

# CMSC 33250: Graduate Computer Security

Grant Ho, Fall 2023

(CMSC 23200 is **not** offered this quarter)

(Some slides borrowed from Stefan Savage and Frank Li)

# Today's Class

- Course info & background
- Whirlwind overview: security areas & research approaches
- Course structure: class format and coursework

# This is **NOT** CMSC 23200 (Intro to Security)

- Due to registrar policy and historical decisions, CMSC 33250 is always co-listed with CMSC 23200.
- But this course is **not** an Introductory course for computer security (CMSC 23200 will only be offered in the Winter)
- This is a graduate course focused on security research  
The class assumes you have background in security (e.g., Intro to Computer Security) and/or experience in CS research

# CMSC 33250: Course Goals

1. Gain a broad understanding of the problems that researchers are studying across many areas of computer security
2. Understand the different mental models & methods used in security research
3. Identify & address security issues in your own work
4. Learn to conduct & communicate good security research

# Course Info

Instructor: Grant Ho

- **Email:** [grantho@uchicago.edu](mailto:grantho@uchicago.edu)
- **Office hours:** Wed, 3 - 4pm in JCL 255
- **Course website:** Details about schedule, assignments, project (Work-in-progress):  
<https://classes.cs.uchicago.edu/archive/2023/fall/33250-1/>
- Discussion forum (Ed Discussion) & Announcements on Canvas:  
<https://canvas.uchicago.edu/courses/52391>

# Background about me

- Assistant professor in CS
  - Prev: PhD & Postdoc at UC Berkeley & UCSD
- Research focus:
  - Understanding & improving how we secure organizations / enterprises
  - Generally: how can we use data to advance computer security?
- My work is often informed by collaborations with real-world security teams (e.g., Google, Facebook, Dropbox, Barracuda Networks, LBNL)



# Today's Class

- Course info & background
- Whirlwind overview: security areas & research approaches

# Computer Security

- Most of computer science focuses on building new *functionality*:
  - UX / UI
  - Software & Hardware Architecture
  - Algorithms
  - Operating Systems / Networking / Databases
  - Compilers / Programming Languages
  - AI / ML
- Computer security is *not* about developing functionality
- It is about how technology actually behaves in the *presence of an adversary*



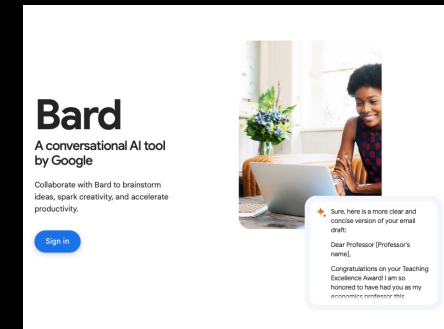
# Classical Security Principles

On a technical level, security research studies whether an attacker can violate one (or more) of the following properties:

- **Confidentiality**
  - Can an attacker learn information they shouldn't from a system or dataset?
- **Integrity**
  - Will a system work exactly as intended, or can an attacker cause the system to behave in unintentional ways?
- **Authenticity**
  - How do I know the true identity of who/what I am interacting with online?
- **Availability**
  - Can an attacker prevent users from accessing a system and its functionality?

# Security Research Landscape

## Application




- How can an attacker exploit AI/ML applications to operate incorrectly?
- Can an attacker learn private or sensitive information from an ML model?

# Security Research Landscape


## Application

- Can an attacker view videos or calls that they were not authorized to see?
- How can an attacker compromise or hijack my account?
  - Could an attacker use this application to stalk or harass a victim?




zoom

# Security Research Landscape




# Security Research Landscape




How can an attacker prevent users from connecting to our services?




# Security Research Landscape




How does a user know if they're connecting to an attacker's machine or the real destination?




