Conducting an Ethical Study of Web Traffic

John F. Duncan
johfdunc@Indiana.edu

L. Jean Camp
ljcamp@Indiana.edu

Abstract

We conducted a study of student web browsing habits at Indiana University's Bloomington campus, in which we examined the web page requests of over 1,000 students during a period of two months. In this paper, we discuss the details of the study development and implementation from the point of view of ethical design. Concerns with stakeholder privacy, the quality of study data collection, human subjects research protocols, and unexpected data anomalies are presented in order to illustrate the many difficulties and ethical pitfalls confronting network researchers even at this small scale. Success and failures to meet the principles of ethical design are highlighted. A secondary contribution is the evolution of the instruments that were developed through the human subjects process. Finally, we discuss the impact of the Menlo Report (DHS-2011-0074) and similar documents on the future directions of network and security research.

1 – Introduction[*]

We sought to evaluate the inadvertent wisdom of crowds in avoiding dangerous web locales. If users are informed that a particular site is seldom visited, this might alter their behavior and trust decisions. The associated research prototype (Tsow et al, 2009), which motivated the experiment, used sub-communities for reasons of privacy and resilience to Sybil attacks. We previously illustrated the efficacy of information in informing individual trust decisions (Camp & Genkina, 2006). If users exhibit similarities in browsing habits (e.g. visit the same sites), they can better collaborate to combat online threats such as phishing.

In order to establish and test metrics for user browsing homophily and session behavior, we conducted a study of student web browsing habits in an academic dormitory at Indiana University's Bloomington campus. Under the aegis of this study, we examined the web page requests of over 1,000 students during a period of two months. The process of conducting such a study ethically has become relevant for the CSET community. Thus, in this paper we discuss the details of the study development and implementation, **not the results of the original study**. For details on results, readers are directed to papers on the properties of browsing sessions (Meiss et al, 2009), the value of information disclosure in web resource evaluation and discovery (Dong & Camp, 2010) and the development of a system informed by this research (Tsow et al, 2009).

Working in tandem with Indiana University's Human Subjects Committee (HSC) and Institutional Review Board (IRB), we designed our study in such a way as to minimize student impact. Our secondary goal was providing education by informing students as to the possible ways in which they might make their web browsing information difficult to detect by malicious snooping agents. To do this, we provided the students with information regarding the ways in which they might opt out of the study (while also increasing their security) through the use of a VPN. While we had desired to also provide students with information about systems for secure anonymous browsing, such as Tor (Dingledine, Matthewson, & Syverson, 2004), we did not receive approval to distribute this information.

To implement this study, we were first required to complete an exam on how to conduct human subjects research (Indiana University, 2011). This was geared primarily towards researchers who would be conducting face-to-face interviews or handing out physical questionnaires. Details on the questions asked by Human Subjects, and our responses to them, can be found in Section 7. Section 8 includes discussion of the progression of the flyer from a more

[*] We would like to acknowledge the individuals who made this study a possibility, including Mary Beth Schmucker, Robert Weith, Residential Programs and Services (RPS), University Information Technology Services (UITS), the Human Subjects Committee (HSC), the Institutional Review Board (IRB), the staff in Residential Services, and IU research librarians. Special acknowledgements to Mark Meiss, without whom the technical implementation of this work would not have been possible. This material is based in part upon work supported by the National Science Foundation under Grant Numbers CNS 0943382 and IIS 1036918. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

technical document to an arguably more informative one. We enumerate failures, successes, and one ethically-challenging surprise. In Section 9 we present clear steps as recommendations to researchers intending to follow on the path of ethical research in the current environment. In Section 10, we discuss the impact of the Menlo Report (DHS-2011-0074) and similar documents on our research and that of others studying networks and security.

2 - Technical Design and Anonymization

The dormitory study was implemented by installing a dedicated FreeBSD server located in the central routing facility of the Bloomington campus of Indiana University. This system had a 1 Gbps Ethernet port that received a mirror of all outbound network traffic from one of the undergraduate dormitories. This dormitory is split roughly evenly between men and women.

To obtain information on individual HTTP requests passing over this interface, we used a Berkeley Packet Filter to capture only packets destined for TCP port 80 (the normal port for web browsing). These packets were mirrored to this traffic monitor, which examined the requested URL to determine if it was a web page (.htm, .html, .asp, .php, etc.) or other content (images, video, etc.). This methodology was also employed in (Gribble, 1997). We used a regular expression search against the payload of the packet to determine whether it contains an HTTP GET request. We then analyzed the packet to determine the identity of the virtual host contacted, the path requested, the referring URL, and the advertised identity of the user agent. Raw data files that contained a timestamp, the virtual host, the path requested, the referring URL, and a flag indicating whether the user agent matches a mainstream browser (Internet Explorer, Mozilla/Firefox, Safari, or Opera). The aggregate traffic of the dormitory was sufficiently low so that the sniffing system was able to maintain a full rate of collection without dropping packets. No content was fetched by the packet monitor during the study. No attempts to observe encrypted traffic were made (nor did we capture requests to TCP port 443), although classic man-in-the-middle attacks would have made such observation possible. The traffic sniffing was stateless, so our system did not behave differently for keep-alive and pipelining. Individual requests in either case would be sent in separate TCP packets by popular clients.

