Measurement + Ethics

CS499/579 :: Empirical Computer Security

Zane Ma (he/him/his)
2024.10.02

From last class…

- In order to understand how computer systems actually work, we need to measure them (e.g., performance / security properties)

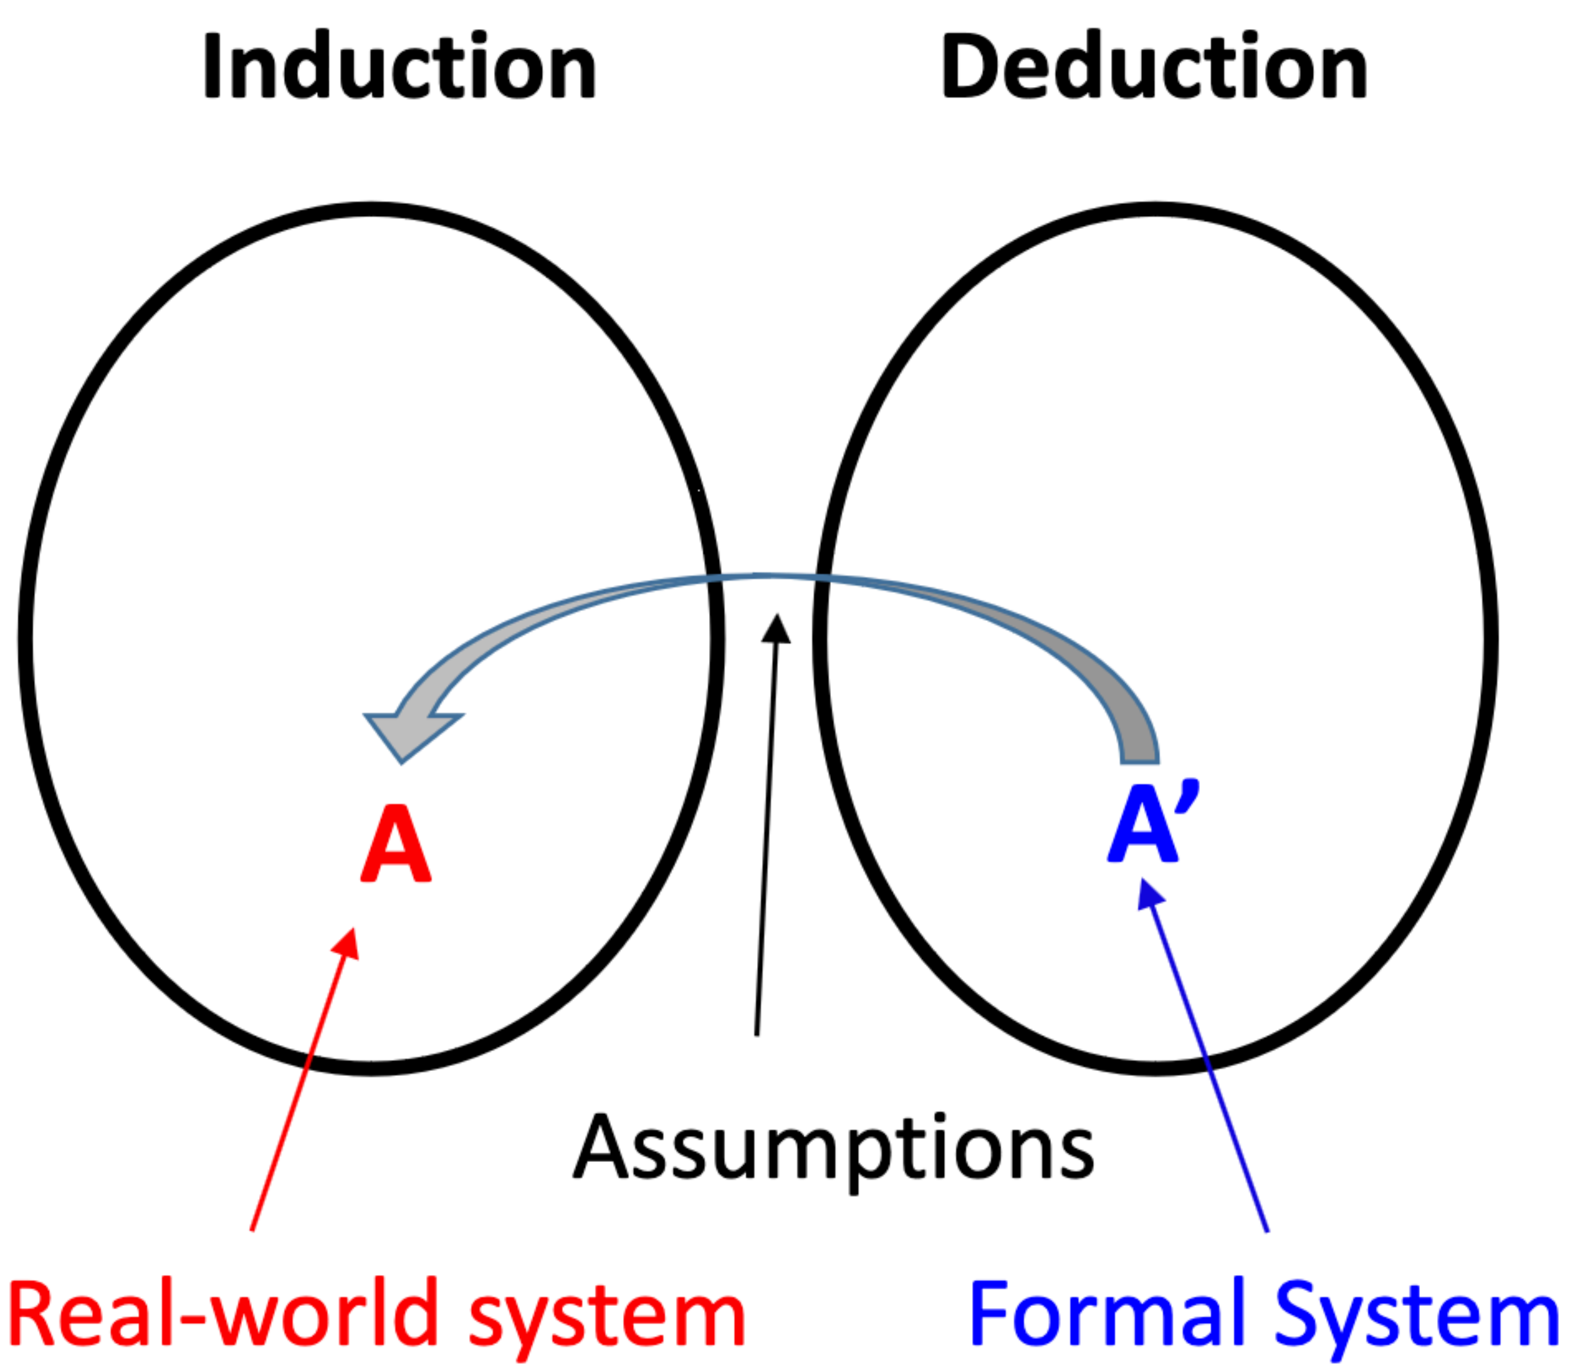

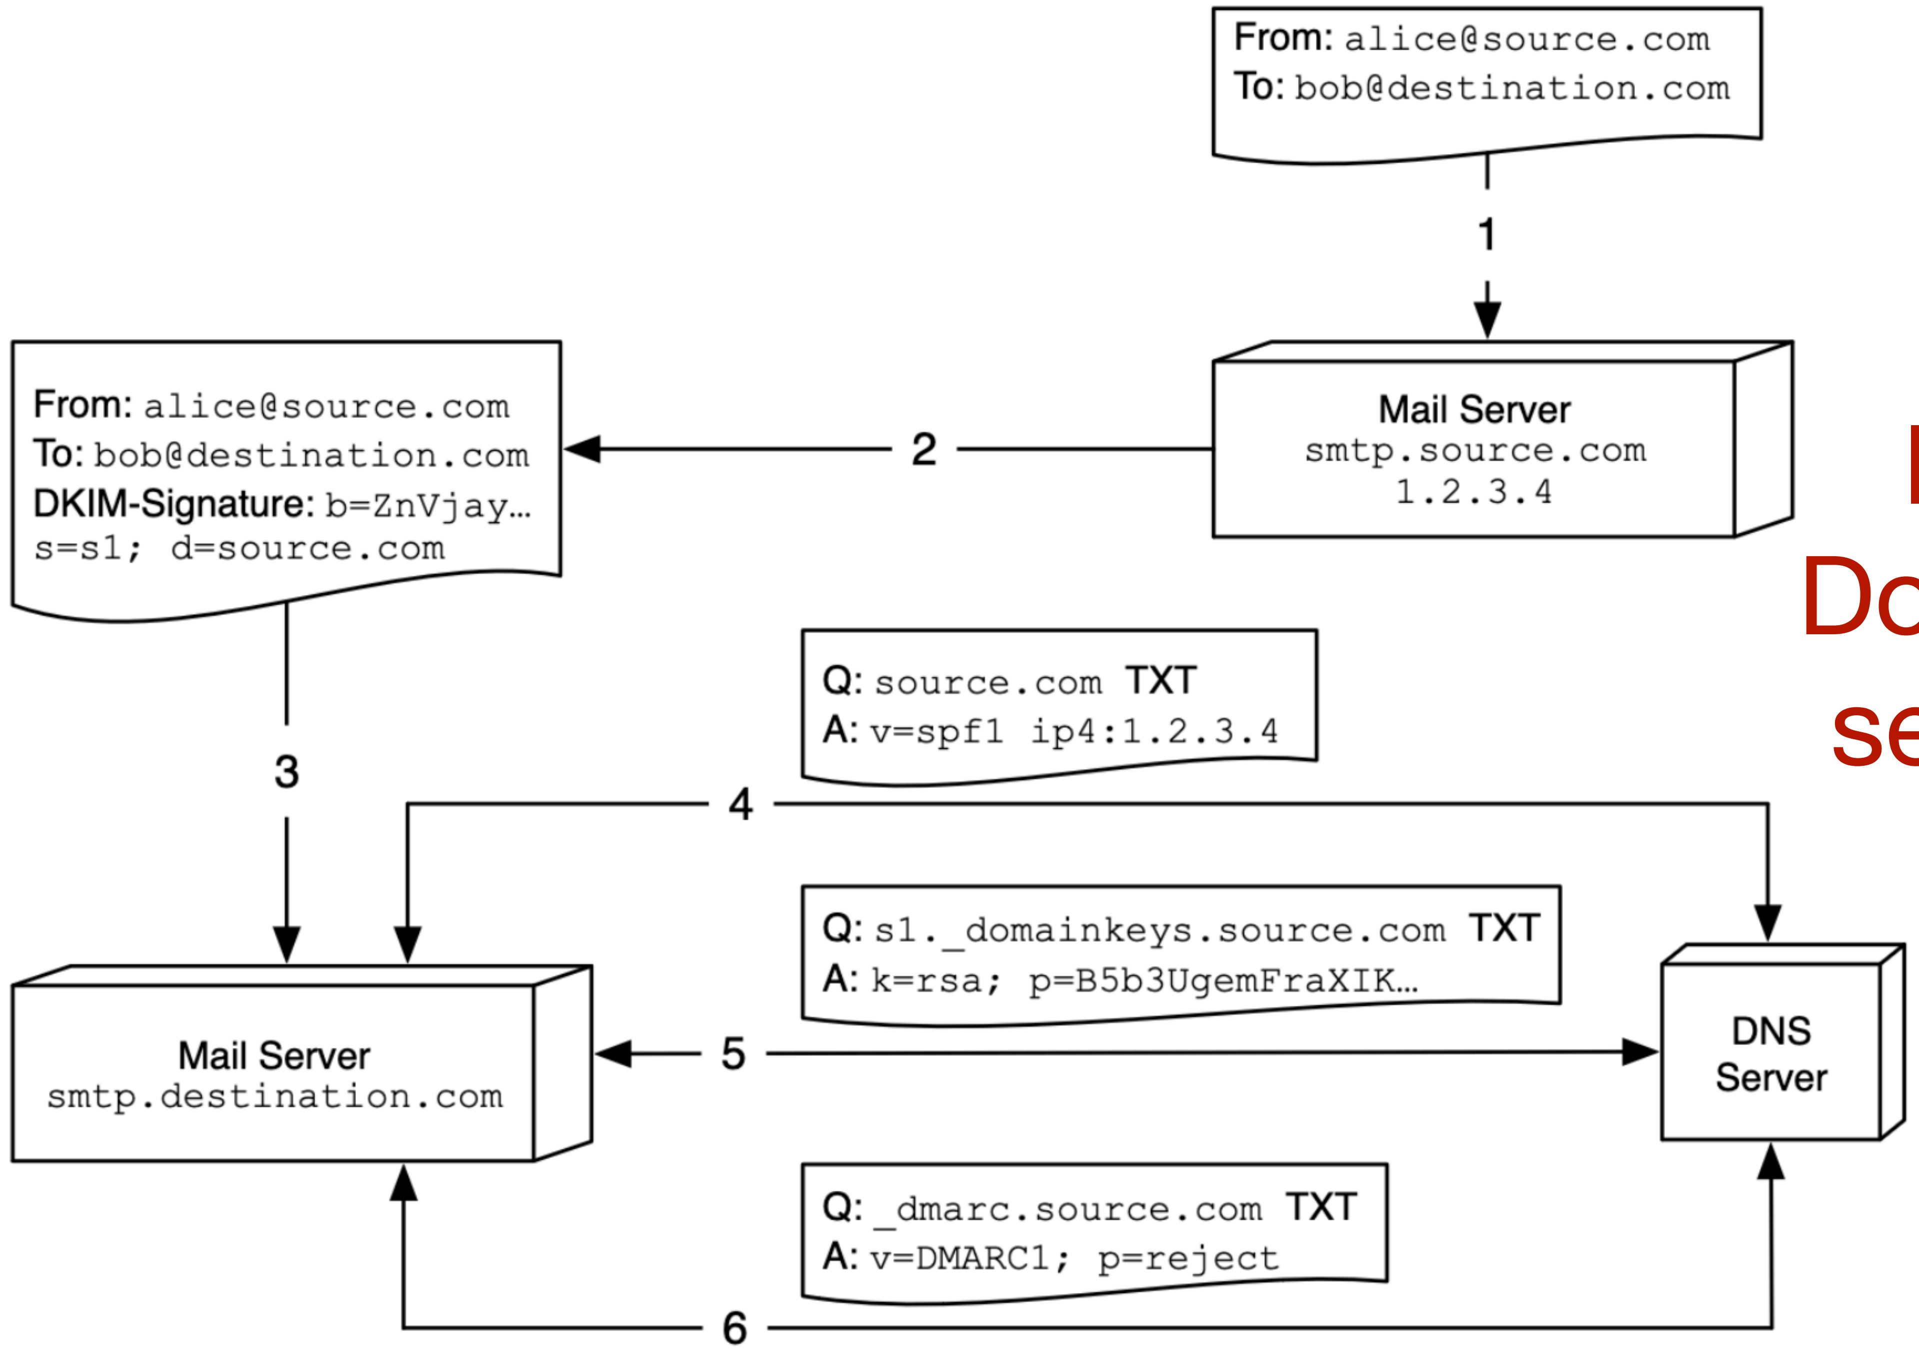

Is email secure?
Do people use email security protocols?
Used securely?