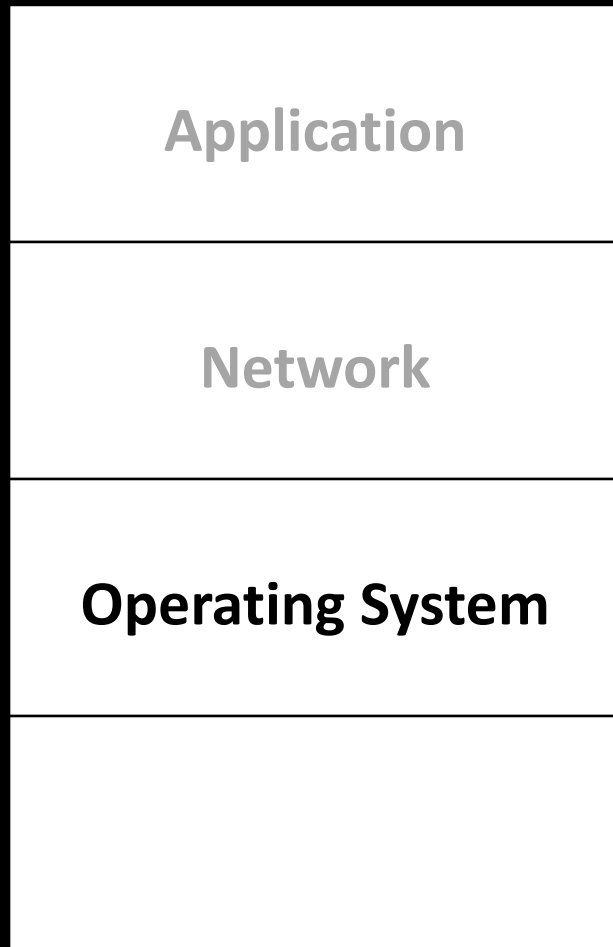
# Security Research Landscape



How can an attacker monitor who a user is communicating with?



# Security Research Landscape


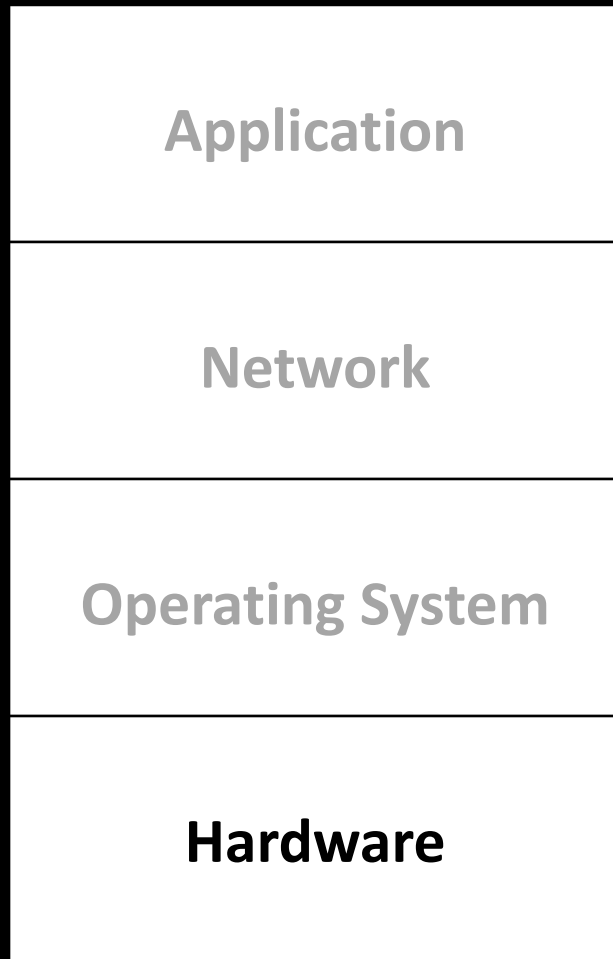


How can an attacker exploit bugs in a program to make it misbehave and/or read secret data?

How can a system ensure that users only access files and data they have proper permissions to view?



