A secure and privacy-preserving web banner system for targeted advertising

Joris Claessens[1], Claudia Díaz, Raimondo Faustinelli, and Bart Preneel

COmputer Security and Industrial Cryptography (COSIC)
Dept. of Electrical Engineering – ESAT
Katholieke Universiteit Leuven, Belgium
http://www.esat.kuleuven.ac.be/cosic/

Abstract

A solution for privacy-preserving targeted advertising through web banners is proposed. The solution allows users to make a balance between the exposure of their privacy and the personalization of advertisements. The general idea of the solution lies in dynamically associating users with profiles according to their interests and/or demographics instead of to individual and unique identifiers. Users have complete control over the content of these profiles. In addition a number of security enhancements are suggested. The solution relies furthermore on an infrastructure for anonymous communication. A proof-of-concept of the web banner system is being developed.

Keywords: privacy, web banners, targeted advertising

[1] Joris Claessens currently works at the European Microsoft Innovation Center. The work presented in this paper was performed while he was at COSIC.

1 Introduction

Anonymity and privacy are of paramount importance within the various electronic services provided in today's expanding digital society. In this paper we focus on the service of targeted advertising. In particular we assume users are potentially interested in receiving web banners with advertisements of their personal interest, while surfing the web. However, they may – and should – not want to give up their privacy in this context. We show that personalized, tailor-made services do not necessarily require true identification, and propose a secure and privacy-preserving web banner system for targeted advertising.

Web banner based targeted advertising has received much negative attention in the past. The most infamous example is probably the DoubleClick case [9]. DoubleClick and many other Internet advertising companies track Internet user behaviour in order to better target banner advertisements. For this purpose they assign a unique identifier to each individual user. This unique identifier is put in a cookie that is sent along with the web banner and that is stored on the user's computer. The cookie is sent back to the banner server with each subsequent banner request. Without any more information, a unique identifier cannot directly be linked to a user's real identity, and hence seems only a harmless pseudonym. However, if the real identity of a user is disclosed for one single web transaction, this identity can be linked to the complete chain of all the other transactions performed by that user. This was exactly the great concern of privacy organizations when DoubleClick merged with Abacus Direct, a giant in off-line marketing information, as merging the collected information of their respective databases allowed linking names with originally pseudonymous user activity on the web, and thus constituted a major threat to the privacy of the individual users. As a consequence, legal complaints were filed alleging among others that DoubleClick violated its own earlier stated privacy policies which claimed that collected information would remain anonymous.

Current web banner systems are without doubt a big risk with respect to the user's privacy. Banner servers deploy mechanisms by which they can keep track of individual users. Since web banners are everywhere, a vast amount of information can be collected and complete user profiles can be maintained. Although individual users are initially only identified with pseudonyms, at some point in time the user profiles will likely to be linked with the real identities of the users. Many users are unfortunately unaware of this. Their privacy is only protected by law, but is not guaranteed by technical means. Anonymity as a technical tool is therefore a good starting point for a correct implementation of data protection legislation (even if in some circumstances there will be overriding rights which do not allow the use of anonymity; note that this is obviously not the case with web banners). From the European Union's point of view, this starting point can be deduced from two principles in the general data protection Directive: personal data must be kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the data were collected or for which they are further processed; and moreover, personal data should not be ‘excessive’ in relation to the purposes for which they are processed. In this paper we state that unique identifiers are excessive for the purpose of targeted advertising. We therefore propose an alternative and privacy-preserving mechanism for ‘personalized’ web banners. Our proposed system meets the interests of both the users and the banner servers. The users can enjoy web banners of their personal interest, while their privacy remains protected. The banner servers can continue their business, while they need to worry less about the protection of personal data as the data they are able to collect has already been anonymized.

Outline of the paper

We start the paper with a discussion of the overall privacy threats on the World Wide Web. Section 3 then explains the concept of targeted advertising with web banners and its privacy and security issues. The secure and privacy-preserving web banner system is presented in Sect. 4. Related and complementary solutions for anonymity on the WWW are described in Sect. 5. Finally we give our conclusions and indicate future work.

