Lecture #2: Beginnings

UCalgary ENSF619

Elements of Software Security

Instructor: Lorenzo De Carli (lorenzo.decarli@ucalgary.ca)

A lightning introduction to security

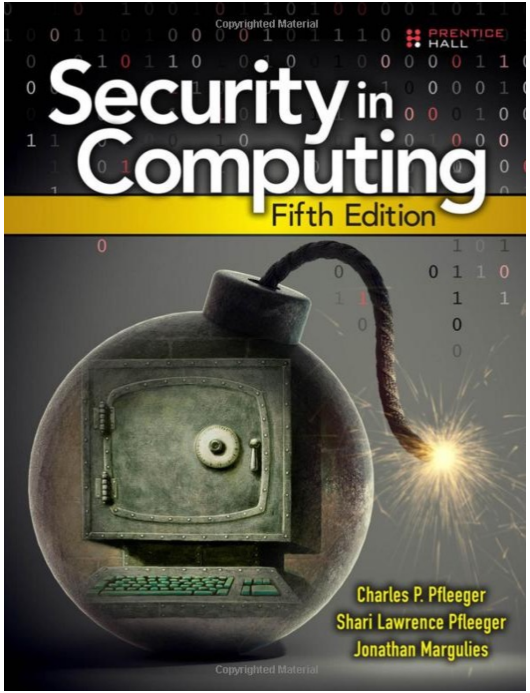

This section is partly based on Pfleeger & Pfleeger, ”Security in Computing”, 4th ed.

Object of study

- Security of software and related problems

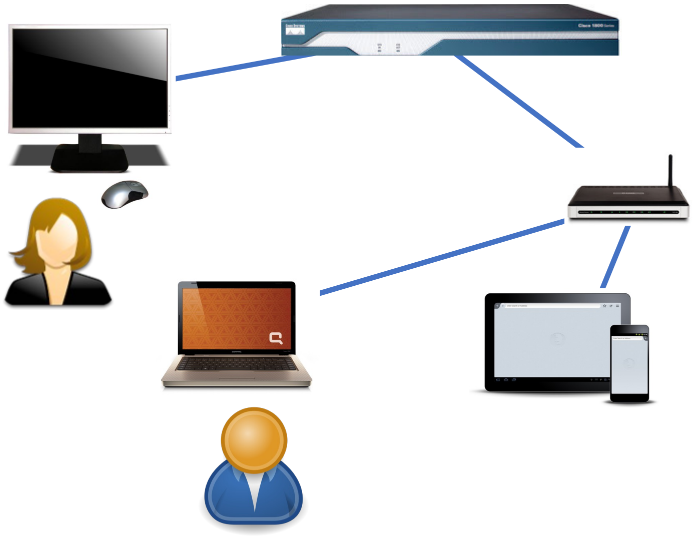

Including:

- **Software artifacts** (source code, executable)
- **Computer systems & devices**
- **Digital data and communications**
- **Users**

Security in a diagram (well, two)

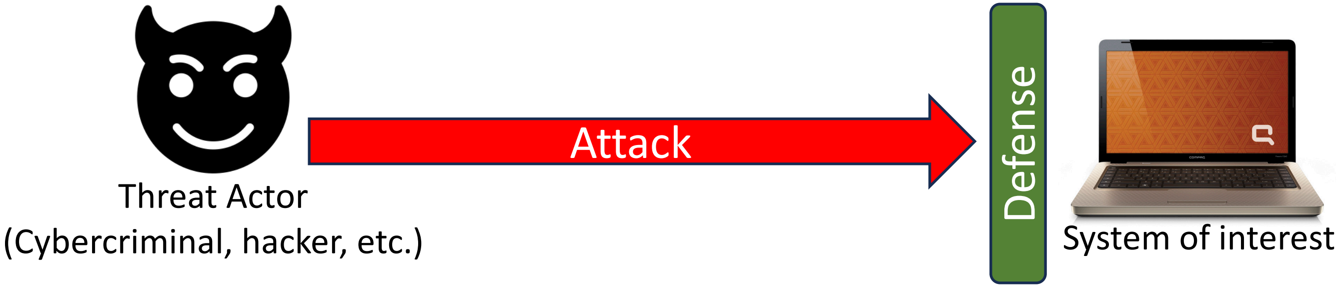

Threat Actor
(Cybercriminal, hacker, etc.)

Attack

Defense

System of interest

Types of security research

Building attacks (how can systems be broken?)

Building defenses (how can systems be protected?)

Measures: characterizing security posture, finding vulnerabilities etc.