Scanning the Internet

- Prior to 2013, scanning the full internet was uncommon
- Why? (Think IPv4)

IPv4 header format

<table>
<tr><th>Offsets</th><th>Octet</th><th colspan="8">0</th><th colspan="8">1</th><th colspan="8">2</th><th colspan="8">3</th></tr>
<tr><th>Octet</th><th>Bit</th><th>0</th><th>1</th><th>2</th><th>3</th><th>4</th><th>5</th><th>6</th><th>7</th><th>8</th><th>9</th><th>10</th><th>11</th><th>12</th><th>13</th><th>14</th><th>15</th><th>16</th><th>17</th><th>18</th><th>19</th><th>20</th><th>21</th><th>22</th><th>23</th><th>24</th><th>25</th><th>26</th><th>27</th><th>28</th><th>29</th><th>30</th><th>31</th></tr>
<tr><td>0</td><td>0</td><td colspan="4">Version</td><td colspan="4">IHL</td><td colspan="6">DSCP</td><td colspan="2">ECN</td><td colspan="16">Total Length</td></tr>
<tr><td>4</td><td>32</td><td colspan="16">Identification</td><td colspan="3">Flags</td><td colspan="13">Fragment Offset</td></tr>
<tr><td>8</td><td>64</td><td colspan="8">Time To Live</td><td colspan="8">Protocol</td><td colspan="16">Header Checksum</td></tr>
<tr><td>12</td><td>96</td><td colspan="32">Source IP Address</td></tr>
<tr><td>16</td><td>128</td><td colspan="32">Destination IP Address</td></tr>
<tr><td>20</td><td>160</td><td colspan="32" rowspan="3">Options (if IHL > 5)</td></tr>
<tr><td>⋮</td><td>⋮</td></tr>
<tr><td>56</td><td>448</td></tr>
</table>

EFF SSL Observatory: A glimpse at the CA ecosystem (2010)

3 months on 3 Linux desktop machines (6500 CPU-hours)

Census and Survey of the Visible Internet (2008)

3 months to complete ICMP census (2200 CPU-hours)

- 32-bit address! 2^{32} = ~4B destination IPs
- Scanning at 100 IPs / second would take 462 days

ZMap: Fast Internet-Wide Scanning and Its Security Applications

Zakir Durumeric
Michigan (now Stanford)

Eric Wustrow
Michigan (now UC Boulder)

Alex Halderman
Michigan

2013 USENIX

Introducing ZMap

An open-source tool that can port scan the entire IPv4 address space from just one machine in under 45 minutes with 98% coverage

With ZMap, an Internet-wide TCP SYN scan on port 443 is as easy as:

```
$ zmap -p 443 -o results.txt
34,132,693 listening hosts
(took 44m12s)
```

97% of gigabit Ethernet linespeed

Weeks / months of scanning —> hours

How does it work?

What are the resource / performance costs?
How would you optimize this?

Naive way of scanning an IP address:

1. Make a randomized stack of all IP addresses
2. Send one packet to random destination (pop off the stack)
3. Wait - if response received, log IP + response payload; otherwise, timeout

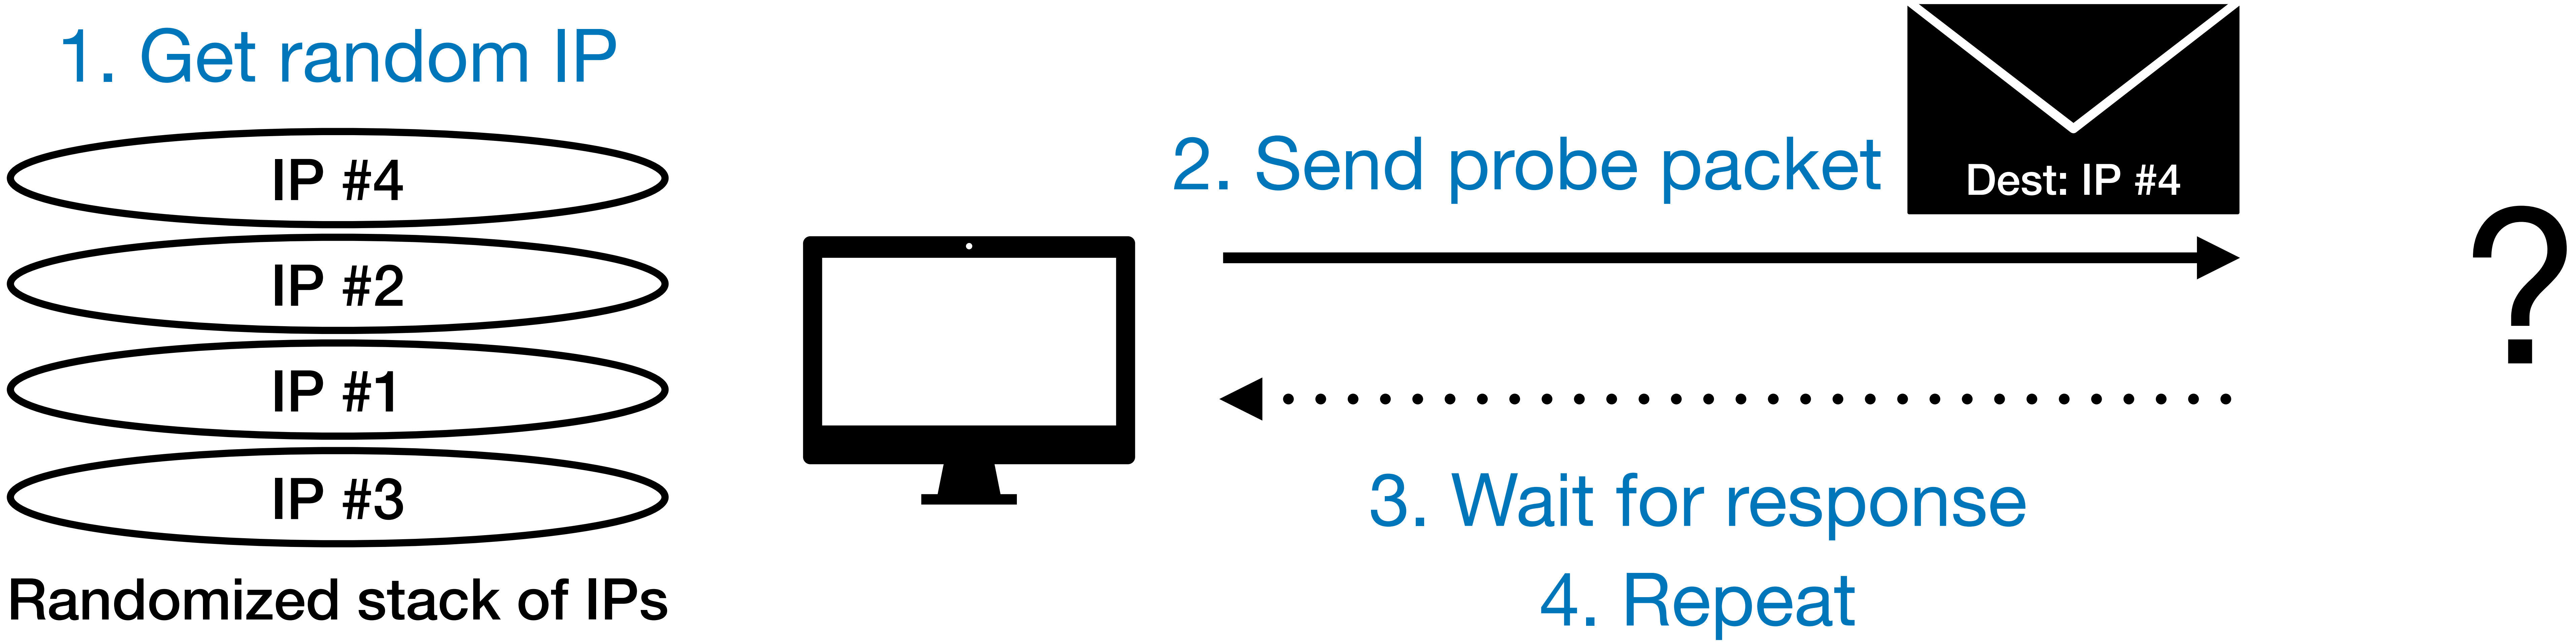

How does it work?

Short answer: reduce / eliminate state associated with scanning!

In other words, reduce how much the scanner has to remember, so you don't need to wait for responses (facilitating parallelization) + you can minimize memory usage

1. Efficient random IP tracking: How can we scan all IPv4 addresses, randomly, without remembering all the ones we have already scanned?

2. Stateless scanning: How can we send out network requests without waiting for a response?

1. Efficient random IP tracking

How can we scan all IPv4 addresses (equivalent to 4-byte unsigned integer), ~~randomly~~, without remembering all the ones we have already scanned?

Order them and keep track of:

1. Current IP address (e.g., 128.193.10.29)
2. Increment size (e.g., 1)
3. Starting point (e.g., 0 = 0.0.0.0)

Randomness is required to reduce the scanning load on individual networks (i.e., adjacent IP addresses).

1. Efficient random IP tracking

How can we scan all IPv4 addresses (equivalent to 4-byte unsigned integer), randomly, without remembering all the ones we have already scanned?

