Privacy policies and their negotiation in distance education
Yee, George; Korba, Larry

National Research Council Canada   Conseil national de recherches Canada

Canada

National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC·CNRC

*Privacy Policies and their Negotiation in Distance Education **

Yee, G., and Korba, L.
2004

Canada

Privacy Policies and their Negotiation in Distance Education[1]

George Yee

Tel : (613) 990-4284

Fax: (613) 952-7151

George.Yee@nrc-cnrc.gc.ca

Larry Korba

Tel : (613) 998-3967

Fax : (613) 952-7151

Larry.Korba@nrc-cnrc.gc.ca

National Research Council Canada

Institute for Information Technology

1200 Montreal Road, Building M-50

Ottawa, Ontario, Canada K1A 0R6

Privacy Policies and their Negotiation in Distance Education[1]

ABSTRACT

This chapter begins by introducing the reader to privacy policies, e-services, and privacy management. It then derives the contents of a privacy policy and explains "policy matching". It next presents an approach for the negotiation of privacy policies for an e-learning service. Both negotiating under certainty and uncertainty are treated. The type of uncertainty discussed is uncertainty of what offers and counter-offers to make during the negotiation. The approach makes use of common interest and reputation to arrive at a list of candidates who have negotiated the same issues in the past, from whom the negotiator can learn the possible offers and counter-offers that could be made. Negotiation in this work is done through human-mediated computer-assisted interaction rather than through autonomous agents. The chapter concludes with a discussion of issues and future research in this area.

INTRODUCTION

Most distance education innovations have focused on course development and delivery, with little or no consideration to privacy and security as required elements. However, it is clear that there will be a growing need for high levels of confidentiality and privacy in e-learning applications, and that security technologies must be put in place to meet these needs. The savvy of consumers regarding their rights to privacy is increasing; new

privacy legislations have recently been introduced by diverse jurisdictions (Canadian Standards Association, Department of Justice). In addition, the move to corporate outsourcing of distance learning will lead to requirements of confidentiality of student information, to protect company sensitive information that might be disclosed if training records were obtained by competitors.

A promising solution to the lack of privacy and security for e-learning systems is to put in place a policy-based management system, i.e. formulate privacy and security policies for the e-learning system and back them up with security mechanisms that ensure that the policies are respected. Policy-based management approaches have been used effectively to manage and control large distributed systems. As in any distributed system, e-learning may also use a policy-based framework to manage the security and privacy aspects of operations. However, policies must reflect the wishes of the e-learning consumer as well as the e-learning provider. In this chapter, we introduce privacy policies, e-services, privacy policy management, and describe an approach for the negotiation of privacy policies between an e-learning consumer and an e-learning provider. We examine negotiation under certainty and uncertainty (where the offers and counter-offers are known or unknown, respectively) and propose a scheme for resolving the uncertainty using the experience of others who have undergone similar negotiation. The choice of whom to call upon for negotiation experience is resolved through the identification of common interest and reputation. The results of this chapter are applicable to all types of e-services, including e-business and e-learning.

The negotiation approach presented in this chapter does not employ autonomous agent negotiation (AAN). We find that: a) AAN is not necessary for our application area, b) current AAN technology would be unable to capture all the nuances and sensitivities involved with privacy policy negotiation, including cultural impacts (Kersten, G. et al, 2002), and c) the level of trust that consumers would have in autonomous agents negotiating privacy policy would be low.

In the literature, most negotiation research is on negotiation via autonomous software agents. This research focuses on methods or models for agent negotiation (Huang, P. and Sycara, K., 2002; Lopes, F. et al, 2001; Benyoucef, M. et al, 2001) and can incorporate techniques from other scientific areas such as game theory (Murakami, Y. et al, 2001), fuzzy logic (Lai, R. and Lin, M., 2002; Kowalczyk, R. and Bui, V., 2000) and genetic algorithms (Tu, M. et al, 2000). The research also extends to autonomous agent negotiation for specific application areas, such as e-commerce (Chung, M. and Honavar, V., 2000; Limthanmaphon, B. et al, 2000) and service level agreements for the Internet (Nguyen, T. et al, 2002). More recently, Chiu et al (2003) and Kim et al (2003) have written on e-negotiation, a systematic web services supported way of negotiating contracts over the Internet for e-commerce. Apart from these works on negotiation, research has also been carried out on support tools for negotiation (Boehm, B. et al, 2001; Druckman, D. et al, 2002), which typically provide support in position communication, voting, documentation communication, and big picture negotiation visualization and navigation.

Regarding privacy policies, there are related works such as P3P (W3C), APPEL (W3C, 2002), and PSP (Carnegie Mellon University), which provide ways of expressing privacy policy and preferences. Service providers use P3P to divulge their privacy policies to consumers. APPEL is a specification language used to describe a consumer's privacy preferences for comparison with the privacy policy of a provider. PSP is a protocol in the research stage that provides a basis for policy negotiation. These works are not necessary for the purposes of this chapter. They only serve as illustrations of what has been done in the related area of capturing privacy preferences in a form amenable to machine processing.

Our work differs from P3P, APPEL, and PSP in that we look at privacy legislation and other regulations in order to derive a core set of privacy attributes that can be used for expressing the content of any consumer privacy policy. P3P, APPEL, and PSP, on the other hand, are not concerned with privacy policy content, but are concerned instead with how policy content can be expressed in machine-processable form, once it is determined what the content should be. We are not aware of any other work that looks at privacy policy content in this way.

We next give more of the background to this chapter by introducing the reader to privacy legislation, e-services (of which distance education is one example), and privacy management.

5

Privacy Legislation and Directives

In Canada, privacy legislation is enacted in the *Personal Information and Electronic Documents Act* (Department of Justice) and is based on the Canadian Standards Association's *Model Code for the Protection of Personal Information* (Canadian Standards Association), recognized as a national standard in 1996. This Code consists of ten Privacy Principles (Canadian Standards Association) that for convenience, we label as CSAPP. In addition, the Canadian healthcare sector has published *Principles for the Privacy Protection of Personal Health Information in Canada* (Privacy Working Group), describing privacy principles specifically for health information. Again for convenience, we label these as CHSPP. We will examine both these sets of principles below. Data privacy in the European Union is governed by a very comprehensive set of regulations called the Data Protection Directive (European Union, 1995). In the United States, privacy protection is achieved through a patchwork of legislation at the federal and state levels. However, privacy has been recognized as a constitutional right and there exists a highly developed system of privacy protection under tort law for the past century (Industry Canada).

E-Service Model

It is useful to describe what we mean by e-learning being an e-service. An e-service is a service that is offered by a provider to a consumer across a computer network. A stock quotation service is often used as an example of an e-service. Here a consumer would logon to the service from a computer, and after appropriate user authentication, would make use of the service to obtain stock quotes. Accessing one's bank account through online banking is another example of an e-service. Here the provider is the bank and the

service consists of allowing the consumer to check the balance, transfer funds, or make bill payments. The network is usually the Internet, but could also in principle be a private enterprise network. At any point in time, one provider may be serving many consumers and many providers may be serving one consumer. For the purposes of this chapter, the business relationship between provider and consumer is always one-to-one, i.e. the service is designed for one consumer and is provided by one provider, payment for the service is expected from one consumer. In addition, service providers may also be service consumers, and service consumers may also be service providers.

