

CONDUCTING AN INTERNET MEASUREMENT PROJECT IN AN INTERDISCIPLINARY CLASS CONTEXT: A CASE STUDY

Alan Nochenson, Jens Grossklags, Kathryn Lambert

The Pennsylvania State University (USA)

Abstract

This paper describes a project that was conducted at a large interdisciplinary college in a Fundamentals of Information Security class. Motivated by the prevalence of behavioural tracking on the Internet, we designed this project for students to garner a better understanding of the related technical issues and policy considerations. The main component of the project was a small-scale Internet measurement study, which was conducted in teams. The learning objectives aimed to foster students' technical abilities, data analytic skills, and their critical thinking. Students developed their technical skills by collecting data with a state-of-the-art measurement tool and their analytic skills by analysing the data and creating effective visualizations. Students engaged in critical thinking exercises through readings and the writing of a reflective essay at the conclusion of the project. This paper discusses the effectiveness of this project in meeting its learning objectives and outlines potential modifications to better serve these objectives. With our contribution, we want to spark a discussion about best practices for conducting comparatively complex measurement projects in interdisciplinary class contexts.

Keywords: Measurement Study, Computer Security, Privacy, Interdisciplinary Education

1 INTRODUCTION

Personal consumer data has been referred to as the new oil of the Internet and the new currency of the digital world [1]. Due to its increasing importance for economic decision-making, individuals need to be aware of the value of data and how to manage their information to avoid falling prey to unfavourable business practices, privacy invasions¹ or even cybercrime. Data about consumers' web browsing, online activities and shopping is at the forefront of the development of new business practices. Particular concerns are related to behavioural tracking and resulting actions such as targeted advertisements and online price discrimination [3] [4].

Motivated by these observations, we developed a semester project that involved students in a multifaceted study about behavioural tracking on the Internet. The project was introduced in a Fundamentals of Information Security second-year class which is a required course in an interdisciplinary bachelor's program on Security and Risk Assessment. This program attracts students with diverse backgrounds and interests, and allows them to specialize in different directions with foci on technology, analysis, or social science. Because of this diversity, our project needed to account for many levels of technological expertise [5] [6]. Further, like a typical sample of Internet users, we expected them to have struggled with security and privacy decisions in the past [7] [8].

With this project, we aimed to address the following intertwined learning objectives to increase *technical abilities*, *analytic skills*, and *critical thinking*:

- Develop technical skills needed for working with a cutting edge computer security and Internet measurement tool
- Utilize the measurement tool on a realistic and practically relevant website dataset to collect evidence on the prevalence of behavioural tracking
- Analyse data to comprehensively answer a number of questions about behavioural tracking
- Create informative and meaningful visualizations of analysed data

¹ For example, the retailer Target came under fire in early 2012 for its "pregnancy prediction" score which predicted a teenage girl's pregnancy and triggered the sending of personalized advertisement materials through the mail. As a result, the girl's father first learned about the pregnancy from the unsolicited mail [2].

- Reflect on the implications of the data analysed, the experience of learning the technical skills, and the public policy implications of the practice of behavioural tracking

To accomplish these objectives, we employed the following instructional strategies. Students were assigned to first acquire, install, and configure a cutting-edge computer security and Internet measurement tool [9]. That is, students were put “in the shoes” of a typical computer security professional who needs to be updating her skills on an everyday basis and familiarize herself with constantly evolving technologies. Second, students were to utilize the tool on a realistic dataset (i.e., a diverse set of popular Internet websites) to measure the extent to which these websites and their third-party affiliates use technologies to track user behaviour. More specifically, we aimed for students to learn what and how many artefacts (e.g., tracking cookies) are used by websites and third-party affiliates to authenticate and to compile behavioural records about users. Third, students were to analyse the collected data to comprehensively answer a number of questions about behavioural tracking and to create meaningful visualizations (for example, to show “which advertisers are most prevalent across the sites visited”). We intended that students would learn about the pervasiveness of user tracking and the deeply complex business relationships that shape, for example, the online advertisement industry. Fourth, students were asked to reflect on their technical and analytic findings with a personal essay.