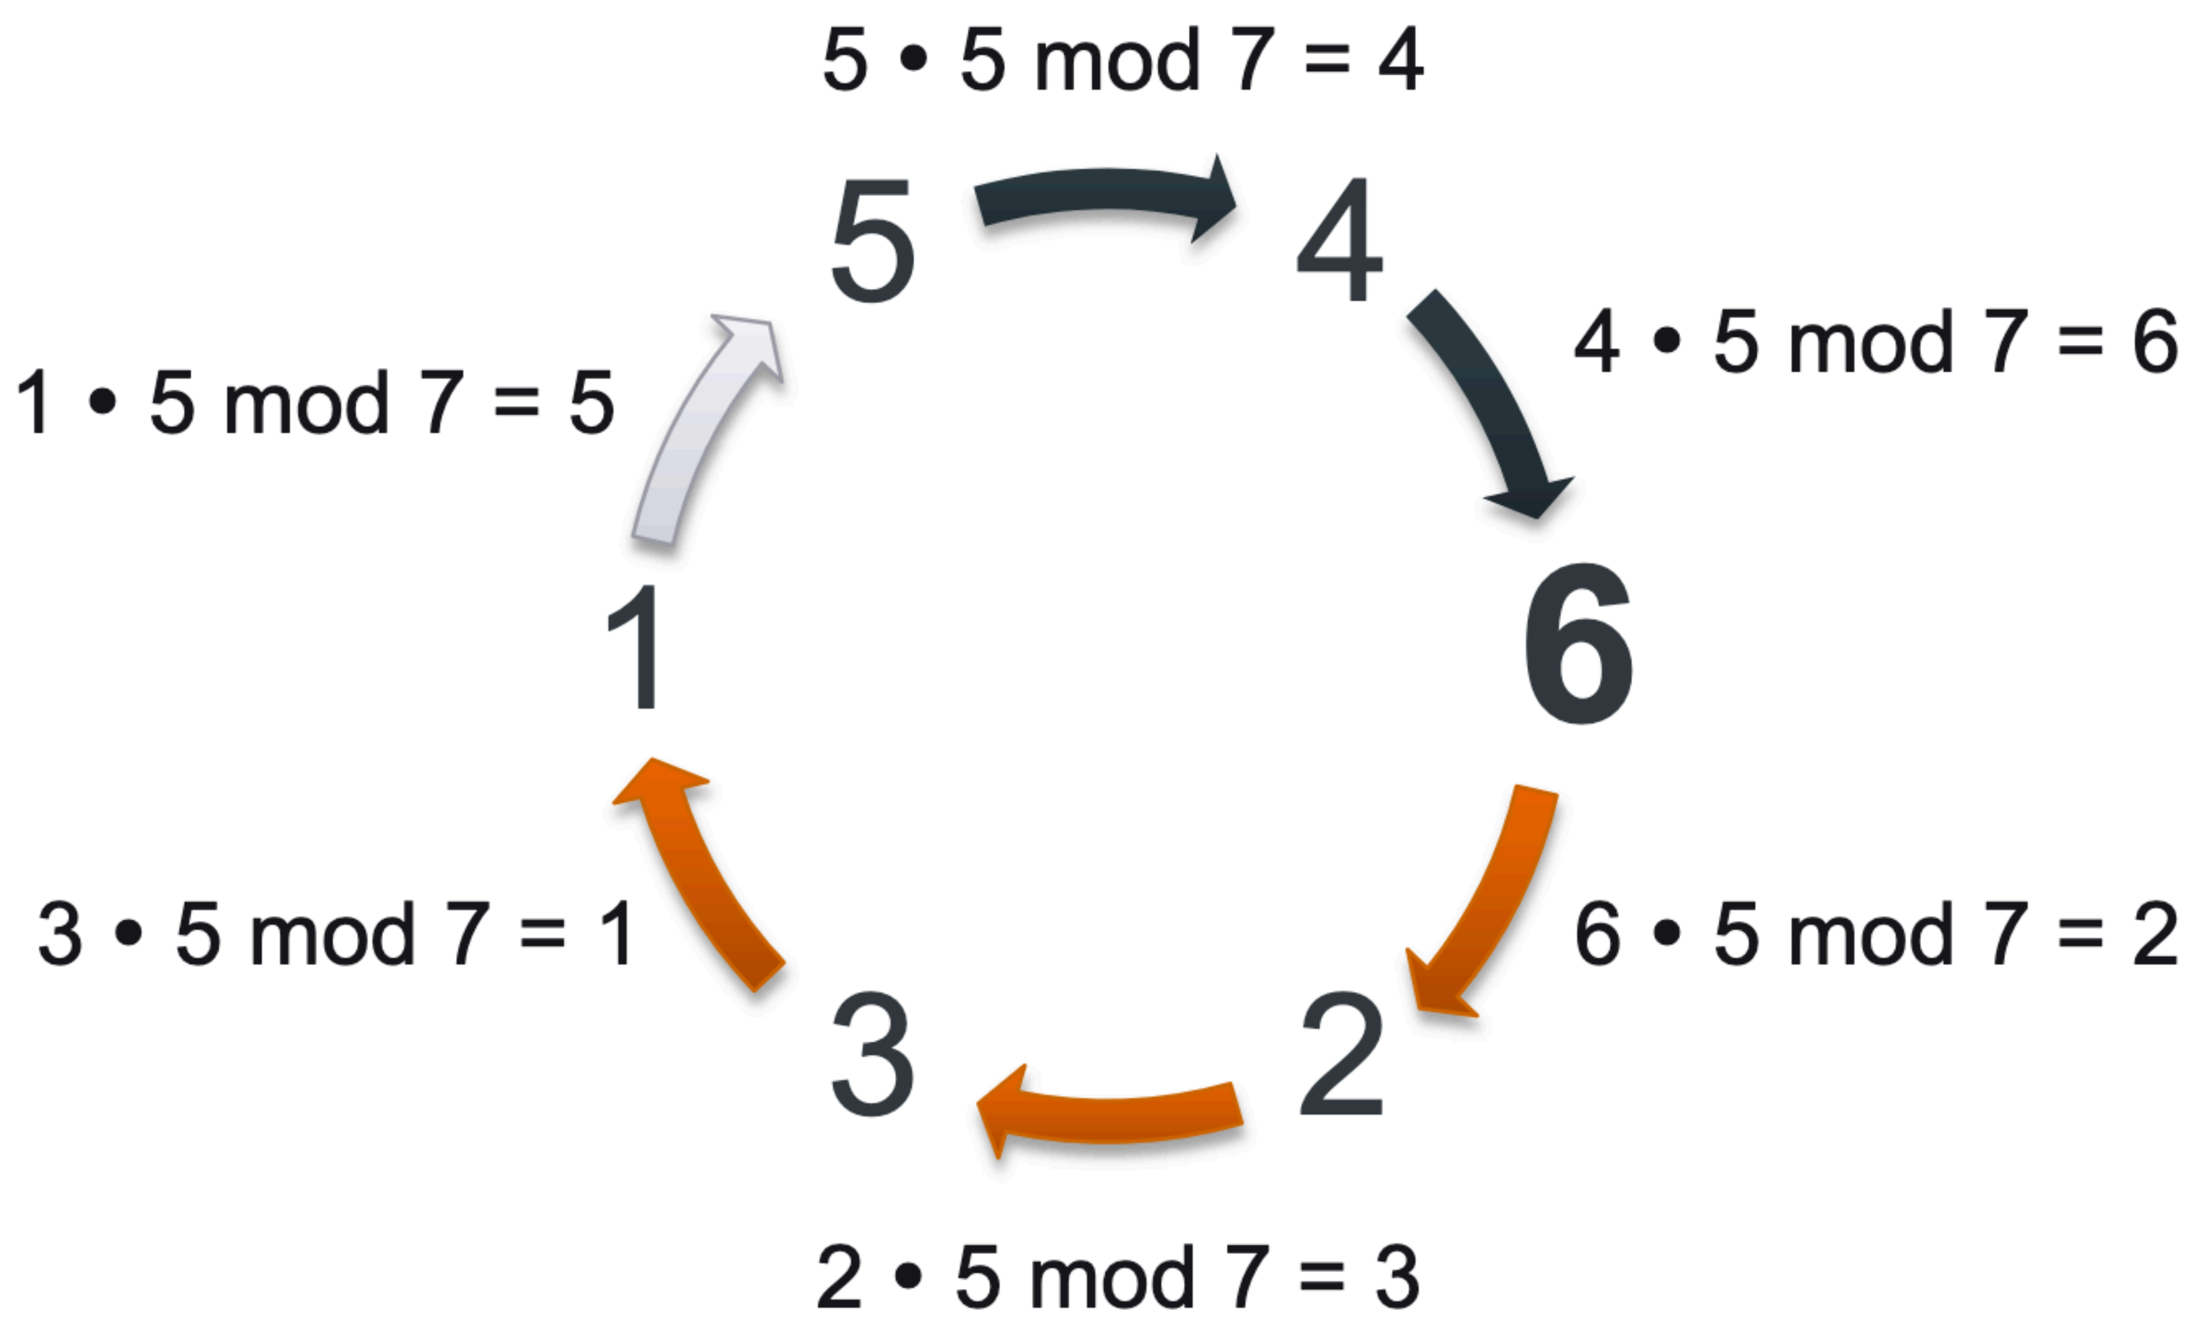

Fancy math ordering = multiplicative group of integers modulo *p*, only track:

1. Current location (current IP)

2. Primitive root (increment size)

3. First address (starting/end point)

Each primitive root is a different ordering

2. Stateless scanning

How can we send out network requests without waiting for a response?

But first: why do we need to wait for responses anyways? Random background noise - unsolicited packets are common

How do we normally distinguish between background noise packets and response packets? Look at response fields predictably related to probe packet

