



CSCI 780:

IoT Security

Prof. Adwait Nadkarni

Lecture 1: Introduction


Lets break it down

- *Internet* of *Things* (IoT) *Security*




The Internet

- Every machine is connected
- Huge, *open*, system
 - No barrier to entry
 - Not just limited to dogs and users
- Built for connectivity, not security (i.e., the “end-to-end” principle)




Things are...





Things are...

Ubiquitous —
7 Billion¹
devices in use!



¹<https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/>



Things are...

**Financially
Critical —**
*\$520 Billion² by
2021*

Expensive —
*Cameras, door
locks cost \$\$\$*




²<https://www.bain.com/insights/unlocking-opportunities-in-the-internet-of-things/>



Things are...



Physical —
*Can view, listen
to, and **modify**
our physical
spaces.*



Some bad news



- We are bad at designing secure systems



Designing secure systems is hard



Fundamental Asymmetry between the attacker and the defender



Functionality is *relatively* easy to measure, but...

Airplane works



Airplane doesn't work



...*security* is almost impossible to measure

Web browser Owned

WILLIAM & MARY

HELP EXIT

User Login

Please enter your WMuserid and Password and click Login.

When you are finished, please Exit and close your browser to protect your privacy.

WMuserid must be lower case

WMuserid

Password

Login

[Click Here for Help with Login?](#)

RELEASE: 8.8

© 2017 Ellucian Company L.P. and its affiliates.

This software contains confidential and proprietary information of Ellucian or its subsidiaries.

Use of this software is limited to Ellucian licensees, and is subject to the terms and conditions of one or more written license agreements between Ellucian and such licensees.

Web browser not Owned

WILLIAM & MARY

HELP EXIT

User Login

Please enter your WMuserid and Password and click Login.

When you are finished, please Exit and close your browser to protect your privacy.

WMuserid must be lower case

WMuserid

Password

Login

[Click Here for Help with Login?](#)

RELEASE: 8.8

© 2017 Ellucian Company L.P. and its affiliates.

This software contains confidential and proprietary information of Ellucian or its subsidiaries.

Use of this software is limited to Ellucian licensees, and is subject to the terms and conditions of one or more written license agreements between Ellucian and such licensees.

Some good news

Computer security is a growth area.