Together these steps aimed to guide students to overcome their lack of knowledge about privacy and security practices in a scenario with high relevance for their own lives. Students studied what websites collect that they use frequently, and where they share deeply personal information or use sensitive credentials. Similarly, we laid out a path for students to not only become sophisticated consumers of data, but even more so producers of knowledge in an area that is still largely unexplored. We report findings about learning assessments and feedback we solicited from the students. We enrich our report with observations about specific challenges we encountered during the semester, and recommendations for conducting measurement projects in an interdisciplinary class context.

The paper proceeds as follows. Section 2 gives additional details about the learning objectives associated with this project. Section 3 describes the project steps in detail and how it was administered. In Section 4, we describe our observations, and in Section 5 we discuss possible strategies for best practices and potential improvements for conducting educational measurement projects. In Section 6, we discuss validity considerations. In Section 7, we discuss related work, and we conclude in Section 8.

2 LEARNING OBJECTIVES

The high-level goal of this project was for students to understand what behavioural tracking is and how someone would go about assessing the extent of the phenomenon on the Internet. To this end, we developed a number of learning objectives for the students. These objectives fell into the categories of technical abilities, data analysis skills, and critical thinking.

2.1 Technical objectives

Work as a security professional involves dealing with an ever-changing problem and threat landscape. As a consequence, security professionals need to constantly update their skillsets to be able to respond to new threats. Typically, this work involves working with numerous pieces of software which are not designed for an average home user, and are thus difficult to use without technical skills and a learned affinity for complex tools.

This project required students to develop their technical abilities by engaging with a real measurement tool, e.g., to develop skills in command-line usage and troubleshooting [9]. With the help of the tool, students were requested to gather a dataset of the prevalence of tracking cookies for a set of popular websites. Utilizing the tool effectively required students to operate like a software developer at the end of a development cycle. Students were not required to do any programming, but were required to run code created by others, in a development environment. Through this exercise, we aimed to give students the confidence needed to use similar tools in the future, in addition to building proficiency at the current task.

In addition to this project, the course also included a number of virtual laboratory exercises, which also had students work with different security tools (e.g. Wireshark²). However, those virtual labs “held the

² <http://www.wireshark.org/>

hands” of students more than we planned for in this project. Motivating factors for the different approach in the course project were the comments of students in previous instances of the course that requested a more open and less controlled technical study environment. Students suggested that the virtual lab instructions were too detailed (e.g., explained every step) and left little room for error or challenging situations. With this in mind, we developed this portion of the project to be at a level somewhat farther beyond what we believed students were capable of *a priori* (compared to the virtual labs) in order to allow them to encounter challenges and to overcome them.

2.2 Data analysis objectives

The project engaged students with a straightforward, but practical, data analysis exercise. Based on the data collected with the help of the measurement tool, students were asked to identify patterns in the data that would shed light on the practice of behavioural tracking, and to create informative and meaningful visualizations.

As professionals, it is a common task to create visualizations of data (e.g., bar charts, line plots and data trends). We allowed students to use any data analysis program to complete the data analysis and visualization tasks, however all groups chose Microsoft Excel. We provided students with the means to create a data file (i.e., by making a conversion script available) which would be compatible with different analysis tools. After preparing the data file, we then asked them to analyse the data and to create visualizations in response to a number of questions about tracking behaviours (e.g. “which advertisers are most prevalent across the sites that you visited”).

To complete the assignment, students needed to interpret what the data they collected means and represents (i.e., how to read and interpret the name and contents of a tracking cookie that was placed when browsing to a particular website) and to choose which data needs to be used to answer a variety of questions.

