Negotiating Exchanges of P3P-labeled Information for Compensation

Scott Buffett, Keping Jia, Sandy Liu, Bruce Spencer, Fang Wang

Institute for Information Technology – e-Business
National Research Council of Canada
46 Dineen Drive, Fredericton, New Brunswick, Canada E3B 9W4
Faculty of Computer Science, University of New Brunswick,
P.O Box 4400, Fredericton, New Brunswick, Canada E3B 5A6
{*Scott.Buffett, Keping.Jia, Sandy.Liu, Bruce.Spencer, Fang.Wang*}*@nrc.gc.ca,*
http://iit.nrc.gc.ca/il.html
Tel: (506) 444-0544, Fax: (506) 444-6114

Abstract. We consider private information a commodity, of value to both the information holder and the information seeker. Hence, a customer can be enticed to trade his/her private information with a business in exchange for compensation. In this paper we propose to apply utility theory to allow each participant to express the value they place on each private datum and, separately, on combinations of data. The PrivacyPact protocol transmits messages that are comprised of possible exchanges. Each participant is prevented from making offers that necessarily have lower utility for the other partner than previous ones. The protocol is complete in that if an exchange exists that is acceptable to both, it will be found as long as neither partner exits the negotiation early. While the space of possible offers grows exponentially on the number of negotiable items, experimentation with simple strategies indicates that negotiations can converge relatively quickly.

Keywords: privacy, multi-attribute utility, negotiating agents, P3P

1 Introduction

People are often reluctant to engage in electronic commerce interaction because they fear loss of private information. Often a website will request personal information such as name, home address, browser type, etc., from a user accessing the site. This information can be used to improve the website's content or organization, or the services offered. However, occasionally a website's policy might include less desirable practices, such as the transmission of private information to third parties. In either case, the information holds some value to the requestor. Consequently, the requestor may be willing to provide incentives such as inclusion in a rewards program, a free product, a discount on a future purchase, etc., in exchange for this private information.

Recent work in privacy economics research shows that people are typically willing to share their private information if they foresee a sufficient reward in return. Chellappa and Sin [2] and Culnan and Bies [7] argue that, when attempting to collect user data in order to personalize web sites, consumers are willing to share preference information in exchange for benefits such as convenience, if the quantified value of services outweighs the quantified loss of privacy. Hann *et al.* [11] show that economic incentives affect users' willingness to share information, and derive consumers' monetary worth of secondary use of personal information. Moreover, Culnan and Armstrong [6] argue that consumers are more willing to share their private information if they believe that fair information practices are in place. Fair information practices are those that 1) reveal why the information is being collected and how it will be used, and 2) give consumers control over its possible uses.

The Platform for Privacy Preferences Project (P3P) [4] enables websites to express their data-collecting practices in a standard format, indicating which data is to be collected, for what purposes, with whom it will be shared, and for how long it will be retained. P3P user agents then allow users to be informed of these website practices and automate the decision on whether or not to share the requested information based on the user's pre-defined privacy policy. While this mechanism constitutes a fair information practice, it is quite rigid. If the website and the user do not completely agree on the terms of the exchange, then the relationship is terminated. To alleviate this, earlier drafts of P3P included a protocol for multi-round negotiation. However, it was believed that this made P3P too complicated, and was thus dropped from the specification. Furthermore, Cranor and Resnick [5] show that under reasonable assumptions of user anonymity, publicly known website strategies, and no negotiation transaction costs for users, take-it-or-leave-it offers yield just as much website profit as any negotiation strategy.

On the other hand, if the negotiation includes rewards in exchange for the private information, the assumption of publicly-known website strategies becomes less reasonable. A business requesting personal information on their website might have several strategies that may be employed, depending on what sort of reward is being requested or offered, and also what type of user is involved in the negotiation. A user who has a long history of dealing with the business may be offered a higher reward, since his/her information may be worth more. This last point can provide incentive to the user to reveal herself, thus removing the assumption of user anonymity.

In this paper, we propose a multi-issue automated negotiation [8, 9, 12] protocol where the owner of the website (the business) can offer a certain level of service (e.g. 10% discount, free delivery) in exchange for information about the user (the customer). If this offer is not acceptable, the user may make a counteroffer, possibly consisting of a subset of the requested information for some higher level reward. Businesses whose privacy policies are less flexible or too complicated to negotiate may choose to negotiate only the reward, while other businesses may choose to negotiate policies or a combination of both.

Multi-attribute utility theory is used to help the parties rank their preferences. The exchange of counteroffers continues until either a deal is reached, or the negotiation is terminated by either participant. A participant might terminate negotiation if it appears impossible or unlikely that the other participant will make an acceptable offer, or if the negotiation is taking too long. We show analytically that our protocol is guaranteed to converge if a mutually acceptable agreement exists and neither party terminates the negotiation prematurely. We also give experimental evidence under normal settings that it can converge in a reasonable amount of time when simple negotiation strategies are used.

The paper is organized as follows. Section 2 provides some background on P3P, while section 3 presents a discussion on the elements of the negotiation, namely the information sought and the types of rewards to be offered in exchange. Section 4 then presents our proposed protocol. Section 5 provides an example negotiation session. Section 6 discusses the protocol's usage and properties. Section 7 describes a working prototype and a set of experiments. Section 8 then discusses conclusions and related work, and finally section 9 discusses future work.