# Security Research Landscape



Courtesy  
Oswald

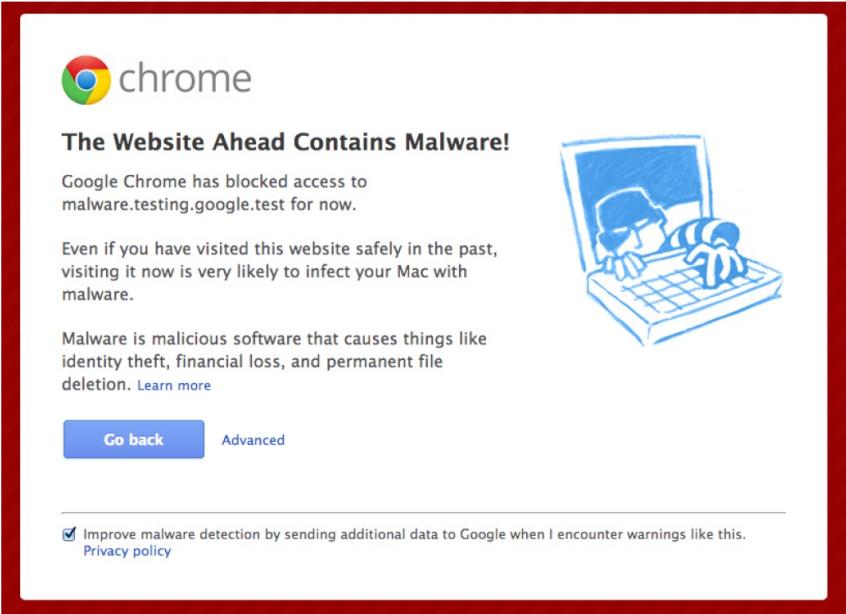
How can an attacker exploit hardware properties (e.g., power consumption) to learn sensitive data about programs running on the machine?

# Security Research Landscape

How do we communicate security information to users effectively?

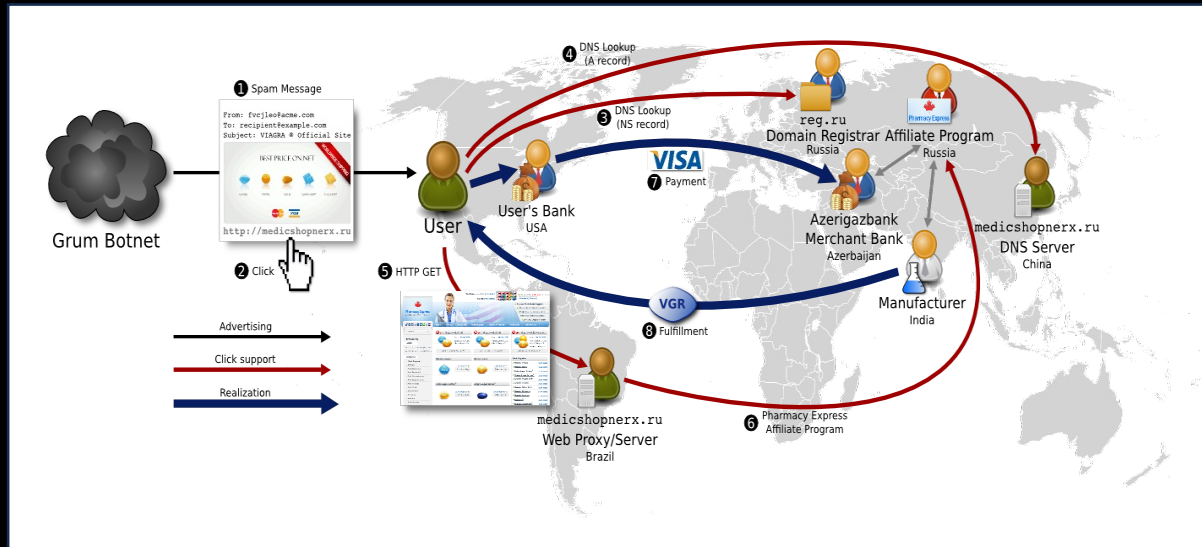
How do we design systems so that the default actions align with users' security & privacy expectations?