2.3 Critical thinking objectives

In response to our overall goal, we aimed to make students aware of the advantages and disadvantages of behavioural tracking (e.g., increased personalization versus potential intrusions of personal privacy). Additionally, we intended students to learn to express a well-reasoned opinion on the issue of behavioural tracking based on their own analyses, and articles in the popular press and research from academia. Further, the essay assigned to students required them to reflect about their learning experience, and to articulate their personal take-away lessons from the project about behavioural tracking.

The project was set up to be of relevance for the everyday behaviours of the students. Specifically, we were interested in having students understand how their own personal browsing habits are subject to data collection and how this information may eventually be used to positively customize their online experience, but also to discriminate (primarily in an economic sense) and to benefit monetarily by selling the data to interested third parties. Students were further encouraged to learn background information and skills to make informed choices about behavioural tracking in their personal lives, such as using anti-tracking mechanisms (e.g., opt-out cookies or browser extensions which limit tracking). Finally, students were required to think about the role of technology policy and public policy in mediating between the interests of the different stakeholders of the online advertisement space (in particular, in the context of self-regulation efforts to implement Do-Not-Track standards).

3 LOGISTICS

3.1 The project

The project was administered in three separate parts. The first and third parts required individual students to complete the work, and the middle part was a team-based effort where students worked in pre-established teams of three or four students.

The first part consisted of each student setting up his or her respective computing environment to run a sample crawl (i.e., visit to a single website) and to log tracking cookies using the measurement tool,

the FourthParty browser extension.³ The tool required the use of the developer version of the Firefox browser which was provided to the students together with the extension [9].

In the second part, students were asked to visit 30 websites of their choosing in order to log cookies with the FourthParty extension. Each group of students was instructed to visit each website in three conditions and to record the cookies that were placed on their computers:

1. Shallow condition: Students were asked to only visit the start homepage of a site (while having the Do Not Track feature of the Firefox browser switched off)
2. Do Not Track condition: Students were asked to switch on the browser's Do Not Track feature and to visit the start homepage of a site
3. Deep condition: Students had to thoroughly browse through different layers of each site (while having the Do Not Track feature of the Firefox browser switched off).

As part of the second part of the project, each group wrote a report (including visualizations) answering a number of questions about the findings of their crawls.

In the third part of the project, students wrote a reflection essay about the advantages and disadvantages of behavioural tracking, their thoughts on the issue and their experience with the project. In all project phases, students were required to provide feedback about potential problems they experienced.

The FourthParty tool was developed by academic researchers and was not created for commercial distribution; however, it has been used for a number of research projects in the past (e.g., [9]). Usage of the tool required installing a number of packages and using the command line to open the developer version of the Firefox browser. Theoretically, it should be possible for the browser itself to keep a complete record of tracking cookies and to display this information to users in a meaningful way. However, by default such features are not enabled in the currently available popular browsers, which poses an impediment to user education and serves as an obstacle to market forces (i.e., to encourage more privacy friendly behaviour by web sites). In particular, browsers typically do not present information about third-party cookies, e.g., those placed by advertising companies, and other intermediaries and third-party affiliates. Since behavioural tracking is mostly associated with advertisements on a website, students could not use a commercially ready product to complete the objectives of the project.

From our point of view, we expected the following aspects to pose the most significant challenges to our student group with diverse technical background knowledge:

- Using command-line (terminal) skills to move from directory to directory, copy and move files, and run scripts: The only script that we required students to run was written by a teaching assistant and extracted data from their database files (SQLite files created by the FourthParty extension) into a plain data file, which students were then able to use for their data analysis.
- Data analysis of the tracking cookie data: We expected students to create visual representations of the data (e.g., bar graphs and line plots), to identify trends in the data, as well as to evaluate qualitative aspects of the data. Sub-skills in this category included the use of a data analysis tool, the creation of graphs (including the use of graph design principles such as data-ink ratio [10]), and working towards a good understanding and interpretation of the data via qualitative analysis.
- Develop a high-level understanding of behavioural tracking and the involved stakeholders: Students were required to connect their findings with additional information about the policy debate around behavioural tracking to form their own opinion about the extent of the problem.

