**Floating Point Precision Vulnerabilities in Differential Privacy Sampling Methods**

**Analyzing Data Collection Policies and Their Transparency to End Users**

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Jack Liu

November 4, 2022

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Caitlin Wylie, Department of Engineering and Society

Tianhao Wang, Department of Computer Science

## Introduction

More than ever before, we are surrounded by devices that are connected to the internet in our daily lives. Many of these devices are equipped with cameras, microphones, and other sensors that companies are willing to use to collect information whether the user is aware or not. Recent studies show that some video conferencing applications continuously monitor microphone input even when on "mute" and even transmit the data to remote telemetry servers (Yang et al., 2022). Moreover, the researchers were able to classify the background activity in the intercepted audio data sent to the servers. This has only been more pertinent as we are now using these conferencing applications in our personal lives and spaces due to the recent pandemic. One bright spot in this area is the development of differential privacy which is a guarantee of using data without revealing the details of any individual. An example of this in practice is a technology developed by Google that can be used to collect aggregate statistics from users of their Chrome Web browser (Erlingsson et al., 2014). However, this field is still new and vulnerabilities may still exist. In my technical report, I will discuss potential attacks on privacy-preserving algorithms based on vulnerabilities introduced by the translation from theory to physical practice. For my STS research project, I will take a step back to analyze the transparency of data collection in our digital world and what improvements can be made to protect the users.

## Technical Topic

Differential privacy is a mathematical definition of privacy that has been a hot area of study in recent years. It describes the property of an algorithm to calculate aggregate statistics on a population while keeping the individual's data hidden. More precisely, it ensures that the probability of any output occurring when one individual contributes to the dataset only differs

from the probability if the individual is removed from the dataset by a small tunable factor. The advantage of defining such a notion of privacy is that individuals can feel protected when their data is collected since it cannot be misused and traced back to a particular individual.

However, there are some nuances related to the finite limitations of computers when translating these theoretical ideas into practice. An insecure implementation can introduce vulnerabilities that destroy the differential privacy guarantee. One paper by Mironov (2012) discusses a possible attack vector for systems using Laplacian noise based on analyzing the least significant bit (LSB) of a calculation. Mironov found that the LSB can serve as an artifact produced by rounding and reveal the original, noiseless, data. This is important as the Laplacian mechanism is one of the standard ideas when creating a differentially private algorithm. In addition, this work has been referenced by other authors in the field as it serves as one of the first to study the effects of finite precision computing on differential privacy.

One such paper that references Mironov's work is Ilvento's (2020) research on implementing a differential privacy algorithm using base-2 arithmetic. The original LSB vulnerability originated from the inability to exactly calculate and represent the natural logarithm on standard computer hardware. However, Ilvento shows that with some modifications to change all calculations to base-2 we can preserve the theoretical guarantee of differential privacy in real-world implementations as modern computers can perfectly calculate base-2 arithmetic. This work is an important step forward in the field of differential privacy, but it only addresses one of the mechanisms used to implement differential privacy. Thus, it can be a reference for future work on whether base-2 arithmetic can be a patch to other precision-based vulnerabilities.

My proposed research in conjunction with Professor Tianhao Wang will attempt to apply this principle of the limited precision of computers to other cryptographic mechanisms used in

ensuring the privacy of data. Many of these mechanisms rely on sampling probabilistic distributions similar to the Laplacian distribution mentioned before, and this is where the research will focus on. By looking at the standard open-source implementations, we hope to produce a proof-of-concept attack through a precision-based vulnerability. This in turn can be used to break the privacy "guarantees" and show that careful consideration must be taken when implementing cryptographic principles.

## STS Topic

While my technical report focuses on the mathematical definitions and theory behind privacy, my STS research will focus on the question of how data collection and data use policies are communicated to users, and how they can be made more transparent. This question has become convoluted with the prolific technologies and services that have been embedded in our daily lives. For example, Pollach (2011) examined the privacy policies and other publicly available information for several IT companies and argues that companies do not yet do enough to protect the privacy of all parties and that there are no clear standards for companies to follow. This can be a problem as without standardization it can be overwhelming for users to understand how each and every service they use manages their own data. While this work was done over 10 years ago, we can investigate to see what progress has or has not been made since.

Privacy policy statements are critical in understanding what data a company collects and how it might be used. However, they can often be full of legalese that are difficult for the average user to read and thus ignored. Hintze (2017) argues that privacy policies should be drafted to be fully comprehensive, as this ultimately promotes more transparency, even at the cost of being left unread for the end user. Instead, other parties like reporters will be more likely to read longer privacy statements and can give this information to the general public as

necessary. While I agree that companies should provide detailed statements about all the various sources of data they collect, it is also important that all users should be able to understand risks and protections to themselves directly rather than through a third party.

Towards this effect, we can turn to a similar situation involving information disclosure agreements for research projects. These agreements are intended to inform and protect participants, but they can be poorly designed which results in participants accepting without fully grasping their rights. Rossi and Lenzini (2020) developed several information design patterns for researchers so they can more effectively and transparently disclose how information may be used to empower the participants. One example is the use of visuals to help retain the attention of the reader, make abstract concepts more tangible, and aid those with lower literacy. Attention retention and the de-abstraction of legal concepts are some of the same factors that previous sources have discussed. Therefore, it may be useful to take the design patterns proposed by the authors of this paper and see if they can be applied to improve transparency around data collection by companies and technology.