The click data was collected over a period of about two months, from March 5, 2008 through May 3, 2008. This period included a week-long vacation during which no students were present in the building. During the full data collection period, we logged nearly 408 million HTTP requests from a total of 1,083 unique MAC addresses. We filtered out a small subset of users with negligible activity; their traffic consisted largely of automated Windows Update requests and did not provide meaningful data about user activity. We also found that some web clients issue duplicate HTTP requests (same referring URL and same target URL) in nearly simultaneous bursts. These bursts occur independently of the type of URL being requested and are less than a single second wide. We conjecture that they may involve checking for updated content, but we are unable to confirm this without access to the original HTTP headers. Because this behavior is so rapid that it cannot reflect deliberate activity of individual users, we also removed the duplicate requests from the data set. While there no doubt remained some automated requests in the data set, we removed as much noise as was feasible. The resulting data set is the basis for all of the discussion in this paper. It contains a total of 29.8 million page requests from 967 unique users. These requests were directed toward a total of over 630,000 distinct Web servers, and nearly 110,000 distinct servers provided referrers for those requests.

3– Related Work and Design Motivations

In the first instantiations of the experiment we sought to implement an ethical and privacy-preserving experiment based in no small part on our self-definition as privacy researchers. The goal was not to implement a publishable experimental design; however, the expanded interest in ethical design indicates that this discussion may now be timely.

Based on the literature, we considered for the purposes of this experiment that ethical systems are ones that respect the rights of all stakeholders in the use of that system (Friedman, 1996; Friedman & Nissenbaum, 1996; Nissenbaum, 1998a, 1998b). The increasing concern with ethical design in the engineering of physical systems suggests that ethical design for software systems and research protocols is also increasingly important (Cummings, 2006). This is especially notable in network research, where as the Director for CAIDA (Cooperative Association for Internet Data Analysis) noted, “No one is investing in technology to learn about networks while minimizing the amount of privacy compromised in the process” (Claffy, 2008). Since that 2008 paper, and thus after our experiment, considerable progress has been made in the development of ethical experiments. In part of this realization, the networking community has been moving to meet the challengg; graduating from panels (Soghoian, 2011), to directed meetings, (Hall, 2011), dedicated workshops (WECSR), and currently the Menlo Report (DHS-2011-0074). There are also specialized discussions in research communities, e.g.

the ethics of research on users of anonymity systems (Soghoian et. al, 2011).

Students, as a population being considered as subjects for research, are stakeholders. Our process was derived from Value-sensitive design. VSD informed both the framing of the study and definitions of data. Risks, in this study, took the form of incidents or data capture policies that would threaten subject privacy. Identifying such risks and eliminating them when possible, while preserving the functionality of the data, was a major component of the study. Where we were unable to uphold the standards of ethical design completely, we discuss what occurred and identify points in which further work is needed. By beginning a dialogue on ethical study design in network research, we hope to strengthen privacy and security for the studied populations, while preserving the ability for researchers to capture the data needed for their work. This potential tension between the needs of stakeholders, researchers, and research review boards is identified by Mosher as a "justice-as-fairness" concern (Mosher, 1988). While he identifies this tension as most visible in the area of controversial human-subjects research on subjects such as sexuality, the controversial nature of traffic studies on real-world networks, involving actual usage data, renders his arguments equally applicable.

During our data collection, requests for URLs were trimmed by discarding both the query string and fragment identifier if present, to prevent the inclusion of any session-based information (such as cgi variables) that would identify the user. Cockburn et al. performed similar trimming, although theirs was primarily designed to correctly register multiple visits to sites that employ search parameters (Cockburn, McKenzie, 2001). When these trimmed URLs were added to the database of browser requests, we attached an identifier constructed by taking the MAC address of the requesting party and passing it through a suitable one-way function to produce a pseudo-anonymous identifier. We maintained a BerkeleyDB which took the MAC address onto a serially assigned integer ID, used during the collection period, and subsequently thrown away. In our final analysis, URLs were also passed through one-way functions as they were not inherently relevant to our study. By choosing the MAC address as the origin of our identifiers, we made the assumption that most computers in the building have a single primary user. Students may certainly have guests present who might use their computers, but the bulk of the traffic will be generated by the primary user.

Removing session identifiers and passing the MAC address through a one-way function are often characterized as sufficient to preserve anonymity. A similar method was used in Bestavros et al., although no analysis of its effectiveness was performed (Bestavros et al., 1995). Gribble employed hash functions more extensively - to conceal source IP, destination IP, and requested URL (Gribble, 1997). From the pattern of their requests, it may or may not be possible to identify users. This increases in likelihood if user traffic is highly dissimilar (i.e. exhibits low homophily). If it is not possible to infer identity directly from traffic patterns, it may be possible to create a profile that allows tracking even without current methods such as tracking cookies. Lasko and Vinterbo suggest that removing identifiers and obfuscating the MAC address are not necessarily sufficient to prevent users from being identified. They describe a tension between privacy protection and analytic utility, suggesting that techniques such as their proposed spectral anonymization must be used to maintain subject privacy in high-dimensional datasets. This technique was not publicized at the time of our study, however, it appears to offer some further subject protections (Lasko, Vinterbo, 2010).

