The Cookie Conundrum: Examining the Dual Nature of Cookies

Vicki Ha
Lina Hdeib
Farah Al-Shaar
Kori Inkpen

Technical Report CS-2005-18

October 12, 2005

Faculty of Computer Science
6050 University Ave., Halifax, Nova Scotia, B3H 1W5, Canada

The Cookie Conundrum: Examining the Dual Nature of Cookies

Vicki Ha, Lina Hdeib, Farah Al Shaar, Kori Inkpen
Faculty of Computer Science, Dalhousie University
6050 University Ave., Halifax, Nova Scotia, Canada B3H 1W5
{vha,lhdeib,alshaar,inkpen}@cs.dal.ca

Abstract

Proper cookie management methods have long been the source of frustration to consumers and researchers alike. This is due to the fact that cookies have a dualism in the way that they can be both beneficial and malicious, unlike other malware which are intrinsically malicious. Because of this duality, cookie management requires a subjective component unlike managing other malware. This in turn requires awareness and control on the part of the user in order to subjectively manage cookies. In this paper, we show the social and technical considerations required because of this duality, and the consequences which result from these considerations. We will also show, through examining the results of focus group sessions, that an increase in awareness is the best partial solution to the privacy problems associated with cookies.

Keywords

World Wide Web, privacy, cookie management, cookies, cookie dualism, privacy awareness, focus group, third-party cookies, , P3P, anonymity.

1. Introduction

HTTP applications on the World Wide Web are naturally stateless and will not store information about users across multiple sessions. Lou Montulli introduced cookies in 1995 [5] to circumvent this problem by allowing persistent state information to exist in the form of cookies. Cookies are small text files, containing information obtained from the user, that are placed on the client-side of a Web connection by the server to which it is connected. These cookies can later be retrieved by the web server, which then extract the information stored in them.

However, cookies have since been used for malicious intent, placing them at the centre of many controversial issues, including invasions of privacy and informed consent violations. In 2002, two lawsuits were filed against DoubleClick Inc. alleging that the company was using cookies without proper consent to track the activities of Internet users and using the gathered information to compile comprehensive profiles of each individual [3]. As a result, cookies were suddenly regarded as a security loophole abused by marketing companies and other entities to infringe on the privacy of Internet users. Cookies were found to be a threat, along with other malware like viruses and worms, to the safety of private information.

In subsequent years, steps have been taken to try to alleviate this problem. Many browsers currently have options to accept, reject, and clear all cookies. Malware management software has been developed for consumer use, providing options to view cookies and their content, and often the ability to change the contents of cookies. Where technological solutions fail, social self-regulation has been adopted by some users who make a point to clear cookies regularly, change passwords often, and maintain an awareness of the consequences of cookie intrusion.

Despite the increase in potential solutions, cookies remain a challenging problem. Often the simplest resolution to proper cookie management is to do nothing, in which case users are at a high risk of potential cookie abuse. The other extreme is for users to manually examine each individual cookie deposited on their computer in order to decide whether it should be removed or retained. This ensures full control over the amount of information being disclosed at any time by retaining only abuse-free cookies and removing abusive cookies. However, this not a feasible expectation in everyday life due to the sheer volume of cookies encountered while browsing the web.

This dual nature of cookies – the ability for them to be either "good" or "bad" – turns them into a bit of a conundrum. It is precisely this duality which has impeded the development of suitable solutions to properly manage cookies.

In this paper, we will illustrate that this duality is the source of many frustrations involved in developing proper cookie management tools and solutions. We begin by defining our motivations for this project, then by surveying the background literature related to cookies.

This will be followed by a description of the focus group sessions which were conducted to aid us in our investigation and a discussion of the results of the focus groups. Finally, we will attempt to identify the necessary considerations required in developing adequate management methods for cookies, and conclude by defining the measures that should be taken to alleviate the privacy concerns associated with cookies.

2. Motivations

Motivation for this project began with a realisation that many friends and colleagues pay little attention to cookie management and the privacy issues surrounding them. Further exploration into the subject led us to four key reasons why this area warrants further investigation.

- As previously mentioned, the dual nature of cookies makes them unlike any other malware on the Internet today. This difference may affect how cookies should be managed.

- Since their inception in 1995, cookies have been a large factor in issues relating to personal privacy. A useable and effective solution to this problem is still lacking and users remain largely unaware of the severity of the cookie problem.