**Human**



The screenshot shows a Chrome browser warning page with a red border. At the top left is the Chrome logo and the word "chrome". The main heading is "The Website Ahead Contains Malware!". Below this, it states "Google Chrome has blocked access to malware.testing.google.test for now." To the right is an illustration of a laptop with a person's head and hands on it, representing malware. The text continues: "Even if you have visited this website safely in the past, visiting it now is very likely to infect your Mac with malware." Below that, it explains: "Malware is malicious software that causes things like identity theft, financial loss, and permanent file deletion. [Learn more](#)". At the bottom, there are two buttons: "Go back" and "Advanced". At the very bottom, there is a checkbox that is checked, with the text "Improve malware detection by sending additional data to Google when I encounter warnings like this." and a link to "Privacy policy".

# Security Research Landscape



Ecosystems

How do adversaries monetize their attacks and abuse?

What is the infrastructure that attackers use to run operations at Internet scale?

# Approaches to Security Research

## SoK: Science, Security, and the Elusive Goal of Security as a Scientific Pursuit

Cormac Herley

Microsoft Research, Redmond, WA, USA  
cormac@microsoft.com

P.C. van Oorschot

Carleton University, Ottawa, ON, Canada  
paulv@scs.carleton.ca

*Abstract*—The past ten years has seen increasing calls to make security research more “scientific”. On the surface, most agree that this is desirable, given universal recognition of “science” as a positive force. However, we find that there is little clarity on what “scientific” means in the context of computer security research.

research) in the light of consensus views of science and scientific methods. We find that aspects from the philosophy of science on which most other communities have reached consensus appear surprisingly little used in security, including

# Security through Formalisms / Proofs

## Deductive

Start

Make a set of assumptions  
about computer systems,  
users, the world



End

Use formal reasoning  
(math, logic, etc.) to derive  
a proof / guarantee about  
the security of a system

# Security through Formalisms / Proofs

- Deductive: given a set of assumptions, derive a logically certain conclusion about the security of a system
- Examples: cryptography, verification, trustworthy systems
  - Strength: Protocol (e.g., Diffie-Hellman key exchange) can be proven secure given specific assumptions
  - Weakness: Abstract design  $\neq$  real computers (e.g., side channels)

644

IEEE TRANSACTIONS ON INFORMATION THEORY, VOL. IT-22, NO. 6, NOVEMBER 1976

## New Directions in Cryptography

*Invited Paper*

WHITFIELD DIFFIE AND MARTIN E. HELLMAN, MEMBER, IEEE

## Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems

Paul C. Kocher

Cryptography Research, Inc.  
607 Market Street, 5th Floor, San Francisco, CA 94105, USA.  
E-mail: paul@cryptography.com

# Security through Empiricism

## Inductive

Start

Make observations & measurements about security in the world



End

Analyze the data to draw generalizable conclusions about the real-world state of security

# Security through Empiricism

- Inductive: make observations and draw general conclusions from the data
- Examples: risk management, experimental life sciences & medicine
  - Strengths: Captures how security actually manifests in practice
  - Weakness: Does not provide future guarantees; things can vary across time and context

## **The Security of Modern Password Expiration: An Algorithmic Framework and Empirical Analysis**

Yinqian Zhang  
University of North Carolina at  
Chapel Hill  
Chapel Hill, NC  
yinqian@cs.unc.edu

Fabian Monrose  
University of North Carolina at  
Chapel Hill  
Chapel Hill, NC  
fabian@cs.unc.edu

Michael K. Reiter  
University of North Carolina at  
Chapel Hill  
Chapel Hill, NC  
reiter@cs.unc.edu


## **Consequences of Connectivity: Characterizing Account Hijacking on Twitter**

Kurt Thomas<sup>†</sup> Frank Li<sup>†</sup> Chris Grier<sup>†\*</sup> Vern Paxson<sup>†\*</sup>  
<sup>†</sup>University of California, Berkeley    <sup>\*</sup>International Computer Science Institute  
{kthomas, frankli, grier, vern}@cs.berkeley.edu



# Complementary Approaches

Produce logically certain conclusions / guarantees about security (assuming the assumptions hold)



Inform and validate formal model assumptions / conclusions in reality

# Today's Class

- Course info & background
- Whirlwind overview: security areas & research approaches
- Course structure: class format and coursework

# Class Format

- This will be a discussion driven course
- There will be 2 papers per class on a topic in security & privacy (~4 papers / week)
- I will give a short overview of the area, and then we will discuss each of the assigned papers.
- Everyone is expected to have thoughtfully read the papers and engage in discussion.