2 Platform for Privacy Preferences Project (P3P)

There are various ways to collect personal information from the Internet: asking the user directly, accessing legitimate user agents, sharing "cookie" files between websites, or illegitimately invading private storage. We do not consider illegitimate uses here, but many users do not understand the extent of access to their private data even via legitimate methods. An important first step is to make policies on collecting data explicit, so W3C has developed the language P3P [4]. P3P is used by websites to express their privacy practice. A computerized agent, acting on behalf of the user, can fetch and read the P3P policy file, can inform the user about the site's privacy practices and can make an automatic or semi-automatic decision on behalf of the user.

The P3P policy file is an XML file that is defined for certain regions of a website or the entire website. Each P3P file contains at least one statement, and each statement describes what data will be collected, with whom it will be shared, for how long it will be retained and for what purpose. Figure 1 shows an example P3P statement requesting an element of the "physical" category (specifically the user's given name), and indicates that this information will be used for telemarketing and website administration purposes, that there are no intended recipients other than the requestor itself and that it will be retained for an indefinite period of time.

As an alternative to P3P, IBM's EPAL [1] can also be used to define a business' privacy policy. Where P3P essentially provides a vocabulary for formalizing human-readable privacy policies into a machine-readable format, EPAL aims at formalizing enterprise-internal privacy policies by giving a fine-grained vocabulary for formalizing the privacy-relevant aspects and the hierarchy of purposes for which the enterprise collects data. While we claim that our methods can

```
<STATEMENT>
   <PURPOSE>
      <telemarketing/> <admin/>
   </PURPOSE>
   <RECIPIENT><ours/></RECIPIENT>
   <RETENTION><indefinitely/></RETENTION>
   <DATA-GROUP>
      <DATA ref="#user.name.given"/>
         <CATEGORIES><physical/></CATEGORIES>
      </DATA>
   </DATA-GROUP>
</STATEMENT>
```

Fig. 1. An example P3P statement

be generalized to work for EPAL or any other privacy policy language, for the remainder of this paper we only consider P3P statements.

3 Negotiation Elements

In our model, there are two issues under negotiation between the website and the user: the P3P statement and the reward. By reaching an agreement, the user consents to releasing the private information specified in the associated statement, and the terms specified therein, and the website commits to compensating the user by providing the associated reward. While a P3P policy can actually contain several statements in practice, for simplicity we discuss the special case where only one statement is negotiated, and comment that the protocol can easily be extended to the general case. In this section, we discuss the participants and elements under negotiation in more detail.

3.1 Participants

We consider a two-participant bilateral negotiation where each participant is self-interested and has incomplete information about the opponent. Information is incomplete in that a participant is unsure not only about the opponent's reserve limits and deadlines, but also about its preference ranking of possible offers. For example, a user may be uncertain at first whether a website values his age or e-mail address more. This makes negotiation convergence more difficult since it is possible that one participant, while believing to be making concessions, could actually be moving farther from the set of mutually agreeable deals.

In this paper, we refer to the website (or the agent working on behalf of the website) as the *business* (denoted by b), and the user (or user agent) as the *customer* (denoted by c).

3.2 P3P Statements

Let D, R and P be sets of allowable values for requested data, intended recipients and purposes, respectively, as given in the P3P specification [4]. Let S be the set of P3P statements, where each element $s = \langle d, r, p, \tau \rangle$ contains a set $d \subseteq D$ of data, a set $r \subseteq R$ of recipients, a set $p \subseteq P$ of purposes and a real-valued retention time $\tau \in \Re$. Each participant has a private utility function $u^s : S \rightarrow \Re$ over the set of statements, indicating his utility of each statement. The business' utility for a statement s represents his utility for obtaining the information as specified by s, while the customer's utility for s specifies his utility for giving away the information. We denote u_b^s to be the business' utility function and u_c^s to be the customer's utility function. Note that these utility functions may vary depending on the partner with which the business or customer is dealing. For example, a customer may be less willing to divulge certain information when dealing with a company with which he is less familiar, or a business may be more eager to obtain certain information from repeat customers.

While the utility functions are private, a partial order of each participant's preference ranking is mutually known. Specifically, we assume that the business necessarily values a statement s no more than another statement s' if the data, recipients and purposes specified in s are subsets of those specified in s' and the retention time specified in s is no longer than that in s'. For the customer, the opposite is true. More formally, we define the partial order operator $\preceq$ over the set S as follows: Let $s, s' \in S$ be two statements with $s = \langle d, r, p, \tau \rangle$ and $s' = \langle d', r', p', \tau' \rangle$,[1]

$$s \preceq s' \quad \Leftrightarrow \quad d \subseteq d' \ \wedge \ r \subseteq r' \ \wedge \ p \subseteq p' \ \wedge \ \tau \leq \tau'$$

For example, let s_1, s_2 and s_3 be P3P statements where $s_1 = \langle$\{e-mail address\},\{ours\},\{admin\}, 1 week$\rangle$, $s_2 = \langle$\{name, e-mail address\},\{ours\},\{admin, telemarketing\}, 3 weeks$\rangle$ and $s_3 = \langle$\{name, phone number\},\{ours\},\{admin, telemarketing\}, 3 weeks$\rangle$. Then $s_1 \preceq s_2$ since the data, recipients and purposes specified in s_1 are subsets of those specified in s_2, and the retention time specified in s_1 is less than that specified in s_2. On the other hand, $s_1 \not\preceq s_3$ since \{e-mail address\} is not a subset of \{name, phone number\}. It is possible that the business could value obtaining \{name, phone number\} more and thus have a higher utility for s_3, or conversely that the customer could prefer giving up \{name, phone number\} less and thus have a lower utility for s_3. However, these preferences would be only privately known.