- The popularity of cookies is increasing. Since 1995, cookie use has only increased, especially for uses of e-commerce and marketing. Many standards have been produced to serve as guidelines as to how cookies should be used (eg. RFCs 2109, 2965, 2964), but these guidelines are not always followed. This is especially true of marketing agencies who find the personal information they gather highly lucrative [20].

- Surprisingly little research has been conducted to examine if users are comfortable with the current system of cookie management. Previous research has demonstrated that users value their personal privacy [1], however little has been done to investigate if people are comfortable with current software managers, and what features might be more important to them then others. Similarly, the privacy settings provided by current web browsers have been evaluated from an accessibility standpoint [17], yet little work has been done to explore the actual use of such privacy functions.

These motivations have lead us to believe that more can be done to fully understand the cookie problem and indicated that a different approach is required in order to gain an understanding of current users' perceptions of cookies.

3. Background

Cookies were first introduced by Netscape Navigator as an extension to the Hyper-text Transfer Protocol (HTTP) [19]. Netscape introduced two new headers, "Set-Cookie" and "Cookie", to add state information to the otherwise stateless HTTP. Every HTTP transaction is treated as an independent entity regardless of the initiator and receiver. One user might initiate several requests to the server but by using HTTP alone, the server has no means to determine if these were multiple requests initiated by the same user. With HTTP being stateless, user specific services over the Web are virtually impossible due to the stateless nature of the Web. A popular example demonstrates that cookies can be used to implement "shopping carts" in online-shopping websites. E-commerce applications which utilise shopping carts use cookies to keep track of customers' activities throughout a session, recording and differentiating the items which have been added to each user's cart.

Cookies have several advantages that make them the one of the best techniques to be used in E-commerce applications. Cookies enable the easy development of stateful web applications, can provide more web interactions with users, and can help provide personalized web applications that suits the needs of the users [15][25].

In order to implement state information into a web application without using cookies, several techniques have been developed:

- IP addresses can be used as session identifications. However, they do not provide enough information to identify users. The reason is simply that one computer can be used by several users; one user can use multiple computers to access the same service; and for users of dial up or wireless connections, they are assigned a different IP address for every new connection sessions [15].

- State information can be embedded into URLs. This can be used to keep track of state information within one session. The problem with this approach is that it is not stable. For instance, in a shopping cart scenario, every new page they visited would be assigned a unique URL. This would result in an increased load on the web caches since several copies of the same content will be saved due to the different URLs that it receives from servers. To use URLs to save information across multiple sessions, the server has to assign a unique URL to every user, and every time the user needs to access their service, they either have to type in that URL or login to verify their identity [25].

- Hidden fields in HTML or dynamic HTML can be used to assign values that are unique to each user. This technique might be helpful to keep track of state information for one session; however, it cannot provide information across multiple sessions.

These techniques are problematic and require more effort on the part of the web developer than cookies to implement state information. Nevertheless cookies also suffer from several disadvantages. These disadvantages result from the fact that cookies are implemented in a user-specific manner. They are used to keep track of important state information to identify a particular user. This reinforces the fact that cookies are not inherently good nor bad. The privacy concerns surrounding cookies stem from the ways in which cookies are used, the poor built-in cookie management techniques in current browsers implementations, and a lack of user control over the data that are being passed to web servers in cookie interactions.

3.1 Cookie Uses

When cookies were initially introduced in 1995, the goal was to add state and maintain user information across sessions. However, cookies can also provide the ability to track users' activities on the Web by recording the users' click movements in the same site or across multiple websites by following hyperlinks from one website to another. The ability to track user behaviour is the focus of privacy concerns. Cookies can store information that users may not expect to be stored and this includes the pages visited in one site, the time of each visit, the login information, the shopping cart details, and, the personal preference attributes. While users give this sort of information both implicitly and explicitly, many users are not aware of the use of cookies to store this information.

Browsing the Web without cookies is generally anonymous but when using cookies, browsing requests become "pseudonymous" [25][9]. These requests can become even identifiable if user profile information from user profiles is added to them. This presents a real threat to users' privacy as this information can be used to infer full user profiles that do not only contain identification information but also browsing activities and online purchases.

