CS 127 Repository Entry
Embedded EthiCS @ Harvard Teaching Lab

Overview

Course:	CS 127/227: Cryptography
Course Level:	Advanced undergraduate
Course Description:	"In this fast-paced course, I plan to start from the very basic notions of cryptography and by the end of the term reach some of the exciting advances that happened in the last few years such as the construction of fully homomorphic encryption, a notion that Brian Hayes called 'one of the most amazing magic tricks in all of computer science,' and indistinguishability obfuscators which are even more amazing. To achieve this, our focus will be on ideas rather than implementations and so we will present cryptographic notions in their pedagogically simplest form– the one that best illustrates the underlying concepts– rather than the one that is most efficient, widely deployed, or conforms to Internet standards. We will discuss some examples of practical systems and attacks, but only when these serve to illustrate a conceptual point."
Module Topic:	Privacy and the Ethics of Client-Side Scanning
Module Author:	Eliza Wells
Semesters Taught:	Fall 2021-2022
Tags:	privacy [phil], security [phil], power [phil], cryptography [CS], surveillance [CS], client-side scanning [CS]
Module Overview:	This module focuses on the relationship between privacy, security, and surveillance. It presents and challenges a standard model of that relationship—that privacy only benefits individuals and so must be sacrificed to security, which benefits communities—by thinking about each concept in terms of different agents' powers. Students are then asked to consider the ways in which particular design decisions impact who has power to do what as a way of determining when sacrifices of privacy or security are justified. This module focuses on design decisions involved in the nascent technology of client-side scanning.
Connection to Course Material:	This module uses as a case study Apple's August 2021 proposal to implement client-side scanning (CSS) to detect and report child sexual abuse material (CSAM). CSS can subvert end-to-end encryption (a standard way of ensuring that content remains private) by scanning content directly on users' devices. Students in this course have learned about different encryption strategies as well as the perceptual hashing technology that Apple proposed to use to target CSAM, so they are well prepared to discuss its technical implementation. As students in a cryptography course, they are also constantly engaged with the concepts of privacy, security, and surveillance that this module explores. The Apple CSS case study was chosen for two reasons: a) it was very current (announced only three months before the module ran) and b) it is less obvious which tradeoffs are justified than in other cases that often feature in the ethics of cryptography (e.g. students tend to already have views about whether the FBI should have access to potential terrorists' phones). However, this module could be run in the same way with a different case study.

Goals

Module Goals:	1. Understand the philosophical conception of privacy as power to control access to your personal information. 2. Consider the relationship between privacy, security, and surveillance in terms of power. 3. Think about how particular design decisions impact who has power to do what. 4. Apply these tools to a case study.	
Key Philosophical Questions:	1. Does client-side scanning pose an unacceptable threat to privacy? 2. What is privacy? What is privacy good for? 3. What is security? What is security good for? 4. What is the relationship between privacy, security, and surveillance? 5. How do we decide when threats to privacy are justified?	The goal of the module is to prepare students to think more carefully about the first question, but not to answer it for them. The key philosophical questions are focused on providing definitions that serve as tools for thinking through what threats to privacy might mean, rather than on providing a theory that determines when privacy ought to be protected.

Materials

Key Philosophical Concepts:	● Power as the ability to do something ● Privacy as the power to control access to your personal information ● Security as it aims to protect powers ● Technology as political in that it influences power	The terms “privacy” and “security” are frequently used, but what we mean by them is not often clear. This module attempts to provide both clarity and conceptual tools by introducing specific philosophical definitions of both of these concepts that students can then work with. Understanding the relationship between privacy, security, and surveillance in terms of power equips students to a) identify threats to privacy by identifying when some agents’ powers to control access to their personal information is limited and b) consider whether sacrifices of privacy are justified by thinking about which agents ought to have which powers to do what. It also helps students to complicate the narrative that privacy and security are inherently in conflict and see how they can both threaten and enhance each other.
Assigned Readings:	● Abelson et al.,“Bugs in Our Pockets: The Risks of Client-Side Scanning,” October 2021	This paper explains the technical details of client-side scanning and argues in depth that it has technical and ethical failings, including risks to users’ privacy and security

		through attacks and government abuse. It also discusses Apple's proposal to implement CSS and argues that it is impermissible. The paper was published by leading computer scientists less than a month before the module ran, so it was a timely discussion. The authors made a number of strong claims about value that we unpacked in the module.

Implementation

Class Agenda:	1. Introduce the big claim that cryptography is political because it influences relations of power. 2. Explore a standard model of the relationship between privacy, security, and surveillance: privacy is good for individuals, security is good for the community, and they are inherently in conflict. Surveillance is justified even though it threatens privacy because it protects security. 3. Challenge the standard model. We can understand privacy as a power and security as something that protects powers. Both can threaten and enhance each other. Whether or not surveillance is justified depends on who ought to have powers to do what. 4. Exercise: apply what we've learned to Apple's CSS proposal.	The big claim and the standard model are drawn from Phillip Rogaway, "The Moral Character of Cryptographic Work," 2015.
Sample Class Activity:	Students were split into groups of 2-3 at the beginning of class and discussed questions together throughout. During the case study, students were asked to discuss the following questions with their small groups in response to different design decisions made in Apple's CSS proposal: ● Whose power is increased by this proposal? Power to do what? ● Whose power is decreased in this proposal? Power to do what? We then discussed groups' answers as a class.	There were two goals to this exercise: first, to build up to the key question of whether or not CSS poses an unacceptable threat to privacy by thinking about how different elements of CSS design impact power; second, to highlight the political nature of cryptography and computer science generally by demonstrating how small design choices can have impacts on agents' power.
Module Assignment:	On a later problem set, students were asked to write a 250-350 word response to the following question: Do you think client-side scanning poses an unacceptable threat to privacy? Why or why not?	This assignment gave students the chance to use the tools they were given in class in order to grapple with a big, difficult question. It also gave them the chance to share their personal opinions without peer pressure. Student responses were thoughtful and nuanced, showing that they were engaged with the module.

Lessons Learned:	1. Students seemed to easily grasp privacy and security in terms of power and were highly engaged in their small groups. Asking about particular design decisions was interesting for students. 2. It was not clear that the “standard model” of the relationship between privacy, security, and surveillance presented in this module was in fact standard for students. Future versions of this module could try to motivate the conception of privacy as power in a different way. 3. Student responses to the assignment indicated that more discussion about who ought to have the power to do what, and so what limitations on privacy are acceptable or unacceptable, would have been helpful.	The course used software that allows students to comment together on reading materials. Students were also highly engaged with the paper beforehand.