3.3 Rewards

In this paper, we discuss the notion of offering rewards in exchange for a customer's private information. Such rewards could include discounts on merchandise, free software or document downloads, air miles, etc. We assume that typically rewards would just be short term offers to be redeemed immediately, rather

[1] For the P3P retention value `<indefinitely/>`, we assume $\tau = \infty$.

than long-term offers such as a lifetime gold membership. Let T denote the set of *tokens*, where each token represents some reward. Each participant has a private value function $u^t : T \rightarrow \Re$ indicating his utility of each token. The business' utility for a token t represents his utility for giving away t, while the customer's utility for t represents his utility for obtaining it. We denote u_b^t to be the business' utility function and u_c^t to be the customer's utility function.

Similar to that for statements, since some rewards are mutually agreeable to be "more" or "better" than others, we define a partial order relation $\preceq$ over T. For example, while the customer's preference over various free software downloads might be subjectively decided upon and therefore be only privately known, it is assumed to be obvious to all that the customer would value a 20% discount more that a 10% discount on the same items. For any two tokens $t, t' \in T$, $t \preceq t'$ if and only if it is mutually known that the reward represented by t is no greater than that represented by t'.

3.4 Offers

Each valid offer $\langle s, t \rangle$ in the negotiation consists of one statement $s \in S$ and one token $t \in T$. Note that this model could be extended to allow sets of statements and sets of tokens in a single offer, but for simplicity we choose to narrow our focus. Each participant $z \in \{b, c\}$ has a private utility function $u_z : S \times T \rightarrow \Re$ over the set of possible offers.

We assume that the two attributes S and T in our model are *mutually utility independent*. That is, each participant's preference relation for statements remains the same regardless of which token is in question, and vice-versa. Given that the attributes are mutually utility independent, and that u_z^s and u_z^t are fully specified over their respective domains, the two-attribute utility function u_z can then be expressed by the bilinear function

$$u_z(s,t) = k_z^s u_z^s(s) + k_z^t u_z^t(t) + k_z^{st} u_z^s(s) u_z^t(t) \tag{1}$$

for all $s \in S$ and $t \in T$, where k_z^s, k_z^t and k_z^{st} are scaling constants which sum to 1 (refer to Keeney and Raiffa [13], for example).

As a result, the assumption of mutual utility independence together with the mutually agreeable partial orderings over statements and tokens induces a partial ordering over offers. For any two offers $\langle s, t \rangle$ and $\langle s', t' \rangle$, if $s \preceq s'$ and $t' \preceq t$ then the customer prefers $\langle s, t \rangle$ at least as much, while the business prefers $\langle s', t' \rangle$. That is, the customer prefers less information and more reward, while the business prefers more information and less reward.

4 PrivacyPact Protocol

Let S be a set of P3P statements and T a set of tokens. Let u_z map elements of $S \times T$ into the normalized range $[0, 1] \subset \Re$ where the participant z is either the business b or the customer c. Each side seeks an exchange such that his

utility is increased. Let $s^* \subseteq S$ be the information to be divulged and $t^* \in T$ be the token granted in such a mutually agreeable exchange, and call $\langle s^*, t^* \rangle$ a target exchange. There may be several such target exchanges. Each participant z will both approve of a target exchange if the utility for each exceeds some acceptability threshold α_z.

$$u_b(s^*, t^*) \geq \alpha_b \text{ and } u_c(s^*, t^*) \geq \alpha_c \tag{2}$$

The protocol gives the rules of the conversation between the business and the customer. In all cases it is assumed that if either partner wants to discontinue, the communication port can be closed. If either partner violates the rules of the protocol, or even if a partner is taking too long to make progress, the other is entitled to close the port.

There are three phases to the protocol.

Initial Phase The participants agree on the subdomains of negotiation, $D' \subseteq D$, $R' \subseteq R$, $P' \subseteq P$, $\Re' \subseteq \Re$, $T' \subseteq T$. The initial messages in the handshake are exchanged, depending on the supporting protocol, such as HTTP, FTPS, SOAP, etc. The first message in the PrivacyPact protocol contains sets D_b, R_b, P_b and $\Re_b$, indicating the data, recipients, purposes and retention times that the business is seeking, and also a set $T_b \subseteq T$ containing the selections from T that the information seeker may be willing to bestow. This is followed by a further specialization of these sets from the customer, with a message containing $D_c \subseteq D_b$, $R_c \subseteq R_b$, $P_c \subseteq P_b$, $\Re_c \subseteq \Re_b$ and $T_c \subseteq T_b$. It is assumed for now that these messages are sufficient to settle on the terms of reference. Thus $D' = D_c$, $R' = R_c$, $P' = P_c$, $\Re' = \Re_c$ and $T' = T_c$. For each participant z, there must exist a target $s_z^* = \langle d^*, r^*, p^*, \tau^* \rangle$ (where $d^* \in D'$, $r^* \in R'$, $p^* \in P'$ and $\tau^* \in \Re'$) and $t_z^* \in T'$ that represents an acceptable exchange. That is, $u_z(s_z^*, t_z^*) > \alpha_z$. If either partner finds no such elements, the conversation is discontinued.