As Cranor [9] discusses, consumers on the Web often reveal their personal information to websites in order to benefit from services which the websites offer. This is because they will not usually have access to these services otherwise. Some of these services include purchasing products online, typically with better prices than normal retail price, and getting access to services and information that require registration to obtain (e.g. name, email, address). However, users in many situations would prefer to stay anonymous. Cranor [8] lists several products that help users stay anonymous while browsing the Web. Products like Anonymizer, Crowds, Onion Routing, LPWA, and P3P-enabled browsers all help users to stay anonymous while browsing the Web [8]. Nevertheless, these products need to be adopted by most users and vendors to make sure that web users have more control over their personal information that are being passed through cookies or any other state management techniques.

Accessing services through the Web is still a valuable resource for online customers. But these customers have to be assured that their privacy is being protected while they are browsing or shopping online. Anonymity tools do provide privacy to users while browsing, but they do not protect the data that is passed through web transactions. The data that users provide while browsing or that is transmitted through cookies, needs to be protected through legal regulations and self-regulation practices. Cranor and Reagle [8] described the Platform for Privacy Preference Project (P3P) which helps provide resources for service providers on the Web to explain their privacy and information policies, and for users to identify their privacy needs and preferences. Other third party agencies such as TRUSTe and the Better Business Bureau's BBBOnline Seal programme provides businesses with a digital or a visual "assurance seal" or a "trustmark" which assures web customers that these businesses' practices are compliant with what is written in their privacy policies, and that the users' personal information will not be used in a manner that is not described in the privacy policies [8][9].

3.2 Browser cookie management techniques

Privacy concerns that are associated with cookies are partially due to the fact that users have no control over the data that is being transmitted through cookies and the lack of feedback provided by the browsers built-in cookie management techniques.

Previous research has shown that the built-in cookie management tools in web browsers can not be used alone to manage cookies. Most current browsers are configured to accept cookies by default. Some browsers are too conservative in their default privacy policies; they block all cookies by default and the results are often inconvenient to most of their users. Other browsers, such as IE6 and Mozilla, manage cookies based on the P3P policy of the websites visited [25]. P3P-based cookie management tools are often very complex and for ordinary web users this can be very cumbersome especially since it does not provide direct feedback on their bases for blocking cookies as shown in [25].

Still, most built-in techniques provide limited options for managing cookies and users have to depend on filtering rules and trust that these rules will function properly and filter cookies according to the users' preferences.

3.3 Cookie management software

As discussed in the previous section, cookie management tools that are built into browsers do not provide the user with enough control over the cookies that are passed or accepted during web browsing. To solve this problem, several cookie management tools have been developed and made available for web users either as a freeware or commercial software.

Cookie management software is largely divided into two categories: (1) software that primarily manage cookies (standalone cookie management software), and (2) software that manages malware in general including cookies (bundled malware management software).

The more widely available cookie management software is the standalone managers. These tools are widely available for download as freeware and as commercial software. Cookie management products can be categorized depending on the set of functions they offer the users. Some of these products provide information in general about the cookies being passed but do not provide the user with any tools to control cookies. Other products (such as the ones used in our focus group) prompt the user to decide what to do with cookie either per cookie or per site (by providing the "Accept All" or "Reject All" options for the same website).

In bundled software, cookie management is normally sidelined. This category of tools normally comes with a default global policy that overrides the browser's policy. These cookie management tools can be configured by the users to control cookies by applying several filtering rules. However, these tools do not provide the user with enough information about the types or content of cookies that are passed. It can globally control cookies by allowing or rejecting first or third party cookies depending on the default policy or the users' preferences.

Some types of commercial cookie management software have the added functionality of detecting third party cookies. These cookies are either blocked entirely or users are prompted for a decision. Most cookie managers maintain a list of cookies accepted per session or in general. Many of these managers allow the user to create filtering rules by creating "Allowed" and "Rejected" lists where the user can list the websites that they want to accept cookies from unconditionally or reject all cookies coming from websites in the reject list.

Few cookie managers give the user more information about the actual content or category of the cookie. From the users' point of view, cookie contents (i.e. the contents of the value field) mean nothing without their context of use by the company who disperse these cookies. For example, a cookie from http://www.google.com contains the following string, "PREF ID=92d9daea01a52931:LD=en:TM=1105900308:LM=1105900308:S=ljrllmeUaddbKePG... ." It can be inferred that "ID" indicates an identification number, "TM," probably a time reference, and "LD" probably refers to the language preference. For the normal web user, there is no way to confirm this information without contacting Google. Furthermore, it is not obvious what the "S" and "LM" fields contain. This information is from a single cookie and to individually dissect each and every incoming cookie is not feasible. Most cookie managers simply display to the user this information without extra explanation, leaving the onus of analysis entirely on the user.

