Sistemi di Elaborazione dell'informazione II

Corso di Laurea Specialistica in Ingegneria Telematica
II anno – 4 CFU
Università Kore – Enna – A.A. 2008-2009

Alessandro Longheu
http://www.diit.unict.it/users/alongheu
alessandro.longheu@diit.unict.it

Il problema del Trust

Trusting - Fiducia

- Trust is an **integral component in many kinds of human interaction**, allowing people to act under uncertainty and with the risk of negative consequences.
- In computer science, trust is a widely used term whose definition differs among researchers and application areas. Trust is an **essential component of the vision for the Semantic Web**, where both new problems and new applications of trust are being studied.
- Trust often refers to **mechanisms to verify that the source of information is really who the source claims to be**. Signatures and encryption mechanisms should allow to check the sources of that information.
- The web motto "**Anyone can say anything about anything**" makes the web a unique source of information, but we need to be able to understand where we are placing our trust.

Trusting - Fiducia

- Trust has another important role in the Semantic Web, as **agents and automated reasoners need to make trust judgements** when alternative sources of information are available.
- These trust judgements are now made by humans based on their prior knowledge about a source's perceived reputation, or past personal experience about its quality relative to other alternative sources they may consider.
- **Trust judgments are currently in the hands of humans. This will not be possible in the Semantic Web**, where humans will not be the only consumers of information. Agents will need to automatically make trust judgments to choose a service or information source

Definizioni di Trust

- **Three general definitions** from existing research:
- **Reputation-based**: "[Trust is] a subjective expectation an agent has about another's future behavior based on the history of their encounters."
- **Competence-based**: "[Trust is] the firm belief in the competence of an entity to act dependably, securely, and reliably within a specified context."
- **Action-based**: "Trust of a party A to a party B for a service X is the measurable belief of A in that B behaves dependably for a specified period within a specified context (in relation to service X)."

Meccanismi per la valutazione del trust

- Two common ways of determining trust are through using policies or reputation:
  - **Policies** describe the conditions necessary to obtain trust, and can also prescribe actions and outcomes if certain conditions are met. Policies frequently involve the exchange or verification of **credentials**, hence it is based on "hard-evidence" owned by the entity
  - **Reputation** is an assessment based on the history of interactions with or observations of an entity, either directly with the evaluator (**personal experience**) or as reported by others (**recommendations** or third party verification). How these histories are combined can vary.

Policies in trust

- Evolving work in policies highlight a more complex problem in trust: how much to trust another entity to see your own credentials when you wish to earn that entity's trust
- An important problem in establishing trust is that revealing a credential may incur a loss of privacy or control of information. In some systems (e.g. TrustBuilder), **trust is earned when sufficient credentials are revealed** (but not too many to sacrifice privacy).
- In a more specific view, (Gandon and Sadeh, 2004) have proposed using ontologies to enable **context-aware applications** on the Semantic Web. Contextaware applications will only reveal credentials in the correct context

Policies in trust

- **several current policy languages**, designed for use in the Semantic Web, have been proposed, as KAoS and Rei; they address security and privacy issues in the semantic Web, while allowing each entity to specify their own policy. Other works uses ontologies to flexibly represent trust negotiation policies (rules used to negotiate trust). Ontologies have more flexibility than set standards, they simplify policy specification, and they enable more information to be specified to control privacy during trust negotiation.
- The well-known **Kerberos protocol** (Kohl and Neuman, 1993) is used to exchange credentials. The Kerberos system uses a third party to facilitate the exchange of credentials (digital signatures) between a user and a computer. Kerberos does not determine access rights, but instead enables two parties to securely exchange verifiable credentials.

Reputation in trust

- Reputation-based trust uses personal experience or the experiences of others, possibly combined, to make a trust decision about an entity.
- Just as in policy-based trust, one solution to obtaining trustworthy reputation information is to consult a **central**, trusted third party that has had prior experience with the entity in question and can provide an assessment of its reputation. The majority of existing work avoids this solution, and **most research focuses explicitly on decentralization** for reputation management. Reputation management avoids a hard security approach by distributing reputation information, allowing individuals to make trust decisions instead of a single, centralized trust system making the decisions for them.
- Another question is that **trust changes over time**; some approach uses statistical analysis to characterize trust and reputation so that computation remains scalable for long time

Reputation in trust

- A key recent example of the reputation-based approach describes how trust is computed for the application TrustMail. Reputation is defined as a measure of trust, and each entity maintains reputation information on other entities, thus creating a "web", that is called a **web of trust**. The work uses ontologies to express trust and reputation information, which then allows a quantification of trust for use in algorithms to make a trust decision about any two entities. The quantification of this trust and associated algorithms are called trust metrics.

Reputation in trust

- Some questions:
  - **local vs global** reputation (personalized or not)
  - **controversial users**; some works shows that the globally computed trust value (in a web of trust) for a controversial user may not be as accurate as a locally computed value due to the global disagreement on trust for that user
  - **empowering** (an user is supported maliciously by others)
  - **whitewashing** (bad users get new identities)
  - **feedback**, e.g. propagation of distrust (if A distrusts B and B distrusts C, we cannot say if A trusts C, this is a problem!)
  - objective vs subjective trust evaluation (e.g. based on maximum network flow)
  - **context**: some works considers the domain of knowledge (context). This work enumerates several kinds of **referral** (trust in ability to recommend) and **associative** (two agents being similar) trust: domain expert (trust in an agent's domain knowledge), recommendation expert (trust in an agent's ability to refer others), similar trusting (two agents having similar trust in other agents), and similar cited (two agents being similarly trusted by others)

Modelli per il trust

- **Four qualities** are important when making a trust decision:
  - competence (ability to give accurate information)
  - benevolence (willingness to expend the effort)
  - integrity (adherence to honest behavior)
  - predictability (evidence to support that the desired outcome will occur).
- One of the first works on trust proposed a set of (subjectively set) variables, an a way to combine them to arrive at one **continuous value** of trust in the range [-1,1], being -1 complete distrust, 1 complete trust and 0 the indifference

Trust in the Semantic Web

- Any statements contained in the Semantic Web must be considered as claims rather than facts until trust can be established, hence there is more to trust than simply reputation
- Noting that "trust is at the heart of the Semantic Web vision", some name **five trust strategies for agents using the Semantic Web**: optimism, pessimism, centralized, investigation, and transitivity.
  - Optimism is to assume trust,
  - pessimism is to assume distrust,
  - centralized is to trust through a single third party,
  - investigation is to collect trust information from others,
  - transitivity is to use a web of trust.

Trust in the Semantic Web

- Several works evaluate trust through hyperlinks:
  - some assume that **all Web links are positive endorsements** (and indications of trust).
  - Others propose a **minor addition to HTML**, enabling the author to specify whether a link is positive, negative, or neither
  - Others describes the concepts of **a hub and an authority**, the former being a page that points to many authorities, and the latter being a page that is pointed to by many hubs. The PageRank algorithm exploits Kleinberg's ideas of using links as human encoded judgments of relevance and uses the concept of authorities to compute a heuristic of popularity.
- The details regarding the sources and origins of information (e.g., author, publisher, citations, etc.) are referred to as **provenance**, and they serve as a means to evaluate trust.
- Trust on the Web is needed to make decisions when information conflicts or is non-authoritative.