However, the policies that a company state is only half of the story. Since the passing of the General Data Protection Regulation (GDPR), websites have begun asking visitors to accept certain tracking cookies that the site uses. Researchers surprisingly found that sometimes when users chose to reject all cookies, tracking activity actually increased (Papadogiannakis et al., 2021). I agree with the authors that these behaviors by the websites are highly problematic as the cookie consent banner gives users a false sense of power and obfuscates the real actions happening behind the scenes.

Yet another method proposed to improve transparency regarding data use is the development of software tools aimed at providing information to users. One such tool was

developed by the researchers that took disclosures mandated by the EU General Data Protection Regulation (GDPR) and showed them to users through a graphical user interface (Fischer-Hübner et al., 2016). This allowed consumers to see which companies have what information about them. They found that users appreciated the transparency granted by the tool, but were also concerned regarding the security of the tool itself. This is an interesting approach to promoting transparency which places the onus on third parties rather than the company collecting the data itself. The tools can be connected to some of the design patterns discussed by Rossi and Lenzini such as providing visuals to aid explanation. However, there is also a paradox as any tool similar to the ones created by the researchers will need to collect the user's data to display and thus would need to make clear to what extent the data is used.

For my proposed STS research, I hope to answer this open question of how users can be made aware of what data they are agreeing to give as to protect their individual privacy. Toward this goal, I will employ the ethical framework of deontology to analyze what companies see as their duty when it comes to protecting the privacy of their users and where their intentions may differ from that of the general public. Moreover, this framework can help guide what should be done when designing privacy disclosures assuming transparency as the underlying duty. My methods will include investigating the data collection policy of companies as well as their actual actions to see if they align. In turn, this can be used to pose improvements to improve transparency to end users.

## Conclusion

My proposed capstone research project intends to study one facet of the theoretical foundation for modern cryptographic methods. The hope of this research is to identify the limitations of theory when applied to the real world and how to be conscious of this as designers

and developers. Meanwhile, the STS research asks the broader question of how data collection and use can be communicated to the end users so that they can make informed decisions. Together, the two pieces of research may work to aid individuals on how to navigate privacy concerns within our technology-filled world. Users can be more confident in knowing what data is being collected and that they are used by secured algorithms designed to protect their privacy.

## References


Erlingsson, Ú., Pihur, V., & Korolova, A. (2014). RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 1054–1067. https://doi.org/10.1145/2660267.2660348

Fischer-Hübner, S., Angulo, J., Karegar, F., & Pulls, T. (2016). Transparency, Privacy and Trust – Technology for Tracking and Controlling My Data Disclosures: Does This Work? In S. M. Habib, J. Vassileva, S. Mauw, & M. Mühlhäuser (Eds.), *Trust Management X* (pp. 3–14). Springer International Publishing. https://doi.org/10.1007/978-3-319-41354-9_1

Hintze, M. (2017). In Defense of the Long Privacy Statement. *Maryland Law Review*, *76*(4), 1044–1084.

Ilvento, C. (2020). Implementing the Exponential Mechanism with Base-2 Differential Privacy. *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 717–742. https://doi.org/10.1145/3372297.3417269

Kudina, O., & Verbeek, P.-P. (2019). Ethics from within: Google Glass, the Collingridge dilemma, and the mediated value of privacy. *Science, Technology, & Human Values*, *44*(2), 291–314. https://doi.org/10.1177/0162243918793711

Mironov, I. (2012). On significance of the least significant bits for differential privacy. *Proceedings of the 2012 ACM Conference on Computer and Communications Security - CCS '12*, 650. https://doi.org/10.1145/2382196.2382264

Papadogiannakis, E., Papadopoulos, P., Kourtellis, N., & Markatos, E. P. (2021). User Tracking in the Post-cookie Era: How Websites Bypass GDPR Consent to Track Users. *Proceedings of the Web Conference 2021*, 2130–2141. https://doi.org/10.1145/3442381.3450056

Pollach, I. (2011). Online privacy as a corporate social responsibility: An empirical study. *Business Ethics: A European Review*, *20*(1), 88–102. https://doi.org/10.1111/j.1467-8608.2010.01611.x

Rossi, A., & Lenzini, G. (2020). Transparency by design in data-informed research: A collection of information design patterns. *Computer Law & Security Review*, *37*, 105402. https://doi.org/10.1016/j.clsr.2020.105402

Vimalkumar, M., Sharma, S. K., Singh, J. B., & Dwivedi, Y. K. (2021). 'Okay google, what about my privacy?': User's privacy perceptions and acceptance of voice based digital assistants. *Computers in Human Behavior*, *120*, 106763. https://doi.org/10.1016/j.chb.2021.106763

Yang, Y., West, J., Thiruvathukal, G. K., Klingensmith, N., & Fawaz, K. (2022). Are You Really Muted?: A Privacy Analysis of Mute Buttons in Video Conferencing Apps. *Proceedings*

*on Privacy Enhancing Technologies*, *2022*(3), 373–393. https://doi.org/10.56553/popets-2022-0077