

Personalized Services with Negotiable Privacy Policies

Sören Preibusch

German Institute for Economic Research (DIW)
Königin-Luise-Str. 5, 14195 Berlin, Germany
spreibusch@diw.de

ABSTRACT

This paper examines how negotiation techniques can resolve the trade-off between service providers' personalization efforts and users' individual privacy concerns, how they lead to efficient contracts, and how they can be integrated into existing technologies to overcome the shortcomings of static privacy policies. The analysis includes the identification of relevant and negotiable privacy dimensions. A detailed privacy negotiation scenario from multi-channel retailing is examined. Based on a formalization of the user's privacy revelation problem, we can solve the selection of the efficient privacy level as an optimization problem. Finally an extension to P3P is proposed that allows a simple expression and implementation of negotiation processes. Support for this extension has been integrated in the Mozilla browser.

Author Keywords

Privacy Negotiations, Privacy, Personalization, P3P, E-Commerce.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous. K4.4. Electronic Commerce, K.4.1 Public Policy Issues

INTRODUCTION

Online users are facing a large and increasing complexity of services offered over the Web, due to their quantity and their diversity. In online retailing, stores are constantly expanding their assortments in width, depth and quality levels. Individuals are confronted with a number of product alternatives that makes their exhaustive comparison impossible [8]. They appreciate being offered effective guidance through automated recommender systems [13].

Moreover, the commodization of digital services increases price competition among service providers and successful customer value extraction requires attracting and binding customers by new means. Personalized services, individually tailored for a single consumer, create lock-in effects. A recent survey concluded that knowledge from, for and about customers is a mission-critical factor [16].

Yet Customer Relationship Management (CRM) typically relies on large data amounts to be collected and kept over time. Careless data collection activities and data misuse are nowadays discussed in mass media and remember customers to care about their Privacy.

The depicted situation is known as the Privacy-Personalization trade-off. Users get immediate better service quality when revealing personal information but may experience negative consequences due to sensible information disclosure. A common way for websites to communicate their data-handling practices is to post "privacy policies" on their website. Though, this approach is too rigid and comparing present benefits to expected values of future costs constitutes a major hurdle for most online users. The empirical proof of users' stated privacy preferences diverging from their actual behaviour is a symptom of this burden [19].

Our contribution is to depict how negotiation techniques can overcome current drawbacks of static privacy policies, and reconcile privacy and personalization. We explain the economic benefits of Privacy Negotiation Techniques (PRINT) and how they support the customer. We investigate how negotiations can be implemented using existing technologies. The following two sections examine negotiable privacy dimensions and present the optimization calculi of the user and the service provider respectively, based on a formalization of privacy negotiations. The implementation of PRINT using W3C's Platform for Privacy Preferences (P3P) is portrayed. Before concluding with a summary and outlook, a detailed negotiation scenario in multi-channel retailing is provided.

RELATED WORK

The privacy-personalization trade-off as presented above has led to several technical approaches both in research and in practice. These technologies can be classified according to the market structure and the targeted transaction phase, each of them being briefly portrayed in the following paragraphs: (a) Service providers publish P3P Policies that are retrieved by a user agent (UA) acting on the user's behalf. (b) The UA checks if the P3P Policy is compatible with the user's privacy preferences. Latter can be coded using the privacy preference languages APPEL or XPref. (c) Intra- and inter-organizational guidelines governing the handling of collected data can be expressed using EPAL.

(a) P3P is an XML-based language developed by the World Wide Web Consortium (W3C) [24]. It became a recommendation in 2002 and aims "to inform Web users about the data-collection practices of Web sites" [25]. P3P has become widely adopted by service providers but it remains restricted to the "take-it-or-leave-it" principle: The

service provider offers a privacy policy; the potential customer has to accept as a whole if she wants to use the service. A negotiation process between the involved parties is not intended. Although the first drafts of the P3P specification included multi-round negotiation mechanisms, these parts had been removed in favour of easy implementation and early and wide adoption of the protocol. The latest version of the P3P 1.1 specification [25] does not mention negotiations either.

(b) In addition to the P3P specification, the W3C conceived APPEL1.0, A P3P Preference Exchange Language 1.0 [23]. APPEL is a language “for describing collections of preferences regarding P3P policies between P3P agents”. It is primarily intended as a transmission format and a machine-readable expression of a user’s preferences. Given a P3P privacy policy, it may be evaluated against a user-defined ruleset to determine if her preferences are compatible with the service provider’s intentions for data. Though standard behaviours and basic matching operations are supported by APPEL, its applications are still limited and the capability of expressing negotiation strategies is explicitly excluded from the language’s scope. Using APPEL as a negotiation protocol is neither supported by its semantics nor is the language designed for this purpose.