2 Anonymity on the World Wide Web ?

When surfing the World Wide Web (WWW), users are subjected to many potential privacy threats. Although many are concerned about these threats, some users are still not aware of them, while others do not seem to care about on-line privacy, as they may not know the potential consequences.

Information about a user’s identity can be revealed in two ways: at the application level, that is, through the content itself that is exchanged, and at the network level, that is, via the network address of the machine from which the user is surfing.

Threats to anonymity at the application layer

A user’s browser discloses a substantial amount of information within the *HTTP headers*, that can help identifying a particular user, and that can help building a complete profile of the user’s interests and behaviour. For example, the Referrer header informs the server about the page that contained the link the user is requesting at that moment (this is how a banner server knows on which page a banner is placed). The User-Agent header informs the server about which browser on what platform is being used. Last but not least, *cookies* [13] constitute the key tool to maintain user profiles. Cookies are little pieces of information that a web server can request the browser to store on the user’s machine. When the user returns to that server, the cookie is retrieved by the browser and transmitted again to the server. Cookies can contain any type of information. They are mainly used to establish (authenticated) sessions in an e-service, or to store a user's personal preferences. However, in many current banner systems they contain a unique identifier which allows the banner server to link the information collected in different sessions to the same user, and thus to maintain and keep complete user profiles.

Besides data that is transparently sent without the user really noticing it, personal information is often requested explicitly to the user. The user should trust the web server to take proper care of this personal information. Many web sites require users to provide a username, password, and an e-mail address. This allows the web site to offer a personalized service. Unfortunately, users will mostly choose easy-to-remember usernames that can be associated with the real identity of the user. The same e-mail address will often be used, making it very easy to link different usernames to each other. Note that users will frequently choose the same passwords too, so if one password gets compromised, this may give an adversary access to other (more sensitive) services too. One cannot expect that ordinary users have different usernames, passwords and email addresses for all of the web sites that they visit.

Finally, besides the browser itself, third-party browser add-ons (e.g., that show comments posted by other users of the web site the user is currently visiting, or that show an updated list of related sites, etc.) are potentially dangerous for the privacy of a user.

Threats to anonymity at the network layer

More fundamentally, at the network level, *IP addresses* are required for establishing communication between browser and web server. In many cases, these addresses can be linked to a limited set of persons, if not one person. An IP address may thus identify an individual to some extent, or at least may significantly help identification when combined with other information.

3 Targeted advertising with web banners

This section explains the concept of targeted advertising with web banners and investigates its privacy and security issues.

3.1 Model and participants

In a web banner system for targeted advertising, there are essentially four different parties:

- *Users*. The users surf to various web sites. We assume that they are potentially interested in receiving advertisements that are related to their personal interests. Accordingly, the banner server will try to include the appropriate web banners with personalized advertisements into the web pages the users surf to.

- *Web sites*. The different web sites want to include web banners in their pages, as they can earn money, per banner that is shown to the users via their pages. In addition, if banners are shown that are interesting for the user, it might make the web site more attractive.

- *Banner Server*. The Banner Server's business model is to make money by serving banners to the appropriate category of users. For this purpose, the users' profiles are of crucial importance. Tracking of users is also desired for choosing the appropriate banner (e.g., the same banner should be shown 3-7 consecutive times), and for marketing reports.

- *Banner Payer.* The Banner Payers will pay for the service offered by the Banner Server. Via the banners, they hope to attract users to their web sites, products and services.

From a brief technical point of view, the current banner systems for targeted advertisement can be outlined as follows:

- Banner requests are associated with a unique identifier. This can be done based on the IP address of the user and/or based on storing a unique identifier in the cookie that is delivered to the user together with the banner.

- The unique identifier is associated with a user profile, a database which contains information related to the user's activities on the web. While the unique identifier allows to link all actions of the same user, it should not directly be related to the user's real identity. However, the risk exists that sooner or later the real identity of the user will be revealed to the banner server for one single action, hereby linking all other actions to this user.