Negotiation Phase The messages alternate from each participant to the other, starting with the first message from the business. Each message from participant z consists of an offer $\langle s_z^i, t_z^i \rangle$. It is a reasonable demand that the negotiation protocol allows the business and the customer to freely choose their own secret utility values on items exchanged. Negotiation strategies are also secrets kept from the other side. The protocol constrains the messages so that the interaction makes progress toward an acceptable exchange. Each message from one partner is not allowed to be less interesting than a previous one to the other partner, where less interesting for the business means exchanging less information for a higher token, and for the customer it means exchanging more information for a lower token. The following constraints control where the interaction goes and when it needs to stop:

Constraint 1-b Business makes progress Let $\langle s_b^n, t_b^n \rangle$ be the n-th message sent by the business. Suppose that this n-th message is the current

message and the customer is determining whether the message is legal. $\langle s_b^n, t_b^n \rangle$ must satisfy

$$\forall i = 1, ..., n-1 \quad s_b^i \not\preceq s_b^n \quad \text{or} \quad t_b^n \not\preceq t_b^i \tag{3}$$

That is, no offer to the customer may be worse than a previous offer. For every previous offer, either it does not ask for more information or it offers more in return than every previous offer.

Constraint 1-c Customer makes progress Now consider $\langle s_c^n, t_c^n \rangle$ the n-th message sent by the customer. This offer must be no worse for the business than any previous offer.

$$\forall i = 1, ..., n-1 \quad s_c^n \not\preceq s_c^i \quad \text{or} \quad t_c^i \not\preceq t_c^n \tag{4}$$

Acceptance Message During the negotiation phase, one of the following two conditions may arise. If so, the negotiation phase is over, as one of the participants has made an offer that is more beneficial to the opponent than one of that opponent's offers. It must therefore be accepted. In other words all previously made offers are still "on the table".

Condition 2-b Termination If the business' current offer $\langle s_b^n, t_b^n \rangle$ is the same as, or an improvement from the customer's point of view over one of the customer's previous offers, then this offer must be accepted by the customer. Such an offer ends the negotiation.

$$\begin{array}{l} \text{if } \exists \ \langle s_c^i, t_c^i \rangle \text{ where } i = 1, ..., n-1 \\ \text{such that } s_b^n \preceq s_c^i \text{ and } t_c^i \preceq t_b^n \\ \text{then the current offer } \langle s_b^n, t_b^n \rangle \text{ must be accepted by the customer.} \end{array} \tag{5}$$

Condition 2-c Termination If the customer's current offer $\langle s_c^n, t_c^n \rangle$ is the same as, or an improvement from the business' point of view over one of the business' previous offers, then this offer must be accepted by the business. Such an offer ends the negotiation.

$$\begin{array}{l} \text{if } \exists \ \langle s_b^i, t_b^i \rangle \text{ where } i = 1, ..., n \\ \text{such that } s_b^i \preceq s_c^n \text{ and } t_c^n \preceq t_b^i \\ \text{then the current offer } \langle s_b^n, t_b^n \rangle \text{ must be accepted by the business.} \end{array} \tag{6}$$

Last Chance Message One participant z may instead choose to terminate the negotiation at time n, perhaps because of time constraints or because the protocol disallows any further offers that he deems satisfactory. At this time, z can offer a take-it-or-leave-it message over the set of previous offers $\langle s_z^i, t_z^i \rangle, \quad i = 1, ..., n$. In this case, z's opponent can choose to accept one of z's previous offers or decline. In either instance, the negotiation is terminated.

Certification Phase The negotiated private information and a certificate entitling the customer to the negotiated token level of service are exchanged.

Theorem 1 (Convergence) A mutually acceptable offer can always be reached under this protocol if one exists, provided that neither participant prematurely terminates the negotiation process.

Proof. Let M be a message representing a mutually acceptable offer. If M is still allowed by the protocol then a solution still exists. Then assume without loss of generality that as a result of previous offers made by participant b in the negotiation, the protocol does not allow b to offer M. Then there exists an offer M' previously proposed by b that is as good as or better for the customer than M (i.e. less information for more reward). So M' is acceptable to c, and since it was offered by b, it is acceptable to b.

Let it be c's turn to offer. Since negotiation has not yet terminated under condition 2-c, then c has not made any offer that is better for the business than M' (i.e. more information for less reward). Therefore, c can offer M' without violating constraint 1-c, and an agreement will be reached under condition 2-c.

Let it be b's turn to offer. At any time, b can give a last-chance message and offer a take-it-or-leave-it message. As above c can then choose to accept M' and an agreement is reached. Thus if a mutually acceptable deal exists, then one can be reached using this protocol. □

5 An Example

To illustrate the point, this section demonstrates a simple example negotiation, where the business is interested in obtaining the customer's salary, name, phone number and email address. In practice, at least one purpose and one recipient as well as a retention time are needed in a P3P statement, but for simplicity we omit these in the example. After establishing these desired data elements, the customer then responds with a subset of these items, perhaps omitting the salary element. This indicates that he is not interested in divulging any information on salary, no matter what the business is offering, and effectively takes it off the bargaining table. Also, consider the set T of reward tokens presented to the customer to contain three levels of service: Silver, Gold and Platinum, where Silver $\preceq$ Gold $\preceq$ Platinum. This forms the initial domain of negotiation and concludes the initial phase.