Cookie Crusher was used in our study because it provides extra information about the type of cookies encountered and whether these cookies are secure or not. However, the software does nto provide any feedback as to its basis of consideration or categorisation. The Cookie Crusher categorises the cookies into four categories, namely, "Site-Tracking," "Advertising," "E-commerce," and "Unknown." Though these classifications provide some insight as to the nature of each cookie, there is no indication by the programme as to how to how each cookie is classified. This lack of feedback means that users have to trust the software, even though they have no information about the exact content of each cookie, when deciding whether to remove or keep a cookie.

4. Methodology

During the Spring of 2005 we developed a series of focus group sessions to provide us with insights to users' knowledge about cookies. We were interested in several issues:

- We wanted to know how much people were aware of cookies, how they are used, and what kind of information they provide to companies.
- We wanted to examine current strategies for cookie management.
- We wanted to tease out ideas about potential cookie management tools which will help people become more aware and in control when managing cookies.

To achieve our goals, we conducted a series of focus group studies to examine each of these issues.

4.1. Focus Group Study

We recruited 16 participants to take part in a focus group. Participants were classified according to age (below 22 year and above 28 years of age) as well as technical familiarity. We ended up with the following four groups: technical/younger, non-technical/younger, technical/older, and non-technical/older. The reason behind this was to get a better distribution of knowledge, and to ensure that each group shared similar experiences in terms of technical know-how. The focus group sessions lasted two hours each with a short break midway during the session, and took place in the Usability Lab in the Computer Science Building.

We chose to conduct a focus group instead of a survey primarily because of the limitations inherent in survey data. According to Singleton and Harper [21], surveys are difficult for obtaining information about personal privacy because they tend to group concepts such as identity fraud, spam, and other security threats together under the general umbrella notion of "privacy." This leads to misleading results based on assumptions and personal experiences on the part of respondents. Additionally, because surveys do not inform or educate respondents about the variety of issues involved, they are unlikely to elicit the same responses as in real-world situations. Moreover, survey questions may be manipulative in order to draw out certain types of responses from respondents, based on the interests of the surveyor. Based on these suggestions, we used a focus group to enable a more intimate discussion about topics and to educe a true reflection of current attitudes and responses towards cookies.

During the focus group sessions, participants were asked to discuss perceptions of cookies, current ways in which cookies are being managed, as well as various other issues related to cookies. They were also asked to test out two commercial standalone cookie management applications, namely the Cookie Pal (fig. 1) [7] and the Cookie Crusher (fig. 2) [6], while browsing the Web. The two applications were chosen because of the different ways in which they approached cookie management. Both used pop-ups as the notification tool for incoming cookie alerts. However while Cookie Pal was highly simplified, providing one pop-up per cookie, the Cookie Crusher had many more functions and options in each pop-up and collated cookies into multiple cookies per pop-up for each website. The Cookie Crusher also provided cookie classifications (advertising, site tracking, e-commerce and unknown), which were derived from the contents of the cookies, and presented this classification together with the contents of each cookie. The user then had to click on an option to accept or reject the cookie to varying degrees.

After testing out the two applications, participants were engaged in more discussions about their experiences with the software, and were asked to envision what an ideal cookie manager might look and perform like.

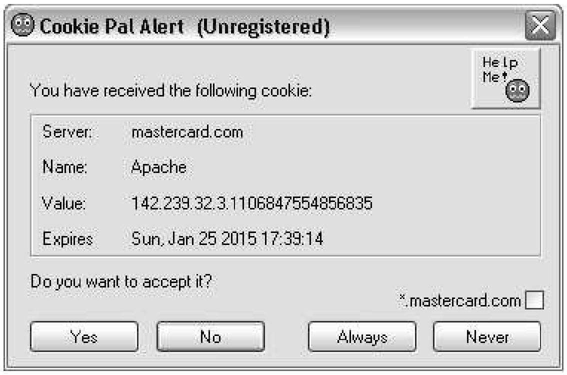

You have received the following cookie:

Server: mastercard.com
Name: Apache
Value: 142.239.32.3.1106847554856835
Expires Sun, Jan 25 2015 17:39:14

Do you want to accept it?