Privacy Management Model

An effective and flexible way to protect privacy is to manage it using privacy policies. A provider has a privacy policy stating what private information it requires from a consumer and how the information will be used. A consumer has a privacy policy stating what private information the consumer is willing to share, with whom it may be shared, and under what circumstances it may be shared. An entity that is both a provider and a consumer has separate privacy policies for these two roles. A privacy policy is attached to a software agent that acts for a consumer or a provider as the case may be. Prior to the activation of a particular service, the agent for the consumer and the agent for the provider undergo a privacy policy exchange, in which the policies are examined for compatibility. The service is only activated if the policies are compatible (i.e. there are no conflicts), in which case we say that there is a "match" between the two policies (Section "Privacy Policies" defines a "match"). In addition, *we assume that in general, the provider always asks for more private information from the consumer than the consumer is willing to give up.* Figure 1 illustrates our privacy management model. For the purposes

of this chapter, it is not necessary to consider the details of service operation. However, the provider is expected to comply with the consumer's privacy policy if service is initiated. We discuss policy compliance in Section "Privacy Policies".
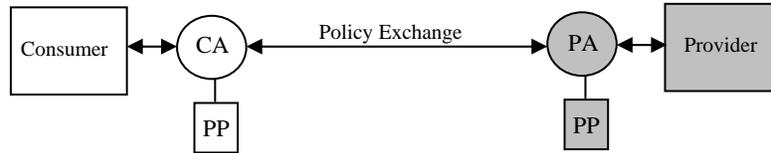
Figure 1- Exchange of Privacy Policies (PP) Between Consumer Agent (CA) and Provider Agent (PA)

The remainder of this chapter is organized as follows. Section "Privacy Policies" examines the specification of privacy policies by identifying some attributes of private information collection, using the privacy principles from the Canadian Standards Association and the Canadian health sector as guides. Section "Negotiation – Structure and Representation" considers the mathematical structure of negotiation. Section "Negotiation in Certainty and Uncertainty" examines negotiation under certainty and uncertainty. For the latter case, we explore using the experience of others in making decisions. Section "Scheme for Negotiating Privacy Policy Under Uncertainty" gives a scheme for negotiating a privacy policy under uncertainty. Section "Conclusions, Discussion, and Future Research" presents a discussion of issues and future research together with our conclusions.

PRIVACY POLICIES

Requirements from Privacy Principles

In this section, we identify some attributes of private information collection using CSAPP and CHSPP as guides. We will then apply these attributes to the specification of privacy policy contents. We choose to use CSAPP because it is representative of principles behind privacy legislation in many countries, including the European Union. As a result, our requirements will be applicable in many countries. We include CHSPP to reflect the fact that there is great concern over the privacy of healthcare information (Privacy Working Group, Kumekawa). In fact, many consumers consider their health information as more "private" than other personal information such as financial status. This may be due to the fact that involuntary disclosure of certain critical health conditions may affect their careers or their health insurance (Kumekawa). Healthcare information is so sensitive and personal that it merits special protection by healthcare providers and legislators (Privacy Working Group). Tables 1 and 2 show CSAPP and CHSPP respectively.

We interpret "organization" as "provider" and "individual" as "consumer". In the following, we use CSAPP.n or CHSPP.n to denote Principle n of CSAPP or CHSPP respectively. Principle CSAPP.2 implies that there could be different providers requesting the information, thus implying a *collector* attribute. Since on a provider policy, the collector is always the provider, *collector* is only used in consumer privacy policies. Principle CSAPP.4 implies that there is a *what* attribute, i.e. what private information is being collected. Principles CSAPP.2, CSAPP.4, and CSAPP.5 state that there are *purposes* for which the private information is being collected. Principles CSAPP.3,

CSAPP.5 and CSAPP.9 imply that the private information can be disclosed to other parties, giving a *disclose-to* attribute. Principle CSAPP.5 implies a *retention time* attribute for the retention of private information. Thus, from the CSAPP we derive 5 attributes of private information collection, namely *collector*, *what*, *purposes*, *disclose-to*, and *retention time*.

The CHSPP provides much of the same privacy protection as the CSAPP. However, the CHSPP adds the notions of "proxy" (CHSPP.3) and "identifiable" (CHSPP.10, CHSPP.11) due to the nature of healthcare. We treat "proxy" by adding a proxy field to the privacy policy. For a consumer policy, the proxy field holds the name of the proxy if a proxy provided the information. Otherwise, this field has the default content of "no". For a provider policy, the proxy field has a default value of "yes" to indicate that the private information may be provided by a proxy. Otherwise, this field has the value "no". The address and telephone contact of the proxy may be specified as an informational item in the privacy policy itself. We treat "identifiable" by requiring the name, address, and telephone number of the healthcare information owner to be placed as an informational item in the policy.

Table 1 - CSAPP - The Ten Privacy Principles from the Canadian Standards Association

Principle	*Description*
1. Accountability	An organization is responsible for personal information under its control and shall designate an individual or individuals accountable for the organization's compliance with the privacy principles.
2. Identifying Purposes	The purposes for which personal information is collected shall be identified by the organization at or before the time the information is collected.
3. Consent	The knowledge and consent of the individual are required for the collection, use or disclosure of personal information, except when inappropriate.
4. Limiting Collection	The collection of personal information shall be limited to that which is necessary for the purposes identified by the organization. Information shall be collected by fair and lawful means.
5. Limiting Use, Disclosure, and Retention	Personal information shall not be used or disclosed for purposes other than those for which it was collected, except with the consent of the individual or as required by the law. In addition, personal information shall be retained only as long as necessary for fulfillment of those purposes.
6. Accuracy	Personal information shall be as accurate, complete, and up-to-date as is necessary for the purposes for which it is to be used.
7. Safeguards	Security safeguards appropriate to the sensitivity of the information shall be used to protect personal information.
8. Openness	An organization shall make readily available to individuals specific information about its policies and practices relating to the management of personal information.
9. Individual Access	Upon request, an individual shall be informed of the existence, use and disclosure of his or her personal information and shall be given access to that information. An individual shall be able to challenge the accuracy and completeness of the information and have it amended as appropriate.
10. Challenging Compliance	An individual shall be able to address a challenge concerning compliance with the above principles to the designated individual or individuals accountable for the organization's compliance.

The Privacy Principles also prescribe certain operational requirements that must be satisfied between provider and consumer, such as identifying purpose and acquiring consent. Our service model and the exchange of privacy policies automatically satisfy some of these requirements, namely Principles CSAPP.2, CSAPP.3, CSAPP.8, CHSPP.2, CHSPP.3, and CHSPP.10. The satisfaction of the remaining operational requirements depends on compliance mechanisms (Principles CSAPP.1, CSAPP.4,

Table 2 - CHSPP - The Privacy Principles from the Canadian Health Sector

Principle	Description
1. Overarching (affects all principles)	An individual's right of privacy of personal health information is paramount; however, it is not absolute. This right is subject to reasonable limits, *prescribed by law, as can be demonstrably justified in a free and democratic society.*
	Principles for Individuals
2. Privacy	Individuals have a right of privacy with respect to their personal health information.
3. Consent	Individuals have the right to provide or withhold consent with respect to the collection, use, disclosure, or access of their personal health information. Consent may be from the individual's designated proxy, where the individual is incompetent or lacking in decision-making capacity.
4. Knowledge	Individuals have a right of knowledge with respect to their personal health information.
5. Individual Access	Individuals have the right to access their own personal health information.
6. Accuracy	Individuals have the right to have their personal health information recorded as accurately as possible and to review and amend their health records to ensure accuracy.
7. Recourse	Individuals have the right to recourse when they suspect a breach in the privacy of their health information.
	Principles for Providers and Organizations
8.Confidentiality	Providers and organizations have an obligation to treat personal health information as confidential.
9. Trusteeship and Accountability	Providers and organizations entrusted with personal health information have an obligation to safeguard the privacy of individuals and the confidentiality of this information.

