

What is Trust? My Own Point of View

Michel Deriaz

Position Paper

Abstract. What is trust? From an intuitive global definition on how humans perceive trust, to a real and practical implementation, the way is long. Most people stop their quest of an answer at midway, in the field of formal, mathematical and theoretical models. And the path is not unique, each researcher taking a different way. This paper presents the author's own way. Trying to stay as close as possible to the human notion of trust, we discuss also how to take into account the time component, since people give intuitively more importance to recent or scattered events.

1 Introduction

The Merriam-Webster English dictionary defines the trust as “assured reliance on the character, ability, strength, or truth of someone or something”. But everyone has its own definition of the word “trust”. While in general trust refers to an aspect of the relationship between individuals, the term has a completely different meaning depending on the domain it is used. In sociology, for example, we say that a person trusts someone else if he accepts to rely on him. In law, trust is a legal arrangement in which one person manages the property on behalf of another. A trust company is a financial institute, most often a kind of bank. In computer hardware security, trusted computing refers to chips or devices which the user is forced to trust; vulnerabilities in such a component can compromise the overall security of the computer. In a higher level of computer security, trust is a mechanism which aims to connect together entities that behave as expected. In cryptography, trust is the level of certainty that a public key belongs to the correct owner.

All the above definitions are in no way “official” ones. Even if we reduce our scope to computer security, we notice that different authors have also different definitions of what trust is. This paper gives the author's point of view, according to the context in which he works. The aim is to be as close as possible to the human notion of trust, or as close as possible to what the most people would intuitively describe as being trust.

2 Differences between computers and humans

One trusts someone or something if he or it behaves as expected. Computers often use a very simple trust policy. Anyone that is able to give a username and the

corresponding password is considered as trusted. A swindler that gets your password can enjoy freely of all your privileges, as long as he wants. A strange or unusual behavior, like installing key loggers or deleting system files (instead of typically checking for email as soon as the computer starts), won't alarm the system. Once you are considered as trustful (correct password), you remain trustful until you log out.

However, humans use a more sophisticated trust policy when they interact with each others. When you go to a new doctor, you probably trust his competences just because you see the letters "Dr." on his badge. But if your health deteriorates after some visits, you become suspicious and you want to reevaluate the trust value you previously accorded to this doctor. You need a stronger proof of his competences. Perhaps you will search information about its global reputation, by asking a doctors association (do they know him?), or by checking if his name appears somewhere in the yellow pages. Or you will ask your friends if they know about him. We see clearly that the intuitive human trust policy is much more complex than the static and one-time-check policy typically used by computers. Humans use dynamic trust values, computed according to their own observations, past experiences, global reputation, and recommendations given by trusted parties, like friends. Context is also important. You can have a high trust value for a friend in a specific topic, like his know-how in cars, and a very low value for the same person in a different topic, like cooking.

Another important point is the notion of risk. Actually the notions of trust and risk are indissociable. You do not need to trust someone or something if there is no risk. More precisely, the amount of trust you need in a specific situation is directly related to the amount of risk that is involved. If your neighbor asks you 10 € while explaining that he forgot to go to the bank, you will probably lend him the sum. The probability that you get your money back is high, and anyway, the maximum risk is low. But if the same neighbor asks you one million euros, while explaining that he wants to buy a house overseas, you will probably refuse. The risk is clearly too high.

3 Trust policies in practice

Apart from the very basic trust policy for computers, described above, we find in practice mainly two ways to handle trust information: centralized and decentralized.

Centralized is the easiest one. A well-known example is the eBay site [ebay] in which buyer and seller rate each other after every transaction. Information is stored on the server and informs about the global reputation of each entity. The trust model is very simple. It consists just in comparing positive outcomes with negative ones.

The decentralized trust policy is more complex. A well-known example is file sharing in a peer-to-peer network. Peers rate each other and the combination of all the values informs about the reputation of the peer. The challenge here is where to store trust values, as there is no central server. The answer is given by algorithms like Eigentrust [eigentrust], which works even if the peers are anonymous. The idea behind Eigentrust is that each peer has a set of mother peers responsible for storing its trust value, and therefore each peer acts also as a mother peer for others. Eigentrust excludes malevolent as well as collusion of malevolent peers, even when up to 70% of the peers are trying to subvert the system.