While our data was removed post-analysis, and by the very design of the study, data was never intended to be shared, there is some argument for the benefit to all stakeholders that might arise from increased data sharing among researchers. Kenneally and Claffy, specifically, argue that data sets such as ours, with first-order and second-order identifiers removed are less likely to pose privacy risks for subjects (Kenneally, Claffy, 2010). Where researchers have defaulted to not sharing data sets, this is often due to unclear policies on data sharing permissibility, but also increasingly on an inherent concern for stakeholder privacy.

4– Difficulties in Ethical Research

Despite the precautions and design of the study, we encountered an ethical quandary for which we were unprepared. In the course of the data analysis for this study, we also discovered the presence of a poorly-written anonymization service that was attempting to obscure traffic to an adult chat site by spoofing requests from hundreds of uninvolved clients. Adjusting the technical parameters of the study was a simple matter: these requests were removed from the data set because they were not genuine. However, this behavior was identified with enough time left in the study to have altered the code collecting our data to retrieve and store the actual MAC address of the interface card making the requests. From there, with the assistance of university technology professionals, we could have located the actual student machine responsible, and presumably, its owner. In doing so we might have informed this student of the poorly-coded nature of this software and directed them to specific other anonymization resources, such as Tor, that provide more robust anonymity (Dingledine, Matthewson, 2006). While network usage policies requiring only work-related content are both uniquitous and ubiquitously ignored, an official identification can be problematic. That the individual was using privacy-protecting software was taken as an expressed preference that the user sought anonymity. Thus identification, no matter how well-intentioned, would violate this express preference even if the result were long-term decrease in technical risk. Also, given the sites identified orientation, the social loss of privacy would arguable override the technical increase in privacy.

Ultimately, we decided not to make an effort to contact the student responsible for these anomalies. The chance that they might avoid future embarrassment from the detection of their behavior simply did not outweigh the immediate costs. No matter how carefully delineated the data collection policy, there can always be a user that locates that window of uncertainty. This incident is an excellent example of the tension between security, privacy and trust issues in network research.

Had we discovered a user who was visiting child pornography sites, we might have been legally obligated to report this occurrence. Steinberg et al. identify this concern in the context of child abuse and clinical studies, but this dilemma is equally applicable in network research. The central question, as they put it, is "whether researchers have a moral duty to place the health and safety of children above concerns about confidentiality and the benefits of obtaining new knowledge" (Steinberg et. al, 1999).

Solutions to this problem require further exploration of the ethical responsibilities of researchers who are conducting studies in which data is being anonymized in real-time.

5– Ethical Design Failures

In several situations, the execution of this study was not able to wholly comply with the philosophy of ethical design. To begin with, the physical structure of the network architecture imposed severe limits on our ability to conduct the study with an opt-in recruitment model, which more fully preserves the rights of stakeholders. While contacting individual students was possible, and it is certainly possible that these students might have been able to successfully identify and transmit their MAC addresses to us for the purposes of the study, the packet monitor could only have been attached at a central data collection point. There, each packet could have been compared to a list of approved MAC addresses, but this would have involved viewing unauthorized packets simply for the purposes of determining whether or not they were originated by valid participants, thus violating the opt-in model entirely. The assumption of one student to one MAC address would also need to be made in this case. As well, during particularly high traffic times this approach risks discarding packets by adding several additional processing steps. Finally, concern about the number of students who would be willing to undertake the steps necessary to identify a MAC address and provide it prompted us to decide that the study would be much more accurate with a larger population. In this case, technological limits drove study choices, even when ethical design advocated different methods.

While notifications were provided to the students by mail, we had no way of accurately assessing how many of them read or understood the notices, outside of the small number of emailed questions which were received about the study. Generally, the students appeared to exhibit a lack of concern, but this is not the same as full understanding and consent. Similarly, while a method to opt out was provided, opting out is fundamentally less friendly to a study population than opting in, whether or not the method of opting out increased user security. Since we were not, in the end, permitted to conduct an educational panel at the dormitory, we cannot be sure of student response or understanding of the study truly was. We acknowledge that, had we been seeking research on ethical study design at the time, such a survey would have been valuable.

6– Ethical Design Success

While the above section details the difficulties we had in conducting this study in a manner fully in accord with ethical design, there were also several notable areas in which we succeeded in doing so. During design, data collection, and analysis, when the study generated or identified risk all data were purged of identifiers or other elements of identifiable risk to the subjects, whether or not this was initially in accordance with our design methodology.

Because opt-in was not technologically friendly nor likely to generate sufficient data for the study, we took careful steps to ensure that no identifiable data were recorded on any individual, in accordance with current best practices for aggregate data anonymization at that time. The data set was retained only for the period of the data analysis. In the two cases in which stakeholders contacted us with concerns, we communicated with each of them to discuss how these concerns might be mitigated in a technological and procedural context. Aside from the subject who had concerns about student government involvement, the other student had a simple question about study dates. Limited student inquiries make it difficult to assess how successful this approach was in informing students of their rights and any potential risks, as we could not conduct any follow-up studies with individuals.

7– Experiment Proposal & HSC Review

This section has two components. In the first component, we present a selection of the questions asked by the Human Subjects Committee, and the answers we gave to them, in order to illustrate the concerns that were addressed in the course of our receiving approval to proceed. Questions are presented inset in italic font, and our answers follow. While this framework is geared towards research performed in a laboratory setting, it still assisted us somewhat in our study design. We include this here both as a (admittedly highly imperfect) model for other researchers, but also as background for the development of the student communication (i.e., the flyer) describe in the section component. The changes in the research flyer illustrate the evolution of the communication to the student population.