3.2 Student composition

Out of 41 students involved in this project, 30 were male (73%) and 11 were female (27%). The students came from different majors, predominantly Security and Risk Analysis (18 students) and Information Sciences and Technology (12 students).⁴ The average age of the students was 21 years

³ <http://fourthparty.info/>

⁴ Students do not declare a major at Penn State until the end of their sophomore year. For those students in the course who were sophomores, we report their "intended" major as their major.

(standard deviation of 1.85 years). Six students were older than 22 years. A single student had been enrolled in the course in a previous semester, but did not complete it. There was no first-year student enrolled in the class. Nine students were in their second year at the university. 18 students were in their third year, 12 students in their fourth year. One fifth-year student was part of the class.

3.3 Data collection

We collected data in different stages of the project and in multiple ways. Following the classification offered by [11], we collected naturally occurring data, research specific data, and reflection data. The *naturally occurring data* which became part of our analysis are grades, problem reports for the first two parts of the project (the second part is a group submission), and a final essay about the project. We obtained approval from our University's Institutional Review Board (IRB) to use this data for research purposes. The *research specific data* that we collected consisted of a survey administered immediately after the second part of the project.

The survey, which was anonymous, asked for demographic information and had two reflection sections, each consisting of eight 5-point Likert-style questions. One section asked students to answer the questions as if they were administered prior to the project's start, while the other asked them to answer the questions from the point of view of their current perspective. Questions were identical between the two sections except for slight wording changes to make tenses consistent.⁵ In addition to course materials and surveys, the instructor and teaching assistant conducted debriefing sessions. These took place after office hours were held and after completed parts of the project were turned in. The data gathered from these sessions was integrated into the *reflection data*.

3.4 Data analysis

The analysis that was performed on the data described above was largely qualitative. One of the authors of this paper examined the naturally occurring data (problem statements and final essay) in depth, identifying common themes within each part of the project and across the various parts of the project. The other authors corroborated these themes with the documents they had read and with sentiments expressed during the project. We also performed selected quantitative analysis on data from the survey.

4 RESULTS

In our discussion of the results, we follow the same categories as outlined in the learning objectives in Section 2.

4.1 Technical results

The primary technical learning objective of this project was for students to effectively utilize a non-production-ready tool to gather real-world data. To help them with this objective, we provided sample command-line commands to run (for each respective operating system platform), and a sample walkthrough of how to use the tools for a single crawl. Executing this crawl correctly was the central task of the first part of the project.

Through qualitative analysis of problem statements from the first phase of the project, we found that several students did not understand the technical meaning of the sample commands provided to them, and were not able to search independently for information to resolve their initial problems. Therefore, some students encountered difficulties when a command did not work when directly copied and pasted from the instruction sheet, e.g., because they forgot to fill placeholders for directory names. Some students also misread critical parts of the directions for this portion of the project. Errors resulting from this included using the wrong commands (i.e., using commands for a Windows system on a Mac machine) and downloading an incorrect version of needed software.

As the project progressed, students had a number of purely technical difficulties with the project that we did not anticipate. For example, when loading certain websites in the modified browser, the browser would "freeze" without an apparent reason. As each unique difficulty arose, we compiled a growing list of Frequently Asked Questions (FAQs) over the course of the project. Eventually, all

⁵ While it would have been better from a methodological standpoint to actually administer multiple surveys, this was not initially discussed. As such, the inclusion of the first section of questions aimed to take the place of the pre-survey.

student teams were able to overcome their initial technical challenges. To facilitate this outcome, we employed different instructional strategies. Students were able to draw from the compiled list of FAQs, we encouraged collaboration between classmates and teams to resolve problems, and we offered office hour sessions (which were held numerous times by the teaching assistant at the request of students).