Decentralized trust systems are a highly active research topic. We cite here some readings that interest us particularly according to our future research plan.

- The Secure project [secure] aimed to describe in a formal way what trust is, staying as close as possible to the human notion of trust.
- Kinateder and Rothermel [Kinateder2003] present a peer-to-peer system that provides trust and recommendations about different categories of topics. Similar than sites like www.epinions.com or the rating system that we find in eBay, but peer-to-peer.
- The TrustMe protocol [trustMe] builds trust in peer-to-peer network. The trust value of a specific peer is anonymously stored on another peer. Communications are encrypted using sets of private/public keys. The drawback is that all peers have to connect to a bootstrap server when they join and when they leave the network (in order to transmit the hosted trusted values to another peer).
- An interesting system that is similar to Eigentrust, but in which peers stores their own trust value locally, can be found at [p2pRep]. The Elicitation-Storage protocol is used to protect cryptographically the trust value. The requester gets the IP address of the former requesters and checks with them the authenticity of their vote.
- [webOfTrust] presents how P-Grid can be used to implement a distributed PKI infrastructure, enabling c2c (customer to customer) services like eBay but without any centralized system. Unlike PGP that uses the web of trust approach to access a particular public key, this system uses a statistical method; many peers are queried, and the information is rejected if a quorum of peers cannot be obtained.
- A paper that presents a mathematical framework for expressing trust management systems can be found at [trustManSyst].

4 Trust policies in the future

People use more and more electronic devices to assist them in their tasks or for entertainment, but no network is built between devices without human interaction. A transfer of a business card from one PDA (Personal Digital Assistant) to another necessarily requires a manipulation from both end users, even if the operation is actually very simple. The emergence of mobile computing in the next few years will probably bring some changes in the way that computer networks are built and evolve with the time. Trust policies will be affected by these changes. The following hypothetical scenario, taking place in the nearby future, gives a first approach of how we imagine day-to-day live modified by the contribution of ambient intelligence.

4.1 Scenario

Like a majority of his friends, Martin is wearing at his wrist a small electronic device, called an IPDA (Intelligent Personal Digital Assistant). Not bigger than a conventional watch, it is the kind of device that people are wearing all the time. It records as many information as possible about its owner, like preferences and behaviors adopted in the different experiences of his life. After a while, the IPDA knows its owner very much, and is therefore able to autonomously take decisions.

Martin is working as a security consultant and has to travel a lot in order to make conferences and organize classes. His diary has to be very flexible, because some emergencies (like a successful hacker's attack in a company) can loom up suddenly and oblige Martin to reorganize his planning. It is a time consuming task that would borrow Martin's mind (for example not to forget to rent a car, make sure that the conference room will be fitted out with all the needed devices, ...) if he could not just rely on his IPDA to do this job according to his requirements and preferences.

After a phone call, Martin asks verbally his IPDA to organize him a conference in Berlin. He indicates also that it is a high priority request (80/100) and that he feels quite tired. The IPDA re-organizes the planning, cancels the low priority tasks, and because Martin is tired, will book a flight not too early in the morning even if the price is a little bit higher. The IPDA organizes the whole trip (travel, hotel, conference room, cars, restaurants, information mail for his colleagues...) and asks Martin's attention only for crucial decisions that cannot be computed without a reasonable chance of choosing the best solution.

The main idea is that Ambient Intelligence Devices (AID), like IPDA, flight booking systems, or hotel reservations, are able to communicate with each other. For instance, a hotel will not only provide a web site for online booking, but also an AID interface allowing others AID to communicate with it. Every object that is able to communicate in an Ambient Intelligence System is called an AID. We can also imagine that a radiator asks about the preferred temperature of a client directly to his IPDA.

After landing in Berlin, Martin's IPDA guides him to his car, opens it (the AID of the car recognizes the digital signature of the client's IPDA), and computes the best itinerary to the hotel. Once arrived, Martin decides he wants actually eat Chinese tonight. His IPDA finds in its database that a good friend of Martin, Jordan, who loves Asian food, often travels in Berlin. It asks this friend's IPDA for advice. The "Peking" seems to be one of the bests. Martin's IPDA then connects to the restaurant in order to get the opening hours and a price list. It trusts more Martin's friend about the quality of the food, but trusts more the restaurant about administrative information that can change over the time. The last pieces of information are gathered from a web site that registers traveler comments about good and bad experiences they lived during their journey. After computing a different trust value for each source of information, a final calculation confirms that the "Peking" will probably be the best restaurant for Martin at this time.