Briefly describe, in lay terms, the general nature and purpose of the proposed research, and where the study will take place. If student research, indicate whether for a course, thesis, dissertation, or independent research. If the study is only for a course, please review the Student Research Policy to ascertain if this project requires HSC review.

The purpose of this study is to measure the degree of homophily (same-ness) in web browsing traffic, and the effectiveness of current standard practice for de-identifying individuals. The first part (the degree of same-ness) is important because it governs the degree to which users in social networks can gain a benefit by voluntarily cooperating to protect themselves against hostile agents. The second part (current practices for anonymization) is important, because the degree to which these practices are effective has not been properly evaluated. This research will allow us to better protect users from malicious agents on the web.

To implement this study, we plan to record certain parts of the requests for web pages made by web browsers. In particular, we plan to record the URLs in requests made by students in an IU Bloomington residence hall. By following currently accepted anonymization protocols, it is believed that no individual student will be identifiable during this process. However, testing this assumption is also part of this research. The anonymized data will be collected by a secure server, and then placed onto a CD for processing. This CD will be stored in a locked room.

Specifically, we will be recording traffic to port 80 (the port web browsers use to make their requests). Non-browser traffic to this port will be ignored. We will not actually be fetching the web pages requested, or any content thereof. Any traffic that is secured (i.e. encrypted or otherwise secure traffic, such as email, bank traffic, or records) cannot be monitored by our system and will not be recorded. URLs (web page addresses) that contain any sensitive information that could be used to identify individuals will have this information removed prior to storage. To monitor this information, special hardware will need to be installed in the network closets of the residence hall. We have been coordinating this request with UITS.

Describe the process by which subjects will be recruited, how many (or estimate) subjects will be involved in the research, and how much time will be required of them. List specific eligibility requirements for subjects (or describe screening procedures), including those criteria that would exclude otherwise acceptable subjects. If your study uses only male or female subjects, explain why. For NIH-funded research only, address the inclusion of women, minorities and children in the research. Disclose any relationship between researcher and subjects - such as, teacher/student; superintendent/principal/teacher; employer/employee.

A particular residence hall on the IU Bloomington campus will be selected based on its technical suitability for this research, as decided by UITS. No monetary gain or reward will be offered for this study.. No investment of time or resources is required by the users. There is no relationship between the researchers and the subjects. No particular sub-population of students is being targeted for this research.

We plan to collect 2 months worth of traffic data. As a student holiday such as Spring Break may fall during this period, and other delays may be involved in deploying the system, we would like to request 10 weeks to perform the data collection period of this study.

List all procedures to be used on human subjects or describe what subjects will do. If done during regular class time, explain what non-participants will do. If you are taping, explain that here (see item 13 on page 11). Asterisk those you consider experimental. For those asterisked procedures, describe the usual method(s), if any, that were considered and why they were not used. (See item F on page 2 for more information.)

Subjects will simply continue to use web browsers as normal. Any page requests they make that are unencrypted will have their URLs recorded and be anonymized relative to the user, according to current best practices.

Describe methods for preserving confidentiality. How will data be recorded and stored, with or without identifiers? If identifiers are used describe the type: names, job titles, number code, etc. How long are identifiers kept? If coding system is used, is there a link back to the subject's ID? If yes, where is the code list stored in relation to data and when is the code list destroyed? How will reports will be written, in aggregate terms, or will individual responses be described? Will subjects be identified in reports (see item 5 on page 10)? Describe disposition of tapes/films at the end of the study. If tapes are to be kept, indicate for how long and describe future uses of tapes.

We will be using a one-way function to convert the hardware address of each student's computer (an identifier) to a number in a manner that is non-reversible. This number will be the only tag attached to the page requests we record. This is currently believed to be sufficient for anonymity, and there is no link back to the hardware address.

Reports about the data gathered will be presented in aggregate form – no individual history will be presented for any user. While it is believed that no subject can be identified from our data, if we should succeed in doing so, this result will not be presented in conjunction with any individual identity. After the study, data will be retained only on a single encrypted CD stored in Dr. Camp's room in a locked cabinet. This CD will not be made publicly available.

We will attempt to remove all information in the truncated URLs that could be used to identify individuals prior to storage in our research dataset, and will be keeping track of the number and kind of these instances (without any personally identifiable information associated with those instances) as an important research finding. In order to perform this removal, we will need to personally examine the dataset, as no automated process exists to perform this operation.

What, if any, benefit is to be gained by the subject? In the event of monetary gain, include all payment arrangements (amount of payment and the proposed method of disbursement), including reimbursement of expenses. If class credit will be given, list the amount and the value as it relates to the total points needed for an A. List alternative ways to earn the same amount of credit. If merchandise or a service is given, indicate the value. Explain the amount of partial payment/class credit if the subject withdraws prior to completion of the study.

No individual rewards will be offered to subjects. However, users will be presented in the attached notice with easy steps to take to increase their security while browsing the web. Even after the study has been completed, if they continue to follow these steps they will continue to receive increased security.

What information may accrue to science or society in general as a result of this work?