In the second part of the project, students worked in groups to complete a total of 90 visits to 30 websites. Some students encountered a subset of the same technical difficulties that they experienced while working on project part 1; that is, the training phase was not always fully effective. However, the students were able to draw on their previous experiences to resolve these reoccurring difficulties. We also utilized the same set of instructional strategies to assist the student teams. Overall, all student teams were able to submit a full set of measurement results. We conclude that the project successfully engaged students with a cutting-edge security tool; however, students encountered more difficulties than originally intended. We plan to revise this project part to account for the unpredictability associated with the heterogeneity of computing resources, and to provide students with a tutorial of basic command-line computing usage.

4.2 Data analysis results

The data analysis task took place during the second part of the project, after teams had completed their respective 90 visits. The size of the sample was chosen to enable students to collect a wide variety of sites that vary across genre and tracking practices. Once students overcame their initial technical measurement difficulties, students were engaged in the relatively time-consuming, but according to their feedback comparatively uncomplicated data analysis task. We expected that students would closely collaborate on the analysis of the complete dataset. However, based on our observations the task was not always distributed evenly, due to group self-management issues including time-constraints or different levels of expertise of the team members.

Part of the data analysis learning objective was for students to get a qualitative and quantitative understanding of the data. From student comments during office hours and in the problem statements they submitted, some aspects of the task proved more difficult than we had expected. For example, we made a deliberate decision while developing this project to keep the unmodified data files with raw headings (as given by the FourthParty extension) such as "originSite" and "baseDomain" and to leave in columns that were not needed for their analysis, in order to have students experience the typical output of measurement tools. While we did instruct them on what most of the columns meant, several students were still confused about the naming. Similar to the technical challenges observed in part 1, we partly attribute this to an incomplete and hasty reading of the instructions. Such problems were easily resolved when students engaged with the teaching team. In Fig. 1, we provide an example of a data file and a visualization created by students.

Another challenge during the data analysis was to identify the part of the data that would be most suitable to answer a particular question such as "which sites respect Do Not Track" and "how many more cookies do sites store when Do Not Track is disabled". During office hours, some students attempted to get the teaching assistant to tell them which data to use to answer each question to save them the effort to reason through the tasks themselves. At the end, however, nearly every team was able to correctly identify the data needed to answer the questions.

Following from the learning objectives, students were further tasked to create effective visualizations. Successfully completing this part required the collection of valid data and an initial exploration of the data as required by the previously posed analytic questions. From our conversations with students as well as from their reports, we learned that prior to this project many students had not created graphs in any computerized program. For this reason, several students asked us how to use Microsoft Excel to make a graph from data. We provided these students with an informal tutorial, but also encouraged them to find solutions for such relatively simple tasks via web search. We also explained to them the importance of being able to find answers to their own questions as part of their future professional practice. From their write-ups, all of the groups succeeded in creating visualizations. However, some followed good practices more closely than others (e.g., appropriate labels, including titles or captions, using distinct colours).

Taken as a whole, students did accomplish the data analysis learning objectives we defined at the project start. By the end of the second phase students had thoroughly explored the data, identified appropriate parts of the data to answer the different questions, and created somewhat effective visualizations to present their results.

