COSC 366 Lecture 1

Course Logistics

Fall 2024

Dr. Suya

Announcements, materials, assignments, exams will be posted on Canvas

Instructors

- Professor: Dr. Suya
  - Office: 344 Min Kao
  - Office hours: Tue 10:30 – 11:30 AM
  - Email: suya@utk.edu
- TA: Brandon Marth
  - Office hours: Mon 11:30-12:30, Wed 12:30-1:30, Fri 1:30-2:30
  - Location: Min Kao 235
  - Email: bmarth@vols.utk.edu
  - Comments are welcome to improve the course!

A bit about me…

- Full name: Fnu Suya
  - Fnu: **F**irst **N**ame **U**nknown
  - Suya, Dr. Suya, Prof. Suya are all fine.
- PhD (2023) in Computer Science from the University of Virginia (~5h drive from UT)
- Postdoc (2023--2024) in the Maryland Cybersecurity Center (MC2) at the University of Maryland, College Park

A bit about me…

- Assistant professor in EECS
  - Actively looking for undergraduate students for research
    - I have some ideas for you to try, also encourage you to come up your ideas.
    - Send me an email with your CV.
  - Primary research areas:
    - AI/ML for security applications (e.g., apply ML for malware detection)
    - Trustworthy ML (e.g., how ChatGPT produces harmful content when misused)

What is computer security?

Normally, we are concerned with functionality and correctness.

functionality

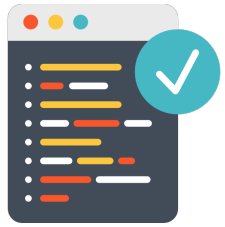

Security is also a form of correctness: correctness even against some “bad” guys.

Attacker

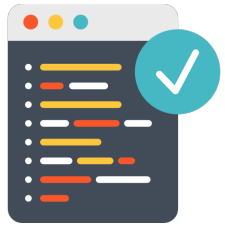

Difference between you and attackers

Normal users: hate bugs/flaws, want to use without interruptions in normal settings

Attackers: actively find the bugs/flaws out and to exploit them; if there is a defense, also circumvent the defenses

What does it mean to be secure?

- “Security” itself is not a measurement: we only measure the degree of insecurity.

- Goal: raise the bar for the attacker
  - Too difficult
  - Too expensive
  - Impossible to attack

- Ultimately, we want to mitigate undesired behavior while maximally preserving the normal functionally

COSC 366: Intro to Cybersecurity

- The Goals:
  - Teach you about security threats and solutions when computers are plugged into a network.
  - Teach you to make and break secure systems
- Outcomes:
  - Learn what we mean when we say secure
  - Learn what the challenges are to security in different settings

COSC 366: Intro to Cybersecurity

- Outcomes (cont.):
  - Learn to think like an adversary (**Security Mindset**)
  - Learn to identify problems in existing systems
  - Learn to avoid flaws in your design
- Get exposed to current security research
  - IEEE S&P, ACM CCS, USENIX Security, NDSS, etc.
  - Machine learning security extends further to AI conferences: ICML, ICLR, NeurIPS, CVPR, ICCV, etc.

Course Logistics Overview

- Course Format
- Textbooks
- Grading
- Assignments/Projects/Exams
- Policies
- Rough Schedule

- All information is on Canvas
  - Please regularly check the Canvas, enable email notifications.

Course Disclaimer

- **My promise to you:** If anything on the schedule changes, especially assignments, due to whatever circumstances, you will NOT be harmed if it was outside your control (i.e., things will work in your favor w.r.t. grades)
  - Examples: change to exam schedule or assignments, etc.

Course Format

- Classes: in-person
- Office hours:
  - Dr. Suya: Tue 10:30-11:30AM.
  - TA: Brandon Marth
- Four assignments will be turned in on Canvas
- Exams:
  - Format TBD (do not worry!)
- If you would like to meet with me one-on-one about something outside the scope of this class, like research, career paths in security, industry vs. academia, etc., just email me!