Data from this study will help lead to better systems for protecting users from malicious agents on the web. Additionally, our research about current anonymization techniques will help establish whether these methods, commonly used today, are actually accomplishing their goals. If not, new techniques need to be developed to protect user privacy. This finding will be extremely important either way.

Is modification of the required elements for informed consent requested?

For technical reasons, as discovered by UITS, it is not feasible to obtain individual consent from each potential subject. Instead, we are attaching a notice, which would be disseminated in hardcopy to every student living in the targeted residence hall. This notice outlines the purpose and practice of the research, and presents ways to opt out of it (which also increase each individual student's security), as well as the information that no penalty or harm comes from opting out.

Explain how this research involves no more than minimal risk to the subject. Loss of confidentiality is, under most circumstances, more than minimal risk. However, contact by primary care givers or others who by the nature of their involvement with the subject already have access to the data, will be considered no further loss of confidentiality and, therefore, may be less of a risk to subject confidentiality. Risk may also vary with the type of information being collected.

Users are being de-indentified according to current best practices for doing so. As well, all data collected is non-encrypted, and is not technically protected from eavesdropping. Confidential information such as web browsing in regards to email, financial services, or academic records should be protected by encryption and thus explicitly non-recordable for this research.

Explain how the research could not practically be carried out without waiver of informed consent.

Technically, it is not possible to opt-in subjects for this experiment. Because of the way that networking functions, an area is targeted. For this reason, we are presenting students with a way to opt-out instead, which is technically feasible. No deception is involved in this research.

Explain how, if appropriate, subjects will be informed of pertinent information after participation.

Students will be informed of all pertinent information by the attached notice. Students will receive the notice as a flyer in their mailbox. We will also be working with Dorm management to disseminate this information through any channels they are willing to make available to us.

8– Student Communication Design

It required six iterations of the experimental notice. These iterations and other supplemental materials are maintained online as Appendices, and can be found here: http://www.ljean.com/research_ethics.html. Subtle variations of the experiment were required, but the most vexing difficulty was communicating the nature of the study to the students and, no less problematic, the IRB. Often suggestions by the IRB were rejected by the researchers as conveying information that was technically incorrect. We arrived at a consensus between the research team, the IRB, and the Human Subjects Committee. Specific changes included the alteration of phrases that were viewed as potentially unclear to students, such as "publically available data" to "a truncated version of the URLs". This allowed students to better understand the scope of the research, and to decide whether or not they wished to opt out. For example, the most contested paragraph began in this form:

If you decide you don't want to participate in this experiment, several tools are available that convert all public information transmitted by your web browser into private data, which will then not be recorded in this experiment. Tor (http://www.torproject.org/) is one of these tools, as is IndianaU's Virtual Private Network. An installer for the VPN can be obtained from IUWare (http://iuware.Indiana.edu/list.aspx?id=134). These links also contain information about how to install and use these tools. Students who decide not to participate in this research will not be penalized in any way for doing so, nor can you be identified as someone who opted out.

After significant contesting by the researchers, the reference to Tor was removed as being potentially in conflict with University policy. Indiana University also has a policy against Tor exit nodes operating on the campus network.

If you decide you don't want to participate in this experiment, several tools are available that convert all public information transmitted by your web browser into private data. Remember that private data will not be recorded in this experiment. Indiana's Virtual Private Network is an easy way to improve your security while removing yourself from this study. An installer for the VPN can be obtained from IndianaUWare here: http://iuware.Indiana.edu/list.aspx?id=134 . Students who decide not to participate in this research will not be penalized in any way for doing so, nor can you even be identified as someone who opted out.

The final version included information about what was not included in the study. The information which began as a paragraph, pointing students to Tor, became a page.

What information are we collecting as part of this study?
We will be recording a truncated version of the URLs (web addresses) of the web pages that residents request as the requests are transmitted over the IU network from your dorm to the Internet. We have chosen to study web browsing from a single dorm because the residents are by definition geographically similar, and they also tend to be similar demographically.

What information are we NOT collecting as part of this study?
We will NOT record any of the following:

- ***names, demographic information, or other personally identifiable information.***
- ***the full URL or any content of web pages any resident browses. We will only record a truncated URL. We do not expect to encounter any personally identifiable information in the truncated URLs we are collecting, but if we do, we will delete it promptly.***

We will NOT know the identity of a resident browsing the Internet and will NOT be viewing the content of the web pages a resident visits.

What do I need to do if I don't want to participate in the study?
Any web traffic that is encrypted – i.e., traffic that is sent via https:// or through a VPN connection – cannot be captured or recorded as part of this study. So if you decide you don't want any of your web browsing information to be recorded as part of this study, you can encrypt your web traffic in one of two ways:

- ***ensure that the website you are using has a URL starting with https:// instead of http://***
- ***use Indiana U's Virtual Private Network (VPN) service to do all of your web browsing. For more information about how to use the VPN, see http://kb.Indiana.edu/data/ajrq.html, or call the UITS Support Center at x5-6789.***

There is no penalty for not participating in this study. Since no personal information will be recorded, we won't be able to tell whether a given resident even participated or not.

9- Recommendations to Researchers

Each organization conducting this type of research is currently subject to its own code of participant rights and researcher responsibilities. While divided by these differences, there are nonetheless certain principles likely to be common among R1 institutions. Examining the difficulties we experienced in the course of this research suggests a number of recommendations which may assist researchers embarking on similar investigations.