In order to negotiate, each participant determines his own utility for each potential offer. Since utility is typically normalized from 0 to 1, we set utility of the best option to 1 and the worst to 0, and the utilities for all other offers to sensible values in between. In the customer's utility table (Table 1(a)), we see that giving out the name, phone number, and email in exchange for Gold service has a utility value of .3. On the other hand, the business values this offer with utility .7 (Table 1(b)). Note that each participant's utility table is not visible to the other participant in a negotiation. Also, let the utility acceptability threshold be $\alpha_c = .55$ for the customer and $\alpha_b = .5$ for the business. The areas that are

outlined by dark black are the acceptable offers. For instance, the customer may accept an offer from the business asking for the phone number in exchange for Gold service. The cells with darker gray background are the (unknown) mutually acceptable deals. For example, one such deal involves the customer divulging his name and the business offers gold service in return.

	{n}	***{p}***	***{n,p}***	***{e}***	***{p,e}***	***{n,e}***	***{n,p,e}***
Silver	.4	.35	.3	.25	.1	.05	0
Gold	**.7**	.65	**.6**	.55	.4	.35	.3
Platinum	1	.85	.8	.75	.6	.55	.5

(a) Customer utility table (n = name, p = phone, e = email, α=.55)

	{n,p,e}	***{n,p}***	***{n,e}***	***{n}***	***{p,e}***	***{p}***	***{e}***
Platinum	.5	.45	.35	.3	.2	.15	0
Gold	.7	**.65**	.55	**.5**	.4	.35	.3
Silver	1	.95	.85	.8	.7	.65	.6

(a) Business utility table (n = name, p = phone, e = email, α=.5)

Table 1. Utility values for each participant in the example negotiation.

A sample negotiation session is listed in Figure 2. Private utility values for each participant are given for each offer. In step 7 the business' offer indicates acceptance since condition 2-b is satisfied. Finally, the certification phase is reached and the transaction is completed.

Business	u_b	u_c	Customer	u_b	u_c
1. $\langle \{n, p, e\}, Silver \rangle$	1	0	2. $\langle \{n\}, Platinum \rangle$	.3	1
3. $\langle \{n, p\}, Silver \rangle$	.95	.3	4. $\langle \{p\}, Platinum \rangle$	.15	.85
5. $\langle \{n\}, Silver \rangle$	.8	.4	6. $\langle \{n\}, Gold \rangle$	.5	.7
7. $\langle \{n\}, Gold \rangle$ DEAL!	.5	.7			

Fig. 2. A sample negotiation session using the PrivacyPact protocol after the initial phase

Note that in this example the protocol would not allow the business to make an offer at step 7 of, say, $\langle \{n, e\}, Silver \rangle$, since this is necessarily worse to the customer than a previous offer (given in step 5) and therefore violates constraint 1-b. This prevents a participant from purposely making offers that wear on the patience of the opponent, thus possibly persuading him to accept the last

reasonable offer. However, keeping in the spirit of negotiation, a participant is in no way bound to accept the first acceptable offer received. For example, the business could have countered in step 7 with $\langle\{n, p\}, Gold\rangle$, and possibly reached a better deal. So our protocol does not restrict any reasonable bargaining strategies.

6 Discussion

A computerized agent will run on behalf of each of the participants to create a conversation that adheres to the protocol described. Each agent is directed by the interests of the participant it serves, but according to the protocol each agent and its participant is free to communicate messages that may or may not be closely related to those of the other participant. The offers of one participant are required only to not be the same or less attractive than previous offers by the same participant.

The PrivacyPact protocol is guaranteed to allow the participants to find an exchange, assuming they search for it exhaustively. But it is open to abuses. If D, R, P and T are the sets of data, recipients, purposes and tokens, respectively, then for any given retention time, there are $(2^{|D|} - 1)(2^{|R|} - 1)(2^{|P|} - 1)|T|$ different possible offers (assuming there is always at least one element from each of D, R and P and exactly one from T in any offer). All of these could be exchanged in a session that adheres to the protocol in this way. The first n messages for the business would request to exchange all items for the token of lowest value, then next lowest, etc. It would repeat this series of offers removing one item from the list, then repeat again removing another item. In this way it could traverse, from top to bottom in level order, the lattice of subsets of S with n offers at each subset. The customer would offer to exchange a singleton set of information for the highest valued token, then the next highest, etc., and produce offers in accordance with a bottom to top traversal of the subset lattice. This much traffic is clearly not feasible for a busy e-commerce system.

Given that this exhaustive strategy is to be avoided, we advocate that each participant calculate its next offer according to the following criteria:

1. For the business, initially suggest all the items wanted in the exchange; this is the only chance to put items "on the table". For the customer, reduce this initial set by removing any items from S not to be divulged and any tokens from T not of interest.
2. The business's initial offer should be fair. Offering to buy everything in S' for the lowest token is likely to make the customer lose interest and close the port. Similarly the first counteroffer from the customer should not offer to exchange some singleton set for the highest token.
3. Every offer a participant sends to its partner should be acceptable to itself, as per Equation 2.
4. If the partner's offer meets a participant's acceptability criteria in Equation 2, it is permissible to try to improve the exchange by continuing with

a counteroffer, but this has to be weighed against the risk of the partner closing the port.
5. Before an acceptable offer has been received, a participant may pursue almost any offer still allowed by the protocol. After an acceptable offer has been received, the participant may look for ways to improve that offer, but should limit the length of the conversation.
6. Depending on a participant's utility function, a directed search might create counteroffers more quickly. As in standard AI practice, the search would consider a set of offers that are near to some selected previous offer. Before an acceptable offer has been received, this selected previous offer would be chosen from among those given by this participant, seeking to make one of them less useful but still acceptable to itself. After one or more acceptable offers is received from the partner, the directed search would proceed from among these; the participant would seek to make it increase its utility.