10. Access and Use - Identifiable Health Information	a) To provide direct care to individuals, providers and health care organizations should have access to identifiable health information.
	b) Identifiable health information shall only be used with the consent of the individual, except in extraordinary circumstances where there is:
	  • a demonstrated legal requirement; or
	  • compelling evidence for individual or societal good and a privacy impact assessment that are adjudicated by an independent body according to strict protocols.

11. Access and Use - De-Identified Health Information	a) Access to and use of *de-identified* information should be available to improve population health status and to improve the effectiveness and efficiency of the health system. b) Disclosure, collection and use of personal health information for purposes such as billing, research, evaluation and quality assurance activities should be restricted to de-identified information unless the user can demonstrate why identifiable information is required.
12. Security	Security safeguards must be in place to protect the integrity and confidentiality of health information.
13. Implementation and Enforcement	Providers and organizations should implement policies, procedures and practices to achieve privacy protection.

CSAPP.5, CSAPP.6, CSAPP.9, CSAPP.10, CHSPP.4, CHSPP.5, CHAPP.6, CHSPP.7, CHSPP.8, CHSPP.9, CHSPP.11, and CHSPP.13) and security mechanisms (Principles CSAPP.7 and CHSPP.12). We discuss compliance mechanisms below. Security mechanisms are outside the scope of this chapter.

Privacy Policy Specification

Based on the above exploration, the contents of a privacy policy should, for each item of private information, identify a) *collector* - who wishes to collect the information (for consumer policies only), b) *what* - the nature of the information, c) *purposes* - the purposes for which the information is being collected, d) *disclose-to* – the parties to whom the information will be disclosed, and e) *retention time* – the amount of time for the provider to keep the information. Since a privacy policy may change over time, we add a *valid* field to hold the time period during which the policy is valid. Privacy policies across different types of e-services (e.g. e-business, e-learning, e-health) are specified using these attributes (see examples below) and principally differ from one another according to the values of the attributes *what* and *purposes*. For example, an e-commerce privacy policy might specify credit card number as *what* and payment as *purposes* whereas an e-learning privacy policy might specify marks as *what* and student assessment as *purposes*.

Figure 2 (top) gives examples of provider privacy policies from 3 types of providers: an e-learning provider, an e-commerce provider, and a nursing practitioner who uses the Internet to obtain referrals. Figure 2 (bottom) gives corresponding example

consumer privacy policies. These policies need to be expressed in a machine-readable

policy language such as APPEL (W3C, 2002) (XML implementation).

The Matching of Privacy Policies

When consumer and provider agents exchange privacy policies, each agent examines the

other's policy to determine if there is a match between the two policies. If either agent

Privacy Policy: *E-learning* Owner: *E-learning Unlimited* Proxy: *No* Valid: *unlimited*	Privacy Policy: *Book Seller* Owner: *All Books Online* Proxy: *No* Valid: *unlimited*	Privacy Policy*: Medical Help* Owner: *Nursing Online* Proxy: *Yes* Valid: *unlimited*
What: name, address, tel *Purposes:* identification *Retention Time:* unlimited *Disclose-To*: none *What:* Course Marks *Purposes:* Records *Time:* 1 year *Disclose-To*: none	*What:* name, address, tel *Purposes:* identification *Retention Time:* unlimited *Disclose-To*: none *What:* credit card *Purposes:* payment *Time:* until payment complete *Disclose-To*: none	*What:* name, address, tel *Purposes:* contact *Retention Time:* unlimited *Disclose-To*: pharmacy *What:* medical condition *Purposes:* treatment *Time:* 1 year *Disclose-To*: pharmacy

Privacy Policy: *E-learning* Owner: *Alice Consumer* Proxy: *No* Valid: *unlimited*	Privacy Policy: *Book Seller* *Owner:* Alice Consumer Proxy: *No* Valid: *June 2003*	Privacy Policy: *Medical Help* *Owner:* Alice Consumer Proxy: *No* Valid: *July 2003*
Collector: Any *What:* name, address, tel *Purposes:* identification *Retention Time:* unlimited *Disclose-To*: none *Collector:* Any *What:* Course Marks *Purposes:* Records *Retention Time:* 2 years *Disclose-To*: none	*Collector:* Any *What:* name, address, tel *Purposes:* identification *Retention Time:* unlimited *Disclose-To*: none	*Collector:* Any *What:* name, address, tel *Purposes:* contact *Retention Time:* unlimited *Disclose-To*: pharmacy *Collector:* Dr. A. Smith *What:* medical condition *Purposes:* treatment *Retention Time:* unlimited *Disclose-To*: pharmacy

Figure 2 - Example Provider Privacy Policies (top) and Corresponding Consumer Privacy Policies (bottom)

finds a match, that agent signals a match to the other agent, and service is initiated. If either agent fails to find a match, that agent would signal a mismatch to the other agent and service would then not be initiated; this mismatch also signals the start of negotiation to resolve the mismatch. We next define some needed terminology, including the meaning of a match between two privacy policies.

Since a privacy policy is made up of informational items that have attributes *what*, *purposes*, *retention time*, and *disclose-to*, let these attributes be represented by $w_{i,c}$, $p_{i,c}$, $r_{i,c}$, and $d_{i,c}$ respectively for consumer informational item *i*. Similarly, we have $w_{i,p}$, $p_{i,p}$, $r_{i,p}$, and $d_{i,p}$ for provider informational item *i*. We wish to ascribe a function *pr* that returns a numerical degree of privacy from the consumer's point of view when applied to the attributes *what*, *purposes*, *retention time*, and *disclose-to.* A high *pr* means a high degree of privacy; a low *pr* means a low degree of privacy – from the consumer's point of view. It is difficult to define *pr* universally because privacy is a subjective notion, and one consumer's view of degree of privacy may be different from another consumer's view. We therefore require consumer input at policy exchange time to determine which informational item is more private. In other words, we ask the consumer to assign a numerical value to *pr*, say from a scale of 1 to 3 (the scale may be increased – it depends on how acceptable it is to the consumer), corresponding to low, medium, high respectively for each attribute *what*, *purposes*, *retention time*, and *disclose-to* for all information items in both policies. After the consumer completes this input, we will have almost all the *pr* values needed for the definition of a match. We need a set of *pr* values for *collector* in the consumer's policy in order to distinguish the case where the consumer

does not care who collects the information and the case where she does care. Let c_i represent the collector of informational item i in the consumer's policy. We set $pr(c_i) = 0$ if c_i has value "any", otherwise $pr(c_i)=1$.