- When the banner server receives a banner request, it updates the profile associated with the unique identifier that is sent along with the request, according to the visited web page, and sends back the appropriate banner. Current banner systems operate in a so-called "server-side measurement" setting.

- There are two ways of inserting banners into web pages. In the “invasive” way, the web site includes a link to the banner in the web page; the link possibly contains extra information regarding the category of the page; the actual banner request is automatically performed by the user’s browser. There is a direct interaction between the user and the banner server, allowing the banner server to track the user with a persistent and unique identifier.

  The other way is “automatic insertion of code at web site”, whereby banners are transparently added to the web pages at the web server’s side, before sending the pages to the user; the actual banner request is performed by the web server, and the banner is sent to the user together with the web page. There is no direct interaction between the user and the banner server, but this does not prevent the banner server to track the user, indirectly via the web server, with a persistent and unique identifier. Thus, the same privacy threats may also apply in this situation.

  Note that our proposed solution for privacy-preserving web banners deploys the invasive way of inserting banners, as will follow from the description below. It should however be easily deployed in the other way as well.

3.2 Privacy and security issues

Some important privacy and security issues can be identified in current web banner systems for targeted advertising.

- *Privacy.* As explained before, banner requests are normally associated with a unique identifier that corresponds to a specific individual user. Banner servers can therefore gather an incredible amount of personal information with which they can build up complete profiles of the users' activities on the web. This is (or should be) an important concern for the users.

- *Security.* Banner server and/or web sites possibly want to cheat in order to gain more money. For example, malicious web sites could perform numerous requests of their own web pages, so that it seems that many users received banners through these pages. The web sites can then claim more money from the banner server. The banner server could even cooperate itself, and claim more money from the banner payers.

Privacy is an important concern of the users. Security is a concern of the other three parties. The solution that is presented in the next section mainly intends to address the privacy issue. However, as the privacy-preserving solution makes the security threats worse, specific security enhancements are proposed as well.

4 Privacy-preserving web banners

A trivial solution for guaranteeing a user's privacy is simply blocking banners (and cookies). However, we assume that users are potentially interested in receiving banners with advertisements of their interest. Obviously, banners can be made more personalized if more personal information is exposed. The privacy of the user should however still be guaranteed. More precisely, users should be able to make a balance between the exposure of their privacy and the personalization of the advertisement.

4.1 Main solution

The general idea of our solution consists of directly associating users with profiles according to their interests and/or demographics instead of to individual identifiers. Consequently, the profile will be maintained by the user instead of by the banner server. Users with similar behaviour will have the same profile. Furthermore, a user's profile can change over time. In other words, an individual user is not directly associated with a unique profile that is maintained by the banner server, but can dynamically belong to a group of users that are all associated with the same profile. In this way, users do not expose their individual actions on the web, but can hide within the group of users who have the same profile at a specific time. This idea adopts the definition given by Pfitzmann and Köhntopp in [14]: *“Anonymity is the state of being not identifiable within a set of subjects, the anonymity set.”* Note that this definition indicates that anonymity is always relative to a certain set of subjects: in an application an individual will always be identifiable as being a member of a certain group, though not identifiable as being that particular individual. Therefore, the level of anonymity that a system can provide depends on the mechanisms used and on the number of individuals that are using the system.

Architecture and protocol

Figure 1 illustrates the architecture and protocol of the proposed privacy-preserving web banner system. The four different parties and the interactions between them are indicated. The web banner system consists of the following steps:

[Figure 1]

1. The user requests a web page from a particular web site.

2. The web site delivers the requested web page. The web page contains a link to a banner. It is possible for the web site to include extra information about the web page with the link (i.e., the type of information on the page – the interest category – can be indicated, for example ‘sports’, ‘movies’, ‘news’, ‘computers’, etc).
3. Upon receipt of the web page, the user’s browser will automatically follow the link and request a banner from the banner server.

   This request is intercepted by a client proxy running on the user’s machine, which adds the user’s profile to a cookie in the request.