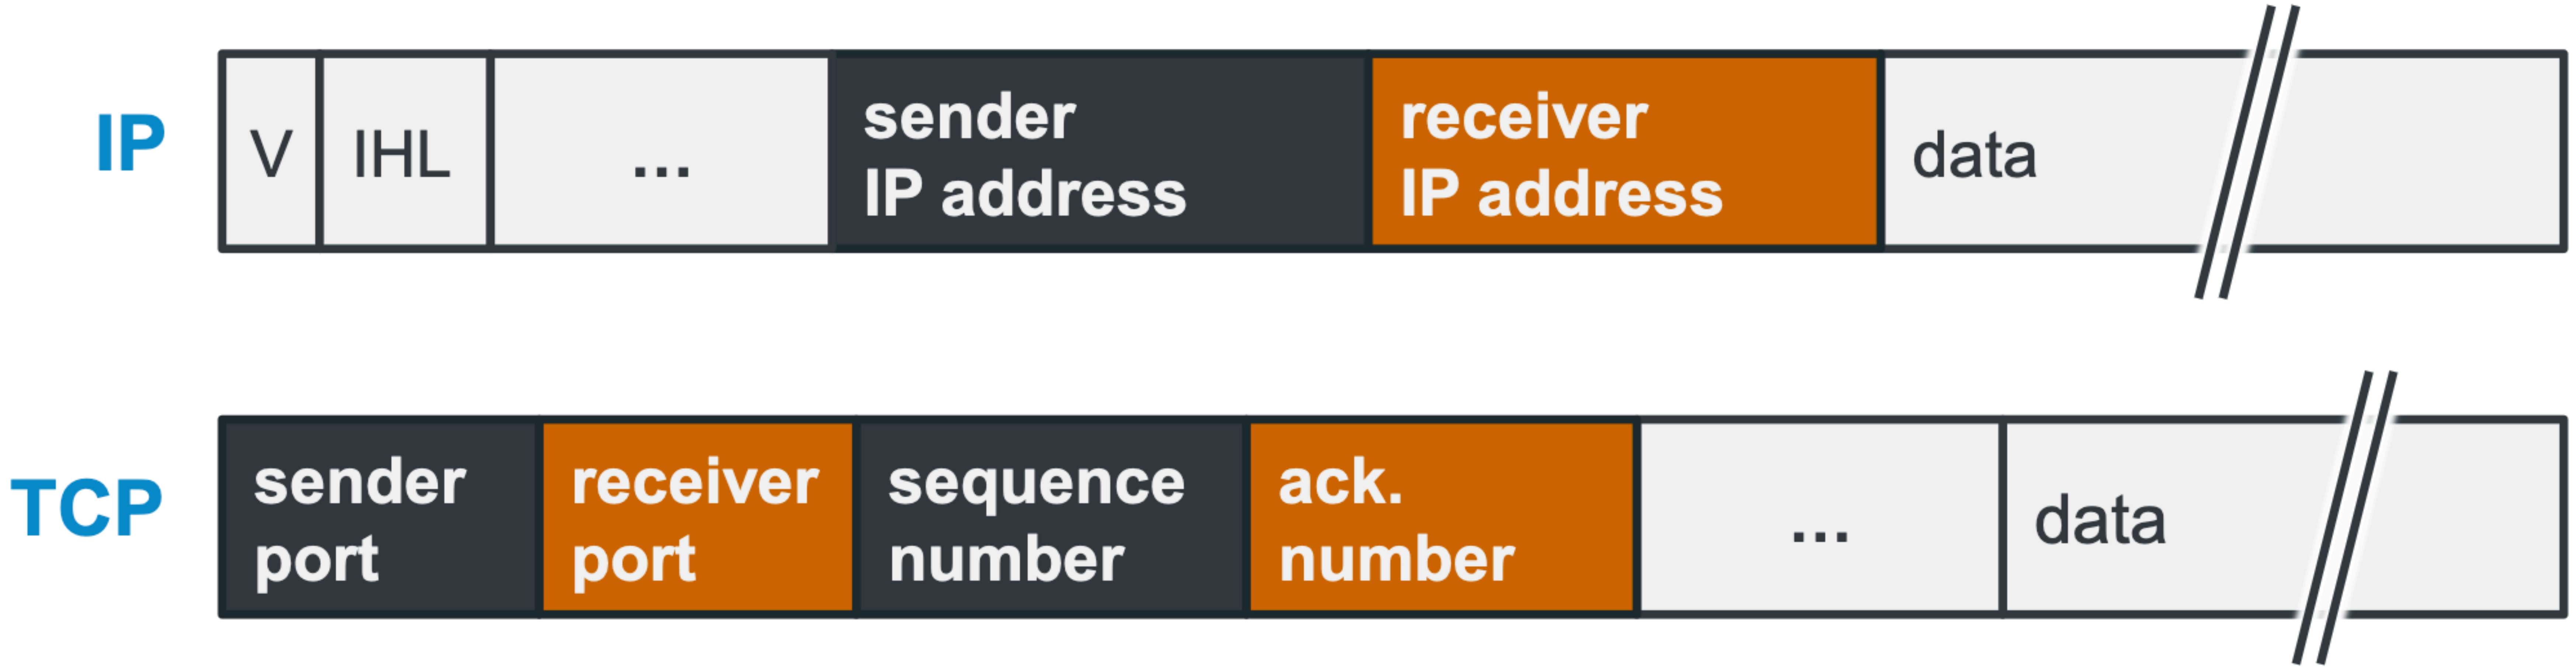

2. Stateless scanning

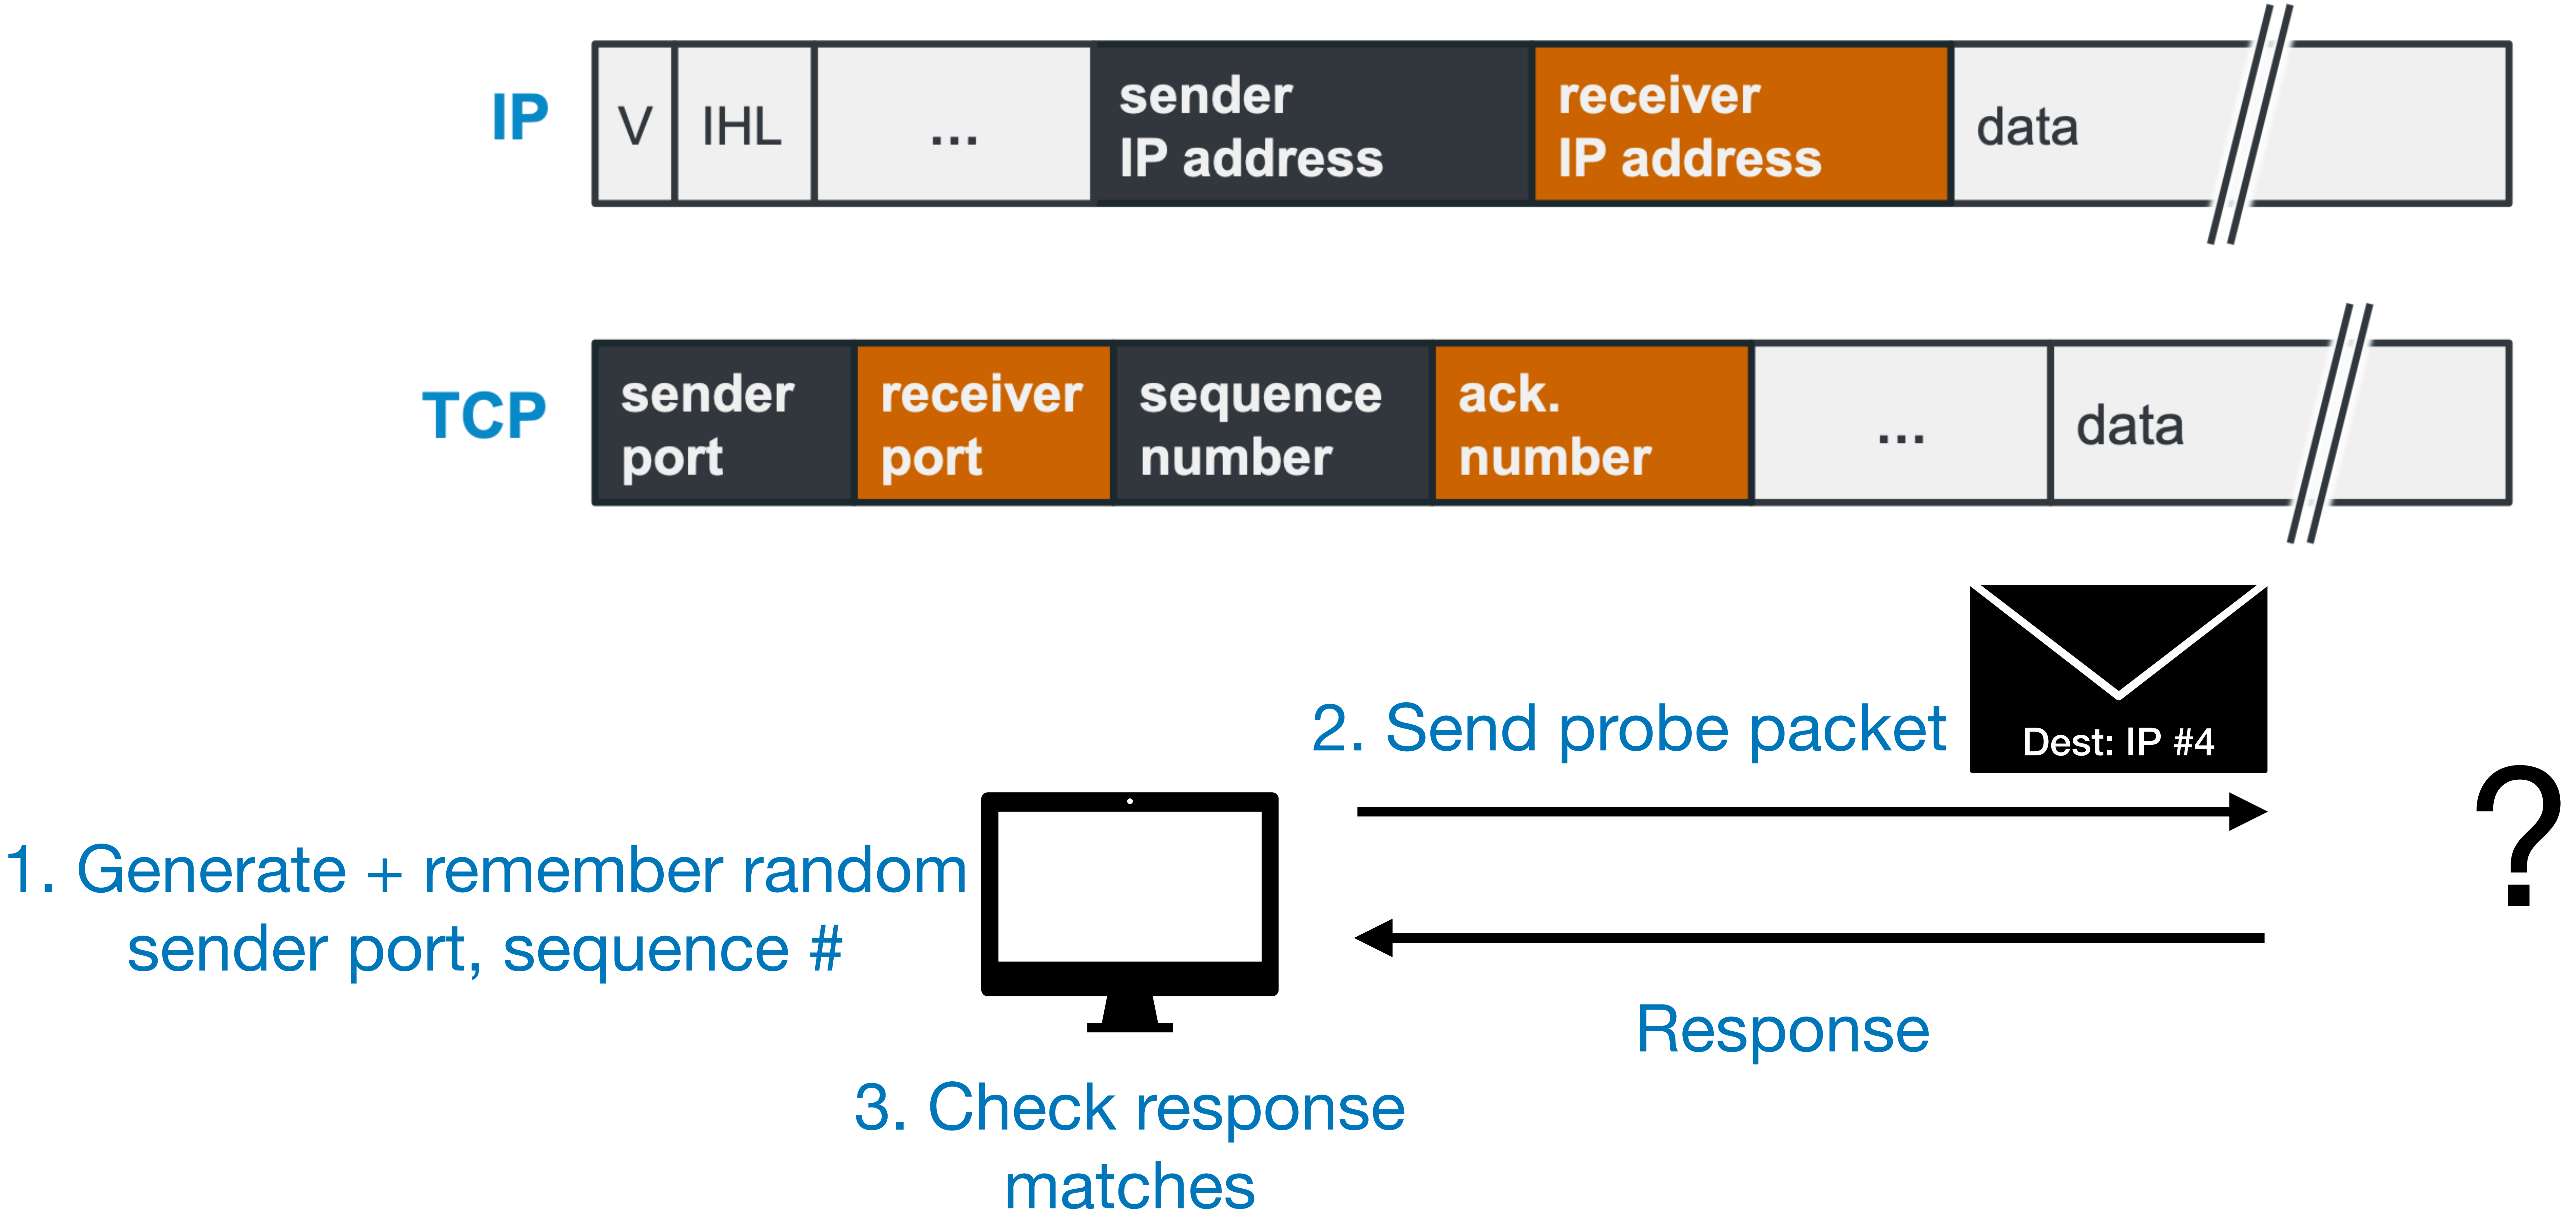

2. Stateless scanning

1. Use the same sender port and initial sequence number every time

    2^16 (16-bit sender port) * 2^32 (32-bit sequence number) uniqueness

2. Per-probe uniqueness: Set the port + sequence number based on the target IP address

    2^16 * 2^32 * 2^32 (32-bit target IP) uniqueness

Downside: can’t distinguish between responses triggered by previous scans

3. Per-probe + per-scan uniqueness (what ZMap does): set port + sequence number based on Message Authentication Code (MAC) computed over the target IP address, using a per-scan key

Scanning Performance

How fast is too fast?

No correlation between hit-rate and scan-rate. Slower scanning does not reveal additional hosts

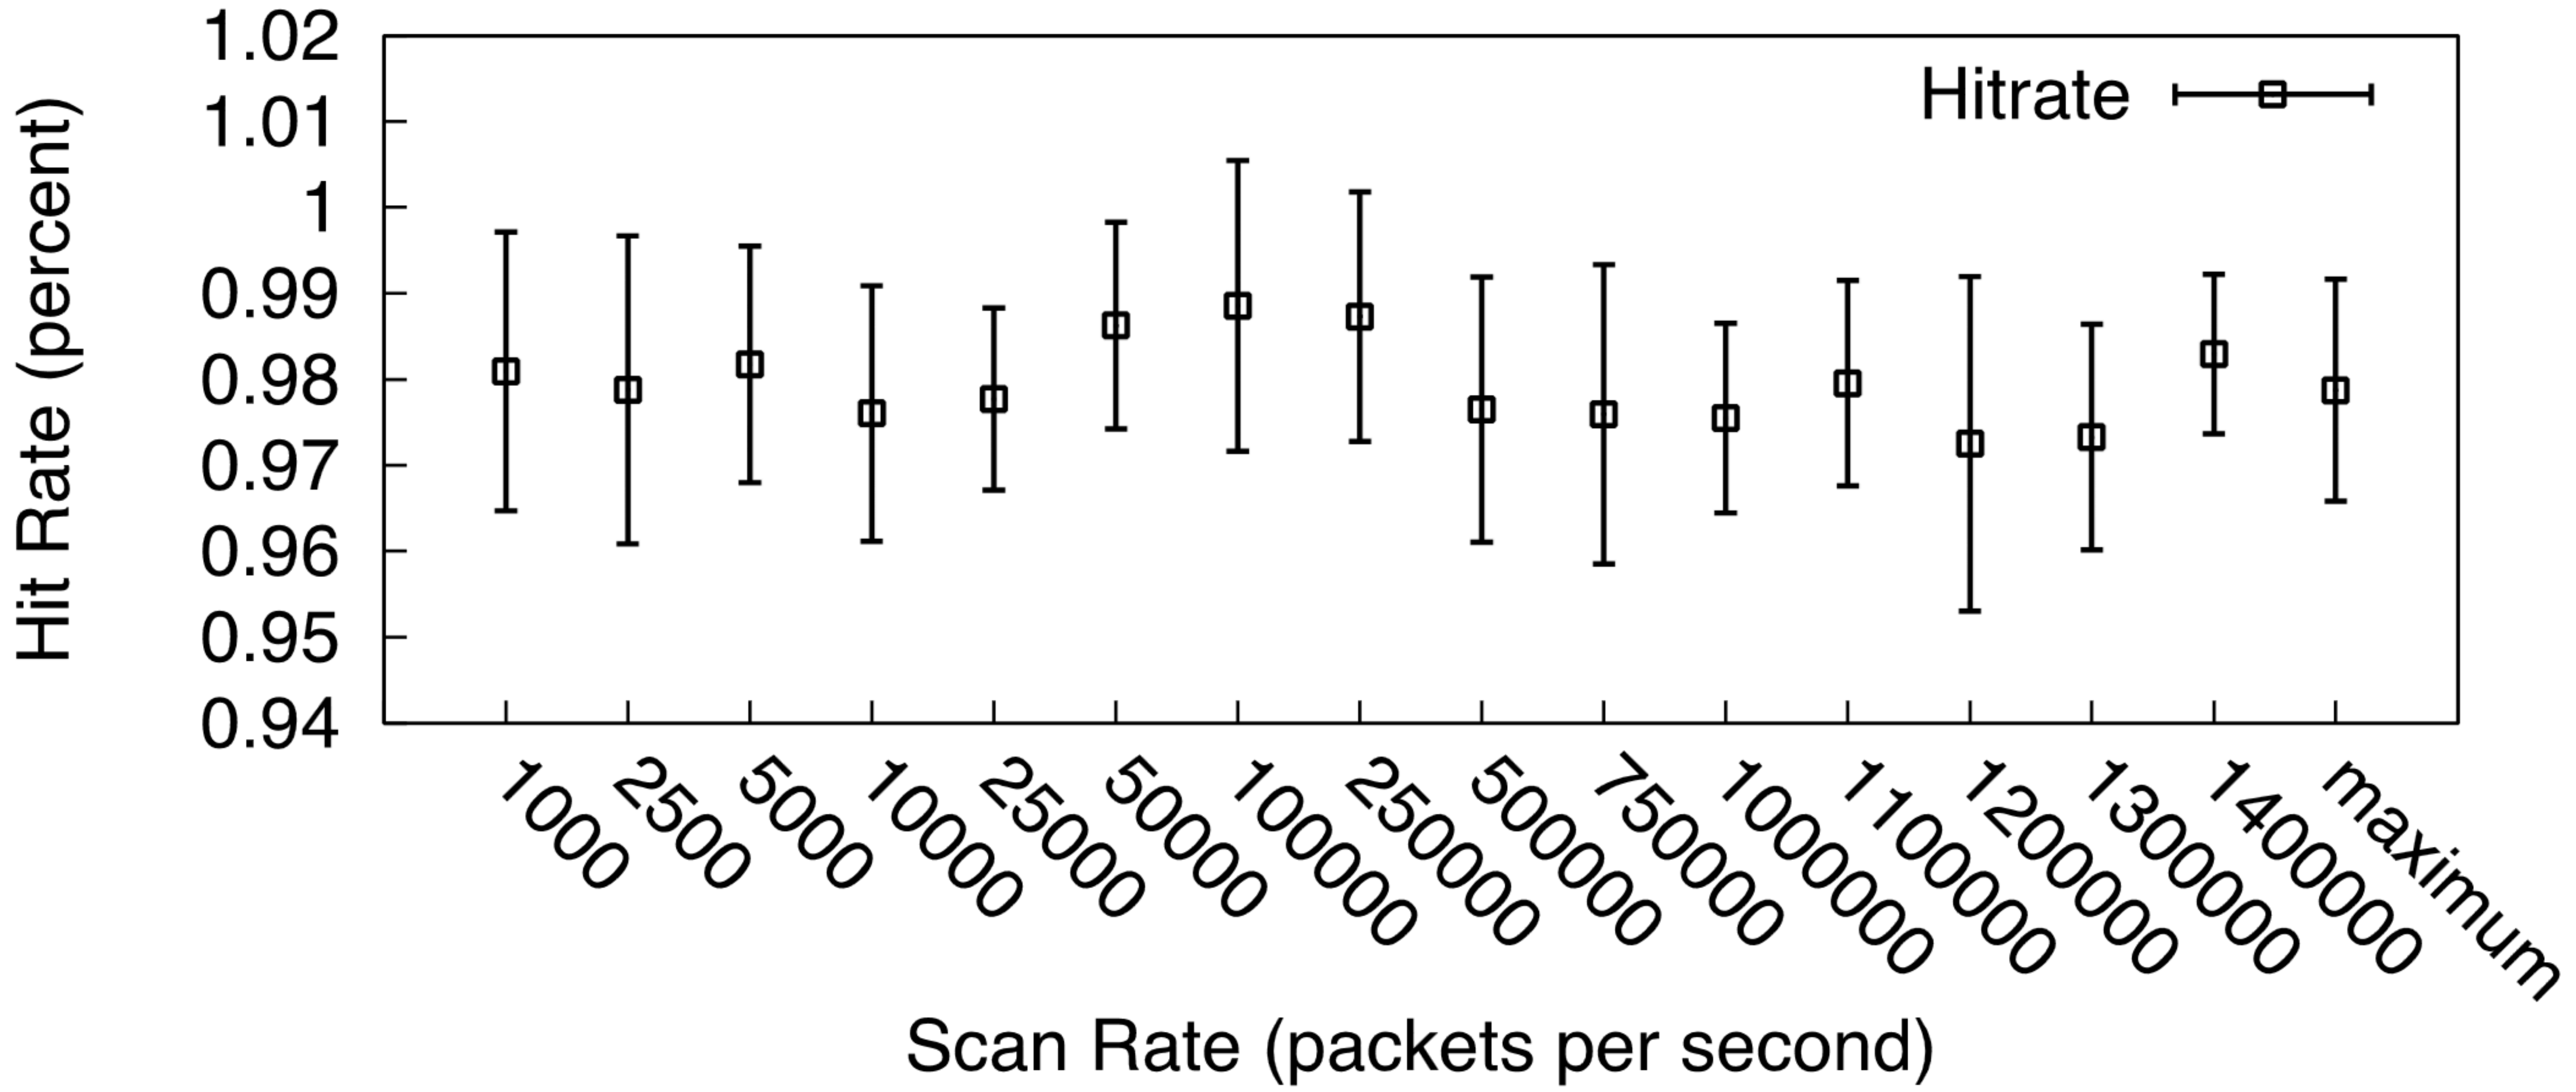

Oregon State University

Scanning Coverage

Is one probe packet per destination IP sufficient?