More formal security concepts

- A **vulnerability** is a weakness in the security of a system that can be exploited to cause loss or harm
- A **threat** is a set of circumstances that has the potential to cause loss or harm
- A human who (directly or indirectly) exploits a vulnerability perpetrates an **attack** on the system
- **Control** is the act of using an action, device, procedure, or technique to remove or reduce a vulnerability
- **A threat is blocked by control of a vulnerability**

General security goals

- **Confidentiality:**
  - Computer-related assets are accessed only by authorized parties (a very narrow interpretation of the concept of **privacy**)
- **Integrity:**
  - Assets are modified only by authorized parties and in authorized ways
- **Availability:**
  - Assets are accessible to authorized parties at appropriate times

What are these “assets” anyway?

- Typically, **information**
- Oftentimes, also **components of a networked computer system** (e.g. software)
- Also, especially in recent times, physical entities: **industrial machinery & smart home devices**

More on privacy

- For much of the history of computer security people have used fairly simple definition of privacy which generally **equates privacy to secrecy**
- Nowadays, **the need is recognized to go beyond this simple definition**
- E.g.: **contextual integrity** (Helen Nissenbaum):
  - Privacy is provided by **appropriate flows of information**
  - “Appropriate” means “conforming to contextual information norms”
  - E.g. it is acceptable to disclose a high-schooler’s grades if the recipient is a parent of the student

Vulnerabilities

- **Hardware vulnerabilities:** interesting, but largely outside the scope of this class
- **Software vulnerabilities:**
  - A software may be vulnerable to modifications that may cause the software to fail, malfunction, or allow introduction of malicious behavior
    - Via programming bugs, backdoors
  - A software may leak information
- **Human vulnerabilities:**
  - Humans may be convinced to misuse the system, causing violation of security properties even if the system itself does not malfunction (e.g., phishing)

Threats

- **Amateurs:** various categories of people with low motivation and skills – example of attacks including stealing and publishing personal information, preventing victim from using an online service
- “**Garden-variety” cybercriminals:** part of criminal organization typically interested in performing attacks for various types of financial gain (e.g. infiltrating a payment processor to steal CC numbers)
- **Advanced Persistent Threats:** highly organized entities, staffed by skilled professionals, which typically work to foster the long-term goals of a nation-state (e.g. infiltrating an embassy to discover undercover enemy operatives, disrupt financial/industrial activity, etc.)

Attribution and forensic

- **Attribution:**
  - The act of attributing an attack to a particular person/organization
  - Complicated by the indirect nature of network attacks, and the fact that online identify is easily altered/hidden
- **Forensics:**
  - The process of understanding a successful attack: how it happened, what was accomplished by the attackers (potentially performing attribution too)
  - Complicated by the need of storing logs/historical data, the fact that attackers may “cover their tracks” or even still be present when forensics begins

A note on retaliation

“Hacking back is a terrible idea that just will not die”
(Bruce Schneier)

Is security important anyway?

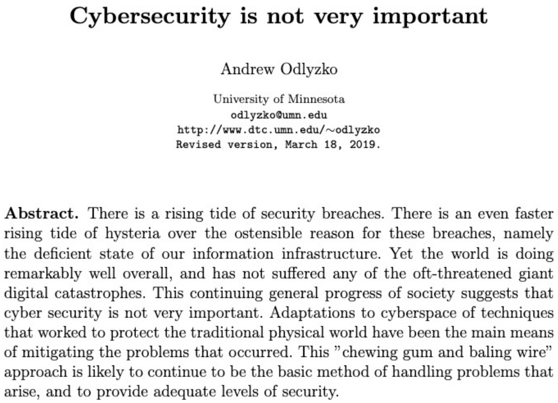

Cybersecurity is not very important

Andrew Odlyzko

University of Minnesota
odlyzko@umn.edu
http://www.dtc.umn.edu/~odlyzko
Revised version, March 18, 2019.

Abstract. There is a rising tide of security breaches. There is an even faster rising tide of hysteria over the ostensible reason for these breaches, namely the deficient state of our information infrastructure. Yet the world is doing remarkably well overall, and has not suffered any of the oft-threatened giant digital catastrophes. This continuing general progress of society suggests that cyber security is not very important. Adaptations to cyberspace of techniques that worked to protect the traditional physical world have been the main means of mitigating the problems that occurred. This "chewing gum and baling wire" approach is likely to continue to be the basic method of handling problems that arise, and to provide adequate levels of security.