   The request thus includes the user’s profile (in a cookie), the referring address of the web page (‘referrer’ header), and extra information about the web page (in the link).
4. Based on the user’s profile, the address of the web page, and the extra information provided by the web site, the banner server chooses the appropriate banner and sends it back to the user. Together with the banner, the banner server will transmit a new cookie containing information with which the user’s profile can be updated.
5. The banner is shown to the user.

   The profile update is performed by the client proxy. The new profile will be sent with the next banner request. A user’s profile will thus change over time, which means that users can dynamically belong to different groups of users with the same profile. Note that the profile can also depend on static data (e.g., demographics and user preferences) which can be altered by the user at any time.

6. At some point in time, the banner server will receive money from the banner payer for having served the banner. The web site will receive its share of the money.

The client proxy is a key component in our solution. It is running on the user's machine, acts in between the user's browser and the WWW, and takes care of the user's profile.

Additionally, all requests should be performed over an anonymous network. Otherwise the user's IP address would be disclosed to the banner server. This address in itself (or at least a range of or a limited set of addresses) would be a unique identifier with which different sessions of the same user could be linked to each other. Moreover, the user's IP address constitutes information that may help to disclose the user's identity. Thus, in addition to maintaining the user's profile and controlling the requests of banners, the client proxy is also the access point to the anonymous network.

Information theory and the degree of anonymity

In [6] we propose an information theoretic model with which we can quantify the degree of anonymity in applications, such as privacy-preserving web banners, where users are divided into groups in which they are indistinguishable from each other. A more generic model for measuring the degree of anonymity in communications has been proposed in [7].

The information theoretic model is particularly suited for measuring the quality of our solution. The model will help to define appropriate user profiles, i.e., user profiles that allow a proper balance between targeted, personalized advertising, and exposure of individual privacy. In other words, the more different possible profiles, the more functionality, while the more users with the same profile, the more privacy.

4.2 Server-side versus client-side web banner systems

A web banner system for targeted advertising essentially consists of two different processes. On the one hand the system should provide a way of maintaining profiles and presenting the right profile with a banner request. On the other hand the system should select appropriate banners for presented profiles. These processes can be performed at the client-side (i.e., on the user’s machine) or at the server-side (i.e., by the banner server). Table 1 situates three different solutions in this respect.

[Table 1]

Traditional solution

In the traditional solution, both processes are performed on the server-side. Users are associated with a unique identifier which allows the banner server to maintain a complete profile. This is a clear threat to privacy. The appropriate banner is chosen by the banner server.

Juels' solution

Juels [12] proposes a solution in which both processes are performed on the client-side. The user's preferences and profile are maintained at the client side, and an appropriate banner is selected by software (the 'negotiant'), running on the user's machine, and typically originating from the banner server. The banner is requested through an anonymous communication network.

Juels' solution is excellent in terms of privacy as individual users are not identified nor is their profile disclosed. However the solution seems to lack flexibility and may require the banner selection software to be periodically updated. This seems not very practical. The solution also requires measures which ensure that the selection software cannot leak the identity or other personal information of the user (problem of 'malicious negotiants').

Our solution

Our solution has both client-side and server-side properties. On the client-side, the profile is maintained and chosen by the user. The profile of a user is based on the complete history of the user's actions, and on static information that can optionally be declared by the user (e.g., age range, country, interests and gender). The users have full control over the information that is sent to the banner server. They may change or delete their dynamic and static profile information whenever they desire. The speed with which their profile can change can also be controlled by the users, depending on the relative weight of the most recent actions in the history. The appropriate banner for a particular profile is chosen on the server-side.

As an individual user can hide within a group of users with the same profile, our solution is still very good in terms of privacy. The actual level of privacy depends on the level of detail in the profile, and can be measured by using the model developed in [6]. Our solution is also very flexible with respect to the banner selection process. The banner server can completely control the choice of banners without needing any trust of the user with respect to privacy.