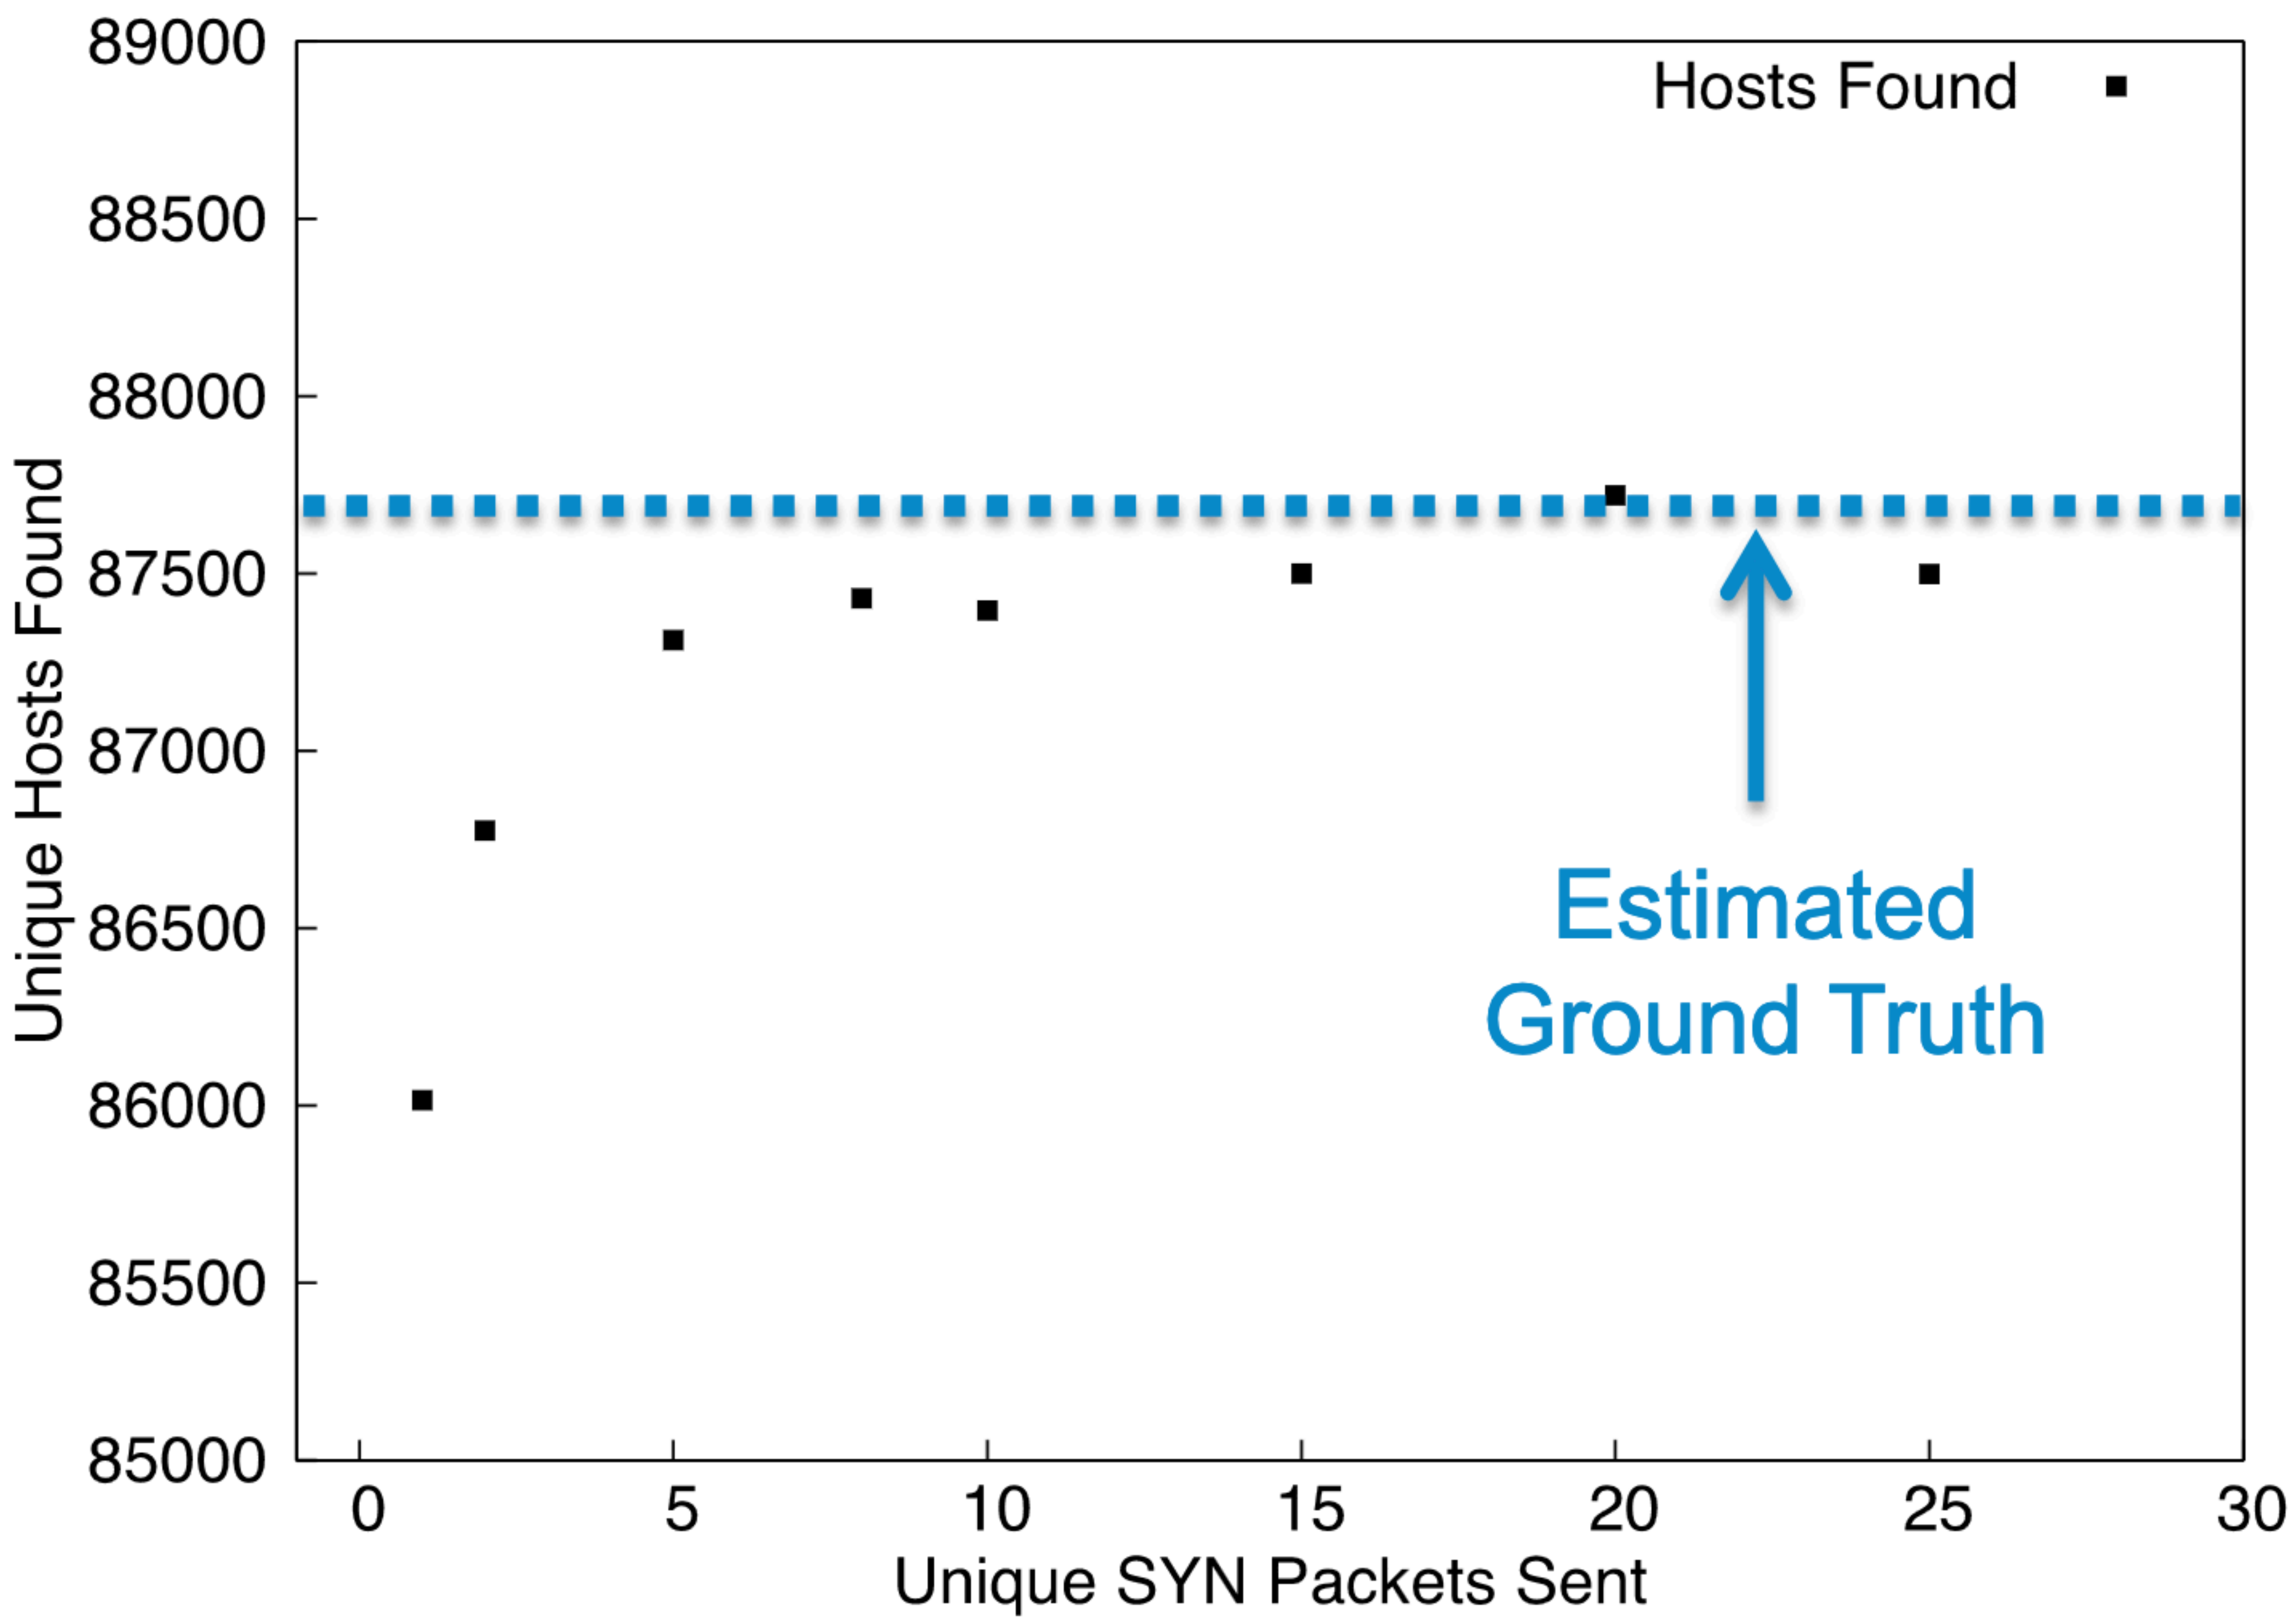

We expect an eventual plateau in responsive hosts, regardless of additional probes.

Scan Coverage

1 Packet:	97.9%
2 Packets:	98.8%
3 Packets:	99.4%

Comparison with Nmap

Scan of 1 million hosts

	Normalized Coverage	Duration (mm:ss)	Est. Internet Wide Scan
Nmap (1 probe)	81.4%	24:12	62.5 days
Nmap (2 probes)	97.8%	45:03	116.3 days
ZMap (1 probe)	98.7%	00:10	1:09:35
ZMap (2 probes)	100.0%	00:11	2:12:35

ZMap is capable of scanning more than 1300 times faster than the most aggressive Nmap default configuration (“insane”)

Surprisingly, ZMap also finds more results than Nmap

Probe Response Times

Why does ZMap find more hosts than Nmap?