Critics have argued that APPEL is difficult to use effectively and have proposed XPref as an APPEL replacement [2]. APPEL and XPref address matching issues antecedent to the approval of the service provider’s privacy policy. Though privacy preference languages are not suited for coding privacy negotiations, user agents may support the user during the negotiation process in reaching an agreement favourable for him.

(c) The Enterprise Privacy Authorization Language (EPAL) is developed by IBM [9]. EPAL allows enterprises to express data handling practices in IT systems. The developed policies are intended “for exchanging privacy policy in a structured format between applications or enterprises” [9]. The language focuses on the internal business perspective, and is not intended for customers to express their own privacy preferences. Although EPAL is not suited for use at the customer interface – which is needed for negotiation – privacy guarantees towards customers can sometimes be deduced from the stated internal procedures and then be expressed in P3P policies.

In parallel to the development of privacy-related technologies and research both in online and offline IT-based transactions, negotiation has been studied in various disciplines. The bases had been set up in game theory, where negotiation is modelled as a bargaining game [10, 20]. Recent influences have arisen with the increasing importance of autonomous agents and collaborative computing [5]. Frameworks for carrying out negotiations have been developed [15]. The rapid development of the Grid and service-based IT-architectures on the technical side, and the enduring process outsourcing to third parties

on the economic side, combined with mobile and ubiquitous computing will broaden the coverage of Privacy Negotiation Technologies in the near future [11, 26].

PRIVACY NEGOTIATIONS

Thompson defines in her book that negotiations are an “interpersonal decision-making process necessary whenever we cannot achieve our objectives single-handedly” [21]. Especially in the case of integrative negotiations, negotiations can unleash the integrative potential that lies in conflicting interests and preferences and turn it into efficient contracts. Two major shortcomings of current online privacy handling mechanisms can be overcome if PRINT are implemented during the transaction between the service provider and the user:

The first shortcoming is the “one-size-fits-all” principle: once the service provider has designed its privacy policy, it will be proposed to all interested users – no matter what their individual preferences are. There may be users who would have accepted offers with less privacy protection and would have agreed to the provider’s proposal even if more personal data would have been asked. Thus, the provider fails to tap the users’ full potential.

The second shortcoming is the “take-it-or-leave-it” principle, i.e. the user can only accept or refuse the provider’s proposal as a whole. The provider is always the one who moves first, he makes the initial offer; the user cannot take the initiative. As demonstrated in the next section, the fundamental inadequacy of the take-it-or-leave-it principle persists even if more than one static privacy policies are offered.

Individualized Privacy Policies

Adopting a broader view and extending the analysis from a single service provider to the whole market, providers might specialize on different privacy levels. Since the amount of service providers (as discrete units) is much smaller than the amount of potential privacy preferences, which can be seen as quasi-continuous due to the large number of gradations for all considerable privacy dimensions, a specialization is not trivial.

Consider n service providers and $m \gg n$ users having different privacy levels with a known distribution. Hence, a given service provider will target more than one privacy level. This may be implemented by giving the users the choice between a set of usage scenarios corresponding to different amounts of personal data to be collected. As the differences between these usage scenarios have to be clearly communicated and the maintenance of one scenario induces costs for the service provider, the set of scenarios will be limited in size to a few possibilities.

The notable difference between the offered privacy levels is part of the service provider’s user discrimination strategy and aims at a successful self-selection of the potential users.

Thence, even under market-driven specialization and alternative usage scenarios, the user still faces fixed policies and a dynamic matching is not carried out.

Negotiable Privacy Dimensions

As observed in the previous sub-section, neither a market-driven segmentation between services providers offering different privacy levels, nor a mechanism based on choices between different static usage scenarios turn out to be adequate solutions, so that negotiation is the remaining approach. Apparently, as it is not feasible to negotiate the entire privacy policy, one important aspect is to identify relevant and negotiable privacy dimensions. We define a *privacy dimension* as one facet of the multi-dimensional concept ‘user privacy’. For each dimension, different discrete revelation levels exist, monotonously associated with the user’s willingness to reveal the data. Privacy dimensions can be identified at different degrees of granularity.

The four top-level privacy dimensions are the recipient of the data, the purpose for which the data are collected, the period they will be stored, and the kind of data. These four dimensions (recipient, purpose, retention time, and data) are in accordance with European privacy legislation [6, 7]. Legislation sometimes identifies the retention time with the purpose, such that the data has to be wiped off in case the intended purpose has been fulfilled or will be at reach in the future anymore (see § 48 of “Allgemeines Gesetz zum Schutz der öffentlichen Sicherheit und Ordnung in Berlin” for an example).¹

It is obvious, that the importance of each of the four dimensions as perceived by the users as well as their respective willingness to provide information depends on the thematic domain of the service. Some recent work proposed to negotiate the recipient of the data in different application scenarios, among them are medical help [26], distance education [27], and online retailing [5]. We will focus on negotiating the amount of data to be revealed.