4.3 Security enhancement

The privacy-preserving web banner system as described above, can be enhanced to improve security.

Malicious parties could perform multiple fake banner requests in order to gain more money. Current banner systems try to tackle this by for example checking if there are not too many requests with the same unique identifier or originating from the same IP address, etc. Thus, fraud *detection* is already included in these systems to some extent. Although these techniques are relatively easy to circumvent, they seem to be good enough in today's banner systems.

As our proposed system does not rely on unique identifiers nor discloses the users' IP addresses, the fraud detection mechanisms are rendered useless. Thus, without any further precautions fraud on a large scale would be possible.

The proposed privacy-preserving web banner system can therefore be enhanced with fraud *prevention* techniques. Note that these techniques would obviously also constitute a security improvement in the current banner systems. Fraud prevention techniques make multiple fake banner requests hard in the first place, instead of (or in addition to) checking against it afterwards. More precisely, the security enhancement aims at allowing banner requests in the scope of normal usage, and preventing banner requests at a higher rate. Two mechanisms can be used for this purpose.

Users can be required to present a valid cryptographic ticket with each banner request. Blinded tickets can be very similar to anonymous cash [4] but with no financial value, and can be obtained by authenticated users. Only a limited amount of tickets would be periodically issued by the banner server (or by a separate trusted party) to each user.

Alternatively, certain ‘client puzzles’ [5] or ‘pricing functions’ [8] could be required to be solved before being able to request a new banner. I.e., the user’s machine should solve some computationally intensive problem (e.g., finding a collision of a hash function) for each new banner request; the problem should still be easy enough to allow normal browsing (e.g., one banner per 5 seconds), but should be difficult enough to prevent large-scale fraud.

4.4 Proof of concept

We have developed a proof of concept to demonstrate our solution for privacy-preserving web banners. More technical details can be found at [2]. The proof of concept consists of a banner server for which a simple banner selection procedure is implemented in PHP, and client software implemented in Java. We rely on JAP [11] to achieve anonymity at the communications level.

The client software is in essence an HTTP proxy which acts between the user's browser and the Internet directly or indirectly via the JAP anonymity proxy. The client software includes functionality to create a profile based on the user's history and declared information, to add this profile to a cookie in a banner request, and to update and maintain the user's history. A graphical user interface allows the user to view the complete history and the declared information and to alter or delete this data whenever desired.

The client proxy does not disclose to the browser the content of the profile cookie sent to the banner server nor the update information received from the banner server. The user profile management is solely performed by the client proxy. This has as an advantage that a malicious browser cannot build a profile of the user. Thus, a user could surf the web with a browser on an untrusted machine, and still use the proxy on remote trusted machine. A mechanism for authenticating the user to the remote banner proxy and storing the user’s history and declared information in encrypted form should therefore be added in the future.

5 Related and complementary work

As indicated in Sect. 2, the user’s privacy is at stake in many ways. A number of privacy-enhancing technologies exist that can be seen as complementary to our work.

The W3C developed the *Platform for Privacy Preferences (P3P)* [17]. P3P provides an automated way for users to gain more control over the use of personal information they send to web sites. In particular, P3P is a standardized way of formulating privacy policies. These privacy policies are presented by servers, can be understood by browsers, and can also be interpreted directly by the user.

A web proxy can hide the user’s IP address from a web server (e.g., the banner server). In addition, a web proxy can provide anonymity services at the application level too, by for example stripping identifying information and/or providing secure management of sensitive information such as cookies and usernames/passwords. An example of such a web proxy is the Anonymizer [1]. Another example is the Lucent Personalized Web Assistant (LPWA) [10] which provides users with a *different*, *anonymous*, and *unlinkable username/password* and *e-mail address* for each different web site, while users only have to remember one secret. Before browsing the WWW, users have to login into the LPWA by giving their identity and their secret. From then on, the LPWA is used as an intermediate web proxy. The LPWA transforms the identity, the secret, and the URL of the web site into a username, a password, and an e-mail address that will be used for that web site.