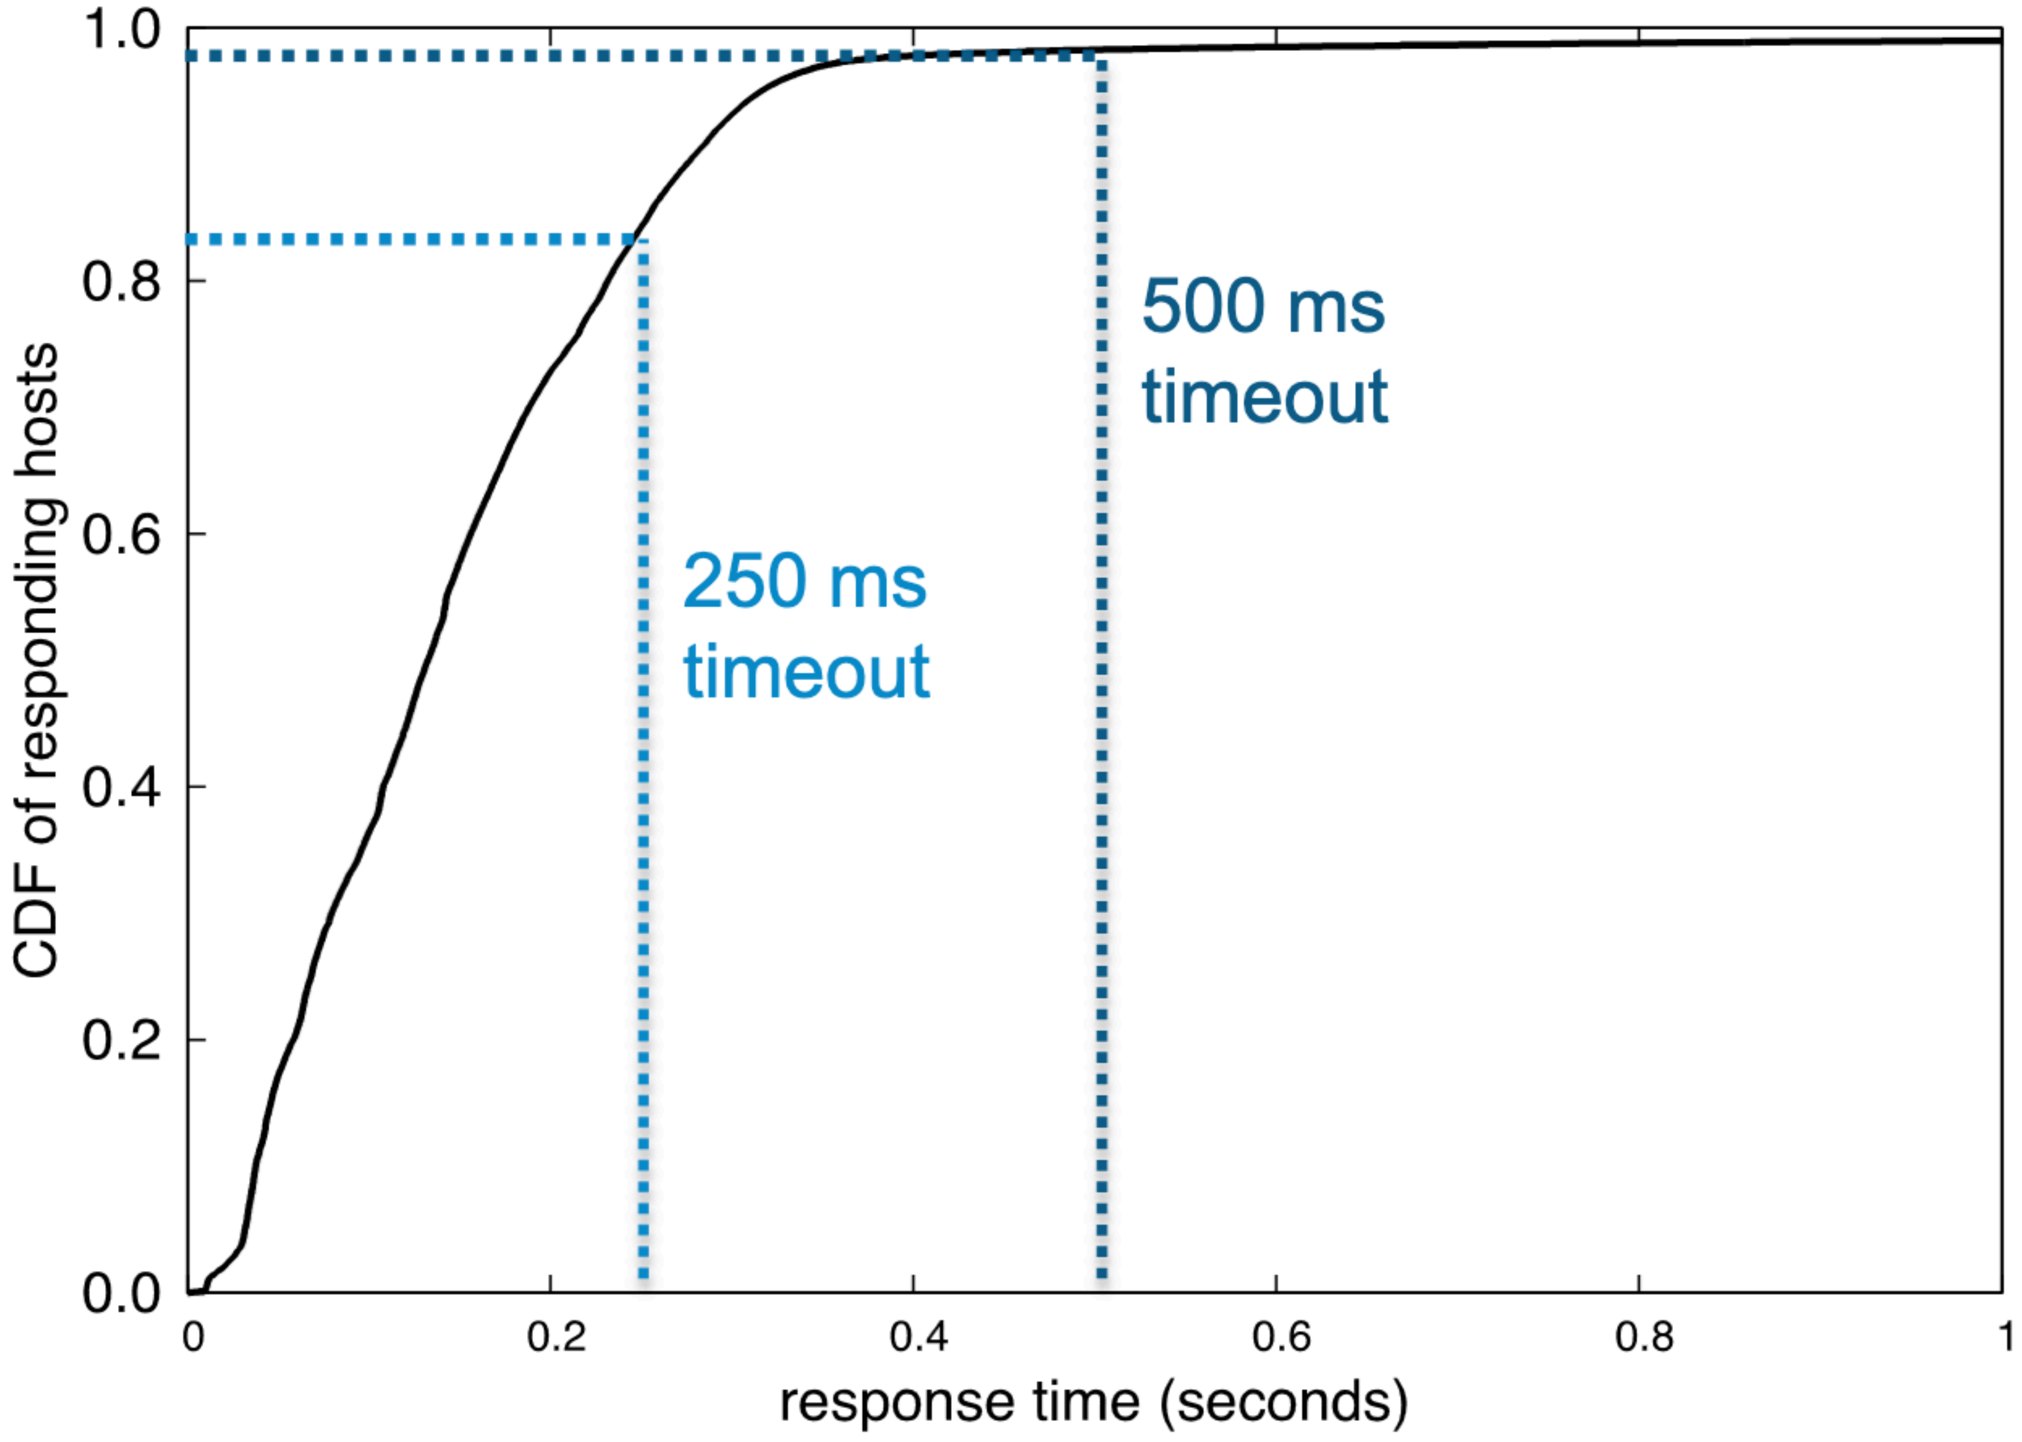

Response Times

250 ms:	**< 85%**
500 ms:	**98.2%**
1.0 s:	**99.0%**
8.2 s:	**99.9%**

Statelessness leads to both higher performance ***and*** increased coverage.

Ethics of Active Scanning

Ethics requires the balancing of harms with benefits

What are potential negative consequences of scanning? Potential mitigations?

- Overwhelming traffic that slows down / takes down network
  - Randomize / spread out probes to a given network
- Sysadmins believe they are under attack + waste resources responding
  - Signal benign nature over HTTP, reverse DNS entries
- Access or modify sensitive or private user data
  - Test locally beforehand; only collect what is needed; remove sensitive data
- Other unforeseen / unknown issues
  - Provide contact info and honor requests to be excluded from future scans

Meta: Do we need to scan the full internet?

- Depends what we are trying to find

When we don't need to scan everything

Determining what percent of websites use HTTPS

Collecting different types of phishing websites to categorize strategies

Make sure to get a random or representative sample!

When we do need to scan everything

Finding really rare (but possibly very impactful) phenomenon

Notifying insecure websites about how to patch vulnerabilities

When we don't feel like doing statistics

BREAK

Computer Security + Ethics

Computer Security + Ethics

- Computers: technology that can easily amplify benefits and harms
- Computer security: evaluation / prototyping of cyberattacks targeting important systems to access to sensitive information; privileged, abusable capabilities for defense
- Ethics is what separates security practitioners (white-hats) from cybercriminals (black hats)

Ethical Frameworks and Computer Security Trolley Problems: Foundations for Conversations

Tadayoshi Kohno
University of Washington

Yasemin Acar
George Washington University

Wulf Loh
Universität Tübingen

2023 USENIX

Scenario: Medical Device Vulnerability

- You are studying the security of a **wireless implantable medical device** – a device that is known to extend the lives of patients by at least 10 years
- You find a vulnerability that, **if exploited**, could cause **significant harm**
- The **company** that made the medical device **no longer exists** (it went bankrupt) ⇒ it is **impossible to patch** the vulnerability
- **Many patients** have the device in their bodies; the device is still being implanted in new patients
- You must choose between disclosing the vulnerability to **everyone** or **no one at all**
- The **likelihood** of an adversary **exploiting** the vulnerability is extremely **low** (**assume zero** for ease of analysis) regardless of whether or how you disclose the vulnerability