Definition 1

Let $I_{i,c}$, $I_{i,p}$ represent informational item i on the consumer's and provider's policies respectively.

$$pr(I_{i,c}) = pr(w_{i,c}) + pr(p_{i,c}) + pr(r_{i,c}) + pr(d_{i,c}),$$

$$pr(I_{i,p}) = pr(w_{i,p}) + pr(p_{i,p}) + pr(r_{i,p}) + pr(d_{i,p}) \; .$$

Definition 2

Given the Privacy Management Model of Section "Introduction", there is a *match* between consumer and provider privacy policies if

$$\sum_i pr(I_{i,p}) \geq \sum_i pr(I_{i,c}) + \sum_i pr(c_i) \; \; .$$

In other words, a match means the degree of privacy in the provider's policy is greater than the degree of privacy in the consumer's policy (the provider is demanding less information than the consumer is willing to offer). Otherwise, there is a *mismatch*.

Definition 3

A privacy policy is considered *upgraded* if the new version represents more privacy than the prior version. A privacy policy is considered *downgraded* if the new version represents less privacy than the prior version.

In comparing policies, it is not always necessary to carry out the above evaluation. We mention three shortcuts here.

Shortcut 1

Both policies are the same except one policy has fewer information items than the other policy. Then the policy with fewer information items always has a *higher* degree of privacy. According to Definition 2, there is a match if the policy with fewer information items belongs to the provider. Otherwise, there is a mismatch.

Shortcut 2

Both policies are the same except one policy has one or more informational items with less retention time than the other policy. The policy with one or more informational items with less retention time always has a *higher* degree of privacy. According to Definition 2, there is a match if the policy with one or more informational items with less retention time belongs to the provider. Otherwise, there is a mismatch.

Shortcut 3

Both policies are the same except one policy has one or more attributes that clearly lead to a higher degree of privacy for itself. According to Definition 3, there is a match if the policy with the resulting higher degree of privacy belongs to the provider.

Thus in the policies above, there is a match for e-learning according to Shortcut 2, since the policy with lower retention time belongs to the provider. There is a mismatch for book seller according to Shortcut 1 since the policy with fewer informational items belongs to the consumer. There is a mismatch for medical help according to Shortcut 3, since the policy with the higher degree of privacy is the one specifying a particular collector (Dr. Smith), and this policy belongs to the consumer.

Privacy Policy Compliance

As we have seen, the above Privacy Principles require a provider to be accountable for complying with the Privacy Principles (CSAPP.1) and the privacy wishes of the consumer (CHSPP.9). In practice, a provider is required to appoint someone in its organization to be in charge for its compliance to Privacy Principles (CSAPP.1). In some organizations, this responsibility may be added to those of the Chief Information Officer (CIO), in others a Chief Privacy Officer (CPO) has this responsibility. The CIO/CPO is to put in place a procedure for receiving and responding to complaints or inquiries about the privacy policy and the practice of handling personal information. This procedure should be easily accessible and simple to use. The procedure should also refer to the dispute resolution process that the organization has adopted. Other responsibilities of the CIO/CPO include auditing the current privacy practices of the organization, formulating the organization's privacy policy, and implementing and maintaining this policy. *We propose that the CIO/CPO's duties be extended to include auditing the provider's compliance to the consumer's privacy policy.*

A weakness of having the CIO/CPO responsible for protecting consumer privacy is that the CIO/CPO belongs to the provider's organization. Will she be truly diligent about her task to protect the consumer's privacy? To get around this question, the CIO/CPO can make use of secure logs to answer any challenges doubting her organization's compliance. Secure logs automatically record all the organization's use of the consumer's private information, both during and after the data collection. Cryptographic techniques (Schneier and Kelsey, 1999) provide assurance that any modification of the secure log is detectable. In addition, database technology such as Oracle9i can tag the data with its privacy policy to evaluate the policy every time data is accessed (Yee et al, 2003). The system can be set up so that any policy violation can trigger a warning to the CIO/CPO.

NEGOTIATION – STRUCTURE AND REPRESENTATION

Negotiation Example

This example illustrates negotiation to produce a privacy policy for a person (consumer) taking a course from an e-learning provider. Suppose the item for negotiation is the privacy of examination results. The employer would like to know how well the person performed on the course in order to assign the person appropriate tasks at work. Moreover, management (Bob, David and Suzanne) would like to share the results with management of other divisions, in case they could use the person's newly acquired skills. The negotiation dialogue can be expressed in terms of offers, counter-offers, and choices, as follows (read from left to right and down):

PROVIDER	*CONSUMER*
OK for your exam results to be seen by your management?	*Yes, but only David and Suzanne can see them.*
OK if only David and Bob see them?	*No, only David and Suzanne can see them.*
OK. Can management from Divisions B and C also see your exam results?	*OK for management from Division C but not Division B.*
How about letting Divisions C and D see your results?	*That is acceptable.*

As seen in this example, negotiation is a process between two parties, wherein each party presents the other with offers and counter-offers until either an agreement is reached or no agreement is possible. Each party chooses to make a particular offer based on the value that the choice represents to that party. Each party chooses a particular offer because that offer represents the maximum value among the alternatives.

Each party in a negotiation shares a list of items to be negotiated. For each party and each item to be negotiated, there is a set of alternative positions with corresponding values. This set of alternatives is explored as new alternatives are considered at each step of the negotiation. Similarly, the values can change (or become apparent), based upon these new alternatives and the other party's last offer.

Let R be the set of items r_i to be negotiated, $R = \{r_1, r_2,...,r_n\}$. Let $A_{1,r,k}$ be the set of alternatives for party 1 and negotiation item r at step k, $k=0,1,2,…,$ in the negotiation. $A_{1,r,0}$ is party 1's possible opening positions. Let $O_{1,r,k}$ be the alternative $a \in A_{1,r,k}$ that party 1 chooses to offer party 2 at step k. $O_{1,r,0}$ is party 1's chosen opening position. For example, for the first negotiation above, the provider's opening position is "exam results can be seen by management". Then for each alternative $a \in A_{1,r,k}$, $V_k(a)$ is the value function of alternative a for party 1 at step k, $k>0$, and

$$V_k(a) = f(I, O_{1,r,k-1}, O_{2,r,k-1}, \cdots)$$

where *I* is the common interest or purpose of the negotiation (e.g. negotiating privacy

policy for "Psychology 101"), $O_{1,r,k-1}$ is the offer of party 1 at step *k*-1, $O_{2,r,k-1}$ is the offer

of party 2 at step *k*-1, plus other factors which could include available alternatives,

culture, sex, age, income level, and so on. These other factors are not required here, but

their existence is without doubt since how an individual derives value can be very

complex. Let $a_m \in A_{1,r,k}$ such that $V_k(a_m) = max \{V_k(a), a \in A_{1,r,k}\}$. Then at step *k*, *k*>0 in the

negotiation process, party 1 makes party 2 an offer $O_{1,r,k}$ where

$$
\begin{aligned}
O_{1,r,k} &= a_m & \text{if} \quad V_k(a_m) > V_k(O_{2,r,k-1}), & \quad (1) \\
&= O_{2,r,k-1} & \text{if} \quad V_k(a_m) \leq V_k(O_{2,r,k-1}). & \quad (2)
\end{aligned}
$$

Equation 1 represents the case where party 1 makes a counter-offer to party 2's offer.