Privacy vs. Personalization – User’s Individual Utility Calculus

In order to model the user’s individual trade-off between personalization and privacy, we present it as a utility maximization problem, taking into account different overall sensitivity levels towards privacy and different importance one may assign to a specific privacy dimension. The formalization allows solving the negotiation game presented in section 4, giving the service provider the opportunity to choose its optimal strategy.

¹ The case that the purpose intended at collection time will not be achievable anymore happens in cases like that volatile preference data is kept unanalyzed for an excessively long time. Hence the data is not accurate anymore and thus the planned purpose is out of reach. The data will have to be removed.

We denote the user’s utility by U , using the following notations:

D^n is a n -dimensional privacy space and $d_i \in D$ are its privacy dimensions

The vectors

$$\begin{aligned} \vec{a} &= (a_1, \dots, a_n) \text{ and} \\ \vec{a}^T &= (a_1^T, \dots, a_n^T) \text{ and} \\ \underline{\alpha} &= (\alpha_1, \dots, \alpha_n) \end{aligned}$$

indicate for each dimension d_i of the privacy space the user’s data revelation level, and the revelation threshold as the required minimum to be revealed, and the weighting of each dimension respectively.

γ indicates the user’s global privacy sensitivity

R is the discount provided by the service provider
 P are other non-monetary personalization benefits
 B is the base utility of the (executed) contract

Revealing personal data reduces the user’s utility.

U_{DD} is the disclosure-induced disutility

Using this notation, the user’s utility can be expressed by:

$$U(\vec{a}) = U_{DD}(\vec{a}) + P(\vec{a}) + R(\vec{a}) + B \quad (1)$$

The user maximizes her utility over her decision variables; those are exclusively the revelation levels a_i . The variables \vec{a}^T , $\underline{\alpha}$ and γ cannot be influenced by her (however not implying that they will be constant over time). Hence, the optimization problem is:

$$\max_{a_1, \dots, a_n} U(.) \quad (2)$$

In case that the user is not willing to provide sufficient data for the contract to be executed, the base utility B and the discount R will be zero (3). The user gets some personalization benefits P even if the involved parties do not conclude on a contract. In case P is less than the negative utility the user gets from providing the necessary data, the user will prefer unpersonalized usage of the services (4).

Finally, as the transaction is carried out on a voluntary basis, the user will not agree to a contract generating him a negative net utility. The contract to be executed must provide him at least as much utility as she had before.

In summary, the optimization problem expressed in (2) is subject to the following *participation constraints* (3) - (5):

$$\vec{a} < \vec{a}^T \Rightarrow R(\vec{a}) = 0 \wedge B = 0 \quad (3)$$

$$U_{DD}(\vec{a}) + P(\vec{a}) + R(\vec{a}) < 0 \Rightarrow \text{usage without personalization preferred} \quad (4)$$

$$U(\vec{a}) \geq 0 \quad (5)$$

As the ability to identify a user individually (identity inference, also known as triangulation) does not increase linearly when more data is provided, we use a *Cobb-Douglas utility function* instead of an additive composition for the user's disutility of data revelation. Two other important characteristics of this utility expression in the context of privacy awareness are discussed at the end of this section.

$$U_{DD}(\vec{a}) = -\gamma \cdot \prod_{i=1}^n a_i^{\alpha_i} \quad (6)$$

The *thresholds* a_i^T are set by the service provider and are usually openly communicated. In implementations, hints like 'required field', 'required information' or form fields marked by an asterisk are common practice. The necessity can be deduced from the nature of the transaction: It is obvious that an online bookstore cannot achieve postal delivery if the user refuses to provide her shipping address. It is to note that in this model, the kind of privacy dimensions is not fixed: The purpose as well as the recipient can be privacy dimensions. In the case of shipping, the threshold for the recipient dimension may be the company itself (no third-party logistics company used) and the minimum purpose the user has to agree upon may be postal delivery.