Scenario: Medical Device Vulnerability

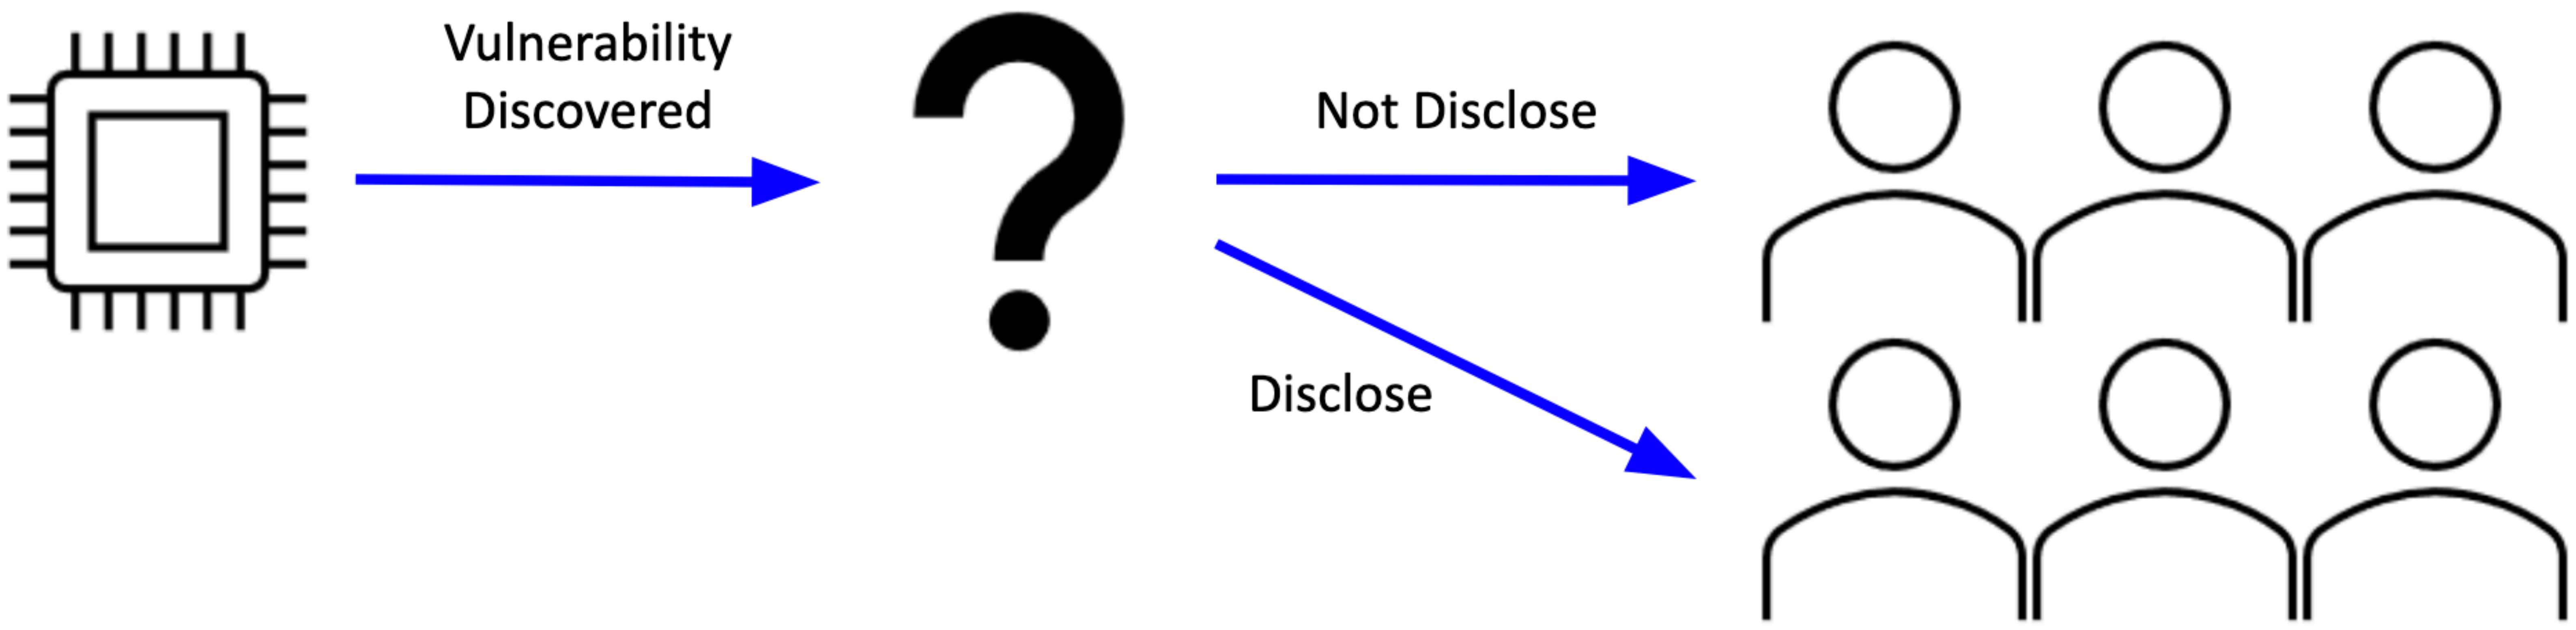

- **If not disclose**: **Patients** have **no awareness** that their device is vulnerable; **patients** keep and/or proceed with obtaining device and **receive** significant **health benefits**

- **If disclose**: **Patients** have the **choice** to not receive or to remove the device; **risk** of **psychological harm** if patients know they have a vulnerable device (even if chance of exploitation is zero); **risk** of **health harm** if patients do not receive / remove the device

Classic Dilemma: The Trolley Problem

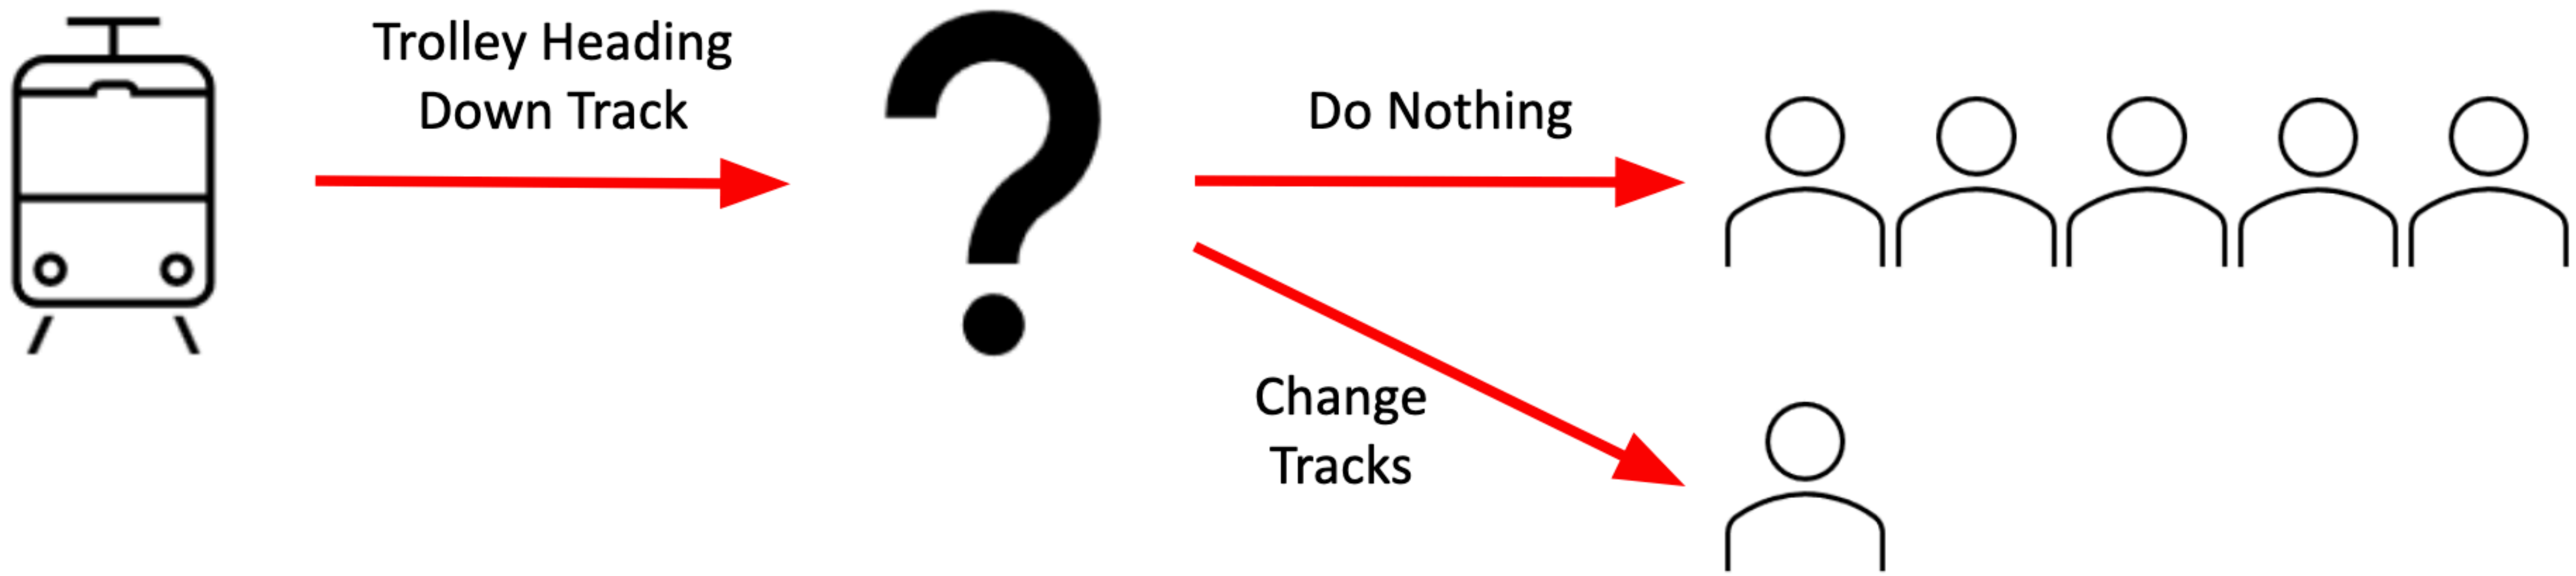

- A runaway trolley with no brakes is heading straight. **Five people** are tied to those tracks. **One person** is tied to an alternate set of tracks. A track operator observes this situation.
- **Should the track operator do nothing** (five people die) **or change the path** of the trolley (one person dies)?

Oregon State University

Consequentialist & Deontological Ethics

- **Consequentialist** and **deontological ethics** are two of today's most common ethical frameworks in computer security, can be found in:
  - Menlo Report: 17-page 2012 Dept. of Homeland Security report on ethical framework for research involving Information and Communications Technologies
  - Conference calls / ethics sections for research papers
- **These frameworks have limitations**, e.g., center **Western** approaches; there is no objectively "correct" framework
- It is not uncommon for people – including modern ethicists – to include elements of multiple frameworks as they reason through decisions

Consequentialist Ethics

- **Consequentialist ethics:** Focuses on **consequences** of actions, policies, institutions
- **Utilitarianism:** Example of consequentialism in which consequences are measured with respect to well-being
- Consequentialists **count numbers** and weigh **benefits / harms**

- **Example**: One death is better than five —> change the trolley's tracks

Deontological Ethics

- **Deontological ethics**: **People** have **fundamental rights**; moral actors have a **duty** to **respect** those **rights**

- Example rights: The **right** to **privacy**, the right to **self-agency**, the right to **informed consent**

- **Kantian deontological ethics**: One should **not violate any single person's rights in order to accomplish another objective**; human beings should be treated as "ends and never purely as means"

- **Example**: **Changing** the trolley **tracks** would **violate one person's right** (their right to live) **in order to accomplish** the **saving** of **five other** lives; changing the track would **use** that **single person** as an "**means**", **not** as an "**ends**"; under Kantian deontological ethics → **do not change the trolley's tracks**

Scenario: Medical Device Vulnerability

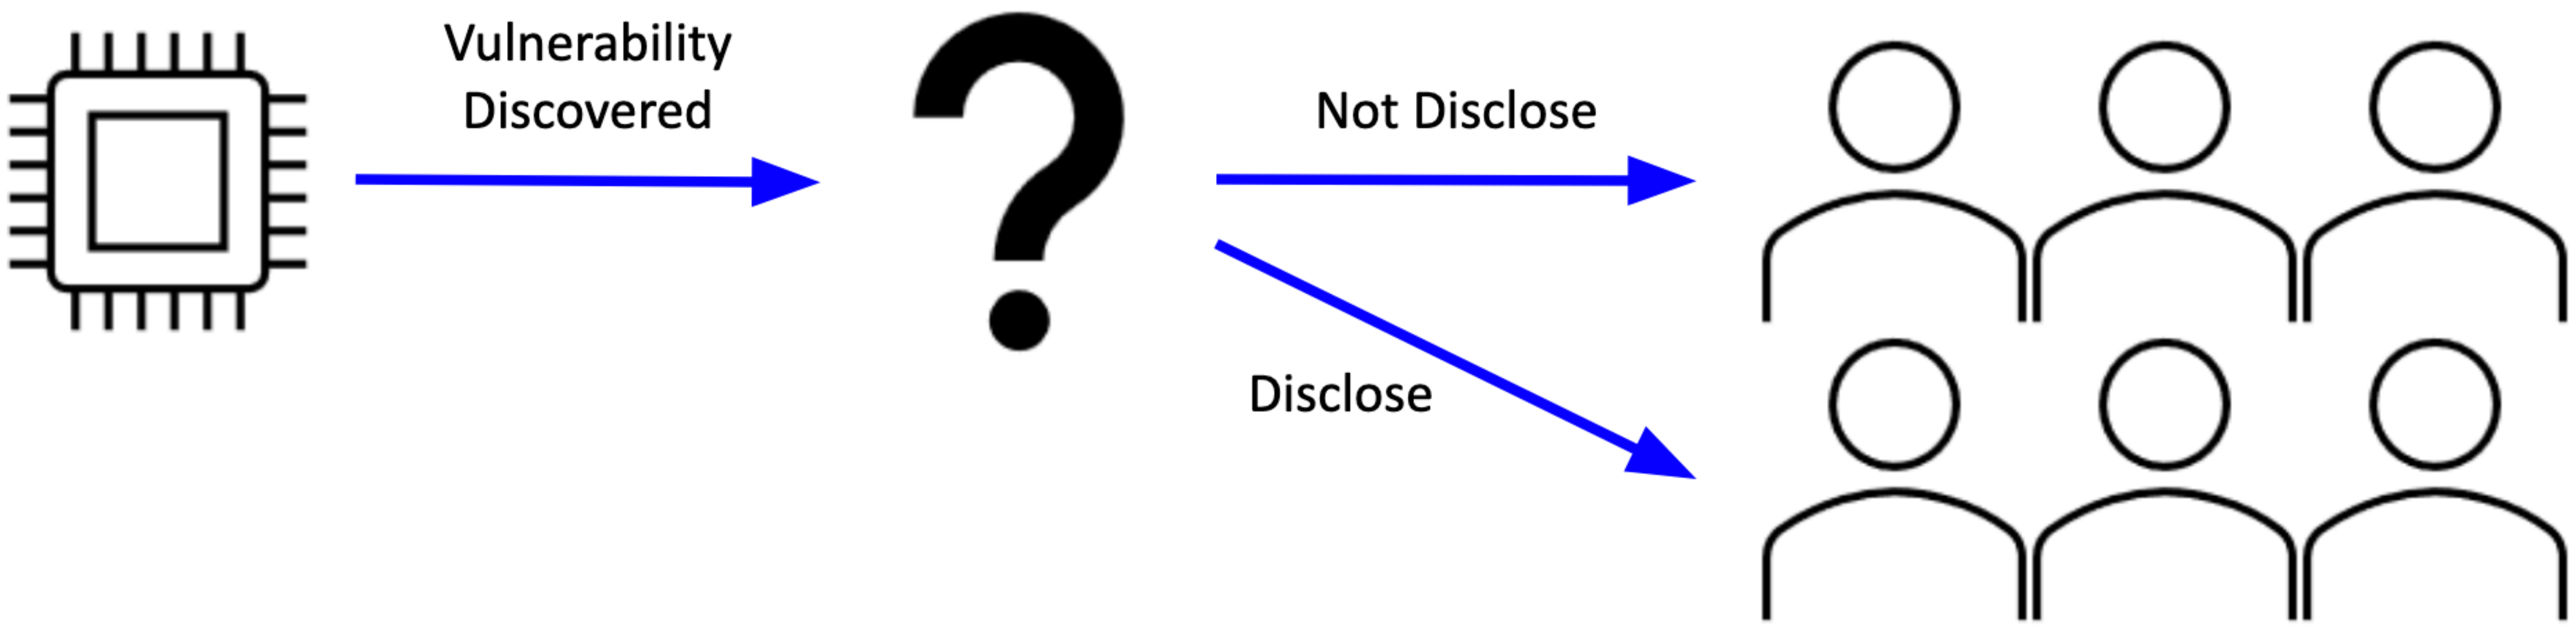

- **If not disclose**: **Patients** have **no awareness** that their device is vulnerable; **patients** keep and/or proceed with obtaining device and **receive** significant **health benefits**