Equation 2 represents the case where party 1 accepts party 2's offer and agreement is

reached! A similar development can be done for party 2. Thus, there is a negotiation tree

$\vec{r}$ corresponding to each item *r* to be negotiated, with 2 main branches extending from *r*

at the root (Figure 1). The 2 main branches correspond to the 2 negotiating parties. Each

main branch has leaves representing the alternatives at each step. At each step, including

the opening positions at step 0, each party's offer is visible to the other for comparison.

As negotiation proceeds, each party does a traversal of its corresponding main branch. If

the negotiation is successful, the traversals converge at the successful alternative (one of

the parties adopts the other's offer as his own, equation 2 above) and the negotiation tree

is said to be *complete*. Each party may choose to terminate the negotiation if the party

feels no progress is being made; the negotiation tree is then said to be *incomplete*.

In Figure 3, the influences arrows show that a particular alternative offered by the

other party at step k will influence the alternatives of the first party at step $k+1$. Figure 4

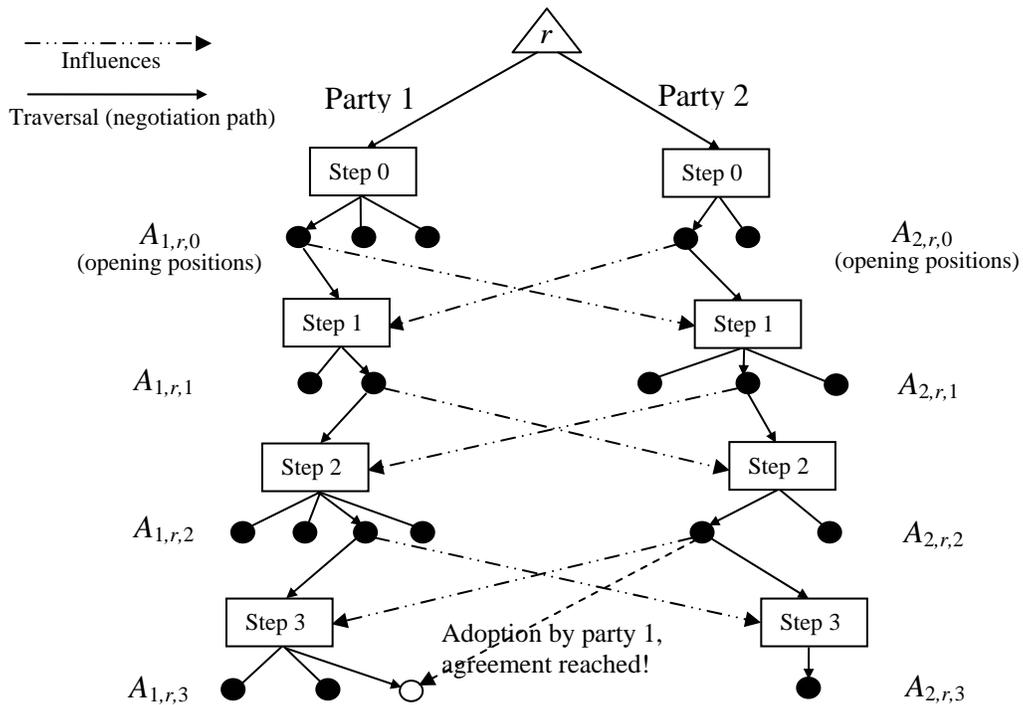illustrates the negotiation tree using the first negotiation above.

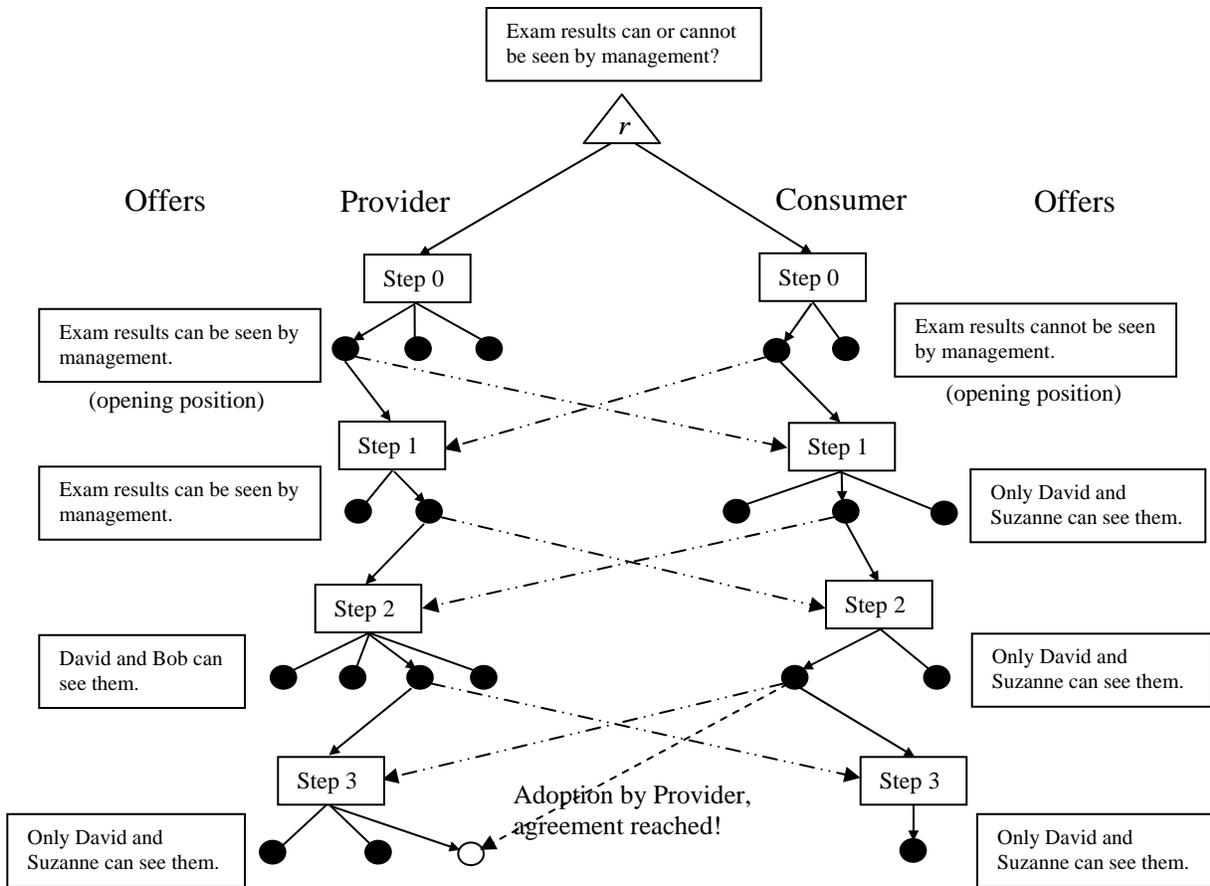Figure 3 – Negotiation tree $\vec{r}$ for a policy negotiation.