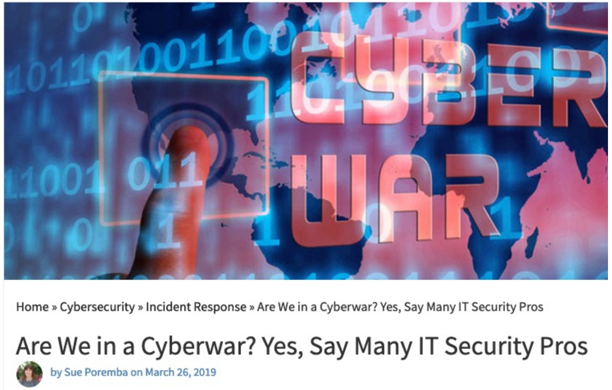

Home » Cybersecurity » Incident Response » Are We in a Cyberwar? Yes, Say Many IT Security Pros

Are We in a Cyberwar? Yes, Say Many IT Security Pros

by Sue Poremba on March 26, 2019

(the real state of things is probably somewhere in between)

Threat Modeling

What is threat modeling?

- Informally, we commonly assess “threats” in everyday life
- E.g.: how early should I leave home to make it on time for lecture?
    - Empirically, try to find best approach to (i) maximize sleep, and (ii) minimize change of arriving late
- **Threat modeling is the same idea, just more formal and for computing systems**

What is a threat model?

- Given a system, it is a **definition** of:
  - What are the **risks** we care about
  - What are the **capabilities** of the **attacker**
  - Which risks we want to **mitigate** and to what **extent**

Threat model example?

Why do we care about threat modeling?

- Just “making a system secure” is a poor, fuzzy goal
- Oftentimes:
  - There are risks **we don’t care about or don’t apply**
  - There are risks **we cannot mitigate**
  - There is only a **finite amount of resources** to invest in security

Threat modeling matters for research too!

- Most self-respecting papers must have a **clear threat model**
- Without a threat model it is **impossible to determine**:
  - Whether an attack is **realistic/relevant**
  - Whether a defense is **actually useful**
  - Whether a defense **actually protects against the attack**

How detailed a threat model should be?

- **It depends** on the task!
- For example, a paper detailing an attack against a cryptographic algorithm may make **very specific assumptions** about attacker computational capabilities, access to information etc.
- A paper presenting a detector for malicious URLs may use a more **empirical, generic mod**el that just outlines the type of attacker

Is there some **guidance** for creating one?

- For academic research, oftentimes people pick it up from reading papers
- When working in industry, threat models may need to be more elaborate and specific processes are followed
- **Shostak's four questions:**
  - What are we working on?
  - What can go wrong?
  - What are we going to do about it?
  - Did we do a good enough job?

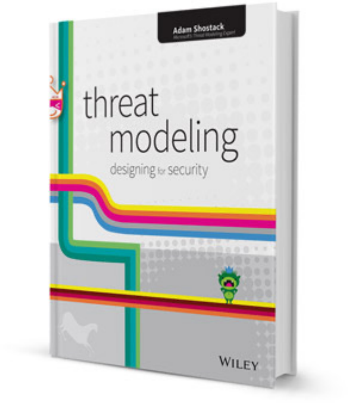

Aside: an informal guide to academic security research

Academic security research

- Academic research may be funded by **public** or **private grants**
  - **Private grants** typically comes with more **specific expectations**
  - **Public grants** may be for **specific projects** or for more **open-ended investigations**
- **Outcomes** of academic research
  - **Publications** (always expected)
  - Oftentimes, the release of some type of **tool or proof-of-concept implementation**

Publications and peer-review

- Researchers typically draft up a report on the research they carried – background, motivation, technical aspects, results, etc.
- Prior to be disseminated through conferences/journals (more on this later), the work must be **peer-reviewed**
- Idea: other members of the community review the work and decide whether it is sound enough to be published
- Goal: prevent “bad science” and identify mistakes
- Does it work? So-so (but better than nothing)

Where is security research published?

- **Highly influential research** typically follow the Computer Science community practice of **conferences over journals**
- **“Tier-1” conferences:**
  - ACM CCS
  - IEEE S&P
  - USENIX Security Symposium
  - Sometimes NDSS is also considered part of this set
  - There are many other great conferences too, which are oftentimes only slightly less selective than the top ones!
  - Several pages with informal conference rankings exist

Security is inter-disciplinary

- …so, oftentimes security-relevant papers may end up being published in **top conferences in other fields** (e.g., software engineering, networking, programing languages, HCI, AI/ML, etc.)

What about journals?

- While journal publications may receive less attention, some journals are **very selective** and **closely followed by the community**
  - Anything with “ACM Transactions” or “IEEE Transactions” is probably good
  - There are some other decent journals from reputable publishers
  - There are also a lot of scammy/predatory journals (avoid publishing there, it will taint your resume!)

Is this gatekeeping?