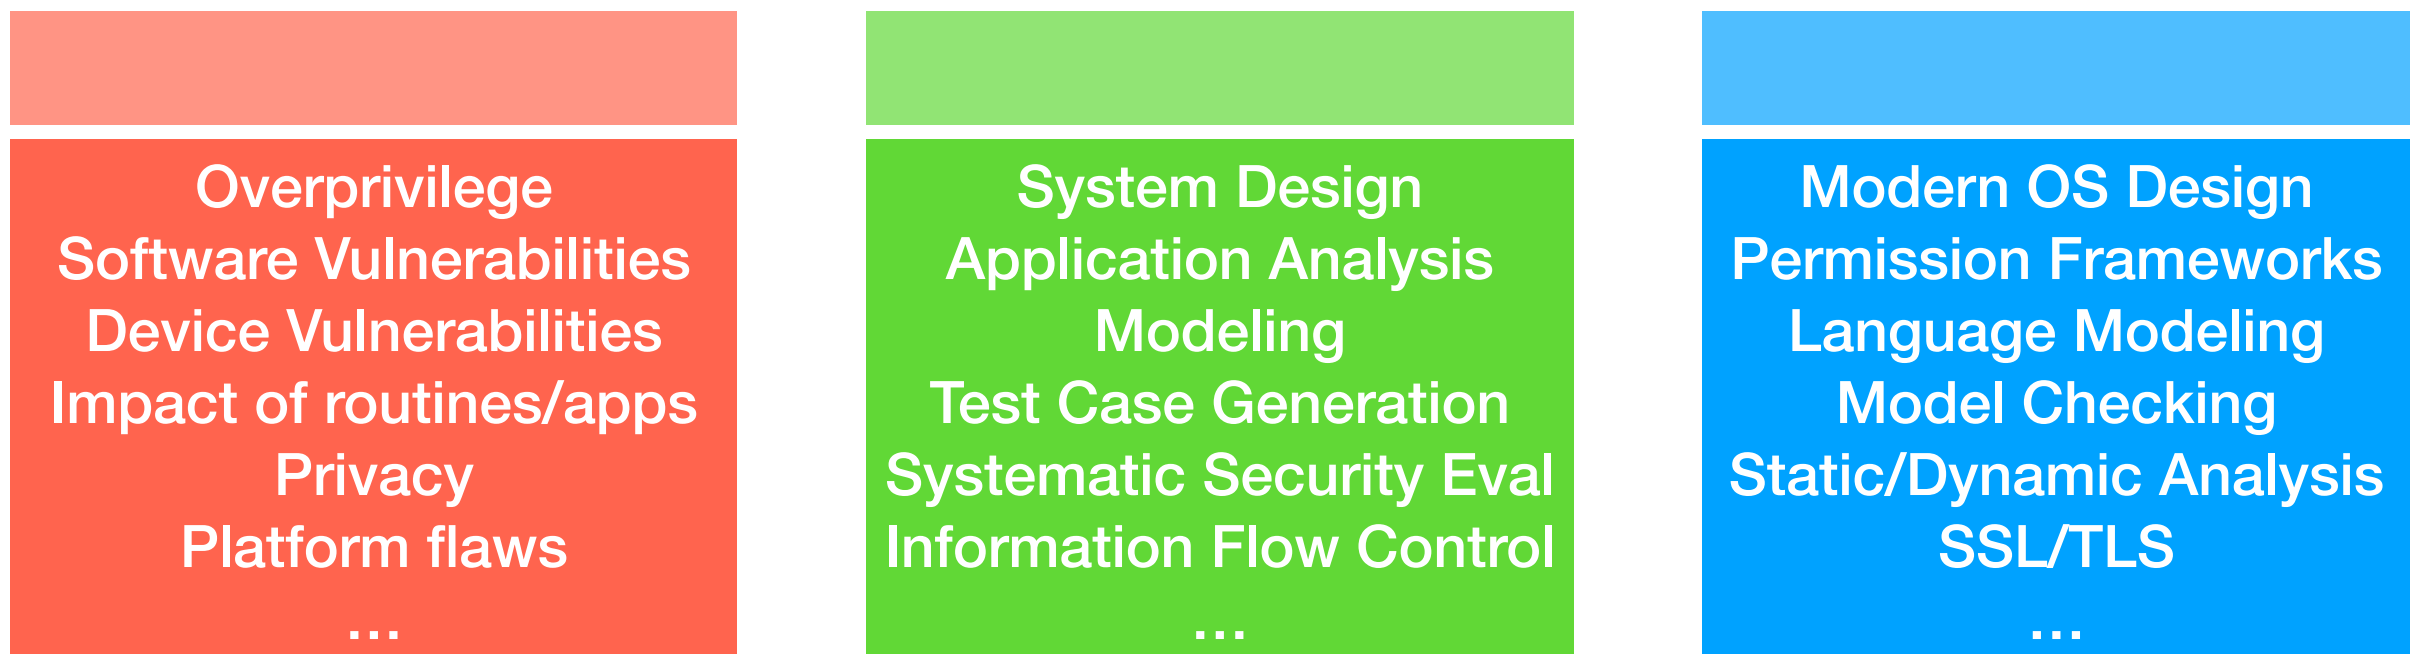


Awesome



Learning Goals

- **My Goal:** To provide you with the tools to (1) *understand*, (2) *evaluate* and (3) *perform* research in IoT Security.



- **What to expect in class:** We will learn the concepts in the context of the papers we read.
 - We will refine our understanding of the core concepts, as well as the space, as we go.
- **Key Activities to ensure learning:** Reading papers, Writing reviews, Participating in class.

Prerequisites

- No hard prerequisites
- However...
 - **Basic knowledge of the following will come handy:** OS Design Principles, Computer Networks, and Algorithm Design.
 - *Please do not hesitate to clarify even the smallest details*
 - Simple questions are often the most difficult to answer

Course Policies & Expectations

Course Website

<https://www.adwaitnadhkarni.com/teaching/csci780>

- **Discussions:** Piazza (<https://piazza.com/wm/fall2025/csci780>)
- **Submissions:** Blackboard (<https://blackboard.wm.edu>)
- **In-class Exercises, Q&A:** Discord (invite emailed to you all)



Office Hours

- **Time:** Tuesday/Thursday after class (11:00 AM - 12:30 PM), *and by appointment*
- **The Zoom link will be shared on Discord**

Textbook

- No required textbook.
- We will heavily rely on *paper readings*
- For specific concepts, we will refer to the following (online) textbooks:
 - Security Engineering, Ross Anderson (Available online: <http://www.cl.cam.ac.uk/~rja14/book.html>)
 - Operating System Security, Trent Jaeger (*Available online via* <https://libraries.wm.edu/>)