7 A PrivacyPact Experiment

The PrivacyPact protocol has been implemented for two reasons, first as a demonstration of its feasibility, and second to experiment with various strategies for dealing with the possibly exponentially long negotiations. Our initial intention is not to study how a highly effective negotiation strategy can be built – we feel that doing that would depend on some real world experience with the protocol, which is not currently available – but rather to demonstrate that simple strategies can be effective for reducing the lengths of the conversations, and thus provide assurance that the protocol is not without merit.

7.1 Assessing the utility of an offer

The first task is to create a utility value for each exchange for each of the two participants in the negotiation. This function maps each P3P statement and token to a real value from $[0, 1]$ representing the utility of that exchange for that participant. This function should meet two criteria: transparency and smoothness. We consider each of these in turn.

It should be transparent to a participant that the utility assignment accurately reflects his opinions about the relative importances of the various aspects of the offer and combinations of these aspects. Thus it is essential that the participant's opinions about these importances are expressed in a simple way. We provide two types of statements that allow the participant to express these importances: for individual items, and for combinations of items. Recall that each offer exchanges a tuple $\langle d, r, p, \tau \rangle$ for some token t, where d is a subset of the data D, r is a subset of recipients R, p is a subset of purposes P, and these are the multi-valued attributes. Also, τ is a real-valued duration and t is a token selected from T, and these are the single-valued attributes. In the first type of statement, the participant gives each item in $D, R, P, \Re$ and T a number in the range $[0, 1]$ that represents his opinion on the importance of this item. For each such item e

let $l(e)$ be this assigned value. These numbers are directly translated into utility values as follows, where e is from an attribute Y and this participant is on the receiving end for this item (the business recieves data, recipients, purposes and retention times while the customer recieves tokens):

$$u(e) = \begin{cases} l(e)/\Sigma_{i \in Y} l(i) & \text{if } Y \text{ is a multivalued attibute} \\ \\ l(e)/\max_{i \in Y} l(i) & \text{if } Y \text{ is single-valued} \end{cases} \tag{7}$$

These values for the importance numbers can come from any distribution, and this equation scales them to a number in $[0, 1]$. If the participant is the holder of this attribute (i.e. the business holds tokens while the customer holds data, recipients, purposes and retention times), then the utility value is one minus the value of u computed by Equation 7. There is less utility in giving away more important items.

For example, a business that wants to receive a customer's name most and email address least, where the I attribute also contains phone number, might express that $l(\text{name}) = 9, l(\text{email}) = 1$, and $l(\text{phone}) = 5$. Then $u(\text{name}) = 9/15 = 0.6$. while $u(\text{email}) = 1/15$. A customer who shares the same opinions of relative importance would have a utility of $1 - 0.6 = 0.4$ for giving his name.

A combination of several items is given the utility equal to the sum of the utilities of these items, except when there is information from the user that gives such a set special importance, which is considered in the next paragraph. In the absence of any special instructions from the business, the offer of a name and email would be given utility of $9/15 + 1/15 \approx 0.67$.

Also a participant might express that a combination of items has a special importance. For instance the business may want to express that receiving a name and phone number combined has a higher importance, since it may be used to identify that person uniquely. In this case the participant would express that the combination has higher importance, perhaps 18. The utility of receiving a name and phone number would be based on considering the name and phone number combination to be one item, redefining I accordingly, and applying Equation 7. Keeping $l(\text{email}) = 1$, and defining $l(\text{name-and-phone}) = 18$ we have that $u(\text{name-and-phone}) = 18/19 \approx 0.947$.

After the utilities from each of the four data dimensions are considered, the utility $u^s(\langle d, r, p, \tau \rangle)$ of a statement is calculated. In this experiment we multiplied the four utilities together. That is, $u^s(\langle d, r, p, \tau \rangle) = u(d) \times u(r) \times u(p) \times u(\tau)$. Thus each component counts equally and the final utility of a statement is a number in $[0, 1]$. Once the utilities for tokens has been determined, the utility $u(s, t)$ of an offer is calculated using the bilinear function in Equation 1. In our experiments we assigned statement and token utility equal weight (i.e. $k^s = k^t = 0.5$, $k^{st} = 0$) for each participant.

We also allow that a combination of items from different sets be considered specially. For instance a participant might choose to place special importance on an offer that includes the phone number when it is allowed to be used for telemarketing purposes, combining consideration from separate dimensions. While

this could be done in terms of importance measures, we found this difficult to explain to users because the different dimensions may be using different scales, so these numbers are given as utilities, as numbers in $[0, 1]$.