Figure 4 – Negotiation tree for the first part of the above negotiation.

NEGOTIATION IN CERTAINTY AND UNCERTAINTY

The following definition defines the meaning of negotiating in certainty and uncertainty.

Definition: Party *i negotiates in certainty* if for every negotiation step k, party *i* knows

both $A_{i,r,k}$ and $O_{i,r,k}$. Otherwise, party *i negotiates in uncertainty*.

Negotiation in certainty is therefore the type of negotiation illustrated in the example of section "Negotiation – Structure and Representation". At each negotiation step, each party knows the alternatives and knows what offer he is going to make. What is more interesting, however, is negotiating in uncertainty. What if a negotiating party does not know what the alternatives are or what offer or counter offer would be appropriate, at any particular step? This party may arrive at such a state as follows:

a) The other party's last offer may be a surprise (e.g. it is not understood).

b) He does not fully appreciate the value of the item under negotiation.

c) He may not be able to discern the values of his alternatives (not be able to compute $V_k(a)$).

In this case, the negotiating party may make use of the experience or decisions of others who have already negotiated the same item.

Negotiation in Uncertainty Example

Suppose you have been offered new employment and it is time to negotiate your benefits, including your salary. You know what you want in terms of vacation, sick leave, and training. However, when it comes to salary, you find it difficult to know what would be a fair salary, since both the job and the company are new to you. You have to negotiate in uncertainty. In this case, and what you may do naturally, is seek out others who you trust and who have negotiated salaries with this company in the past, for similar types of jobs. You would like to know how they negotiated their salaries, what alternatives they considered, and what counter-offers they made based on offers made by management.

You may not use their figures exactly but you may use their alternatives with different figures.

Reputation

As the previous example shows, negotiating an item in uncertainty may be facilitated through the use of knowledge from other parties who have negotiated the same item in the past. The question now is "Which other parties' negotiations knowledge should be used?" This is where reputation is employed.

Definition: The *reputation* of a provider or consumer is a quality that represents the degree to which he has fulfilled the commitments that he has made, either explicitly or implicitly. The commitments could be in everyday life (e.g. commitment to be faithful to a spouse) or in commerce (e.g. commitment to deliver work on time, commitment to respect a privacy policy, or commitment to pay for goods received).

The idea is to use the relevant knowledge of those having sufficiently high reputations. These parties would need to have a sufficiently high reputation and share your interest or purpose for the negotiation (*I* above). There may be other factors too, such as whether or not you know the party personally or have dealt with the party in the past. For manageability, we do not consider these other factors here.

A party's reputation is built-up over time from transactions with other parties. A particular transaction *t* occurs between 2 parties and has associated reputation factors that

contribute to determining the reputation of either party from the point of view of the other party. So for example, if party 1 purchases a book from party 2, factors contributing to party 2's reputation (from party 1's point of view) include whether or not the book received was the one ordered, whether or not the book was delivered on time, and party 2's performance history with other buyers. Factors contributing to party 1's reputation (from party 2's point of view) include party 1's credit history, the nature of past dealings with party 1, and party 1's performance history with other sellers.

One way to compute reputation is simply to rate the performance of a provider or a consumer on the associated reputation factors for a given transaction t. Let $t_{i,j}$ represent a transaction that party i has with party j. Let $q_1(t_{i,j}), ..., q_n(t_{i,j})$ be the associated n reputation factors for transaction $t_{i,j}$ assigned by party i to party j, where each reputation factor (rating) is ≥ 0 and ≤ 1 (each factor is an assigned score such as 3/5 or 6/7). Then party i assigns party j a reputation component $p(t_{i,j})$ corresponding to transaction $t_{i,j}$, where

$$p(t_{i,j}) = \frac{1}{n} \sum_{k=1}^{n} q_k(t_{i,j}) \quad .$$

Over the course of m transactions $t_{i,j}$, party i assigns party j a reputation $P_{i,j}$, where

$$P_{i,j} = \frac{1}{m} \sum_{t_{i,j}} p(t_{i,j}) \quad .$$

Notice that $0 \leq p(t_{i,j}), P_{i,j} \leq 1$. Suppose now that there are h parties that have had transactions with party j. Then party j has reputation P_j, $0 \leq P_j \leq 1$ where

$$P_j = \frac{1}{h}\sum_i P_{i,j} \quad .$$

Let party *1*, party *2*, …, party *h* be *h* parties other than party *k* that have had transactions with party *j*. Then *for party k*, party *j* has reputation P_j, $0 \le P_j \le 1$ where

$$P_j = w_k P_{k,j} + \frac{1-w_k}{h}\sum_{i=1}^{h} P_{i,j} \,, \quad i \ne j, k \ne j$$

and $0 \le w_k \le 1$ is a weight that party *k* applies depending on whether he had transactions with party *j* $(0 < w_k \le 1)$ or not $(w_k = 0, P_{k,j}$ undefined). A typical value for w_k is $w_k = 0.7$, meaning that party *k* places more emphasis on his own interactions with party *j* than the interactions of others in determining party *j*'s reputation, which is what one would expect.

In calculating the P_j by averaging over the $P_{i,j}$, we are in effect building consensus, so that any bias by a particular party is mitigated to some extent. Of course, the degree of mitigation is greater with increasing numbers of parties averaged.

In the literature, there has been much research done on reputation (Mui, L. et al, 2002). Our formulas are consistent with what other researchers have done. In particular, Zacharia and Maes (1999) have claimed that reputation in an online community can be described in terms of ratings that an agent receives from others. As a well-studied example, eBay client transaction ratings (Dellarocas, C., 2001) are not too unlike our proposal above. As another example, Cornelli et al (2002) used a rating system to allow servents (a servent is an entity that is both a server and a client) to accumulate reliability reputations for other servents from which they download in a P2P network. These

reputations are then use by resource requesters to assess the reliability of a potential provider before initiating a download.

SCHEME FOR NEGOTIATING PRIVACY POLICY UNDER UNCERTAINTY

We now describe an overall scheme on using the experience of others for negotiating privacy policy under uncertainty, as follows:

Every e-learning participant (both providers and consumers) accumulates negotiation experience in the form of negotiation trees (section "Negotiation – Structure and Representation").

1. Every e-learning participant calculates and stores the reputations $P_{i,j}$. A *reputation agent* can access these $P_{i,j}$ to calculate and store the P_j (using the first equation for P_j above). This can be done periodically to keep the P_j fresh.

2. A participant who is negotiating in uncertainty would obtain assistance, in the form of negotiation alternatives and offers made, from other reputable participants who have negotiated the same issue. The participant would:

    a. Identify which parties are reputable by asking the reputation agent for reputations P_j that exceed a reputation threshold H. Call this set of reputable parties J. That is, $J = \{ j : P_j \geq H \}$. The value of H can be set according to the level of reputation desired.

    b. Among the parties in J, search for parties that have the same interests I as the participant. This produces a subset J_s.

    c. Among the participants in J_s, search for negotiated items r that match the item the participant is currently negotiating. This produces a subset $J_r \subseteq J_s$.

d. Retrieve the matched negotiation trees $\vec{r}$ of participants in J_r. Use the alternatives and offers in these retrieved negotiation trees to formulate alternatives and offers. This is a manual step, supported by an effective user interface for displaying (or summarizing) the information to the participant for a decision on the alternatives. Note that the retrieved trees may be complete or incomplete (section "Negotiation – Structure and Representation").

e. Update his current negotiation tree.