*.mastercard.com □

Figure 1: Cookie Pal pop-up window

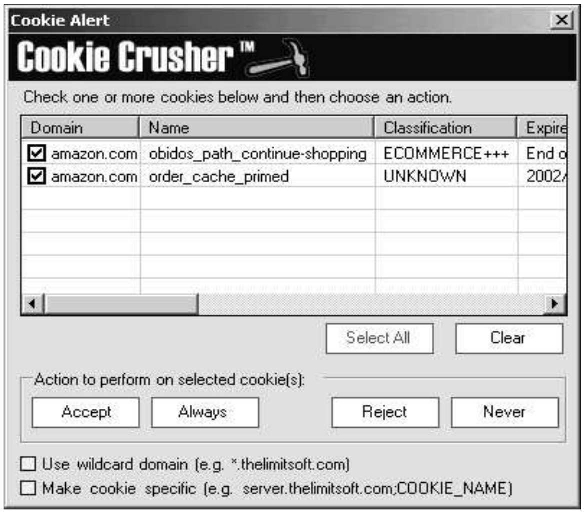

Check one or more cookies below and then choose an action.

Domain	Name	Classification	Expire
☑ amazon.com	obidos_path_continue-shopping	ECOMMERCE+++	End o
☑ amazon.com	order_cache_primed	UNKNOWN	2002/

Action to perform on selected cookie(s):

□ Use wildcard domain (e.g. *.thelimitsoft.com)
□ Make cookie specific (e.g. server.thelimitsoft.com;COOKIE_NAME)

Figure 2: Cookie Crusher pop-up window

5. Results

Overall, the participants were more confident working with the Cookie Crusher because it helped limit the number of pop-ups received per website. Though the focus group discussions were designed to educate and raise awareness as well as elicit responses from participants, participants with a non-technical background found it hard to grasp the idea of what cookies did exactly. Though many found the Cookie Crusher more helpful because of the classification fields, they were still unsure what an advertising or an e-commerce cookie might entail. Many indicated confusion as to how the software was able to classify the cookies in the first place, and found that there was no way to make informed decisions as to the nature of each cookie.

After trying out the software, many people also became concerned with the amount of effort required to manage cookies. They found the pop-ups distracting and commented that they interrupted the flow of their web-browsing. Technical users indicated that they felt the Cookie Pal would be more useful for non-technical users because of its simplicity. However, the non-technical users indicated that they felt more comfortable using the Cookie Crusher instead. They expressed that the Cookie Pal, though simple, seemed to require more technical knowledge in order to make judgement calls to accept or reject each cookie. This was because no explanation as to the nature of the cookies was provided, unlike the Cookie Crusher, which at least gave each cookie a classification.

Another interesting point was the decision by participants to accept the first few cookies from every website based on the perception that rejecting these initial cookies will prevent them from visiting the website. Based on experience, this is true in most cases but not in general. In our study, we noticed that many websites send third-party cookies as part of the few initial cookies. Unknowingly, users under the misconception stated above, accept these cookies without carefully examining the cookies' domain field, which will confirm that these cookies are indeed third-party.

Many non-technical users were also apprehensive about the amount of space cookies might take up on their computers. Moreover, users indicated a preference for bundled security software instead of standalone software (like the Cookie Pal and Crusher) because of the amount of overhead a standalone application might require, both in terms of space and effort.

The older groups of participants were more interested in the privacy issue than younger participants who were more concerned with the amount of added effort a cookie manager might require.

6. Implications and Goals

Due to the dual nature of cookies, their usage is a double-edged sword. Accepting the wrong cookie may elicit malicious attacks on a user's privacy, while deleting non-abusive cookies may lead to decreased functionality and service from a particular website. Current cookie managers (standalone programmes, bundled with browsers, malware handling agents) generally prompt users, whether in real-time or through back-end handling software, to decide how to deal with incoming cookies. The onus is on the user to identify and decide on the appropriate action to take. This may occur with help from the programme, such as the cookie classifications provided by the Cookie Crusher, or with no help at all except for the contents of the cookie itself, which itself may mean absolutely nothing without the context of its use. Even with the help of the classifications, cookies are difficult to distinguish, as shown by the focus group results. Beyond the superficial level of the nature of the cookie contents, arises the more intricate privacy issue of how the information stored in the cookies will be used by its collectors. As a consumer, unless the entity who owns the cookie has stated strict privacy guidelines on the website itself, it is almost impossible to ascertain the actual use of the cookie information. Consequently, subjective judgement calls are required to enable good cookie management.