The *weightings* α_i for each of the privacy dimensions as well as the global privacy sensitivity γ are private information of the user and constitute her type. The same holds for the valuation of the non-monetary *personalization benefits* P and the *base utility* B , but these two components can be neglected in the further analysis: First, users tend to only value additional personalization benefits, known solutions will shortly be seen as a standard service and thus there will be no special appreciation. Nevertheless, some personalization benefits may remain. In case of classical implementations such as active guidance, purchase suggestions based on purchase or service usage history, product highlighting or implicit search criteria, the personalization improves the perceived service quality. Through the active support, the user can save search time and simultaneously the matching quality between her

preferences and the store's offers increases: These savings can be seen as monetary benefits and thus subsumed under the variable R . This is especially appropriate, as increased matching quality only becomes effective in case the product is purchased (and R is zero in case of no contract). The base utility can be neglected as it does not depend on the data revelation levels. Hence, the user's type is determined by α_i and γ and the optimization problem (2) can be simplified to:

$$\max_{a_1, \dots, a_n} U_{DD}(\vec{a}) + R(\vec{a}) \quad (7)$$

As mentioned above, the multiplicative structure of the *Cobb-Douglas utility function* allows a good expression of inference threats. In addition, there are two other interesting characteristics in the context of profile data, related to each other. First, the different privacy dimensions are not perfectly substitutable (e.g. the user's telephone number and her e-mail address constitute two possible ways to contact the user but they are not completely interchangeable). Second, different to an additive composition, the substitution rate between two privacy dimensions (which yields here to the ratio - $\alpha_i a_i / \alpha_j a_j$) is not constant or independent from the current level of revealed data: it decreases with the amount of data already provided.

The influences of the different parts on the user's utility function are described by the partial derivatives and their interpretations shown below:

- $\partial U / \partial a_i \leq 0$: Any privacy infringement reduces the user's utility except in the case where she does not care.
- $\partial U / \partial R > 0$: The user appreciates discounts.
- $\partial R / \partial a_i \geq 0$: But the service provider is only willing to grant discounts in case he gets some personal information in return. The case $\partial R / \partial a_i = 0$ is applicable for a privacy dimension irrelevant in the current transaction scenario or (more restricted) for which the service provider does not honour revelation.
- $\partial P / \partial a_i \geq 0$: The more data the service provider can access, the better the personalization will be.
- $\partial B / \partial a_i = 0$: The contracts base utility is independent of the user's revelation level.

Negotiating the 'data'-Dimension

While the recipient may be the relevant negotiation dimension for distance education or health services, we propose the extent and amount of shared data as negotiation dimension for online retailing. First, the willingness of customers to provide personal information is mainly determined by the service provider's reputation, who is the (nonnegotiable) initial recipient of the data. Second, disclosure practices are often determined business processes (e.g. outsourced billing services or delivery by third-party companies). Third, the relevance of the retention time is rated considerably less important [1]. Finally, all data

carries with it a more or less pronounced intrinsic purpose that cannot be subject to a negotiation (e.g. phone numbers are used for personal contact and telemarketing). Hence, negotiating the kind of data seems appropriate in the case of online retailing.

Generally spoken, for a type of data to become part of the negotiation process, it must at least meet the following criteria:

- the user must be able to provide the data
- the data must not be off-topic; the user should see at least a slight reason for the necessity of providing it
- it must not be indispensable for the execution of the contract, either by its nature or by the level of detail (i.e. no negotiations for $a_i < a_i^T$)
- the service provider must gain the user's favour for collecting the data, i.e. the data is part of an explicit profile [17]

The empirical findings of [1] allow establishing a cardinal ordering of types of data according to the willingness of user's to provide the information. Ackerman et al. found significant differences in comfort level across the various types of information, implying weighting factors α_i in the user's utility function constituting one aspect of the user's type. The other aspect, the global privacy sensitivity expressed by γ , will be examined in the following section.

THE SERVICE PROVIDER'S PERSPECTIVE

Facing Different Types of Users

The service provider is confronted with different types of customers that have various global privacy sensitivity levels, and may rate the importance of one kind of data differently. Efficient customer value extraction is based on a combination of discrimination and negotiation techniques. Discrimination relies on the identification of different groups of customers having the same (or a comparable) type. [1] identified three types: the 'privacy fundamentalists', the 'pragmatics', and the 'marginally concerned' users. [18] distinguishes the pragmatic majority into 'profiling averse' and 'identity concerned' users, hence establishing four user clusters.

The distribution of the four types is assumed to be common knowledge. For marginally concerned users, γ will take values close to zero; pragmatic users will have mid-range γ values. Privacy fundamentalist with γ values close to one may be offered static privacy policies as in most cases, the valuation of hiding personal data will be higher than the discounts the service provider can offer; the inequality (4) becomes binding [14].

Modelling the Negotiation Process

Various methods for modelling negotiation processes exist, some more influenced by computer science (e.g. using

finite state machines), others more influenced by microeconomics. We will adopt a game-theoretic approach, examining two possible negotiation scenarios: a sequential game as framework and a simultaneous game that may be played on every step. [4] has examined negotiation protocols in different contexts: customer anonymity (or not), complete knowledge of the service provider's strategy (or not) and transaction costs for both parties (or not).