- Maybe?
- It’d be great to have less reliance on rankings etc…
- But, having clear community agreement on which venues are the “top” ones help:
    - People have finite time to read paper and discover new publications
    - Top venues guarantee at least some baseline “quality control”

A note about arXiv

- Oftentimes, researchers will put drafts of their paper on arxiv.org **before** the work is peer-reviewed and accepted for publication
- Very common in ML/AI!
- **Reason:** work can start accruing citations, “plant the flag” in a particular research area
- **Not a bad practice, but keep in mind those papers have not been peer-reviewed!**

The typical research process

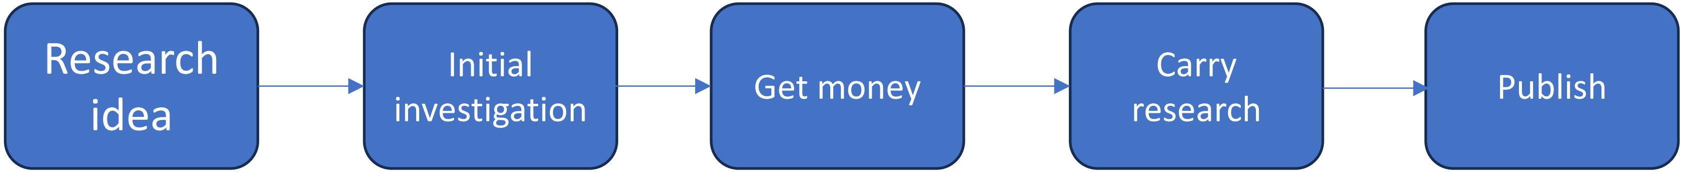

Research idea

Initial investigation

Get money

Carry research

Publish

The cycle of ideas in technical research

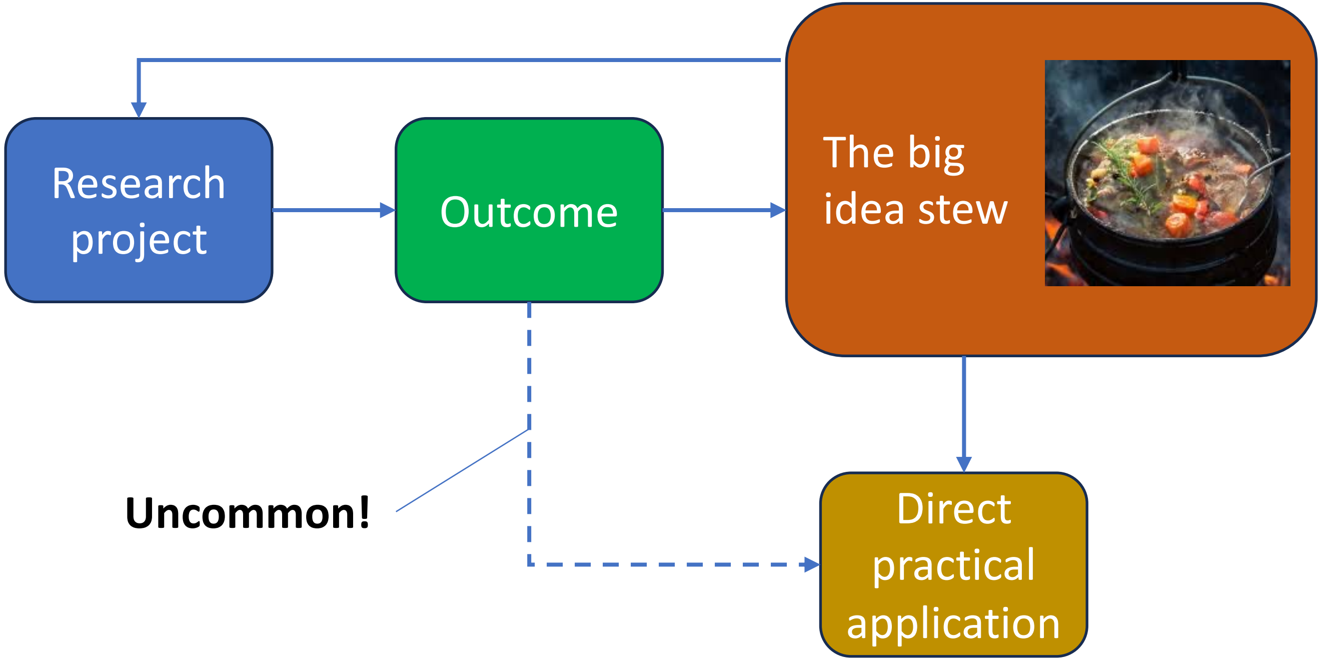

Research project

Outcome

The big idea stew

Direct practical application

Uncommon!

Why am I telling you these things?

- You are going to read a **lot of academic papers** and do a **lot of research** (in this course and otherwise)
- It is important that you can **place it in context**!
  - Know what is the **purpose** of the research
  - Get a sense of **which community** published the work

See you next lecture!