Traveling seems to make Martin very hungry. He ordered the "Peking menu" made of two starters, one main course and a dessert. Once full, he tells his IPDA that he is really happy with this restaurant. This information is recorded as a pleasant

experience and a good tip is computed. Every IPDA runs an e-purse software, and payments are made with e-cash.

The next morning, Martin gets up and, after breakfast, turns on the terminal that is in his room. For security reasons (Trojan horses), there are no hard-disks on such devices, clients use their own one which is more and more often their IPDA. Martin checks his e-mails, the latest news and then makes some final corrections to the slides he wants to present during his conference which begins in two hours. To avoid stress, Martin's best remedy is shopping. He lets his IPDA to direct him towards the nearest shopping centre, and then heads for the music department. After a while, Martin is reminded (from his IPDA) that time is running out. He heads for the exit, catches some chocolate on his way, and walks through the payment gate. Martin does not need to queue up at the traditional till. Every item is labeled with a RFID (Radio Frequency Identification) tag and payment is made automatically by the e-purse when its owner walks through the payment gate.

At the bus stop, Martin gets tempted by a coffee machine. His IPDA consults the virtual notice board attached to it. It is very common for such kind of machines to have a virtual notice board, on which users' IPDA write their experience. In our case, we observe that the machine is working on average only 50% of the time. But we see also that the last four clients enjoyed a positive outcome. According to this and to the limited amount of risk involved, Martin's IPDA accepts the financial transaction without asking its owner's authorization.

Relaxed, Martin joins the conference room, launches the welcome slide of his presentation using the speaker's terminal, and goes to the main door in order to personally greet every participant.

4.2 Scenario discussion

Let us consider the following points (hypothetical):

- There are different trust values for the same person. Martin accords different trust values for his friend Jordan. He gives 80% when they are talking about Chinese food, but only 30% when the subject is politics.
- Trust is transitive. Martin accords 70% of trust (about Chinese food) to the former girlfriend of Jordan, despite the fact he never saw her. He just knows that Jordan accorded 90% of trust to her for that topic, and therefore he computed his own value of trust.
- Trust evolves with time. If Martin is happy about a restaurant, he will increase the trust to the devices that recommended this place, or decrease it if he comes out disappointed.
- Decisions are not only taken regarding trust. The risk of every action is computed and combined to the trust value. The risk represents actually the maximal cost of an operation that fails. The context is very important in the evaluation of these values. For example Martin accepts to go in a restaurant if the chance of being happy is 75%, but he refuses to make a financial transaction if the chance of not being hacked is also 75%.

- More importance is given to recent events. In the example of the coffee machine, Martin's IPDA accepted the transaction despite the fact that the machine works on average only 50% of the time. Intuitively, if it worked well a few minutes ago, then the probability that it will work well ones more is high.

As machines imitate human social behaviors, it seems clear that trust policies will also have to follow this tendency. We need policies that can be used between machines themselves as well as between humans and machines. Like human to human interactions, our devices are meant to meet unknown ones. Trust will be built and constantly updated according to the outcomes of these interactions [secureD3.2][secure].

5 Trust models

Trust is well understood by humans, but seems to be very difficult to model. We can of course take a very simple model in which we compute the trust value T like:

$$T = \frac{m}{m+n} \quad (1)$$

in which we compute m and n as follow:

$$\begin{aligned} m &= |PO| \\ n &= |NO| \end{aligned} \quad (2)$$

where $|PO|$ is the number of positive outcomes and $|NO|$ the number of negative outcomes. But this first model does not take into account time. We will try to modify our simple model in order to give more importance to recent events (like humans would do) and to give more importance to dispersed events (a regrouping of negative outcomes could be a temporally breakdown). A first proposition of our new model computes m and n as follow:

$$\begin{aligned} m &= a \cdot |PO| + \frac{b}{T} \sum_{i=0}^{i=|PO|} t(po_i) + c \cdot |PO| \cdot \delta_p \\ n &= a \cdot |NO| + \frac{b}{T} \sum_{i=0}^{i=|NO|} t(no_i) + c \cdot |NO| \cdot \delta_n \end{aligned} \quad (3)$$

where:

- a , b and c are parameters $\in [0..1]$, with a default value of 1.
- T is the current time.
- $t(po_i)$ is the time when positive outcome i occurred.
- $t(no_i)$ is the time when negative outcome i occurred.
- δ_p is the dispersion of the positive outcomes.