The service provider's strategy is a function that associates discounts to data revelation level vectors ($D^n \rightarrow \text{ran}(R)$). Determining the service provider's best strategy results in solving the following optimization problem: For users being drawn from a known distribution, maximize the total profit. The total profit is the revenue generated by the whole population minus the granted discounts, minus the costs for implementing the personalization, and minus other costs. Latter encompass in particular customers that are lost during the negotiation process by cancelling (e.g. due to psychological reasons or just because they feel overstrained). This maximization is subject to constraint of the users' participation constraints (3) to (5). We deliberately refrain from a detailed solution, as rigorously integrating the service provider's cost structure would go beyond the scope of this paper.

The framework for the negotiation process is a dynamic game where the service provider has high bargaining power: He opens the negotiation with a basic offer, consisting of a small discount and a few personal data (the threshold) to be asked. This constitutes the fallback offer in case the user does not want to enter negotiation. In case the user accepts, she will be presented another offer with a higher discount and more data to be asked. On every step, the user may *cancel* (i.e. no contract or the fallback solution are implemented), *continue* (i.e. reveal more data or switch to another privacy dimension) to the next step or *confirm* (i.e. the reached agreement will be implemented).

This wizard-like structure is strategically equivalent to a set of offers as (data, discount)-tuples from which the user can choose one. However, a sequential implementation allows better guidance, better communication of the benefits in providing the data and instantaneous adaptation of the strategy. Note that for a given offer, the requested data are always a superset of the requested data of the previous offer, even if the customer only enters the additional information (monotonously increasing revelation level for a given dimension). The service provider can also implement more alternatives for one step, so that the user can choose which data she will provide (for example the service provider can ask either for the home address or the office address). This is particularly useful for addressing different weightings of privacy dimensions that are equivalent for the service provider. Implementations may offer the multiple privacy dimensions sequentially. A switch to another dimension is performed in case the user refuses to provide further data or the service provider is not interested in a

higher detail level for the current dimension. A current implementation is described in the next section.

In this basic case, the service provider grants a fixed discount on every single step, which is cumulated along the process. A more sophisticated procedure could also include the service provider's concessions into the negotiation process, e.g. by a simultaneous game on every stage: the user indicates the minimum discount she wants to get for revealing the data and the service provider indicates the maximum discount he wants to grant. Problems will arise as the service provider's maximal willingness can be overt due to the unlimited number of times one or several anonymous users can play this simultaneous game. [3] argues that this problem can be neglected for multi-faceted and individually valued benefits offered to the user.

IMPLEMENTATION

Privacy dimensions in P3P

The four top-level privacy dimensions (recipient, purpose, retention time, and data) identified above can be mapped directly to P3P. P3P policies express the service provider's data processing practices using STATEMENTS; each of those statements having child elements indicating the RECIPIENT of the data, the PURPOSE for which the data will be used, the RETENTION time and what kind of DATA will be collected.

Other optional child elements of a privacy statement, such as the CONSEQUENCE element or possible EXTENSIONS may not be included in the negotiation process: the consequence is only a short summary or a human-readable explanation of the data processing practices described in the (rest of the) statement. The user agent is supposed to show contents of this element to a (human) user. As for possible extensions, the semantics of issuer-defined additions may be ambiguous and one cannot presume that issuer-defined extensions will be understood by all user agents.

Integrating Privacy Negotiations into P3P

The negotiation process as described in the previous section can be implemented using the already mentioned extension mechanism of P3P, which can be used both in a policy reference file and in a single privacy policy. The extensions in the privacy policies will not be optional, but in order to ensure backward compatibility, these extended policies will only be referenced in an optional extension of the policy reference file. Hence, only user agents capable of interpreting the negotiation extension will fetch extended policies.

In a P3P policy, two extensions can be added: a NEGOTIATION-GROUP-DEF in the POLICY element, and a NEGOTIATION-GROUP in the STATEMENT element. The

mechanism is comparable to the tandem of STATEMENT-GROUP-DEF and STATEMENT-GROUP in P3P 1.1 [25].

The STATEMENT-GROUP-DEF extension is used to define an identifier and optionally properties that can be applied to a group of STATEMENT elements using the STATEMENT-GROUP extension. A statement group allows service providers to describe what sections of their P3P policy apply to different user interactions with their site/service. A statement can be associated with a statement group by having at most one STATEMENT-GROUP extension. A STATEMENT-GROUP element can carry at most two attributes: The id-attribute associates a STATEMENT with a certain group of STATEMENTS to cluster them together. The name-attribute associates a name to a certain statement. User agents may use this name to improve the display of the policy to the user in a human readable format.

A NEGOTIATION-GROUP-DEF element defines an abstract pool of alternative usage scenarios. One or several statements (identified by the attribute id) code a possible usage scenario; the pool membership is expressed by the NEGOTIATION-GROUP extension in the statement (attribute groupid), which describes relevant parameters of the given scenario, such as the benefits for the user. The fallback contract can be indicated via the fallback-attribute of the NEGOTIATION-GROUP-DEF element. The standard-attribute indicates which scenario is offered as default.

