Sistemi di Elaborazione dell'informazione II

Corso di Laurea Specialistica in Ingegneria Telematica
II anno – 4 CFU
Università Kore – Enna – A.A. 2009-2010

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

Trust

Trusting - Fiducia

- The Semantic Web makes heavy use of data and meta data, collected from a wide range of distributed sources. Typically we think of that information as a set of statements or assertions, collected into datastores that are somewhat like current databases but with more flexibility, and endowed with more powers of analysis and logical inference.
- **But all statements aren't equal**. For one reason or another, one statement may not be reliable, whereas another may. A piece of information may be wrong due to a mistake, ignorance, a typographical error, or malice.
- We usually think of **data in a database as being authoritatively correct**, but this isn't possible on the Web, and it won't be possible on the Semantic Web either. To complicate things, **a given source may be reliable in one area but not in another**.

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.

Definizioni di Trust

- **Three definitions for trust** 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

Some questions about the policy-based 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. Context aware 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** 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

- Some questions to address within reputation-based trust:
  - **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!)

Reputation in trust

- Other questions:
  - **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.

Trust and Belief

In realta', alcuni autori distinguono fra trust e belief:

Trust:

- Identity: Who are you?
- Why should I trust you?
- Who else trusts you?
- How much should I trust you?
- How can I know that you said what you've claimed to have said?

Belief:

- How much confidence should I place in what you say?
- What should I believe when different "facts" don't agree?
- How much should my prior beliefs influence my confidence in what you say?
- How can I establish the correct degree of belief for a given set of information?

Of course, the word "you" might refer to an agent, to any other source of information or services, or to any entity that vouches for another's identity.

Trust and Belief

Differenza fra trust e belief:

- Trust could be subsumed under Belief, on the grounds that, for example, trust in a person's identity amounts to a belief that the claimed identity is the actual identity. Thus, the one can be seen as a special case of the other.
- Yet there is a qualitative difference, in that Trust tends to be about quasi-official information—identity, responsibility for statements, and the like—while Belief tends to be about meta data, alleged facts, and opinions.
- According to WordNet, trust (as a verb) means "to have confidence in," whereas belief is "cognitive content held as true." Obviously the boundary is fuzzy, and the terminology in use isn't yet consistent.

Trust

Un'altra questione è la **gestione di informazioni contraddittorie**:

- We should detect the appearance of a contradiction and know what new input caused it. It could then request that the **user make a decision** regarding what to do. A more advanced system might be able to **automatically assess the reliability** of the new information relative to the old and automatically take appropriate action (more on this possibility later).
- Now, if the original statement **should be retracted**, all those additional statements would need to be identified and withdrawn as well. For a large datastore, this task could be extremely lengthy and compute-intensive. This kind of activity is sometimes called **truth maintenance**; active research is going on in this area

Trust

- Ancora peggiore è il caso quando la contraddizione non viene rivelata, però **permane un certo livello di incertezza** sulla validità dell'informazione. In general terms, you look for ratings of the degree of confidence one person has in another. The aim is to arrive at a way to assess what confidence one person would (or should) have in a statement by another. This assessment could become the basis for an automated recommendation. Qualche esempio:
  - **Google's PageRank** approach to rating web pages.
  - **eBay**, maintains ratings of the sellers given by those who have bought goods from them in the past.
  - **Amazon** publishes customer reviews of books and other goods that it sells, and many customers are influenced by these reviews.
  - The **Epinions** web site has collected millions of consumer reviews.
  - **FOAF**, the Friend of a Friend network, could form the basis for a network of selfratings for trust assessments.
- Another school of thought advocates digitally signing as many Semantic Web statements as possible.

Uso del Trust

- Trust in the Semantic Web, should allow **agents and automated reasoners 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

Uso del Trust

Un esempio conclusivo:

- suppose I wish to travel to a foreign country on vacation. My PDA is supposed to handle my arrangements. It asks an airline travel agent (a software agent) to check schedules. Then it asks a booking agent to reserve the seats and purchase the tickets. Next it contacts a passport agent and uses the tickets as proof that a trip has been arranged, and it supplies a certificate as proof of my identity.
- Now, by what authority can my assistant be authorized to pay for the tickets? Why should my credit card company allow the transaction? In addition, the agent for the State Department needs to somehow be sure that the ticket reservations have been made for me and authorized by me.
- These are complicated sequences and we're a long way from having this capability today.
- The term Web of Trust currently maps to the PKI infrastructure, so that it includes all the areas of identity, authentication, and belief. Fortunately, this Web of Trust will be able to evolve