A web proxy is only a basic solution that protects against *local observers* (e.g., the web server itself). More advanced solutions will have to protect against *powerful observers*, who are able to overview the *global* network. These solutions should protect against eavesdropping, not for confidentiality purposes, but preventing the content of messages to be traced from destination to source. They should also protect against *traffic analysis*, to prevent messages to be traced based on size and timing measurements. Chaum's *mix* [3] forms the basis for the more advanced solutions. The messages of all parties who want to communicate anonymously are sent through the mix. The mix hides the correspondences between messages in its input and those in its output. The mix hides the order of arrival of the messages by reordering, delaying and padding traffic. As web traffic is real-time and bidirectional, for example, delaying is not really possible. Practical solutions therefore require a *chain of mixes* in order to provide an adequate level of anonymity. Onion Routing [15] and JAP [11] are examples of such a solution. Crowds [16] is an alternative solution for anonymous communication in which users can hide in a crowd by forwarding web requests among each other before sending them to the destination.

6 Conclusions and future work

We have proposed a secure and privacy-preserving web banner system for targeted advertising. We feel that our solution can be very useful to balance the requirements of the data protection legislation with the genuine interests of the user and the commercial interests of web sites, banner servers and banner payers.

As more users are becoming concerned about privacy, many products are being developed that control the information that is sent to and received from the Internet: personal firewalls, password managers, form fillers, cookie managers, banner managers, keyword alerters, etc. The latest browsers also include part of this functionality. Our solution is complementary to these products and provides privacy to a specific personalized web service. It is clear that our solution could therefore be integrated into browsers and/or independent applications to extend their privacy-enhancing functionality.

The proposed solution only prevents tracking based on cookies associated with web banners. Cookies are used in several other circumstances for personalized interaction, and for session identifiers. Individualized services are often needed in these circumstances. A solution based on hiding in a group may thus not be useful in that case.

It may be interesting to explore how our solution can be deployed by an intermediate party as a privacy-enhancing banner service. This party could request banners to existing banner server on behalf of the users. The intermediate party's task is then to map the profile presented by the user, to the unique identifier required by the particular existing banner server. The advantage of this approach is that existing banner systems would not have to change their internal system to be able to work with profiles instead of unique identifiers, even if they want to adopt privacy-preserving techniques. The intermediate banner service does not have to be trusted but to map the profile of the user to a corresponding relevant unique identifier.

Our web banner system has been designed to prevent tracking of individual users, as this is desired for privacy, and allows advertisements to be targeted to users with the same profile. The system however also prevents to show the same banner a number of consecutive times to the same individual user, as required by common marketing rules. A more advanced privacy-preserving web banner system should therefore support tracking of users over a short period of time (similar to session identification) and still prevent tracking of users over a long period of time.

As the degree of privacy depends on the level of detail and the amount of information in the profile, it is interesting to enhance the solution by allowing multiple profiles with different information. This would allow a user to play with different levels of information granularity in the different profiles, hereby providing the flexibility of a more personalized service, while still remaining an equivalent degree of privacy. For example, while one profile of a user may include a generic expression of interest in movies in addition to many other information, another profile could only indicate a specific interest in action movies. This other profile should then only be revealed when dealing with movies, and should be unlinkable to the first profile.

Acknowledgements

This work was supported in part by the IWT STWW project on Anonymity and Privacy in Electronic Services (APES) [2], and by the Concerted Research Action (GOA) Mefisto-2000/06 of the Flemish Government. The authors want to thank Jan Teerlinck of Data4s for the helpful discussion on web banner technology, and Caroline Goemans and Prof. Jos Dumortier of ICRI, K.U.Leuven, for their legal input on web banner systems.

References

[1] Anonymizer. http://www.anonymizer.com/.

[2] APES. Anonymity and Privacy in Electronic Services. https://www.cosic.esat.kuleuven.ac.be/apes/.

[3] Chaum, D. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM 24*, 2 (February 1981), 84–88.