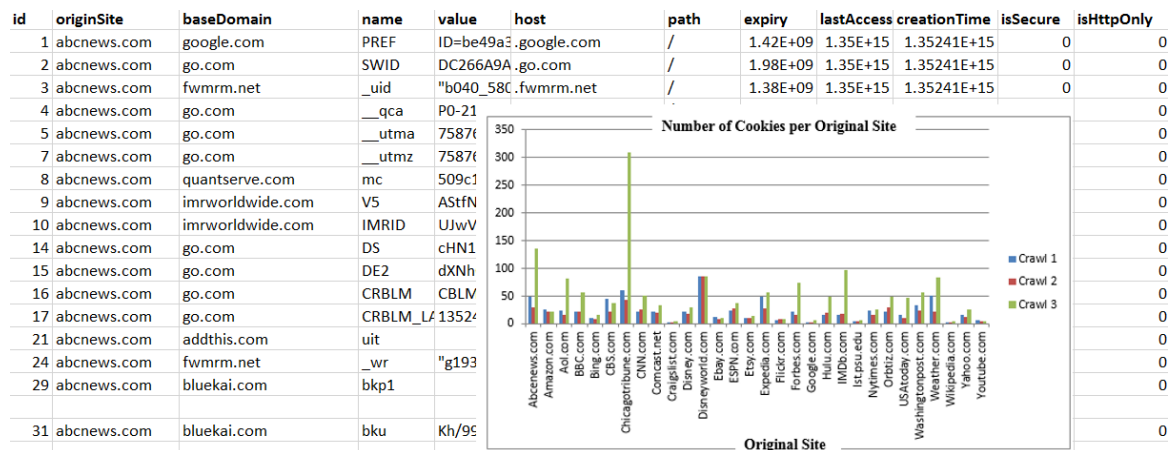


Figure 1: Example data file and visualization to compare differences between number of cookies set by sites when only visited in a shallow manner (crawl 1), visited with the “Do Not Track” setting on (crawl 2) and visited in a deeper manner (crawl 3), respectively. Data and graph above were collected and created by a single team of students. Each row of this data file represents a single cookie. The “originSite” column denotes the site visited and the “baseDomain” column is the name of the site that site the cookie.

4.3 Critical thinking results

The critical thinking learning objectives required students to effectively evaluate the advantages and disadvantages of behavioural tracking and to articulate a well-developed opinion on the issue (informed by their research and other sources). We also solicited students’ reflections on the project to improve the next offering of the course.

Students were able to identify the positive and negative aspects of behavioural tracking, though many students focused very heavily on one side of the problem space without alternative points of view. Student views ranged from “personally, I think that behavioural advertising is great for consumers and online businesses” to “[my] sense of privacy is shattered”. There was also a considerable amount of variability as to the students’ perception about the scale of the problem. Many concluded that sites tend not to respect the “Do Not Track” campaign, while others found that it is mildly successful. The variability in responses can be partly attributed to the fact that each group visited a different set of sites and therefore saw a different sample of tracking behaviours. This variability of results can be seen as a strength but also as a weakness of the project. On the one hand, it is interesting for students to be able to choose the sites that are most relevant to their personal lives. On the other hand, a larger and more predictable sample would have yielded more robust insights. However, it would have been an additional challenge to move from a small-scale measurement study to a more representative set of sites. Several students already remarked that the data collection part was too time-consuming.

Students’ reflections in the final part of the project indicated that they generally had a change of mind over the course of the project about its utility. In the survey and in responses for the first two parts of the project, students voiced doubts about its usefulness and they felt that it was too difficult. In their final reflection essays (part 3), there was a fairly pervasive change of mind. In total, students did not think the project was too hard, and they thought the skills they were learning were more useful than they had previously stated. One student’s words say this better than we can: “I understand that this is college [and not high school] and things don’t just come by hand to us, but even though this project was quite challenging, it helped me learn a lot about things that I knew nothing about. I feel as though I will be able to take the lessons learned from this project and apply them to real world situations.”

5 POSSIBLE IMPROVEMENTS

Given the results we identified above, we developed a number of possible strategies for improvement of future iterations of the project. While these recommendations are based on the analysis of this experience, they will hopefully be transferable to other related situations as well. These improvements include making the process of gathering data more straightforward, giving students additional training

opportunities before the start of the core parts of the project, and holding periodic interviews/forums for discussion with students about issues as they arise.

The process of conducting the actual crawls could be made simpler by “sandboxing” the project or creating a “virtual lab” environment for students to complete the project⁶. Another related option would be to have students use some type of tool which has similar functions and allows them to collect the same data in a more user-friendly fashion (e.g. FireBug⁷ or Google Chrome Developer Tools⁸). This would alleviate many of the technical problems that came out as a result of varied system configurations. However, there would be significant downsides to these strategies as well. Since the act of struggling through the data collection helped students to operate in the role of a security professional, sandboxing the project or letting students use a more production-ready tool would remove some of the technical difficulties and the related expected learning outcome.