- δ_n is the dispersion of the negative outcomes.

Fig. 1 shows an example of positive outcomes (PO) and negative ones (NO).

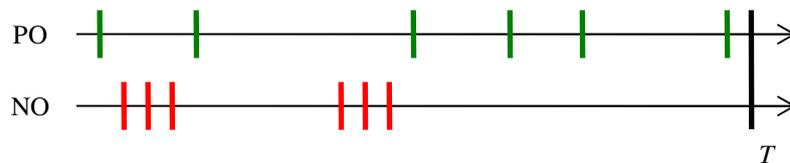


Fig. 1: Timeline with positive and negative outcomes

Since there is the same number of positive outcomes than negative ones, our first basic model gives a trust value of 50% (Fig. 1), But intuitively we would probably give a higher trust value. First because the last outcomes are positive. Fig. 1 represents the virtual notice board of the coffee machine (see the scenario in 4.1). Martin's IPDA accepts the transaction because the last outcomes were positive. If the machine was correctly working a few times ago, so it should probably also work now (at time T). Secondly because the dispersion of negative outcomes is much smaller; it means that the machine sometimes falls down, but that it works quite well otherwise.

Our new model takes care of these observations. In the computation of m and n , the second term gives more importance to events that occurred near T , and the third one gives more importance to events that are scattered.

It is clear that this model is only a simple example. As we wrote in some sections above, the trust policy used by humans is very context dependent. We notice also that this model does not take into account the risk that is involved. However, by choosing correctly the different parameters a , b and c , we can still fit to some every day life situations.

6 Conclusion

We started this paper by presenting an overview of the trust concept. Then we proposed a simple model that computes trust information in a similar way than humans. We underlined that the time component is very important, since recent and scattered events are given more weight during an intuitive human trust computation.

Our future work consists in including trust information in LBS (Location Based Services). Examples of LBS include virtual tags. The idea behind this concept, also called spatial messaging or air graffiti, is to allow a mobile user equipped with a location system to place at his current position or in the neighborhood a virtual tag. Everyone that enters the visibility area of this tag receives it. An interesting question is how to trust the content of these tags, and how to handle the time component, since a tag posted a long time ago can contain outdated data.

References

- [secureD3.2] Jean Bacon, András Belokisztolszki, Daniel Cvrcek, Nathan Dimmock, David Eyers, David Ingram, Ken Moody. Preliminary Definition of a Trust-based Access Control Model. SECURE Deliverable 3.2, 17th February 2004. <http://secure.dsg.cs.tcd.ie/>
- [secure] V. Cahill, et al. Using Trust for Secure Collaboration in Uncertain Environments. IEEE Pervasive Computing Magazine, July-September 2003.
- [ebay] eBay website, <http://www.ebay.com/>.
- [eigenTrust] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The Eigen-Trust Algorithm for Reputation Management in P2P Networks. 2003.
- [Kinateder2003] Michael Kinateder, Kurt Rothermel. Architecture and Algorithms for a Distributed Reputation System. 2003.
- [trustMe] Aameek Singh, Ling Liu. TrustMe: Anonymous Management of Trust Relationships in Decentralized P2P Systems. Proceedings of the Third International Conference on Peer-to-Peer Computing (P2P'03). IEEE.
- [p2pRep] Prashant Dewan. Peer-to-Peer Reputations. Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS'04) IEEE.
- [webOfTrust] Anwitaman Datta, Manfred Hauswirth, Karl Aberer. Beyond “web of trust”: Enabling P2P E-commerce. Proceedings of the IEEE International Conference on E-Commerce (CEC'03).
- [trustManSyst] Stephen Weeks. Understanding Trust Management Systems.