The following example illustrates the usage: users of a website can subscribe to a generic or a personalized newsletter (see next page). The generic newsletter will contain only unpersonalized information; the personalized newsletter includes addressing the subscriber per name and promotions targeted towards her interests. Note the additional DATA elements to be collected as well as the additional PURPOSE. The RECIPIENT and the RETENTION time remain unchanged. See [14] for another example about negotiating delivery details of physical and digital goods.

Note that the benefits given in human-readable format need to be displayed concisely by the user agent. The example above shows that the human-readable privacy policy and other information resources on the site must work hand in hand with the P3P policy. The exhaustive machine-readable coding of the benefits is a remaining challenge – especially for multi-dimensional phenomena other than just a reduced purchase price. ebXML and its sub-standards, e.g. the Core Components Technical Specification by UN/CEFACT by the United Nations Centre for Trade Facilitation and Electronic Business, may be used as the basis for further development [22].

<POLICY>

```

<EXTENSION optional="no"> <NEGOTIATION-GROUP-DEF
  id="newsletter" standard="newsletter_personalized" fallback="newsletter_generic"
  description="Choosing newsletter format" /> </EXTENSION>

<STATEMENT>
  <EXTENSION optional="no"> <NEGOTIATION-GROUP groupid="newsletter" id="newsletter_generic"
    description="Generic newsletter without personalization"
    benefits="You get a standard newsletter and no personal data is collected" /> </EXTENSION>
  <CONSEQUENCE>We use your email address for sending you our newsletter.</CONSEQUENCE>
  <RECIPIENT><ours/></RECIPIENT>
  <PURPOSE><contact/></PURPOSE>
  <RETENTION><stated-purpose/></RETENTION>
  <DATA-GROUP><DATA ref="#user.home-info.online.email"/></DATA-GROUP>
</STATEMENT>

<STATEMENT>
  <EXTENSION optional="no"> <NEGOTIATION-GROUP groupid="newsletter" id="newsletter_personalized"
    description="Personalized newsletter, tailored to your personal preferences"
    benefits="You get a personalized newsletter, promoting only the products you are interested in" /> </EXTENSION>
  <CONSEQUENCE> We use your email address for sending you a newsletter targeted to your interests. </CONSEQUENCE>
  <RECIPIENT><ours/></RECIPIENT>
  <PURPOSE><contact/><individual-decision/></PURPOSE>
  <RETENTION><stated-purpose/></RETENTION>
  <DATA-GROUP>
    <DATA ref="#user.name"/><DATA ref="#user.home-info.online.email"/>
    <DATA ref="#dynamic.miscdata"><CATEGORIES><preference/></CATEGORIES></DATA>
  </DATA-GROUP>
</STATEMENT>

</POLICY>

```

Listing 1. Example of an extended P3P policy, including the proposed elements NEGOTIATION-GROUP-DEF and NEGOTIATION-GROUP (fragment, XML namespaces omitted)

Example: Negotiating User Identifiers in Multi-Channel Retailing

In addition to the introductory example of the previous section, we want to outline a possible privacy negotiation for a multi-channel retailer. The scenario is as follows: The service provider wants to address its customers by name and offer them special promotions on their birthday (supposing that they are in a more lavish mood this day). Moreover, users have to choose a login identifier

Two privacy dimensions can be identified:

- the user's name (d_1), with the revelation levels: {none (0), nickname (N), email (E), first name (F), first name and last name (FL)}.
- the user's birth date (d_2), with the revelation levels: {none (0), year (Y), year and month (YM), year and month and day (YMD)}.

Possible negotiation outcomes are depicted in figure 1:




Figure 1. Possible negotiation outcomes as the product of the revelation levels on both privacy dimensions

The revelation thresholds are $a_1^T = \text{none}$ and $a_2^T = \text{nickname}$. Note that the usage of privacy dimensions allows a

unified modelling of personalized usage, pseudonymous usage (e.g. $a_2 = \text{email}$), and anonymous usage (e.g. $a_2 = \text{nickname}$). Introducing the thresholds restricts the number of possible negotiation outcomes (figure 2)




Figure 2. Negotiation outcomes not fulfilling the revelation thresholds are removed

The service provider is aware of the users' general preference structure, as defined in (6): the user's disutility increases when moving to the upper right corner, as the revelation levels increase. However, the service provider ignores the exact positions of the user's iso-utility curves, as those are based on her type \underline{a} and γ – yet private information.