Two major goals can be defined to help manage cookies. Firstly, the user must be aware of the cookie interactions taking place at any time. They must understand what information is being passed to the web servers via the cookies, as well as what the server owners can do with the information provided. Secondly, the user must be able to control the types of information stored and passed in the cookies, as well as have the appropriate tools to do so.

7. Reflections

During the course of this project, several concerns repeatedly presented themselves illustrating the reasons why the above goals to cookie management are difficult to achieve. From these concerns, we have distilled two major considerations, namely technical considerations and social considerations, and their resulting consequences arising from problems with the current system. These considerations should serve as a basis for judging the feasibility of any implementation of a solution.

7.1. Technical Considerations

General consensus from the focus group study indicates that cookie management software needs to be "smarter." When asked what being "smarter" might entail, participants provided several technical considerations which we have distilled into the following points of interest. In addition, we also inferred several social issues which might arise as a consequence.

- A possible solution to the cookie problem is to discard the use of cookies in favour of some other technology with better privacy considerations. However, to do so would require fundamental changes to internet standards. The use of cookies is so widespread that a major overhaul such as this would be practically impossible, which requires us to find ways to work with the existing technology.

- A major feature of standalone and browser-based cookie managers is a series of alerts to indicate the presence of an incoming cookie while a user is browsing the Web. Although this feature would provide users with a heightened awareness of the cookie interactions taking place, users generally find this form of immediate feedback frustrating and overly intrusive. However, with immediate feedback comes excellent control over the information being sent out as the only way to detect cookie passing is when a server sends a cookie to be stored on users' computers. No detection mechanisms have been set up to indicate when web servers request cookies and the only control users have over the interactions is when the cookie is first set. In the case of no immediate feedback, this becomes a problem because cookies which have already been set will be passed automatically to web servers by browsers, thus leaving users no control over their private information.

- Results from the focus group study indicate that participants appreciate the Cookie Pal for its intuitive and easy-to-use interface, but appreciated the Cookie Crusher for its advanced functions and customisability. In this sense, there exists a trade-off between ease of use and functionality because the more functionality a programme offers, and the more complex its interface becomes. Our results indicated that only the older, non-technical group appreciated the simpler interface of the Cookie Pal, but since our sample size was small, this result may not be indicative of any overreaching trend.

- In addition to ease and functionality, participants also expressed a need for programmes which are convenient to use. Participants expressed an interest in "smarter" programmes which could decipher the contents of cookies to indicate what sort of information was being sent out, as well as what this information was going to be used for. They also expressed a desire for recommendations by the programme as to what course of action to take, as well as the ability of the programme to take this course of action on their behalf, so as not to distract them from web browsing.

- Two categories of cookie management software exist on the market currently, standalone cookie specific software, and general security bundled software (these include the cookie managers built into browsers). Cookie specific software presents an extra layer of overhead which users may find excessive. They can be very intrusive by providing immediate feedback or alternatively non-intrusive with the added requirement for users to sort through the lists of cookies after each browsing session. Bundled software has the advantage of protection from other forms of malware, but has a habit of sidelining the cookie problem by providing only minor support for cookie management. Bundled software often has "set it and forget it" options which provides no awareness of cookie interactions taking place at all.

Design implications which follow from the above issues continually raise the question of awareness and control, which more often than not are contradictory in relation. Often, there is no optimal solution, and compromises have to be made in choosing awareness over control or vice-versa. The result of this is a social reaction which may prompt users to do nothing instead, which in turn leaves users vulnerable to privacy invasions.

7.2. Social Considerations

In this section, we consider the social issues which arise, as well as the technical consequences which follow:

- One of the most important results from the focus group turned out to be the lack of true understanding of the advantages and disadvantages of using cookies. The younger groups tended to have a limited awareness of the functions of cookies, regardless of their technical experience. The older, technical group had the best understanding of cookies, yet hardly any of them used any sort of cookie manager to counteract the problem. The older, non-technical group had a vague understanding that cookies could be used for malicious purposes, but neither understood how this could occur, nor how to deal with the problem. Another interesting result was the shock that many participants received when using the Cookie Pal and Crusher for the first time. They indicated they were unaware that so many cookies were being used on the Web. This trend indicates a striking need for users to be educated regarding the issues surrounding the use of cookies. In particular, illustrating that cookies are not always security threats like viruses and worms, but are useful in the way that they provide state information to websites. Additionally, there is a need for users to be made aware of the types of information that can be passed through these cookies. Many participants had misconceptions that cookies could be used to transmit viruses, which is false.