Textbooks – Nothing required but…

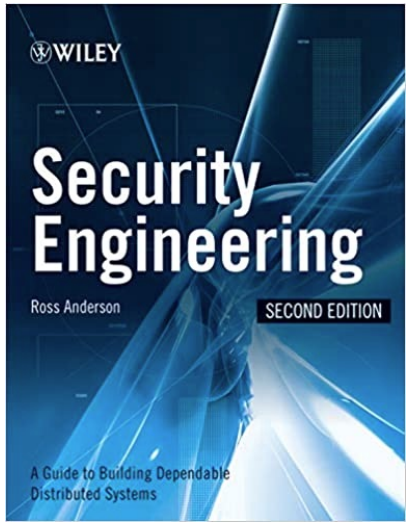

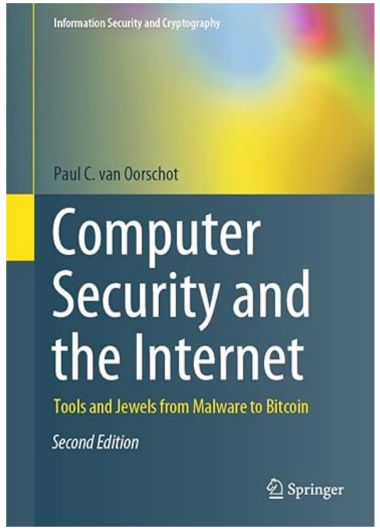

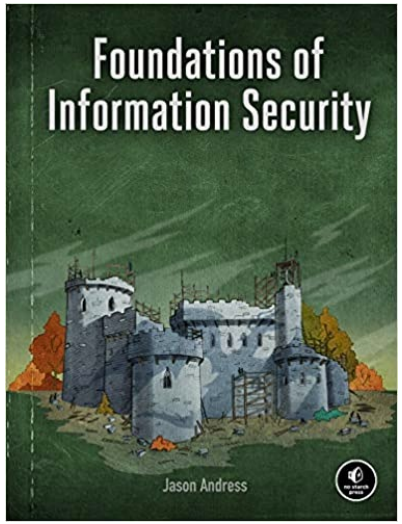

Good baselines

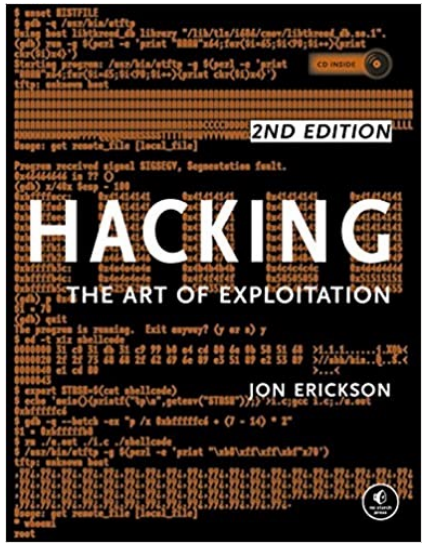

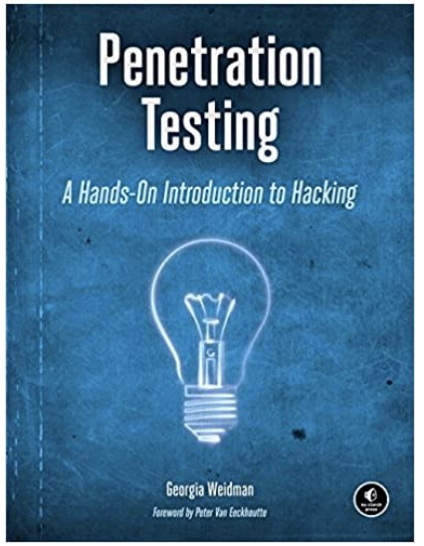

Offensive security techniques

How are you graded?

- Homework (4 in total): 30%
- Projects (4 in total): 40%
- Midterm Exam: 10%
- Final Exam: 20%
- Bonus points: TBD (likely to be associated to class attendance)

Homework

- Four in total, 30% of your grade
- Due dates on Canvas at 23:59:59 of due date!
- Conceptual questions: questions ask you to think conceptually about security
    - Roughly a paragraph or two long answers
    - Ideas should be clearly expressed, but need not be formal

Projects

- Four in total, 40% of your grade
- Due dates on Canvas at 23:59:59 of due date!
- Programing tasks that applies the knowledge you learned in class
  - More detailed instructions will be announced with the project assignments.
  - Highly discouraged to use ChatGPT, UTK Verse or similar LLM to directly complete the assignments for you (without your active thinking).

Exams: TBD

- Midterm and final
  - Likely online or offline with your laptop.
- We will have a prep day for both exams!
- Format and content detail to come at least a week in advance.

Grade Thresholds (ceiling)

>= 93.00	A
90.00 - 92.99	A-
87.00 - 89.99	B+
83.00 - 86.99	B
80.00 - 82.99	B-
77.00 - 79.99	C+
73.00 - 76.99	C
70.00 - 72.99	C-
<= 69.9	F

Attendance

- Attendance is HIGHLY encouraged, but if you must miss, you do not need to tell me
  - Hint: people who participate in class generally get higher grades (and likely to have some bonus points)
- No recordings are uploaded, but students can audio record the class on their own, with approval from the instructor (need justifications).