Our research is grounded in the origins of theoretical ethical design of information systems (Friedman, 1996; Friedman & Nissenbaum, 1996; Nissenbaum, 1998a, 1998b). These theoreitical recommendations are; however, sometimes quite difficult to operationalize (Nissenbaum, 2009). Thus we attempt to translate these into straight-forward guidelines.

First, creating a values statement establishing the goals and constraints of the experiment engenders a clear exposition of the benefits of the research being proposed. This not only serves as a mechanism for data minimization, but also inherently brings forward the subjects as moral participants of the research. This creates a natural flow to the next step, identifying the stakeholders in your proposed research. Once they are identified, the next difficulty is establishing the proper representation for each group. This will depend on the structure of research oversight and is addressed in the next section.

Depending on the research subjects, there are four appropriate models of communication. First, when stakeholders are identifiable, then the best practice is to provide ex ante communication which enables opt-out. Second, inadvertent participants (Levchenko et al., 2011) who could not be notified ex ante can be notified ex post, and in fact can benefit from the research through an educational notification. Researchers should recognize that such notification must be itself subject to the values statement in advance. In our case, the difficulty arose in identification of the orientation of an individual whose technical choices illustrated a desire to remain anonymous. To provide another example (again consider Levchenko et al., 2011) idenficiation of individuals who have endeavored to order controlled substances is likely to cause at least psychological harm. This brings us to the third, which is to make an ethical choice to withhold all communication. Research on vulnerable populations or malicious populations (i.e., criminals) arguably compels researchers to make this choice. Under the fourth (remarkably common) model, subjects simply cannot be identified. For example, any large-scale traffic, darknet or third layer security study falls under th is category.

10- Research through the Lens of Menlo

Two critical challenges noted above (and in our research) are the identification of stakeholder interests and the appropriate communication with stakeholders or their represenativies. This is particularly problematic in security research, as the study of real-world fraud inherently requires deceit, and sometimes the obervation of (potentially organized) criminals to whom the research should explicitly not be identified.

Under the Menlo report, the appropriate mechanisms in all four models (listed in Section 9) would be to consult with your Human Subjects Committee (HSC) and Institutional Review Board (IRB) to apply the policies governing network research at your institution. However, this requires identifiable stakeholders with whom you can readily and ethically communicate. Therefore under the third and fourth models, this is infeasible.

When the IRB model is appropriate (i.e., the first and second models), there remains the difficulty of establishing a common language between research boards and researchers. In many cases, even discussing the experiments is not easily done, and requires dialogue in both directions. Researchers are best served by taking the initiative to present their goals and methods in clear language that is understandable to the non-technical. Participants will also need to be able to understand the experiment to provide meaningful informed consent, and are likely to be equally technological naïve. The previously noted values statement should allow the researchers to present the positive outcomes and need for the research so that participants and administrators can both understand the benefits.

Noting the vast diversity of the IRBs themselves, their policies, and their technical competence was a contribution of the USACM/IEEE-USA response (ACM/IEEE/SIGCHI, 2011) to the Belmont Advanced Notice of Proposed Rule-Making (HHS, 1979, HHS-OPHS-2011-0005). USACM/IEEE-USA also noted the unique features of security research (as above); and the sheer impossibility of notifying all participants in networking research.

To be specific, the USACM proposed the following in response to the Menlo Notice of Proposed Rule-Making (ACM, 2011):

1) Collect and analyze data on current practices before taking action.

2) Evaluate the advantages and disadvantages of a variety of research ethics board models, including national and regional review bodies as well as IRBs.

3) Systematically consider related work and guidance from around the globe.

4) Include specialists in research ethics as part of the process.

Personal comments filed by Camp, also to Menlo, reified these recommendations (Camp, 2011). In addition, Camp identified the following as problematic:

1) IRBs lack technical expertise and are under-resourced.
2) IRBs have structural incentives to deny research, as this is the low-risk choice.

3) Ensuring compliance is beyond the ability or even mission of many IRBs; thus creating a perverse incentive for researchers to avoid the IRB process altogether.

Camp's alternative was a national review body composed of ethicists and technologists that provide consistent review that appreciates both networking research and subject protection. In addition to a Federal body, Camp suggested appropriately ethically-augmented bodies under the purview of IEEE, ACM, IETF, or ARIN for the United States. (Camp, 2011)

Harvard faculty proposed a consistent set of standards instead of a consistent set of reviewers. (Vadhan et al., 2011) The core essence of the comment was that the HIPAA standards for privacy and security were not adequate. While the focus of the Harvard comment was not networking research, the inapplicability of HIPAA appears self-evident. For example, networking researchers obviously cannot obtain signed, informed consent from all observed participants. The proposed "safe-harbor" list would enable research that is designed to limit risk and maximize benefit. This would remove review from standard low-risk research. The Harvard proposal suggested that the safe harbor would be a function of class of data; source of data; method of sharing data; informed consent mechanisms; and class of subjects. This is particularly promising for meta-scale empirical evaluation of network traffic, for example.