Step 3 may be done in real time if reputations and past negotiation trees are all in place. Hence a negotiator can receive help in this manner at any negotiation step, if desired. Figure 5 illustrates the above scheme, using $H = 0.7$.
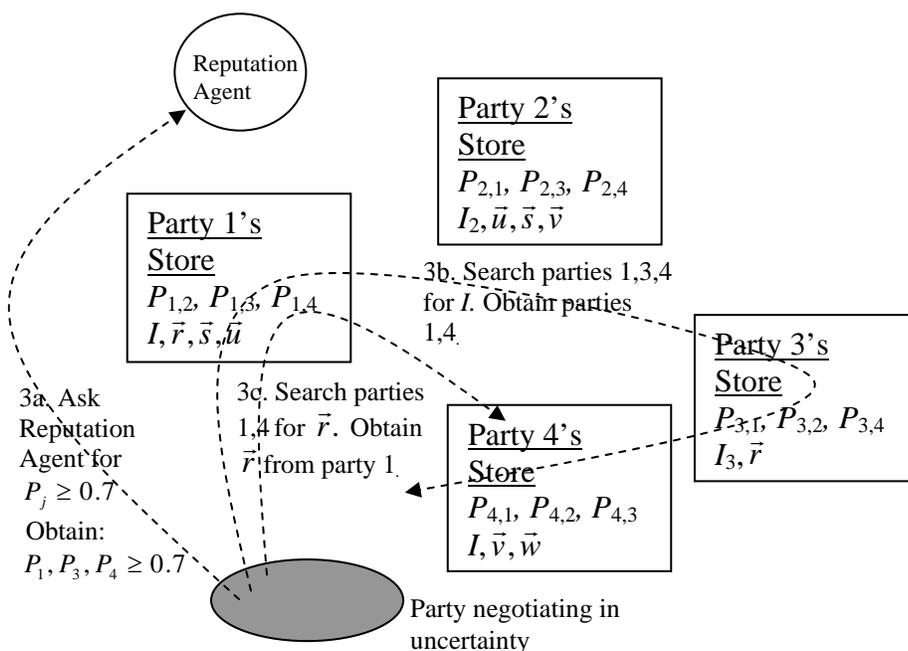

Figure 5 – This schema illustrates the idea of using the negotiating experience of others.

PROTOTYPE

We have implemented the above negotiation scheme in a prototype based on the JADE

platform (Telecom Italia Lab). JADE (Java Agent DEvelopment Framework) is a Java

middleware framework for developing agent-based applications in compliance with the

FIPA specifications for interoperable intelligent multi-agent systems. Our

implementation is totally peer-to-peer except for the use of a Contact Agent, which is

used to put a newly created agent in contact with the rest of the agent community. Figure

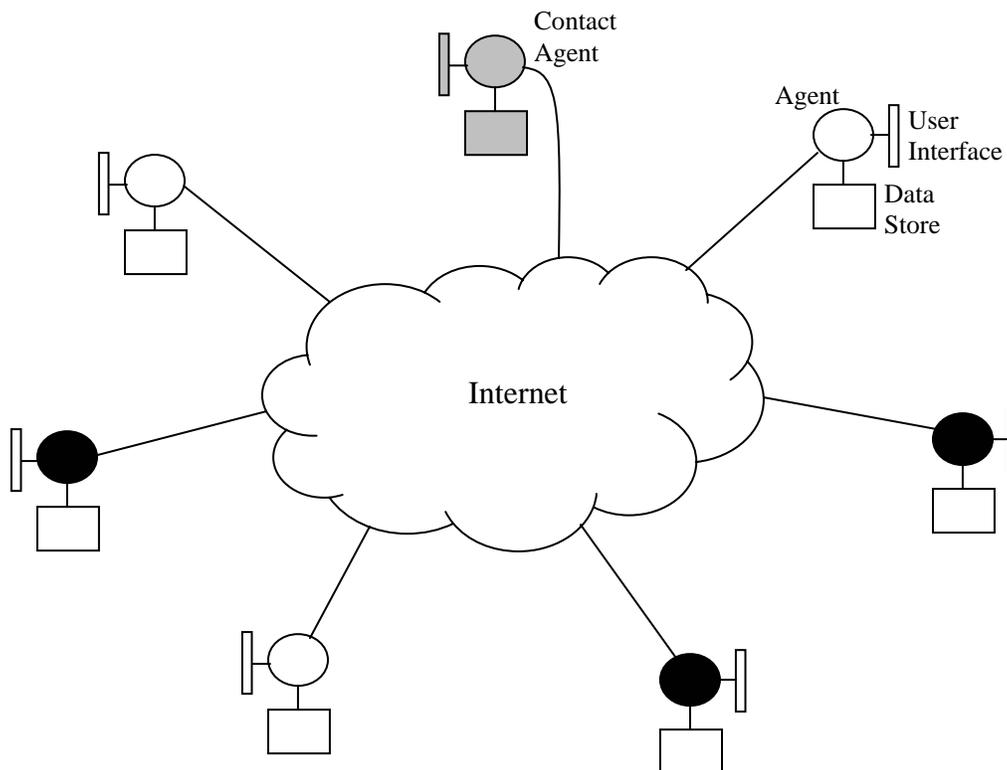6 illustrates the architecture of the prototype.

Figure 6 – Prototype architecture: black agents belong to distance learning providers,
white agents belong to learners; each agent, user interface, data store
combination runs on the learner's or provider's computer.

Here are some highlights of the architecture:

- *Communication*: communication between agents running on different machines uses HTTP MTP.

- *Experience Storage*: the negotiation experience of each learner or provider is stored in the data store of the learner's or provider's agent.

- *Dynamic Agents*: learner and provider agents may enter or leave the distance learning community (corresponding to the learner or provider entering or leaving).

- *Query Reputation or Interest*: taken care of by the agent belonging to the learner or provider making the query.

Full details of this prototype together with the results of experiments using it will be reported in a future paper.

CONCLUSIONS, DISCUSSION, AND FUTURE RESEARCH

This chapter has presented an approach for negotiating privacy policies in distance education using negotiation trees and reputations. The chapter categorized two types of negotiations: negotiating in certainty and uncertainty. The problem of negotiating in uncertainty was discussed and a solution given – that of using the negotiation experiences of reputable people with matching interests as aids in deciding which negotiating alternatives and offers should be employed. A scheme on how this could be done, together with a brief description of a prototype of the scheme, were presented. Our

application of negotiation trees in tandem with a reputation approach to policy negotiation is unique. It should facilitate the implementation of privacy mechanisms, which are key to the wide spread adoption of distance education.

Our work in the negotiations approach presented here is not yet complete. It raises the following questions:

- Security questions: How can negotiations be kept private? How can reputations be protected from malicious tampering? How can reputations and negotiation experiences be authenticated? How can malicious collusion to give a provider or a consumer a bad reputation be avoided?