The service provider develops its strategy: he chooses the discounts he will grant to the customer for each of the remaining possible contracts, “labelling” them with the $R(\cdot)$ function (that maps from D^n to discounts):

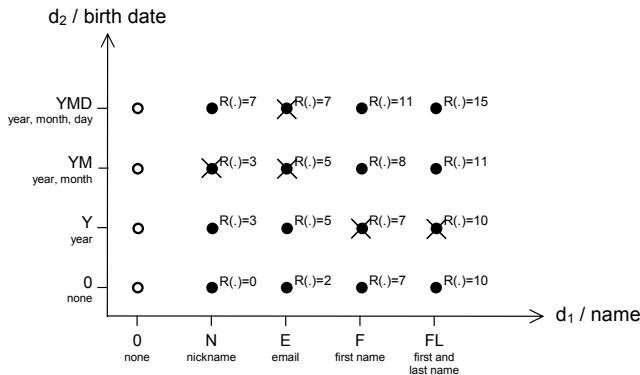


Figure 3. The service provider's strategy $D^n \rightarrow \text{ran}(R)$. Inefficient contracts are crossed

Based on the discount scheme, one can identify inefficient contracts, characterized by revelation levels such that $a^{\text{ineff}} > a^{\text{eff}}$ for the same discount. Figure 4 summarizes the discounts granted by service provider with inefficient contracts removed.




Figure 4. Discounts $R(\cdot)$ offered for efficient contracts

In parallel, the user can determine her disclosure-induced disutility values for the contracts depicted in figure 2. The example is based on a user highly concerned about revealing detailed birth date information as soon as he must provide more than a pseudonymous identifier.




Figure 5. The user's preferences: selected iso-utility curves

For each of the possible contracts shown in figure 2, the disclose-induced disutility values can be determined. Figure 6 shows the disutility values U_{DD} , corresponding to the user's gross utility:




Figure 6. Translating iso-utility curves into disutility levels for each contract (approximate values)

The user's net utility values $U(\cdot)$ for efficient contracts are computed as $U(\cdot) = U_{DD}(\cdot) + R(\cdot)$ and will be used for the optimization problem stated in (7).




Figure 7. The user's net utility values $U(a)$

Users with preferences as outlined will choose the contract (N, YMD), as it is the contract with maximum positive utility value ($U((N, YMD)) = 5$). While revealing birth date at the most detailed level, only pseudonymous data is revealed on the name dimension. Users with other preferences, for example with less reluctance to divulge identifying names, may choose other contracts.

The service provider codes the contracts and the rebate structure shown in figure 4 in a P3P policy using the extensions defined at the beginning of this section. The customer's user agent fetches the policy and serves as a negotiation support system, displaying possible alternatives (a human-readable communication of the data handling practices as coded in the statements along with negotiation benefits) from which the user can choose one.

We have integrated basic negotiation support into the Mozilla Web browser, thence extending its P3P support: a site's privacy policy can be accessed via the "Policy", "Summary" and "Options" buttons in the "Page Info" dialog, directly available from the status bar. Extending the chrome components, we have added a "Negotiate" button: a modal dialog is opened, summarizing the negotiable privacy dimensions (d_i) and the possible realizations (a_i) with drop-down menus. The implementation relies on XUL and JavaScript, uses the Mozilla APIs and integrates seamlessly into the user agent. As the proposed extension to P3P is not restricted to a specific privacy dimensions, neither is the implementation. Any privacy dimension can be negotiated as long it can be expressed using the P3P data scheme. A more sophisticated support relying on an improved XML Schema Definition of our extensions for privacy negotiations is currently under development and will be available by spring 2006 approximately. Installation features will be included to allow easy deployment on multiple devices.

CONCLUSION AND FURTHER WORK

This paper has presented the necessity of negotiation about privacy principles in a relationship between service provider and customer. Negotiating allows a better matching between the seller's needs and the buyer's disclosure restraint and helps to reduce the trade-off between personalization and privacy. Modelling the user's individual utility maximization can take into account the multi-dimensionality of privacy; the service provider may wish to reduce the negotiation space in a way that suits the given business scenario. The incremental revelation of data by the user can be strategically reduced to a choice from a set of alternatives. Using the extension mechanism of P3P, there is no limitation in coding these alternatives even for complex cases involving diverse privacy dimensions: We proposed two new elements that follow the structure of the current P3P 1.1 grouping mechanisms and allow software-supported negotiations in E-Commerce. Software support of the extension was added to the Mozilla browser, integrating privacy negotiations seamlessly into the user agent.

Future work will focus on the practical implementation of privacy negotiation techniques on large scale public websites. We are currently investigating which user interface design best fulfils the usability requirements and how negotiable privacy dimensions are best visualized. Moreover, a taxonomy should be developed to allow a machine-readable coding of the user's benefits for a negotiation alternative. A remaining question is whether users feel more concerned about their privacy when an explicit negotiation process is started. This increasing sensitivity could make take-it-or-leave-it offers more favourable for the service provider.