The research described here followed the shared recommendations of Belmont and Menlo by using a standard IRB model (This model informed Camp's comments to Menlo). By using value-sensitive design and responding to IRB comments, the research complied with the recommendation set forward by the ACM Menlo response. Due to the need for larger-scale data, the research provided a small exemplar for the IEEE/USACM comments. Because of the data compilation method and the degree of anonymization of the data, the research might have been subject to the proposed Harvard Safe Harbor.

11- Conclusions

In this paper, we presented the development and implementation process for a potentially problematic experiment as conducted by researchers at Indiana University. This is intended as a case study. Nonetheless, some of the mechanisms of the study are generalizable, while others are not.

The ethical design approach we applied required identifying and communicating with meaningful stakeholders. Dialog with student stakeholders was difficult due to the restrictions on contacting them. However, having the ability to identify the stakeholders made this research possible. Being able to identify subjects is often a substantive challenge in networking research, much less distributing flyers. Having a formal body capable of representing the stakeholders offered the possibility of informed representation for subjects. We determined that a general use network policy (such as the one students accept in order to gain access to the campus network) was insufficient consent for participation in a specific experiment, even with experimental ethical review.

Due to technological restrictions, we were forced to use an opt-out model, rather than an opt-in model. Users who wished to remove themselves from the study were required to take specific steps to do so. While these steps increased participant network security, it is still not desirable to force users to take action to avoid participation. Yet the ability of individuals to effectively remove themselves from the research before the study was initiated offers a (weak) promise for general network research.

Ethical systems should take steps to reduce data footprint whenever possible, especially when the data derives from stakeholder actions. While the collection of network traffic data at the packet level will never generate a small data footprint, we endeavored to follow the principles of ethical design by observing constraints on data granularity. Pseudonymity was used to protect subjects, and only minimal packet data were stored.

Ethical systems can address incentive misalignment issues should these emerge, engaging all stakeholders in this process. This study was designed to meet the needs of researchers and thus ideally science, but subjects did not experience any difference in their normal activities if they chose to remain in the study.

Developing language that is meaningful to the potential subjects was a difficult task for the researchers. What appeared to be obvious, correct technical notes to the researchers were identified as impenetrable to the subjects. Arguably best technical design (i.e., Tor) conflicted with University policy.

The design and limitations on the study did not prepare us for an ethical quandary: communicating with a subject using an obvious and flawed anonymity tool. Instead of obscuring the traffic, the anonymzing tool made it far more notable. No clear ethical guidance was available from either codes of ethics nor recorded practice. We choose to respect the student's expressed preferences.

Ultimately, this study allowed us to experience the difficulties on a small scale that can emerge from attempting to comply with ethical design principles during network research. This case study illustrates examples of successes, failures, difficulties, and most of all, lessons learned.

The security community has had a well-identified struggle with ethical research. For example, a famed experient used a Tor exit node and de-anonymized the subjects simply because a Tor exit node provides a ready flow of browsing information. Rather than being refused publication, the authors have received both publication and wide citation. The problem of ethical experimentation is not one for the lone reseacher, but instead a challenge to the community as a whole. SOUPS is a venue that has struggled with this by requiring proof of IRB approval for acceptance. The Menlo Report is an important first step; yet best serves as a starting point for dialogue on the future of network research.

12- Future Work

Future work includes a meta-experiment composed of three smaller experiments on student ethical awareness. In this meta-experiment and the smaller experiments it contains, the authors will explore both further examples of research design and the impact of such research on the ethical awareness of stakeholders. We acknowledge that the subject population will make this less than highly representative and will seek a wider population following the initial implementation.

The creation of consensus on the nature or process of ethical network research (whether partially anonymized or not) is not likely to be easily reached. The most common human subject research involves face-to-face (often medical) interaction or surveys of identifiable non-malicious subjects. Despite this, traditions of ethical offline research provide an important source of guidance for network researchers. We hope to contribute to this discourse both as researchers and as active members in the appropriate policy bodies (i.e., USACM and IEEE-USA).

The Menlo Report and the Health and Human Services update of Belmont provide both guidance and an implied warning. The guidance is clear. The implicit risk is that without self-governance or community-driven standards, network and security research might move towards an inappropriate model implemented by a patchwork of sometimes technically naïve, inconsistent, yet official governing bodies.

References

ACM, U.S. Public Policy Council Of The Association For Computing Machinery (2011), Document ID: DHS-2011-0074-0016 ibid DHS-2011-0074

ACM/IEEE/SIGCHI, U.S. Public Policy Council of the Association for Computing Machinery, the Institute for Electrical and Electronics Engineers, Inc. – USA, & the Association for Computing Machinery's Special Interest Group for Computer-Human Interaction (2011), Document ID: HHS-OPHS-2011-0005-0952 ibid HHS-OPHS-2011-0005

Bestavros, A., Carter, R., Crovella, M., Cunha, C., Abdelsalam, H., & Mirdad, S. (1995). Application-level Document Caching in the Internet. *IEEE SDNE.*

Camp, L. Jean (20011), Document ID: DHS-2011-0074-0013 ibid DHS-2011-0074

Camp, L. Jean & Genkina, A., "Social Networks", *Phishing*, Springer-Verlag, eds. M. Jakobsson and S. Myers. (Berlin, DE) 2006.