Course Components and Grading

- This is a *research* and *discussion-driven* class (80% grade)

Research Project
45%

Paper Presentation
20%

Class Participation
15%

Paper Reviews
10%

Readings “bug bounty”
10%

Research Project

- We will learn how to conduct security research, and execute a research project in IoT Security.
- **End Result:** *10-12 page conference-style research paper*
- I will provide sample ideas, but there is significant wiggle room
- Want to do something related to your research?: **talk to me ASAP**
- Grade: For quality of *research*, and *effort*

Project Milestones

- 45% of course grade (100 points for project total)
 1. **Project Proposal (5/100), due 09/11**
 2. Related Work (or preliminary literature survey, in case of an SoK) (10/100), *due 10/14*
 3. Research Plan (20/100), *due 11/04*
 4. Research Artifacts (15/100) and Final Paper (50/100), *due 12/04*
- All submissions (except artifacts) will be in LaTeX.

Paper Presentations

- Each class will be divided into **two phases**: Presentation (first 30-40 minutes) and discussion (the next 30-40 minutes)
- Students will bid on papers marked [Student Presentation] in the class schedule
- The discussion will be guided by the *questions posed at the end of the presentation*, as well in as the reviews
- Points will be awarded for clarity and content (20% of course grade)

Paper Reviews

- Hone your reviewing skills! (*also, 10% of course grade*)
- Only **1 paper review** per week, *even if two are assigned*, submitted *at the beginning of every class*.
 - The presenter does not need to submit a review
- *I will provide a conf-style review form.* Expected contents:
 1. Paper Summary, and Rating
 2. Comments on Technical Correctness
 3. Comments on Contribution and Writing
 4. Detailed comments for authors, justifying the review (1-3 above).
 5. *Three* insightful questions that will generate a discussion
- **Mock Program Committee meeting:** We will conduct a mock PC meeting towards the end of the semester!

Reviews: The good, the bad, and the ugly

- Review others' work like you want yours to be reviewed.
- Appreciate good research, don't nitpick, and be constructive.

Detailed comments for authors

I enjoyed reading this paper, and I think it contains some interesting ideas. Although I do not support acceptance at this time for the reasons specified below, I could imagine a nice publication resulting from this line of work (or perhaps even two publications if you decide to separate out some of your findings on generating sequences).

Paper Focus and Sequence Generation Evaluation

When I was a third of the way through reading this paper, I made a note to myself that I remained uncertain what the paper would actually do. Ultimately, more than half of the paper is devoted to presenting a method for generating likely sequences of events for home automation systems. This feels like far too much space devoted to explaining relatively straightforward concepts (I do not say "straightforward" as a criticism of your approach; my point is that a much more concise presentation seems feasible). The focus of the paper seems to be on sequence generation. To the extent that generating likely sequences has a relationship to security, those considerations feel like an afterthought in the paper.

Furthermore, the authors ultimately skip over details regarding the quality of the resulting sequences. Plenty of promising systems for text generation still occasionally produce humorously bad results (for example, see email auto-response suggestions). I worry about the impact of a system like this on a security-prone system here. I would like to be convinced that your system consistently generates reasonable sequences. One approach would be to have users manually validate the results, which seems reasonable.

I would recommend that you dramatically compress your explanation of the method, add more testing to show that it generates reasonable sequences, and devote more space to discussing and demonstrating its importance for security.

Also note that Section 6 takes an abrupt turn. Earlier sections left me expecting an evaluation focus would be on normal users, but this section focuses on security for vendors.

A

Comments for author

This paper is difficult to follow. The motivation is not quite clear. According to Section 2, Helion could avoid the put of "tremendous amount of time and effort". However, on the other hand, a machine learning-based approach requires "the existence of a set of natural scenarios" as training data, which also needs many efforts. Such motivation is not convincing enough.

Also, Helion does not consider the possibility of personalized settings. It implies that every user has the same need. Further, the output of Helion relies on the quality and quantity of input data.

The evaluations are not convincing. In Section 8.1, no quantitative analysis and horizontal comparison were conducted. The assessment is only based on the judgments of the authors. Such evaluation cannot demonstrate the advancement of Helion, compared with the previous solutions. In Section 8.2, similarly, why the previous solution cannot achieve the same result?

Finally, it seems this paper is not a security paper. The security-related contents/contributions are just a small part of this paper, and these contents could be removed without affecting the logic. Maybe this paper is more suitable for ubiquitous computing or human-computer interaction related conferences, like UbiComp and CHI.