The utility function should also be smooth, so that similar offers that trade items of almost equal importance have similar utilities. If a utility function meets both the smoothness and the transparency criteria, it may be possible for a participant that knows how the competing participant has assessed some offers – e.g. whether some offers were assessed higher than others – to form an opinion about which items are important. We will revisit this point in the next section, and show that this smoothness can help cooperative negotiating partners to discover mutually acceptable offers within short negotiations.

7.2 Computing Negotiations with a Prototype

The goal of our experiment is to show that the PrivacyPact protocol can accommodate simple strategies to give rise to short conversations. We define three negotiation strategies: "miserly" that always makes its next offer according to what is most beneficial to itself, "cooperative" whose next offer is chosen according to its similarity to any of the partner's previous offers, and "hybrid" which is a combination of the previous two.

The miserly negotiator makes a counteroffer by considering all of the valid offers, defined as those admitted by the conditions of the PrivacyPact protocol, and selects the one with maximal utility for itself, without any consideration of the opponent's previous offers. Thus an negotiation involving two miserly participants may require an excessive number of messages to converge, since the space of offers is essentially searched exhaustively to find an agreement.

The cooperative negotiator attempts to overcome this by considering all pairs of offers from two sets: the set of possible counteroffers allowed by the protocol, and the set of previous offers from the opponent. For each pair a similarity measure is determined, and the counteroffer selected is the one most similar to some previous offer of the opponent's. The similarity between a pair of offers is considered according to the similarity of each of the five dimensions. For single valued attributes, the similarity is some defined distance between the values. For instance, retention times are real numbers and the distance can be their difference. For a pair of multivalued attributes chosen from a set S, the distance is calculated according to the number of values from S that they agree upon as a fraction of the size of S. They agree on a value either if they both contain it, or both do not. They disagree if one contains it and the other does not. Once all five dimensions are considered, a linear combination of the five numbers gives the overall similarity; in our case each of data, recipients, purposes, retention time and token similarity counts as one fifth.

The hybrid negotiator combines the other two; it attempts to find good counteroffers quickly by looking first at deals most favourable to itself (i.e. with highest utility), and then choosing from these deals the one that maximizes the similarity measure. In our experiments, the number of such counteroffers was set

at $n = |O|/10$ where O is the set of possible offers at the beginning of the negotiation. Thus at the beginning of the negotiation, the participant only considers the best 10% of the possible offers. Since the value of n remains fixed throughout the negotiation but the number of offers allowed by the protocol decreases, the percentage of valid offers considered increases. Thus the participant becomes more cooperative as the negotiation continues. Since one often wants to make more concessions as time elapses in a negotiation, this is still a very reasonable strategy.

In lieu of real-world examples, which do not (yet) exist, we selected a variety of examples, each with a selection of information items, recipients, purposes, retention times and tokens. Certain goals for the business were set by setting importance of certain subsets high, while goals for the customer were specified by setting low importance for some combinations.

Once the utility functions are defined, and before negotiation can begin, it is necessary to select alpha thresholds for each negotiator. This threshold specifies the lowest utility a partner has for accepting an offer. To make the negotiation as hard as possible, we set the alphas so that a small but non-zero number of offers could be accepted. We do this by considering each offer in turn and the pairs of utilities assigned by of the partners. (Ordinarily no one party would have access to both of these functions.) For each pair, we selected the lower value, and from all these low values we select the highest. The offer associated with this highest low is arguably a hard exchange to find since it represents an offer not much favored by either partner. For each partner, the alpha value is assigned to be that partner's utility of this offer.

We tested the performance of each of the miserly, cooperative and hybrid customer negotiation strategies against a miserly business negotiator. This gives a clear demonstration of how quickly these simple strategies can converge. Note that each strategy was tested against the miserly negotiator rather than against a cooperative or hybrid negotiator since those results, while still much better than miserly versus miserly, were more erratic. For example, in some cases a negotiation involving two cooperative agents would take longer than one involving one cooperative and one miserly. This is because they would both try too hard to please each other and thus make convergence more difficult. Occurrences such as this were completely example-dependent and thus did not warrant consideration in our analysis, simply because our goal is only to show that the protocol can converge quickly under reasonable strategies. Table 2 gives the results. For each run, the number of negotiable items and the number of possible offers are given, as well as the number of messages required for convergence for each of the three strategies. For the sake of simplicity, only elements of D in the statements are negotiated. That is, the purposes, recipients and retention time are agreed upon in the initial phase. Also there are four tokens up for negotiation. Thus for $|D|$ negotiable items there are $2^{|D|} - 1 \times 4$ possible offers. Experiments show that most of the space of offers is searched when the miserly strategy is employed. However, a considerably smaller number of exchanges are required when the hybrid and cooperative strategies are employed. This is a good indication that

protocol has potential to converge quickly when simple but effective strategies are used.

Number of negotiable items	Number of possible offers	Business Strategy	Customer Strategy	Number of messages to converge
3	28	miser	miser	24
			hybrid	18
			co-op	4
4	60	miser	miser	48
			hybrid	12
			co-op	6
5	124	miser	miser	104
			hybrid	16
			co-op	4
6	252	miser	miser	194
			hybrid	16
			co-op	10
7	508	miser	miser	388
			hybrid	20
			co-op	4
8	1020	miser	miser	818
			hybrid	32
			co-op	4
9	2044	miser	miser	1670
			hybrid	52
			co-op	38
10	4092	miser	miser	3178
			hybrid	66
			co-op	56

Table 2. Negotiation lengths for various problems and strategies

8 Conclusions and Related Work

This paper proposes a negotiation protocol that allows users to express the degree of their reluctance to divulge private information, and that this reluctance can be combined in non-additive ways. Our proposal allows users to compare such reluctance to the enticements offered by business, and businesses to decide if suites of information about a specific person are worth the cost of offering the incentive. The protocol is complete in that an acceptable exchange will be found if it exists, barring early exit. It has exponential worst case time complexity, however we give experimental evidence indicating that it can converge in a reasonable amount of time when simple negotiation strategies are used.

