-to-Peer Networks. In Proceedings of the 11th International Conference on Parallel and Distributed Systems (ICPADS'05), pages 244–249, Washington, DC, USA, 2005.
- R. Morselli, J. Katz, and B. Bhattacharjee. A Game-Theoretic Framework for Analyzing Trust-Inference Protocols. In Second Workshop on the Economics of Peer-to-Peer Systems (P2Pecon 2004), Cambridge, MA, USA, June 4-5 2004.

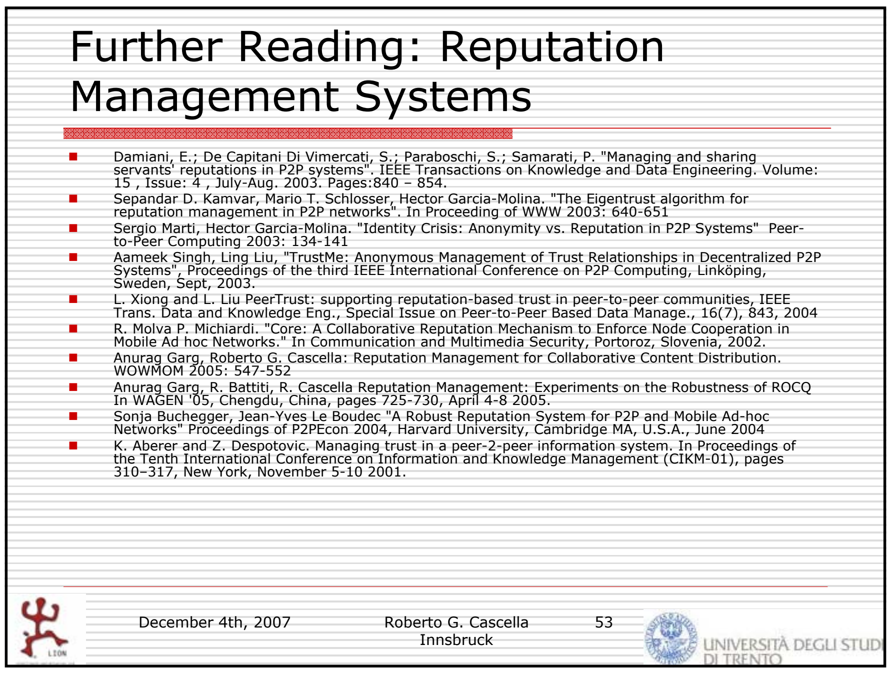

Further Reading: Reputation Management Systems

- Damiani, E.; De Capitani Di Vimercati, S.; Paraboschi, S.; Samarati, P. "Managing and sharing servants' reputations in P2P systems". IEEE Transactions on Knowledge and Data Engineering. Volume: 15 , Issue: 4 , July-Aug. 2003. Pages:840 – 854.
- Sepandar D. Kamvar, Mario T. Schlosser, Hector Garcia-Molina. "The Eigentrust algorithm for reputation management in P2P networks". In Proceeding of WWW 2003: 640-651
- Sergio Marti, Hector Garcia-Molina. "Identity Crisis: Anonymity vs. Reputation in P2P Systems" Peer-to-Peer Computing 2003: 134-141
- Aameek Singh, Ling Liu, "TrustMe: Anonymous Management of Trust Relationships in Decentralized P2P Systems", Proceedings of the third IEEE International Conference on P2P Computing, Linköping, Sweden, Sept, 2003.
- L. Xiong and L. Liu PeerTrust: supporting reputation-based trust in peer-to-peer communities, IEEE Trans. Data and Knowledge Eng., Special Issue on Peer-to-Peer Based Data Manage., 16(7), 843, 2004
- R. Molva P. Michiardi. "Core: A Collaborative Reputation Mechanism to Enforce Node Cooperation in Mobile Ad hoc Networks." In Communication and Multimedia Security, Portoroz, Slovenia, 2002.
- Anurag Garg, Roberto G. Cascella: Reputation Management for Collaborative Content Distribution. WOWMOM 2005: 547-552
- Anurag Garg, R. Battiti, R. Cascella Reputation Management: Experiments on the Robustness of ROCQ In WAGEN '05, Chengdu, China, pages 725-730, April 4-8 2005.
- Sonja Buchegger, Jean-Yves Le Boudec "A Robust Reputation System for P2P and Mobile Ad-hoc Networks" Proceedings of P2PEcon 2004, Harvard University, Cambridge MA, U.S.A., June 2004
- K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM-01), pages 310–317, New York, November 5-10 2001.

Thanks!!!

cascella@dit.unitn.it