- Present day society functions on a basis of trust. To not trust the system in place means that we have to live in constant fear of violation, a state of mind which some participants indicated as futile. Much sentiment shared by the younger groups indicated the

> acceptance that cookies could potentially pose a threat to personal privacy, and yet still choose to do nothing about it, as a sort of resignation to the state of affairs which we live in, a sentiment also observed by Ackerman et. al. [1]. This sign of not caring may be attributed to the lack of solutions or awareness, or may merely be a result of pedantic aloofness.

Managing cookies has become difficult because of the social points of interest indicated above. Consequences of these social considerations result in technology becoming obsolete in the face of such lack of awareness. If users do not know, or do not want to know, about cookies and how to deal with them, then the tools which allow for cookie management become irrelevant if people do not use them.

An example of this is the emergence of the Platform for Privacy Preferences, or P3P project, which allows websites to specify their privacy policies and users to access these policies to determine acceptance. In particular, P3P compact policies, which are indicated in HTTP headers, contain policy information related to cookies, our area of interest. Using P3P, users can easily indicate their privacy preferences, and leave it to user agents to compare these preferences against the privacy policies of visited websites, thus providing a better method to decide the trustworthiness of a website. Though a good answer to the cookie problem, P3P is not used by enough people to be a feasible short term solution. (Granted, P3P has only been around since 1998) According to a survey done by SecuritySpace.com, a web security portal, of 599,019 websites surveyed, only a total of 12,479, or about 0.0208% of the websites used compact privacy policy statements to indicate their privacy policies [12]. This is in contrast to another survey done by the same company, which indicated a low estimate of 20.6% of websites making use of cookies [11]. As a result, although P3P is a good solution to the cookie problem, people's lack of awareness makes it ineffective as a current solution [24].

8. Conclusion

In this paper, we have shown that there are both social and technical considerations required to meet the goals of awareness and control, and we have shown that these technical and social issues do not stand alone, but in fact have in themselves, social and technical ramifications as discussed above.

However, even after all these considerations are met, there exists one point which we have yet to identify – it still remains unknown to what extent companies are using this private information to infringe on consumers' privacy. Guidelines serve only to demonstrate what should be done. Whether these guidelines are actually met are a different matter altogether.

The solution to this particular aspect of the problem is not within the scope of this project, however, an overall increase of awareness on the part of the user, may alleviate the effects of dishonest cookie dispensing websites.

Raising awareness is the key to suppressing the effects of malicious cookies in general. Our study has shown that most people are ill-informed as to the issues surrounding cookies and their management. Because of the dual nature of cookies, they are fundamentally different from other malware such as viruses, which can be protected from by leaving a protection software running in the background, out of sight and out of mind. Cookies on the other hand, have to be subjectively managed, whether actively by the user, or by some artificial intelligence running in the background. Artificial intelligences such as the ones required are impossible under current circumstances because the content of cookies cannot be inferred without the context of the information's use. Thus the only option left is for the user to actively take part in their privacy protection process. This requires additional awareness by the user, which will in turn induce more control to the user when it comes to managing personal privacy through cookie management

This additional awareness can be achieved by conducting focus group sessions such as the ones which we have carried out for this project. Our focus groups served not only for us to gather information from our participants, but also to educate and increase awareness on their part. Even though a definite solution was not found, we believe that our participants left with the additional awareness to continue learning about the various ways in which cookies may be used or abused.

9. Acknowledgements

We would like to thank the members of the Edgelab for their invaluable advice and assistance throughout the course of this project. We would like to especially thank the Computing Research Association's Committee on the Status of Women in Computing Research (CRAW), the Coalition to Diversifying Computing (CDC) and the Collaborative Research Experience for Undergraduates (CREU) programme for funding our research.

10. References

[1] Ackerman, Mark S., Cranor, Lorrie F., Reagle, Joseph, "Privacy in E-commerce: Examining user Scenerios and Privacy Preferences," *E-Commerce, Proceedings of the 1st ACM conference on Electronic Commerce*, Denver, CO, 1999, pp.1-8.

[2] Ackerman. Mark S., Cranor, Lorrie F., “Privacy Critics: UI Components to Safeguard Users’ Privacy,” *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI’99)*, short papers, 1999, pp. 258-259.