REFERENCES

1. Ackerman, M. S., Cranor, L.F., Reagle, J.: Privacy in E-commerce: Examining User Scenarios and Privacy Preferences, First ACM Conference on Electronic Commerce, Denver, CO (1999) 1-8
2. Agrawal, R., Kiernan, J., Srikant, R., and Xu, Y. An XPath-based preference language for P3P. In Proceedings of the Twelfth International Conference on World Wide Web, pages 629–639. ACM Press (2003)
3. Buffett, S., Jia, K., Liu, S., Spencer, B., Wang, F.: Negotiating Exchanges of P3P-Labeled Information for Compensation, Computational Intelligence, Volume 20, Number 4 (2004)
4. Cranor, L. F., Resnick, P.: Protocols for Automated Negotiations with Buyer Anonymity and Seller Reputation, Netnomics, 2(1), 1-23 (2000)
5. El-Khatib, K.: A Privacy Negotiation Protocol for Web Services. Proceedings of the International Workshop on Collaboration Agents: Autonomous Agents for Collaborative Environments (COLA) (2003)

6. European Parliament, Council of the European Union: Directive 2002/58/EC on privacy and electronic communications. Official Journal of the European Communities, 31.7.2002, L 201, 37–47 (2002)
7. European Parliament, Council of the European Union: Regulation (EC) No 45/2001 of the European Parliament and of the Council of 18 December 2000. Official Journal of the European Communities, 12.1.2002, L 8, 1–22 (2002)
8. Gross, P., Gross, Peter (1994): The multi options society (original title: Die Multioptionsgesellschaft), Suhrkamp, Frankfurt am Main (1994)
9. International Business Machines Corporation, Calvin Powers, Matthias Schunter (Editors): Enterprise Privacy Authorization Language (EPAL 1.2), W3C Member Submission 10 November 2003 (2003)
10. Karrass, C. L.: Give and Take: The Complete Guide to Negotiating Strategies and Tactics. HarperCollins Publishers, New York, NY (1993)
11. Kurashima, A., Uematsu, A., Ishii, K., Yoshikawa, M., Matsuda, J.: Mobile Location Services Platform with Policy-Based Privacy Control (2003)
12. Peppers, D., Rogers, M., Dorf, B.: The One to One Fieldbook. New York, Currency Doubleday (1999)
13. Personalization Consortium: Personalization & Privacy Survey (2000)
14. Preibusch, S., Implementing Privacy Negotiations in E-Commerce. in: Frontiers of WWW Research and Development - APWeb 2006: 8th Asia-Pacific Web Conference (APWeb 2006), Harbin, China. LNCS 3841, 604-615 (2006)
15. Rebstock, M., Thun, P., Tafreschi, O.A.: Supporting Interactive Multi-Attribute Electronic Negotiations with ebXML. Group Decision and Negotiation. 12 (2003) 269–286
16. Salomann, H., Dous, M., Kolbe, L., Brenner, W.: Customer Relationship Management Survey, Status Quo and Further Challenges, University of St.Gallen (2005)
17. Schubert, P.: Virtual Virtuelle Transaktionsgemeinschaften im Electronic Commerce, Josef Eul Verlag, Lohmar, Köln (1999)
18. Spiekermann, S.: Online Information Search with Electronic Agents: Drivers, Impediments, and Privacy Issues (2001)
19. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. in EC'01: Third ACM Conference on Electronic Commerce. Tampa, FL, 38-47 (2001)
20. Ståhl, I.: Bargaining Theory. Stockholm: The Economics Research Institute (1972)
21. Thompson, L.L.: The Mind and Heart of the Negotiator. 3rd edn. Pearson Prentice Hall, Upper Saddle River, New Jersey (2005)
22. United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT): Core Components Technical Specification – Part 8 of the ebXML Framework, Version 2.01 (2003)
23. W3C, A P3P Preference Exchange Language 1.0 (APPEL1.0), W3C Working Draft 15 April 2002, <http://www.w3.org/TR/P3P-preferences> (2002)
24. W3C, The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation 16 April 2002, <http://www.w3.org/TR/P3P/> (2002)
25. W3C, The Platform for Privacy Preferences 1.1 (P3P1.1) Specification”, W3C Working Draft 10 February 2006, <http://www.w3.org/TR/2006/WD-P3P11-20060210/> (2006)
26. Yee, G., Korba, L.: Feature Interactions in Policy-Driven Privacy Management. Proceedings from the Seventh International Workshop on Feature Interactions in Telecommunications and Software Systems (FIW'03) (2003)
27. Yee, G., Korba, L.: The Negotiation of Privacy Policies in Distance Education. Proceedings. 4th International IRMA Conference (2003)