Another improvement to this experience would be to offer more related training to students prior to the start of the project (like was done in [12]). This could be accomplished by, for example, holding a class period with an instructional part and exercises related to each of the major skills used in the project (i.e., command-line usage, Excel usage and basic data analysis). A downside to this strategy is that different students have different skill levels, and mandatory sessions may be too elementary for some.

Similarly, we could reserve class time for students to work on the project. In this setting, the instructor and TA would be able to help and guide students in real time. At least one student specifically requested this in order to allow problems to be assessed in a structured environment in timely manner.

Another improvement would be to hold a thorough discussion session after each part of the project was completed, so we could better understand student concerns. However, it is unclear whether the majority of students would engage in such feedback sessions, in particular if they are scheduled after each milestone. For this reason, we propose that a full-class discussion could take place before and after the project, and interviews with a (random) subset of students take place after each project milestone to get a representative sample of the problems facing students. These interviews would largely take the place of the survey that we administered.

6 VALIDITY CONSIDERATIONS

As this paper reports an exploratory case study involving an emerging topic and novel software tools, we do not expect the results to generalize to more traditional technology learning contexts. That being said, there are a number of other possible factors that could affect validity of the study. One threat to validity is that we did not use any specific qualitative coding methods to report our findings (such as is done, for example, in [13]). We believe however that concerns about internal consistency are somewhat ameliorated because we framed outcomes in terms of our original learning objectives. Further, the framing of our learning objectives was done in consultation with a specialist in the fields of education and learning. Nevertheless, we primarily based them on our own evaluation of potential tasks with critical importance for future security professionals.

7 RELATED WORK

McCartney et al. [13] conducted a similar study to ours where they evaluated the effectiveness of using open source software development projects during a software engineering course. The way that they defined objectives is similar to how we did in this paper. They also conducted anonymous surveys and did thematic analyses to tease out interesting results. One key difference is that our course had students of varying technical skills and interests, whereas their course presumed substantial programming skills and software development knowledge. While students in our course were expected to have taken an introductory programming class, many did not seem to have developed the technical maturity to apply the ideas from that course to other tasks, such as this project. The students from [13] were able to better appreciate maintenance and documentation of

⁶ By “sandboxing” we mean creating a controlled environment such as a Virtual Machine that is preloaded and setup to run crawls for this project. With this in place, the technical task would shift from installing and configuring their systems for the crawl task to installing/configuring them for running the sandboxed environment, which is likely less troublesome, but does not have to be trivial.

⁷ <http://getfirebug.com/>

⁸ <https://developers.google.com/chrome-developer-tools>

code, though they did not gain as clear an appreciation for reverse engineering and tool use as the authors had hoped.

Chen and Dong [12] conducted a study of collaborative senior project development for software engineering students. Similar to this paper, they present an exploratory case study in which they used interviews. Also, they collected data from participant logbooks, a source of data which we did not solicit. However, their main evaluation is a quantitative one about how teams work together through meetings. They find that their meetings-flow approach had a positive impact on student performance. Again, the composition of the students in [12] is relatively homogeneous, unlike in our case study.

Our class project is also relevant to the growing number of large-scale measurement studies conducted by academics. For example, Hoofnagle and Good discuss the prevalence of a variety of tracking technologies for the top 25000 websites [14]. Such studies demonstrate how quickly the number of tracking artefacts has grown over a relatively short time span, which raises a number of interesting policy issues [9] [15].

We hope that our students will be able to participate in similar measurement efforts focused on behavioural tracking as part of their professional practice, or apply their knowledge to related measurement issues (such as, for example, privacy and policy issues in conjunction with the growing popularity of third-party applications on social networking sites [16]).