Casassa Mont and Yearworth [15] consider a related problem, that of automating admission to the negotiation. It could be incorporated as a precursor to our work within a system combining the two.

Early drafts of P3P (before August 1999) included a framework for multi-round automated negotiation. This was removed from version 1.0 because it was deemed to be too complicated. It was felt that it was more important for websites to make simple and clear requests and for the users to simply accept or reject [10]. APPEL [3] allows the user to specify a policy. This policy consists of a set of rules specifying certain conditions under which certain types of information may be used. If the requestor meets these conditions then the information is exchanged. There is no facility for negotiation.

Some work has since been done to bring negotiation capabilities back to the exchange, in the setting where the user needs to disclose some information in order to perhaps complete a transaction, receive a free download, etc. Meyer [14] created a protocol for the negotiation of sets of information by constructing rule-based policies for counter-proposals. This can be viewed as the process of the user agent perhaps relaxing the policy slightly with each counter-proposal, as the requestor conforms to the policies a little each time, until an agreement is reached. Cranor and Resnick [5] discuss information negotiation strategies, and the effect that restricted and general protocols have on the effectiveness of those strategies.

There are many privacy issues a user (or agent) could consider, and negotiating over them is a special case of the general task of negotiation over multiple issues. Relatively few papers consider multiple issue negotiation in a distributed environment performed by agents with competing objectives. Faratin, Sierra and Jennings[8] propose such a negotiating system, where a counteroffer is generated in response to an opponent's offer in accordance with three heuristics: make the counteroffer similar to the offer, make the counteroffer have as high a utility as previous offers, and when necessary concede some utility so that negotiation is able to proceed. Given these principles, their hill-climbing system is able to guarantee negotiation will complete with a number of messages proportional to the number of issues to be negotiated. This work does not guarantee that a successful negotiation will result, even if there exists a mutually agreeable exchange. This work does not clearly delineate which utilities are mutually known and which are privately held.

9 Future Work

As we apply the PrivacyPact protocol to general multi-issue negotiation, we will investigate heuristically-based negotiation strategies. The negotiation space we have identified is exponentially large in the number of issues under negotiation, but this does not mean systems that use the protocol must enumerate this entire space. Systems employing heuristics, such as [8], could provide a means of ordering that space by identifying fruitful negotiation paths. If the heuristics do not lead to an acceptable offer, the remainder of the space can still be explored.

Our negotiating constraints allow the negotiating agent to know what offers in the remaining space can be eliminated from consideration because such offers are already guaranteed be unacceptable to the opponent, given previous offers and the mutually-known partial order of utility values.

References

1. P. Ashley, S. Hada, G. Karjoth, C. Powers, and M. Schunter. Enterprise Privacy Authorization Language (EPAL) Specification. http://www.zurich.ibm.com/security/enterprise-privacy/epal/, 03 March 2003.
2. R.K. Chellappa and R. Sin. Personalization versus privacy: An empirical examination of the online consumers dilemma. *Information Technology and Management*. To appear.
3. L. Cranor, M. Langheinrich, and M. Marchiori. A P3P Preference Exchange Language 1.0 (APPEL1.0). http://www.w3.org/TR/P3P-preferences/, 15 April 2002. W3C Working Draft.
4. L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. http://www.w3.org/TR/P3P/, 16 April 2002. W3C Recommendation.
5. L. Cranor and P. Resnick. Protocols for automated negotiations with buyer anonymity and seller reputations. *Netnomics*, 2(1):1–23, 2000.
6. M.J. Culnan and P.K. Armstrong. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1):104–115, 1999.
7. M.J. Culnan and R.J. Bies. Customer privacy: Balancing economic and justice considerations. *Journal of Social Issues*, 59(2):104–115, 2003.
8. P. Faratin, C. Sierra, and N.R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142:205–237, 2002.
9. S.S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
10. R. Grimm and A. Rossnagel. Can p3p help protect privacy worldwide? ACM SIGMM Electronic Proceedings, url = "http://www.acm.org/sigmm/MM2000/ep/grimm", 2000.
11. I. Hann, K. Hui, T.S. Lee, and I.P.L. Png. Online information privacy: Measuring the cost-benefit trade-off. In *23rd International Conference on Information Systems*, 2002.
12. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation*, 10(2):199–215, 2001.
13. R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, Inc., 1976.
14. Jörg Meyer. How to manage, negotiate, and transfer personal information on the web. Written at the IBM Almaden Research Center, submitted as diploma thesis at the University of Applied Sciences Hamburg, Germany, 1999.
15. M. Casassa Mont and M. Yearworth. Negotiated revealing of trader credentials in e-marketplaces mediated by trusted and privacy-aware admittance controllers. http://www.hpl.hp.com/techreports/2001/HPL-2001-216.html, 12 September 2002.