[4] Chaum, D. Blind Signatures for Untraceable Payments. In *Advances in Cryptology – CRYPTO'82* (August 1983), D. Chaum, R. L. Rivest, and A. T. Sherman, Eds., Plenum Press, pp. 199–203.

[5] Dean, D., and Stubblefield, A. Using Client Puzzles to Protect TLS. In *Proceedings of the 10th USENIX Security Symposium* (August 2001).

[6] Díaz, C., Claessens, J., Seys, S., and Preneel, B. Information Theory and Anonymity. In *Proceedings of the 23rd Symposium on Information Theory in the Benelux* (May 2002), B. Macq and J.-J. Quisquater, Eds., pp. 179–186.

[7] Díaz, C., Seys, S., Claessens, J., and Preneel, B. Towards measuring anonymity. In *Proceedings of the 2nd Workshop on Privacy Enhancing Technologies* (February 2003), R. Dingledine and P. Syverson, Eds., Lecture Notes in Computer Science, LNCS 2482, Springer-Verlag.

[8] Dwork, C., and Naor, M. Pricing via Processing or Combatting Junk Mail. In *Advances in Cryptology – CRYPTO'92* (August 1992), E. F. Brickell, Ed.,

Lecture Notes in Computer Science, LNCS 740, Springer-Verlag, pp. 139–147.

[9] Electronic Privacy Information Center. Double Trouble. http://www.epic.org/doubletrouble/.

[10] Gabber, E., Gibbons, P. B., Kristol, D. M., Matias, Y., and Mayer, A. On Secure and Pseudonymous Client-Relationships with Multiple Servers. *ACM Transactions on Information and System Security 2*, 4 (November 1999), 390–415.

[11] JAP Anonymity & Privacy. http://anon.inf.tu-dresden.de/.

[12] Juels, A. Targeted advertising... and privacy too. In *Topics in Cryptology – Proceedings of the Cryptographers' Track at RSA 2001* (April 2001), D. Naccache, Ed., Lecture Notes in Computer Science, LNCS 2020, Springer-Verlag, pp. 408–424.

[13] Kristol, D. M. HTTP Cookies: Standards, Privacy, and Politics. *ACM Transactions on Internet Technology 1*, 2 (November 2001), 151–198.

[14] Pfitzmann, A., and Köhntopp, M. Anonymity, Unobservability and Pseudonymity – A Proposal for Terminology. In *Designing Privacy Enhancing Technologies* (2001), H. Federrath, Ed., Lecture Notes in Computer Science, LNCS 2009, Springer-Verlag, pp. 1–9. Proceedings of the Workshop on Design Issues in Anonymity and Unobservability, July 2000.

[15] Reed, M. G., Syverson, P. F., and Goldschlag, D. M. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in*

Communications 16, 4 (May 1998), 482–494. Special issue on Copyright and Privacy Protection.

[16] Reiter, M. K., and Rubin, A. D. Crowds: Anonymity for Web Transactions. *ACM Transactions on Information and System Security (TISSEC) 1*, 1 (November 1998), 66–92.

[17] World Wide Web Consortium. Platform for Privacy Preferences (P3P). http://www.w3.org/P3P/.