Ethics

- Do not cheat
  - I have to spend time documenting it…
  - No Mercy…
- I will teach you how to attack systems
  - Use this power for good
  - The goal is to foster discovery, experimentation, and exploration

Usage of Generative AI (GenAI)

- Highly encouraged to use GenAI to digest the course materials and seeks answers for course-related questions you have after the class.
  - Real-time feedback with good accuracy.
- Highly discouraged to use these LLMs to complete the assignments, projects, exams for you.
  - For example: prompting "*Hey, this is the assignment, please complete it for me*".
  - As an adult, you are responsible for your own career. You cannot always rely on LLMs to complete the job interviews for you.

Rough Schedule

- Basic Security Goals
- Software Security
- Operating System Security
- Web Security
- Human Factors
- Network Security
- Cryptography
- Machine Learning Security (my research area)

Basic Security Goals

- Confidentiality
  - Only authorized parties should learn certain information.
  - E.g., the contents of a file on a disk, a value in a database, etc.
- Availability
  - Authorized “users” should be able to interact with resources when they wish to in the expected amount of time.
  - E.g., I want to access a website.
- Integrity
  - Everything is as it should be.
  - Authentication integrity - an entity should be who they claim to be

Software Security

- Your code can become vulnerable if you don’t consider security
  - Buffer overflow
- A buffer overflow occurs when a program writes more data to a buffer (a fixed-size memory block) than it can hold, causing data to overflow into adjacent memory locations.
- **Key Points**
  - Buffers are areas of memory allocated to store data temporarily.
  - Overflow happens when the data exceeds the buffer’s boundary.
  - Can lead to unpredictable behavior, crashes, or security vulnerabilities

Buffer Overflow

```
char           A[8] = "";
unsigned short B    = 1979;
```

variable name	A								B	
value	[null string]								1979	
hex value	00	00	00	00	00	00	00	00	07	BB

```
strcpy(A, "excessive");
```

variable name	A								B	
value	'e'	'x'	'c'	'e'	's'	's'	'i'	'v'	25856	
hex	65	78	63	65	73	73	69	76	65	00

Do you know why B has value of 25856 now?

Operating System Security

- Have you taken the OS class?

- Controls access to resources
  - File Permissions (e.g., in Linux)
  - Misconfigured access control can allow attackers to compromise the system (e.g., privilege escalation)