# Coursework Overview

- **Reading Responses (20%)**: short write-up for each assigned paper
- **In-class Discussion (10%)**: actively contribute thoughts and ask questions
- **Discussion Lead (10%)**: each student will lead the discussion of one required paper
- **Course Project (60%)**: conduct, present, and write-up a paper on a research project related to security & privacy

No midterm or final exam

# Discussion: Reading Research Papers

What kinds of things do you pay attention to / look for when reading a research paper?

- Cursory level reading (e.g., is this paper relevant to me?)
  - 10-15min hacks to assess the paper
- Deep read (expectations for this class's assigned papers)
  - Aim to really understand and learn from a paper

# Discussion: Reading Research Papers

1. What is the **problem** and **threat model**?
2. What is the **motivation**?
  - **Societal** (impact on world/people) & **Technical** (why hasn't it been solved?)
3. What are the **contributions**?
  - **Intellectual** (methods/findings) & **Artifacts** (new dataset/software)
4. Can I intuitively **explain the solution** (methods or takeaways)?
5. What was the **evaluation/analysis** and are there any problems?
6. What are the **limitations and future directions**?
7. What **confusions or questions** do I have?

# Reading Responses (20%)

Submit a short, but thoughtful reflection for *each* of the assigned papers **prior to class (due by 11:00am** on Gradescope)

1. What are the paper's main contributions?
2. What parts of the paper are questionable? (e.g., methodology, omissions, relevance, presentation, ethics.)
3. Propose one question for discussion that you will ask or answer/discuss
4. What parts of the paper do you find unclear? (Optional)

I will drop your two lowest scoring assignment sets

# In-Class Discussion & Discussion Lead (20%)

## **In-class Discussion (10%): Not just attendance**

- Actively contribute your thoughts, comments, questions

## **Discussion Lead (10%):** Prepare a 15-20min presentation on one paper (Sign-ups will happen on 10/04)

- Follow my example from the next few classes
- Doesn't need to be pretty, but good communication & structure
- Key parts: Problem & Motivation, Background & Prior work, Methodology, Evaluation/Results, Questions for Discussion



# Course Project (60%)

- Key goal of this course: do a research project in security / privacy
- Project should be done in groups of 2-3 students
  - Post on Ed Discussion (Canvas) for the class to find project partners
- Deadlines & Milestones:
  - **Project Proposal (Oct 13)**: Send me a 1pg proposal for each group
  - **Preliminary Report (Oct 30)**: Preliminary report & Related work outline
  - **Presentation (Nov 29 & Dec 4)**: One presentation per group (15-20min)
  - **Final Report (Dec 8)**: 6-10pg conference-style paper

# Considerations for your Project

- Pick a good **problem**
  - Why is this problem interesting or will become interesting?
    - New technology, new approach, new question, new usage, etc?
  - See the course website for a list of potential project ideas
  - Academic conferences: USENIX Security, ACM CCS, IEEE S&P, NDSS, PETS, SOUPS, etc.
- Pick approaches to problems that are **achievable**
  - What resources would you need to investigate the problem?
- Think about how to **evaluate** your work
- **Start early**: good research takes time and hits many roadblocks

# Next Steps

- If you were planning to take CMSC 23200, unfortunately that is only offered this winter (not this course)
- Make sure you can access Canvas & Gradescope for the course
- Post on Ed Discussion (Canvas) if you are looking for project groupmates
- Read assigned papers for Cybercrime & Ecosystems class
  - Responses for these papers are due at 11:00am on 10/02