Privacy and Self-disclosure in Multiagent Systems

(Extended Abstract)

Jose M. Such
Departament de Sistemes Informàtics i Computació
Universitat Politècnica de València
Camí de Vera s/n, València, Spain
jsuch@dsic.upv.es

ABSTRACT

Agents usually encapsulate their principals' personal data attributes, which can be disclosed to other agents during agent interactions, producing a potential loss of privacy. We propose self-disclosure decision-making mechanisms for agents to decide whether disclosing personal data attributes to other agents is acceptable or not. Moreover, we also propose secure agent infrastructures to protect the information that agents decide to disclose from undesired accesses.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Theory, Design, Experimentation, Security

Keywords

Privacy, Intimacy, Identity, Disclosure

1. INTRODUCTION

Nowadays, in the era of global connectivity (everything is inter-connected anytime and everywhere) with almost 2 billion users with connection to the Internet as of 2010[1], privacy is of great concern. Recent studies show that only 8% of users are unconcerned about privacy [9]. Moreover, almost 95% of web users admitted they have declined to provide personal information to web sites at one time or another when asked [3].

2. MOTIVATION

Autonomous agents play a crucial role to safeguard and preserve their principals' privacy. This is because agents encapsulate personal information of their principal [1]. They usually have a detailed profile of their principal's names, preferences, roles in organizations and institutions, location, transactions performed, and other personal information.

[1] `http://www.internetworldstats.com/stats.htm`

Cite as: Privacy and Self-disclosure in Multiagent Systems (Extended Abstract), Jose M. Such, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1333-1334.
Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Westin [8] defined privacy as a "personal adjustment process" in which individuals balance "the desire for privacy with the desire for disclosure and communication". Westin proposed his definition for privacy long before the explosive growth of the Internet. We consider that it also applies to autonomous agents that engage in online interactions. Agents carry out interactions on behalf of their owners so that they usually exchange personal information of their principals. This may raise privacy concerns, because this exchange of personal information can produce a potential loss of privacy. Thus, agents need self-disclosure decision-making mechanisms to decide whether disclosing personal data attributes to other agents is acceptable or not. Once an agent has decided which information to disclose to what other agent, this information should be protected from undesired accesses. This includes the ability of disclosing information about their principals without disclosing their principals' identities if they decide so.

3. CONTRIBUTIONS

3.1 Self-disclosure Decision Making

Current self-disclosure decision-making mechanisms are based on the privacy-utility tradeoff ([4]). This tradeoff considers the direct benefit of disclosing personal information and the privacy loss it may cause; for instance, the tradeoff between the reduction in time to perform an online search when personal information (e.g. geographical location) is given and the privacy loss due to such disclosure [4].

There are many cases where the direct benefit of disclosing personal information is not known in advance. This is the case in human relationships, where the disclosure of personal information in fact plays a crucial role in the building of these relationships [2]. These relationships may or may not eventually report a direct benefit for an individual. For instance, a close friend tells you what party he voted for. He may disclose this information without knowing (or expecting) the future gain in utility this may cause. Indeed, it might not report him any benefit ever.

We propose a self-disclosure decision-making model based on intimacy and privacy measures to deal with these situations [7]. Our model considers psychological findings regarding how humans disclose personal information in the building of their relationships, such as the well-studied *disclosure reciprocity* phenomenon [2]. This phenomenon is based on the fact that one person's disclosure encourages the disclosure of the other person in the interaction, which in turn, encourages more disclosures from the first person.

Intimacy accounts for the information gain of all the messages received from another agent. Privacy accounts for the information loss caused by sending a message valuated with the sensitivity of the information disclosed. Agents may choose to disclose information that maximizes the estimation of the increase in intimacy while at the same time minimizing the privacy loss. Moreover, they consider how balanced their relationships are, i.e., they may decide not to perform disclosures to agents that do not reciprocate them with more disclosures (following the reciprocity phenomenon).

3.2 Secure Agent Platform

Once an agent has decided which information to disclose to which other agent, this information must be protected from accesses from any other third parties different from the agent to which the information is directed to. This includes parties from their local computer and network but also different locations, even across the Internet. We contribute a secure Agent Platform (AP) that allow agents to interact to each other in a secure fashion [5]. To this aim, our secure AP provides authorization mechanisms based on mandatory access control (agents are confined to access a subset of their principals' permissions), and encryption and decryption of messages exchanged based on Kerberos[2].