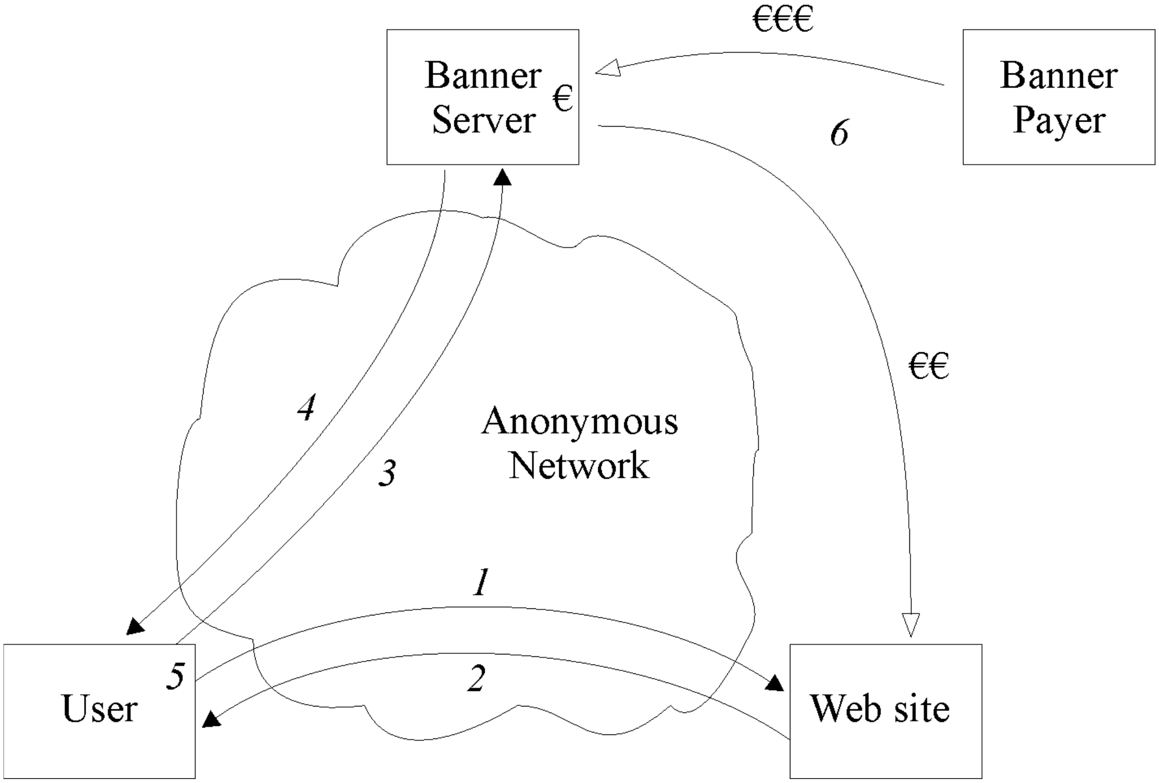

1. User requests web page
2. Web site delivers page; the page contains link to banner and possibly more info about category
3. User transparently requests banner; the request includes the user's profile
4. Banner Server serves banner; information for updating the profile is included together with the banner
5. The banner is shown to the User; the user's profile is updated
6. Banner Payer pays Banner Server; Web site receives part of the money

Figure 1: The privacy-preserving web banner system

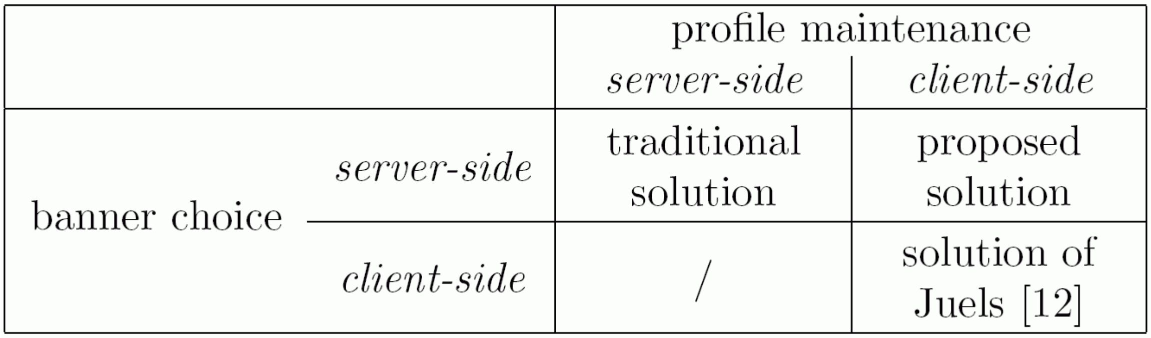

		profile maintenance	
		server-side	*client-side*
banner choice	*server-side*	traditional solution	proposed solution
	client-side	/	solution of Juels [12]

Table 1: Server-side versus client-side web banner systems