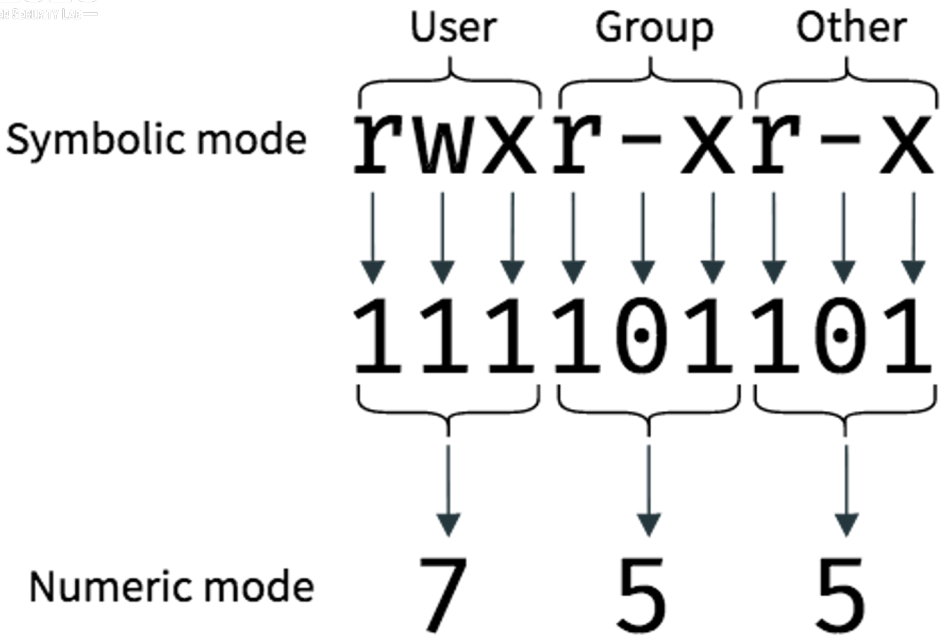

Web Security

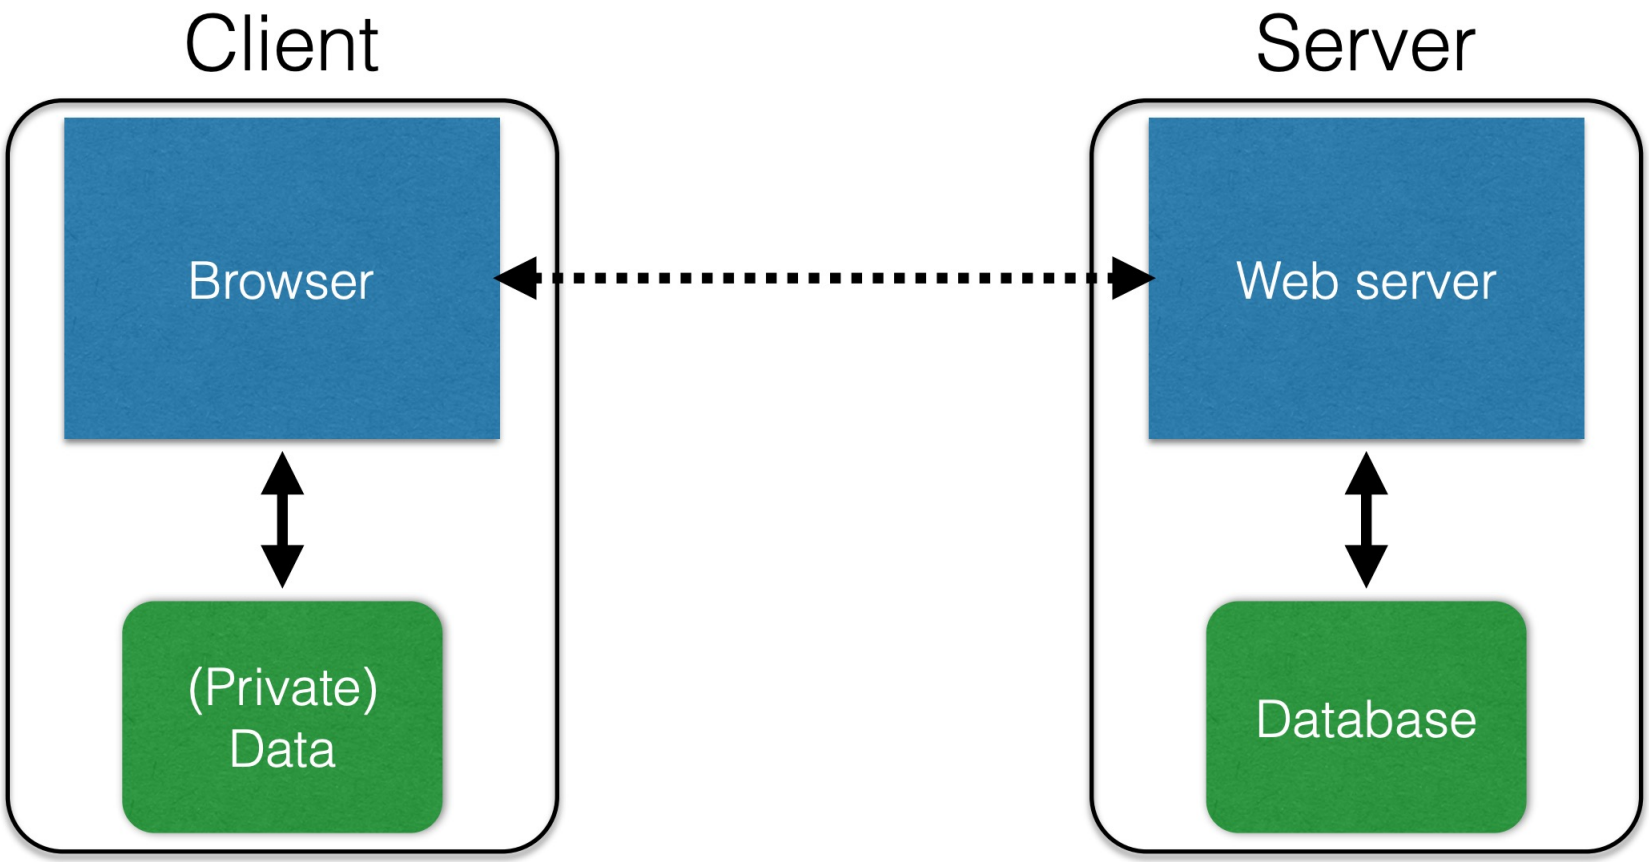

E.g., SQL injection

- ID: suya
- PW: 1234

SELECT * FROM users WHERE username = 'user_input' AND password = 'user_password';

SELECT * FROM users WHERE username = 'john' OR '1'='1' AND password = ‘';

Crypto

- Encrypt/Decrypt
- Hash
- Public Key Infrastructure

Network security

- IP Security & VPNs
- TLS (or SSL)

Machine learning security

Top-5 model accuracy: 98.19%

Top-5 human accuracy: 99.99%

Machine learning security

Adversarial examples