Claffy, KC (2008). Ten Things Lawyers Should Know About the Internet. Ret. March 20, 20011, from CAIDA: http://www.caida.org/publications/papers/2008/lawyers_top_ten/

Cockburn, A., & McKenzie, B. (2001). What Do Web Users Do? An Empirical Analysis of Web Use. *Int. Journal of Human-Computer Studies, Vol 54, Issue 6 .*

Cummings, M. L. (2006). Integrating Ethics in Design through the Value-Sensitive Design Approach. *Science & Engineering Ethics*, 12(4), 701-715.

DHS-2011-0074 (2011) Notice and Request for Comment on "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research" (``Menlo Report") for the Department of Homeland Security (DHS), Science and Technology, Cyber Security Division (CSD), Protected Repository for the Defense of Infrastructure Against Cyber Threats (PREDICT)", *Federal Register v76, (249)*

Dingledine, R., Matthewson, N., & Syverson, P. (2004). Tor: The Second-Generation Onion Router. *Usenix Security, August.*

Dingledine, R., & Matthewson, N. (2006). Anonymity Loves Company. *WEIS,* Cambridge, UK.

Dong, Z. & L. Jean Camp, (2010) "The Decreasing Marginal Value of Evaluation Network Size", SIGCAS Computers and Society, Volume 40, Number 4

Friedman, B. (1996). Value-Sensitive Design. *Interactions, v 3* , 16-23.

Friedman, B., & Nissenbaum, H. (1996). Bias in Computer Systems. *ACM Transactions on Information Systems Vol 14* , 330-347.

Gribble, S. (1997). System Design Issues for Internet Middleware Services: Deductions form a Large Client Trace. *Proceedings of the Usenix Symposium on Internet Technologies and Systems*, (pp. 207-218).

Hall, J. L. (2011) "Responsible Research with Anti-Censorship Technologies", DARPA/NSF Mtg on Ethical, Legal and Social Issues of PII, Arlington, VA

HHS-OPHS-2011-0005 (2011) Advance Notice of Proposed Rulemaking. "Human Subjects Research Protections: Enhancing Protections for Research Subjects and Reducing Burden, Delay, and Ambiguity for Investigators" for the Office of the Secretary of the Department of Health and Human Services (HHS), Office of Science and Technology Policy (OSTP)", *Federal Register v76, (143)*

HHS (1979) Advanced Notice of Proposed Rule Making ("Belmont Report") Ethical Principles and Guidelines for the Protection of Human Subjects of Research, The National Commission for the Protection of

Human Subjects of Biomedical and Behavioral Research, for the US Dept. of Health and Human Services (HHS), from: http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html

Indiana University. (2011, February 26). *IU: Office of Research Administration*. Retrieved February 26, 2011, from IU ORA: http://researchadmin.iu.edu/HumanSubjects/

Kenneally, E., & Claffy, KC (2010). Dialing Privacy and Utility: A Proposed Data-Sharing Framework to Advance Internet Research. *IEEE Security and Privacy*.

Lasko, T. A., & Vinterbo, S. A. (2010). Spectral Anonymization of Data. *IEEE Transactions on Knowledge & Data Engineering*, 22(3), 437-446.

Levchenko, K., Pitsillidis, A., Chachra, N., Enright, B., Felegyhaziz, M., Griery C., Halvorson, T., Kanich, C., Kreibichy, C., Liu, H., McCoy, D., Weavery, N., Paxsony, V., Voelker, G., and Savage, S. (2011), Click Trajectories End-to-end Analysis of the Spam Value Chain, *IEEE Symposium on Computer Security*

Meiss, M., Duncan, J., Goncalves, B., Ramasco, J., and Menczer, F. (2009). What's in a Session: Tracking Individual Behavior on the Web, *Hypertext*

Mosher, D. L. (1988). Balancing the Rights of Subjects, Scientists, and Society: 10 Principles for Human Subject Committees. *Journal of Sex Research*, 24(1-4), 378.

Nissenbaum, H. (1998a). *Values in Computer System Design: Bias and Autonomy*. Delhi: New Academic Press.

Nissenbaum, H. (1998b). Values in the Design of Computer Systems. *Computers in Society* , 38-39.

Nissenbaum, H. (2009) *Privacy in Context*. Standford: Stanford Law Books.

Soghoian, C. (2011) "Enforced Community Standards For Research on Users of the Tor Anonymity Network", *2011 WECSR*.

Soghoian, C., Dingledine, R., McCoy, D., Caspar Bowden, C., and Marcia Hoffman, M., (2011), *"Panel: The Ethics of Research on Tor Users," 11th PETS*, Waterloo, CN, July 27-29, 2011

Steinberg, A. M., Pynoos, R. S., Goenjian, A. K., Sossanabadi, H., & Sherr, L. (1999). Are Researchers Bound by Child Abuse Reporting Laws?. *Child Abuse & Neglect*, 23(8), 771-777.

Tsow, A., Viecco, C., and Camp, L. Jean, (2009) "Privacy-Aware Architecture for Sharing Web Histories", *IBM Journal of Research & Development*, 2009.

Vadhan, S., Abrams, D., Altman, M., Dwork, C., Kominers, S., Kominers, P., Lewis, H., Moran, T., Rothblum, G., & Ullman, J., Harvard, Microsoft Research, University of Chicago, MIT, Herzilya Interdisciplinary Center (2011) , Document ID: HHS-OPHS-2011-0005-1101 ibid HHS-OPHS-2011-0005