Moreover, our secure AP allows agents to authenticate to each other without disclosing their principals' identities. Agents have their own identities that act as pseudonyms for their principals. Our secure Agent Platform keeps track of the association between principal and agent identities. Therefore, principal identities can be obtained for accountability concerns, such as law enforcement.

3.3 Privacy-enhancing Agent Identity Management

Our secure AP keeps track of the agent's principal identity and its association to the agent identity. Thus, the AP itself can be a privacy threat for the principals running agents on top of it. Moreover, agents need to selectively disclose personal data attributes in their identity to other agents following our proposed self-disclosure decision making. This includes the necessity of allowing more than one identity per agent to be used in different disclosures (or different contexts). Thus, different disclosures (in possible different contexts) can remain unlinkable to each other if desired.

We propose an Identity Management Model for Multiagent Systems to enhance the privacy of agent's principals [6]. Our model is based on current Privacy-enhancing Identity Management Systems and uses partial identities as a key concept for identifying entities (agents and principals). In a nutshell and informally speaking, a partial identity can be seen as a pseudonym and a set of attributes attached to it. Our model allow agents to have multiple partial identities and define access control rights for other agents to the attributes in them. Agents can define these rights based on our self-disclosure decision-making model.

In Privacy-enhancing Identity Management Systems, partial identities are issued by Identity Providers (IdPs). In our model, agents must provide their principal's identity, or an existing partial identity to obtain new partial identities. IdPs do not make this association publicly known, but can disclose it if required by a court. Agents can register in an AP using a partial identity. Therefore, agent identity management is decoupled from the system where identities are used, increasing the privacy of principals.

[2] `http://web.mit.edu/kerberos/`

4. FUTURE WORK

We claim that agents following our self-disclosure decision-making model lose less privacy than agents that do not use them when disclosing personal information to other agents. We now want to prove this claim experimentally. To this aim, we are performing experiments comparing agents using these self-disclosure decision-making mechanisms with privacy unconcerned agents that do not use them. We consider environments in which there are different percents of malicious agents, from 0% to 100% of malicious agents. We consider malicious agents to be agents that are only interested in obtaining information from other agents without increasing intimacy, i.e., they do not provide information about themselves or if they do, they lie about themselves.

We are also exploring strategies for agents not to be sincere when disclosing a PDA. This could be useful once these agents detect that they are interacting with malicious agents. They could choose to keep on disclosing PDAs while being insincere instead of not disclosing any other PDA to such malicious agents. Thus, using such strategies agents would be able to lie to liars.

Acknowledgments

This work has been partially supported by CONSOLIDER-INGENIO 2010 under grant CSD2007-00022 and projects TIN2008-04446 and TIN2009-13839-C03-01 of the Ministerio de Ciencia e Innovación de España. The author would also like to thank his thesis advisors Ana Garcia-Fornes and Agustin Espinosa.

5. REFERENCES

[1] M. Fasli. On agent technology for e-commerce: trust, security and legal issues. *Knowl. Eng. Rev.*, 22(1):3–35, 2007.

[2] K. Green, V. J. Derlega, and A. Mathews. *The Cambridge Handbook of Personal Relationships*, chapter Self-Disclosure in Personal Relationships, pages 409–427. Cambridge University Press, 2006.

[3] D. Hoffman, T. Novak, and M. Peralta. Building consumer trust online. *Commun. ACM*, 42(4):80–85, 1999.

[4] A. Krause and E. Horvitz. A utility-theoretic approach to privacy and personalization. In *AAAI*, pages 1181–1188. AAAI Press, 2008.

[5] J. M. Such, J. M. Alberola, A. Espinosa, and A. Garcia-Fornes. A group-oriented secure multiagent platform. *Softw., Pract. Exper.*, page In Press., 2011.

[6] J. M. Such, A. Espinosa, A. Garcia-Fornes, and V. Botti. Partial identities as a foundation for trust and reputation. *Eng. Appl. of AI*, page In Press., 2011.

[7] J. M. Such, A. Espinosa, A. Garcia-Fornes, and C. Sierra. Privacy-intimacy tradeoff in self-disclosure. In *AAMAS*, page In press. IFAAMAS, 2011.

[8] A. Westin. *Privacy and Freedom*. New York Atheneum, 1967.

[9] A. Westin. Social and political dimensions of privacy. *Journal of Social Issues*, 59(2):431–453, 2003.