[3] Chapman, Scott, Dhillon, Gurpreet, “Privacy and the Internet: the case of DoubleClick, Inc.” http://www.unlv.edu/faculty/dhillon/Teaching/Teachingfall02/MBA730-fall02/Cases/Doubleclick.pdf, Last Accessed 20th Jan, 2005.

[4] Chung, Winnie, Paynter, John, “Privacy Issues on the Internet,” *Proceedings of the 35th Hawaii International Conference on System Sciences*, 2002.

[5] Cookie Central 1996-2004. http://www.cookiecentral.com. Last Accessed 8th Aug, 2004.

[6] Cookie Crusher , The Limit Software 1997-2005, http://www.thelimitsoft.com/cookie, Last Accessed 20th Oct, 2004.

[7] Cookie Pal, Kookaburra Software 1996-2005, http://www.kburra.com/cpal.html, Last Accessed 20th Oct, 2004.

[8] Cranor, Lorrie F., “Internet Privacy: Introduction”, *Communication of the ACM, Volume 42, Number 2,* 1999, pp. 28-31.

[9] Cranor, Lorrie F., “Internet Privacy: A Public Concern”. *Networker 2,3*, June, July1998, pp. 13-18.

[10] Cranor, Lorrie F., Arjula, Manjula, Guduru, Praveen, “Use of a P3P User Agent by Early Adopters,” *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, Washington, DC, 2002, pp. 1-10.

[11] E-Soft Inc. Security Space. “Internet Cookie Report,” http://www.securityspace.com/s_survey/data/man.200504/cookieReport.html, 1st May, 2005.

[12] E-Soft Inc. Security Space. “P3P Compact Privacy Policy Report,” http://www.securityspace.com/s_survey/data/man.200504/p3p.html, 1st May, 2005.

[13] Kristol , David, Montulli, Lou, “RFC 2965: HTTP State Management Mechanism,” http://www.faqs.org/rfcs/rfc2965.html, Oct 2000.

[14] Kristol, David, Montulli, Lou, “RFC 2109: HTTP State Management Mechanism,” Superseded by RFC 2965, http://www.rfc-archive.org/getrfc.php?rfc=2109, Feb 1997.

[15] Kristol, David, “HTTP Cookies: Standards, Privacy and Politics,” *ACM Transactions on Internet Technology*, Vol. 1, No. 2, Nov 2001, pp. 151-198.

[16] Lin, Daniel, Loui, Michael C., “Taking the Byte out of Cookies: Privacy, Consent, and the Web,” *Computers and Society,* Policy ’98, Privacy Issues, June 1998, pp. 39-51.

[17] Millet, Lynette I., Friedman, Batya, Felten, Edward, “Cookies and Web Browser Design: Toward Realising Informed Consent Online,” *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI’01)*, 2001, pp. 46-52.

[18] Moore, Keith, Freed, Ned, “RFC 2964: Use of Http State Management,” http://www.rfc-archive.org/getrfc.php?rfc=2964, Oct 2000.

[19] Netscape Communications Corporation. “Persistent Client State: HTTP Cookies.” http://wp.netscape.com/newsref/std/cookie_spec.html

[20] Perkins, Simon, “Internet Cookies: Security Implications,” Student Paper, http://www.cs.uct.ac.za/courses/CS400W/NIS04/resources.html, 2000

[21] Singleton, Solveig M., Harper, James, “With A Grain of Salt: What Consumer Privacy Surveys Don't Tell Us,” *Competitive Enterprise Institute, Issue Analysis,* http://www.cei.org/PDFs/with_a_grain_of_salt.pdf, June 2001.

[22] The Platform for Privacy Preferences (P3P) Project, http://www.w3.org/P3P/, Last Accessed 5th May, 2005.

[23] The Platform for Privacy Preferences 1.1 (P3P1.1) Specification, W3C Working Draft 4th Jan 2005, http://www.w3.org/TR/2005/WD-P3P11-20050104/Overview.html, Last Accessed 28th April, 2005.

[24] Thornberry, Suzanne, “P3P: Big Backer, Slow Pickup,” TechRepublic, http://insight.zdnet.co.uk/hardware/servers/0,39020445,2111631,00.htm, 11th June, 2002.

[25] Yee, Ka-Ping, "A survey of Cookie Management Functionality and Usability in Web browsers," http://zesty.ca/2002/priv/cookie-survey.pdf, 2002.