- Performance questions: Will the machine-readable policy language be sufficiently expressive? How can reputations and negotiation experiences be stored to achieve optimal retrieval times?

- Perhaps negotiation experience should be sought from those who have good reputation as negotiators, not good reputation as online transaction participants. However, it may be more difficult to obtain reputations of negotiators since ratings of negotiation skill are less common.

- The reputations may not be current. For example, if a consumer has not participated in online transactions for a lengthy time, his reputation may not reflect his true nature at the present time. This can be fixed by time stamping when the transactions generating the reputations occurred, and then excluding reputations that are based on transactions that are too old according to a threshold.

- Automation questions: To achieve widespread use, the approach needs to be automated as much as possible. Can the generation of privacy policies be automated? The answer is yes (Yee and Korba, 2003). Can the negotiation be automated? We don't believe full automation is possible (see Section "Introduction") but perhaps some limited automation is possible.

We plan to investigate the above questions as future research. We are currently carrying out experiments on our prototype to answer the automation and performance questions.

REFERENCES

Benyoucef, M. et al (2001). An Infrastructure for Rule-Driven Negotiating Software Agents. In *Proceedings, 12th International Workshop on Database and Expert Systems Applications*, 2001.

Boehm, B. et al (2001). Developing Groupware for Requirements Negotiation: Lessons Learned. In *IEEE Software*, May/June 2001.

Canadian Standards Association. *Privacy Principles*. Retrieved July 3, 2002 from: http://www.csa.ca/standards/privacy/code/Default.asp?language=English

Carnegie Mellon University. *Privacy Server Protocol Project*. Internet Systems Laboratory, Robotics Institute and eCommerce Institute, School of Computer Science. Retrieved August 13, 2002 from: http://yuan.ecom.cmu.edu/psp/

Chiu, D.K.W. et al (2003). Developing e-Negotiation Process Support by Web Services. In *Proceedings of the 2003 International Conference on Web Services (ICWS '03)*, Las Vegas, Nevada, June 23-26, 2003.

Chung, M. and Honavar, V. (2000). A Negotiation Model in Agent-mediated Electronic Commerce. In *Proceedings, International Symposium on Multimedia Software Engineering*, 2000.

Cornelli, F. et al (2002). Choosing Reputable Servents in a P2P Network. In *Proceedings of WWW2002*, Honolulu, Hawaii, May 2002.

Dellarocas, C. (2001). *Analyzing the Economic Efficiency of eBay-like Online Reputation Reporting Mechanisms*. Paper 102, Center for eBusiness@MIT, July 2001. Retrieved October 16, 2002 from:

http://ebusiness.mit.edu/research/papers/102%20Dellarocas,%20eBay.pdf

Department of Justice. *Privacy Provisions Highlights*. Retrieved July 3, 2002 from:

http://canada.justice.gc.ca/en/news/nr/1998/attback2.html

Druckman, D. et al (2002). Artificial Computer-Assisted International Negotiation: A Tool for Research and Practice. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, 2002.

European Union (1995). *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*. Unofficial text retrieved Sept. 5, 2003 from: http://aspe.hhs.gov/datacncl/eudirect.htm

Huang, P. and Sycara, K. (2002). A Computational Model for Online Agent Negotiation. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, 2002.

Industry Canada. *Privacy and the Information Highway, Regulatory Options for Canada*. Chapter 6, retrieved Sept. 5, 2003 from:

http://strategis.ic.gc.ca/SSG/ca00257e.html#6

Kersten, G. et al (2002). The Effects of Culture in Anonymous Negotiations: Experiments in Four Countries. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, 2002.

Kim, J.B. et al (2003). Web Services and BPEL4WS for Dynamic eBusiness Negotiation Processes. In *Proceedings of the 2003 International Conference on Web Services (ICWS '03)*, Las Vegas, Nevada, June 23-26, 2003.

Kowalczyk, R. and Bui, V. (2000). On Fuzzy E-Negotiation Agents: Autonomous Negotiation with Incomplete and Imprecise Information. In *Proceedings, 11th International Workshop on Database and Expert Systems Applications*, 2000.

Kumekawa, J.K. *Health Information Privacy Protection: Crisis or Common Sense?* Retrieved Sept. 7, 2003 from:

http://www.nursingworld.org/ojin/topic16/tpc16_2.htm

Lai, R. and Lin, M. (2002). Agent Negotiation as Fuzzy Constraint Processing. In *Proceedings of the 2002 IEEE International Conference on Fuzzy Systems*, Volume 2, 2002. FUZZ-IEEE'02.

Limthanmaphon, B. et al (2000). An Agent-based Negotiation Model Supporting Transactions in Electronic Commerce. In *Proceedings, 11th International Workshop on Database and Expert Systems Applications*, 2000.

Lopes, F. et al (2001). Negotiation Tactics for Autonomous Agents. In *Proceedings, 12th International Workshop on Database and Expert Systems Applications*, 2001.

Mui, L. et al (2002). Notions of Reputation in Multi-Agents Systems: A Review. In *Proceedings of AAMAS'02*, Bologna, Italy, July 2002.

Murakami, Y. et al (2001). Co-evolution in Negotiation Games. In *Proceedings, Fourth International Conference on Computational Intelligence and Multimedia Applications*, 2001.

Nguyen, T. et al (2002). COPS-SLS: A Service Level Negotiation Protocol for the Internet. In *IEEE Communications Magazine*, Volume 40, Issue 5, May 2002.

Privacy Working Group. *Principles for the Privacy Protection of Personal Health Information in Canada*. Retrieved Sept. 5, 2003 from: http://www.cna-nurses.ca/pages/resources/principles_for_privacy_protection.pdf

Schneier, B. and Kelsey, J. (1999). Secure Audit Logs to Support Computer Forensics. *ACM Transactions on Information and System Security*. Issue 2(2), pp. 159-176, ACM, May 1999.

Telecom Italia Lab. Java Agent DEvelopment Framework (JADE). Retrieved July 20, 2003 from: http://jade.cselt.it/

Tu, M. et al (2000). Genetic Algorithms for Automated Negotiations: A FSM-Based Application Approach. In *Proceedings, 11th International Workshop on Database and Expert Systems Applications*, 2000.

W3C. *The Platform for Privacy Preferences*. Retrieved August 12, 2002 from: http://www.w3.org/P3P/

W3C (2002). A P3P Preference Exchange Language 1.0 (APPEL1.0). *W3C Working Draft 15 April 2002*. Retrieved August 12, 2002 from:

http://www.w3.org/TR/P3P-preferences/

Yee, G. and Korba, L. (2003). Semi-Automated Derivation of Personal Privacy Policies. In *Proceedings of the Information Resources Management Association International Conference 2004 (IRMA 2004)*, New Orleans, May 2004.

Yee, G. et al (2003). Privacy and Trust in E-Government. Book chapter in *Digital Government: Strategies and Implementations in Developed and Developing Countries* (tentative title), co-edited by Wei, K., Siau, K., and Huang, W. To be published by Idea Group, Inc. in 2004.

Zacharia, G. and Maes, P. (1999). Collaborative Reputation Mechanisms in Electronic Marketplaces. In *Proceedings of the 32nd Hawaii International Conference on System Sciences*, 1999.

---

[1] NRC Document No. : NRC-46555