8 CONCLUSIONS

In this paper, we discussed an experience conducting a measurement study in a Fundamentals of Information Security class in an interdisciplinary college. We defined a number of learning objectives including technical, data analytic, and critical thinking objectives. Under these same categories, we presented our findings and discussed difficulties that arose over the course of the project. We proposed a number of possible improvements to the project to address these issues including simplifying the project and increasing the frequency of training sessions. In summation, we believe that this project, while difficult, contributed to the development of skills that help students to become successful security professionals and helped them to gain a better understanding of the challenges related to behavioural tracking.

9 ACKNOWLEDGEMENTS

We would like to thank Lisa Lenze for engaging with us in several discussions during the development of learning objectives, and for her comments on earlier versions of the manuscript. We would also like to thank the students who participated in the course for their time and constructive feedback over the whole duration of the project, and for their permission to use their data in a research setting. Kathryn Lambert gratefully acknowledges support by NASA through the Pennsylvania Space Grant Consortium office as part of the Women in Science and Engineering Research (WISER) program. Further, she is thankful for the support through the Summer 2013 Undergraduate Research Fellowship (SURF) program at the College of Information Sciences and Technology, Pennsylvania State University.

REFERENCES

- [1] World Economic Forum. (2011). Personal data: The emergence of a new asset class. http://www3.weforum.org/docs/WEF_ITTC_PersonalDataNewAsset_Report_2011.pdf.
- [2] Hill, K. (2012). How target figured out a teen girl was pregnant before her father did. <http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/>.
- [3] McDonald, A. and Cranor, L. (2010). Americans' attitudes about internet behavioral advertising practices. In Proceedings of the 9th Workshop on Privacy in the Electronic Society (WPES), pp. 63–72, Chicago, IL.
- [4] Turow, J., King, J., Hoofnagle, C., Bleakley, A. and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it.
- [5] King, J. (2006). Identity in the ischool movement. Bulletin of the American Society for Information Science and Technology, 32(4), April–May.

- [6] Olson, G. and J. Grudin, J. (2009). The information school phenomenon. *interactions*, 16(2), pp.15–19.
- [7] Acquisti, A. and Grossklags, J. (2005). Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1), pp. 26–33.
- [8] McDonald, A. (2010). Cookie confusion: Do browser interfaces undermine understanding? In *Proceedings of the 2010 ACM Conference on Human Factors in Computing Systems, Extended Abstracts*, pp. 4393–4398, Atlanta, GA.
- [9] Mayer, J. and Mitchell, J. (2012). Third-party web tracking: Policy and technology. In *Proceedings of the 33rd IEEE Symposium on Security and Privacy*, pp. 413–427, San Francisco, CA.
- [10] Tufte, E. (2001). *The Visual Display of Quantitative Information (Second Edition)*. Graphics Press, Cheshire, CT.
- [11] Malmi, L., Sheard, J., Simon, J., Bednarik, R., Helminen, J., Korhonen, A., Myller, N., Sorva, J., and Taherkhani, A. (2010). Characterizing research in computing education: A preliminary analysis of the literature. In *Proceedings of the Sixth International Workshop on Computing Education Research, ICER '10*, pp. 3–12, New York, NY.
- [12] Chen, C. and Chong, P. (2011). Software engineering education: A study on conducting collaborative senior project development. *Journal of Systems and Software*, 84(3), pp. 479 – 491.
- [13] McCartney, R., Gokhale, S., and Smith, T. (2012). Evaluating an early software engineering course with projects and tools from open source software. In *Proceedings of the Ninth Annual International Conference on International Computing Education Research, ICER '12*, pp. 5–10, New York, NY.
- [14] Hoofnagle, C., and Good, N. (2012). *The Web Privacy Census*. <http://law.berkeley.edu/privacycensus.htm>
- [15] Grossklags, J. (2011). Do-not-track as a driver for transparency of social networking advertisement practices? In *Proceedings of the W3C Workshop on Web Tracking and User Privacy*, Princeton, NJ.
- [16] Wang, N., Grossklags, J., and Xu, H. (2013). An Online Experiment of Privacy Authorization Dialogues for Social Applications. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2013)*, pp. 261-272, San Antonio, TX.