Some small issues:

(1) What is the meaning of "Helion"? It is a bit confusing.

C

There has been a great degree of interest in security and safety policies for smart homes in recent years. While a number of papers cited by this work have sought to create and enforce security properties for such situations, the evaluation of these properties has always been somewhat incomplete because few large, real-world data sets from actual homes are available. This paper aims to fix that, which is a laudable goal. The core problem is that what this paper accomplishes is not sufficiently real, and also perhaps not sufficiently novel, to achieve that goal.

The paper claims that analysis of Samsung SmartThings apps is the state of the art for home automation data sets. However, this is not quite true. Some recent home automation papers with similar goals, like [33], used large-scale data collected from IFTTT:

- Mi et al. An empirical characterization of IFTTT: ecosystem, usage, and performance. From IMC '17 <https://dl.acm.org/citation.cfm?id=3131369>
- Ur et al. Trigger-Action Programming in the Wild: An Analysis of 200,000 IFTTT Recipes. From CHI '16 <https://dl.acm.org/citation.cfm?id=2858556>

This paper collects IFTTT-like trigger action programs in a small user survey. This paper's technique has the advantage that a collection of apps comes from a single user, whereas those large IFTTT data sets only contain the name of the person who created the app, rather than everyone who is using it. As such, this paper could make a great contribution. Unfortunately, the apps users in this study make don't control anything in the real world. Users don't have the chance to iteratively add to or improve their rules. In addition, the traces are simulated based on a language model based on these rules. Thus, it seems like this artificial approach does not have enough ecological validity or external validity to approximate real traces. Furthermore, based on the

entered manually for these rules, rather than generated from real-world events. In other words, the interface in the study is not like the one used in systems like IFTTT.

at Washington State from a decade ago also <https://www.cse.washington.edu/datasets/>

security and safety violations are very rare. In the end, why not use a formal model? Much of

B

Readings Bug Bounty!

- Reading papers is hard; *reading 22-30 papers in a semester is even harder.*
- New this year: *you do not have to do all the readings.*
- *Instead:* Each student owes Prof. Nadkarni **2 bugs** from published papers assigned in class (5 pt each, **10 course points total**)

Rule 1: You must be the *first to report* the bug

Rule 2: It must be *non-trivial* (e.g., impractical assumption, logical flaw that affects the paper's claims)

Rule 3: You must be able to *explain it*

Cheating Policy

- Cheating is not allowed
- We run tools
- If you cheat, you will probably get caught

- If you get caught, you will get a ~~negative score~~ on the project.

This includes the course project!

All text and figures should be your own.

- **I REFER ALL ACADEMIC DISHONESTY INCIDENTS TO THE OFFICE OF STUDENT CONDUCT, WITHOUT EXCEPTION**
- When in doubt, *ask*

Course Credo

Think like an attacker, but behave like a responsible adult

W&M's computer usage policies apply to this class.

Security course != permission to disrupt or cause harm

Ethics Statement

- This course considers topics involving personal and public privacy and security. **As part of this investigation we will cover technologies whose abuse may infringe on the rights of others.** As an instructor, I rely on the ethical use of these technologies. Unethical use may include circumvention of existing security or privacy measurements for any purpose, or the dissemination, promotion, or exploitation of vulnerabilities of these services. Exceptions to these guidelines may occur in the process of reporting vulnerabilities through public and authoritative channels. **Any activity outside the letter or spirit of these guidelines will be reported to the proper authorities and may result in dismissal from the class and or institution.**
- When in doubt, please contact the instructor for advice. Do not undertake any action which could be perceived as technology misuse anywhere and/or under any circumstances unless you have received explicit permission from Professor Nadkarni.

Other Policies

- Please turn off cell phones during class.
- I will do my best to respond to emails within 24 hours. You will receive faster answers if you post to Piazza.
- Students may appeal to the instructor for reconsideration of a grade, but the appeal must be in writing (i.e., email), and must be sent within 3 weeks (or the close of the semester, whichever is sooner) of receiving the graded assignment.
- Behave civilly: **don't be late for class**; don't read newspapers/blogs/etc. during class; don't solve Sudoku puzzles during class; don't struggle with crossword puzzles during class; **respect others' opinions**, *even if they are clearly wrong*.
- Adhere to good scientific principles and practices, and uphold the W&M Student Code of Conduct.

.

Lecture Notes

- Slides will be released on the course schedule after each class.
- If you are presenting, please email me a PDF of the slides after class.

Good Luck!