- **If disclose**: **Patients** have the **choice** to not receive or to remove the device; **risk** of **psychological harm** if patients know they have a vulnerable device (even if chance of exploitation is zero); **risk** of **health harm** if patients do not receive / remove the device

Scenario: Medical Device Vulnerability

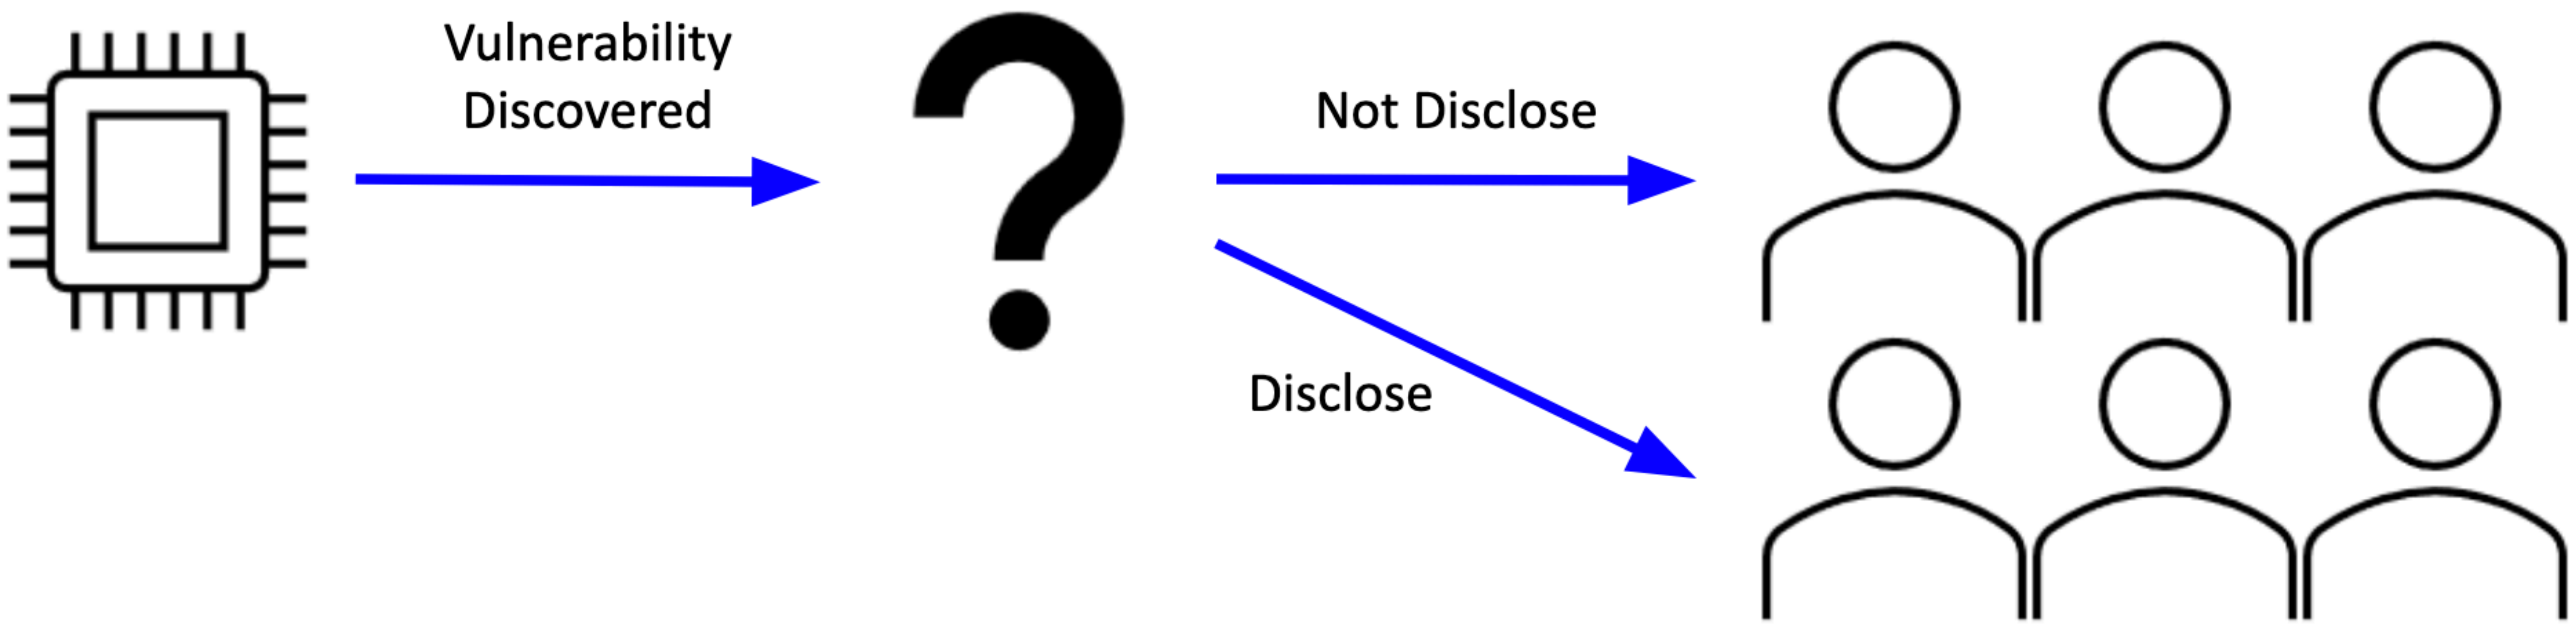

- **Consequentialist** Ethics: **Likelihood** of **exploit** is **zero**; **harms if patients informed** (health: remove device / not get device; happiness: live with knowledge that the device has faults) → **do not disclose vulnerability**

- **Deontological** Ethics: Duty to respect people's **right** to **informed consent** (e.g., warnings on medicine ads) and right to **self-agency** (make their own decisions about what is best for them) → **disclose vulnerability**

Ethical Takeaways

- Different ethical frameworks can lead to different / same conclusion; or can lead to no conclusion
- Deciding what ethical framework to use is a personal choice; however, decision makers should not pick a decision and find the framework that justifies it
- Sometimes the morally correct action is not in the best interest of the decision maker
- Ethical frameworks can provide tools for discussion and help ensure that everyone is speaking the same language
- Historically, security community has adopted a blend of consequentialist / deontological ethics

BLOOKET

Machine Learning

- Step 1: Collect lots of data
- Step 2: Analyze data to see current state of security
- Step 3: Use ML for prediction: perform attacks, automate defenses, etc.
- Step 4: …

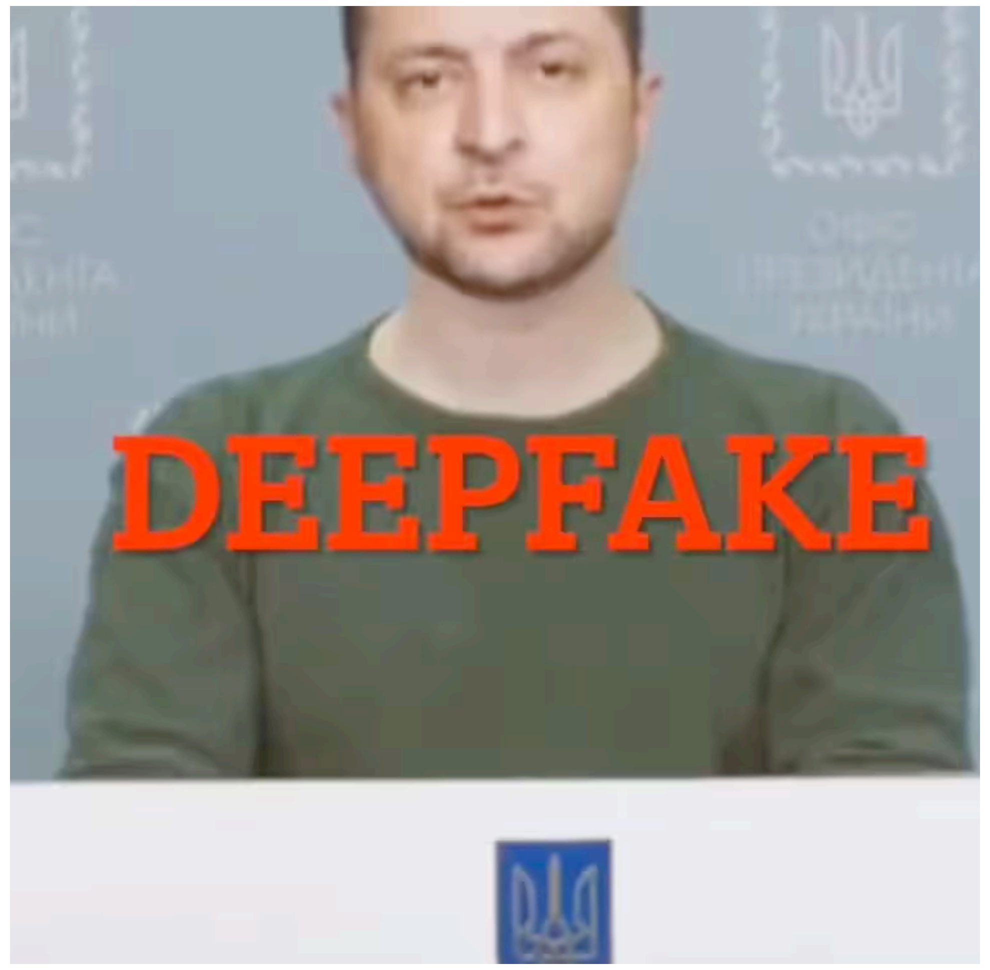

Dos and Don’ts of Machine Learning in Computer Security

Daniel Arp (Technische Universität Berlin) et al.

2022 USENIX

Machine Learning Workflow

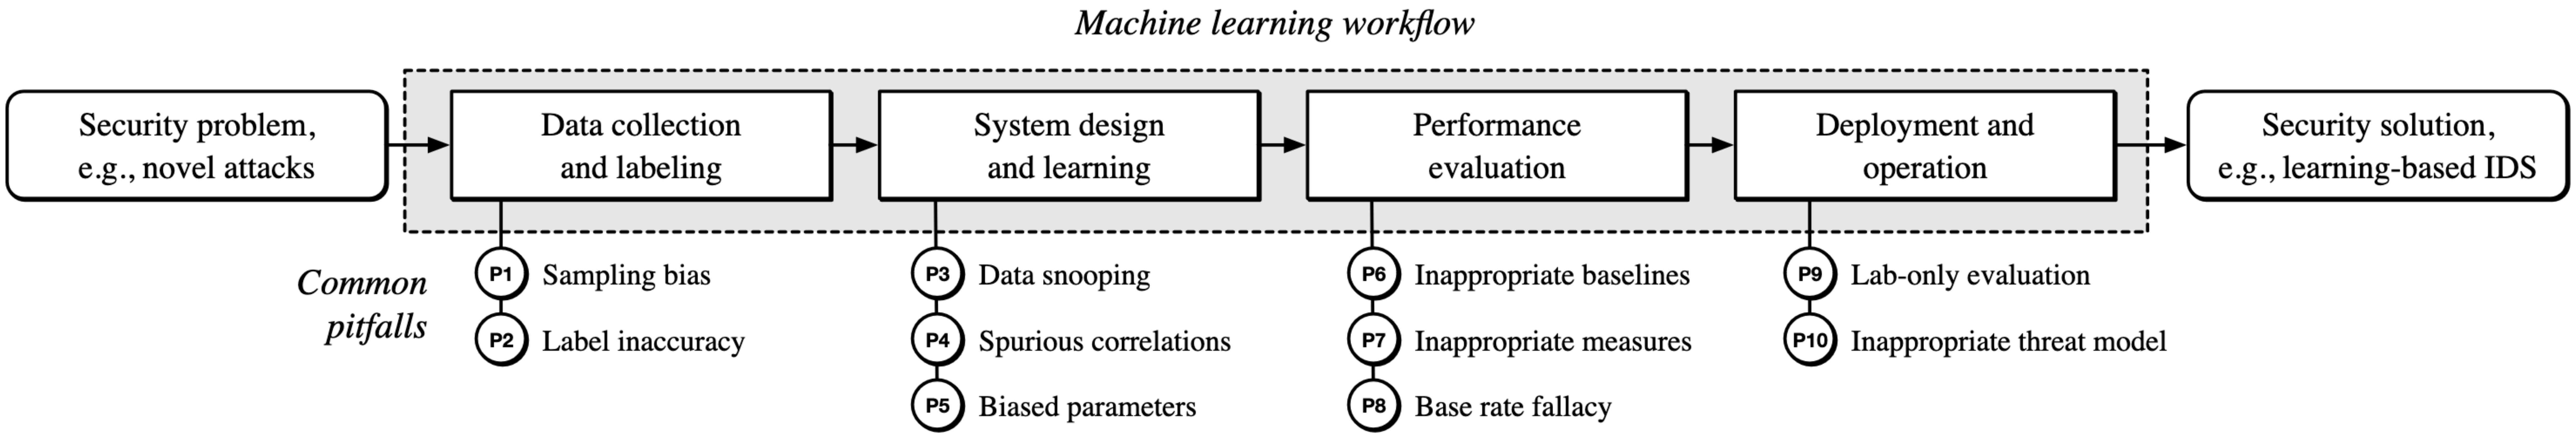

Machine Learning Flaws

Measured 30 top security papers

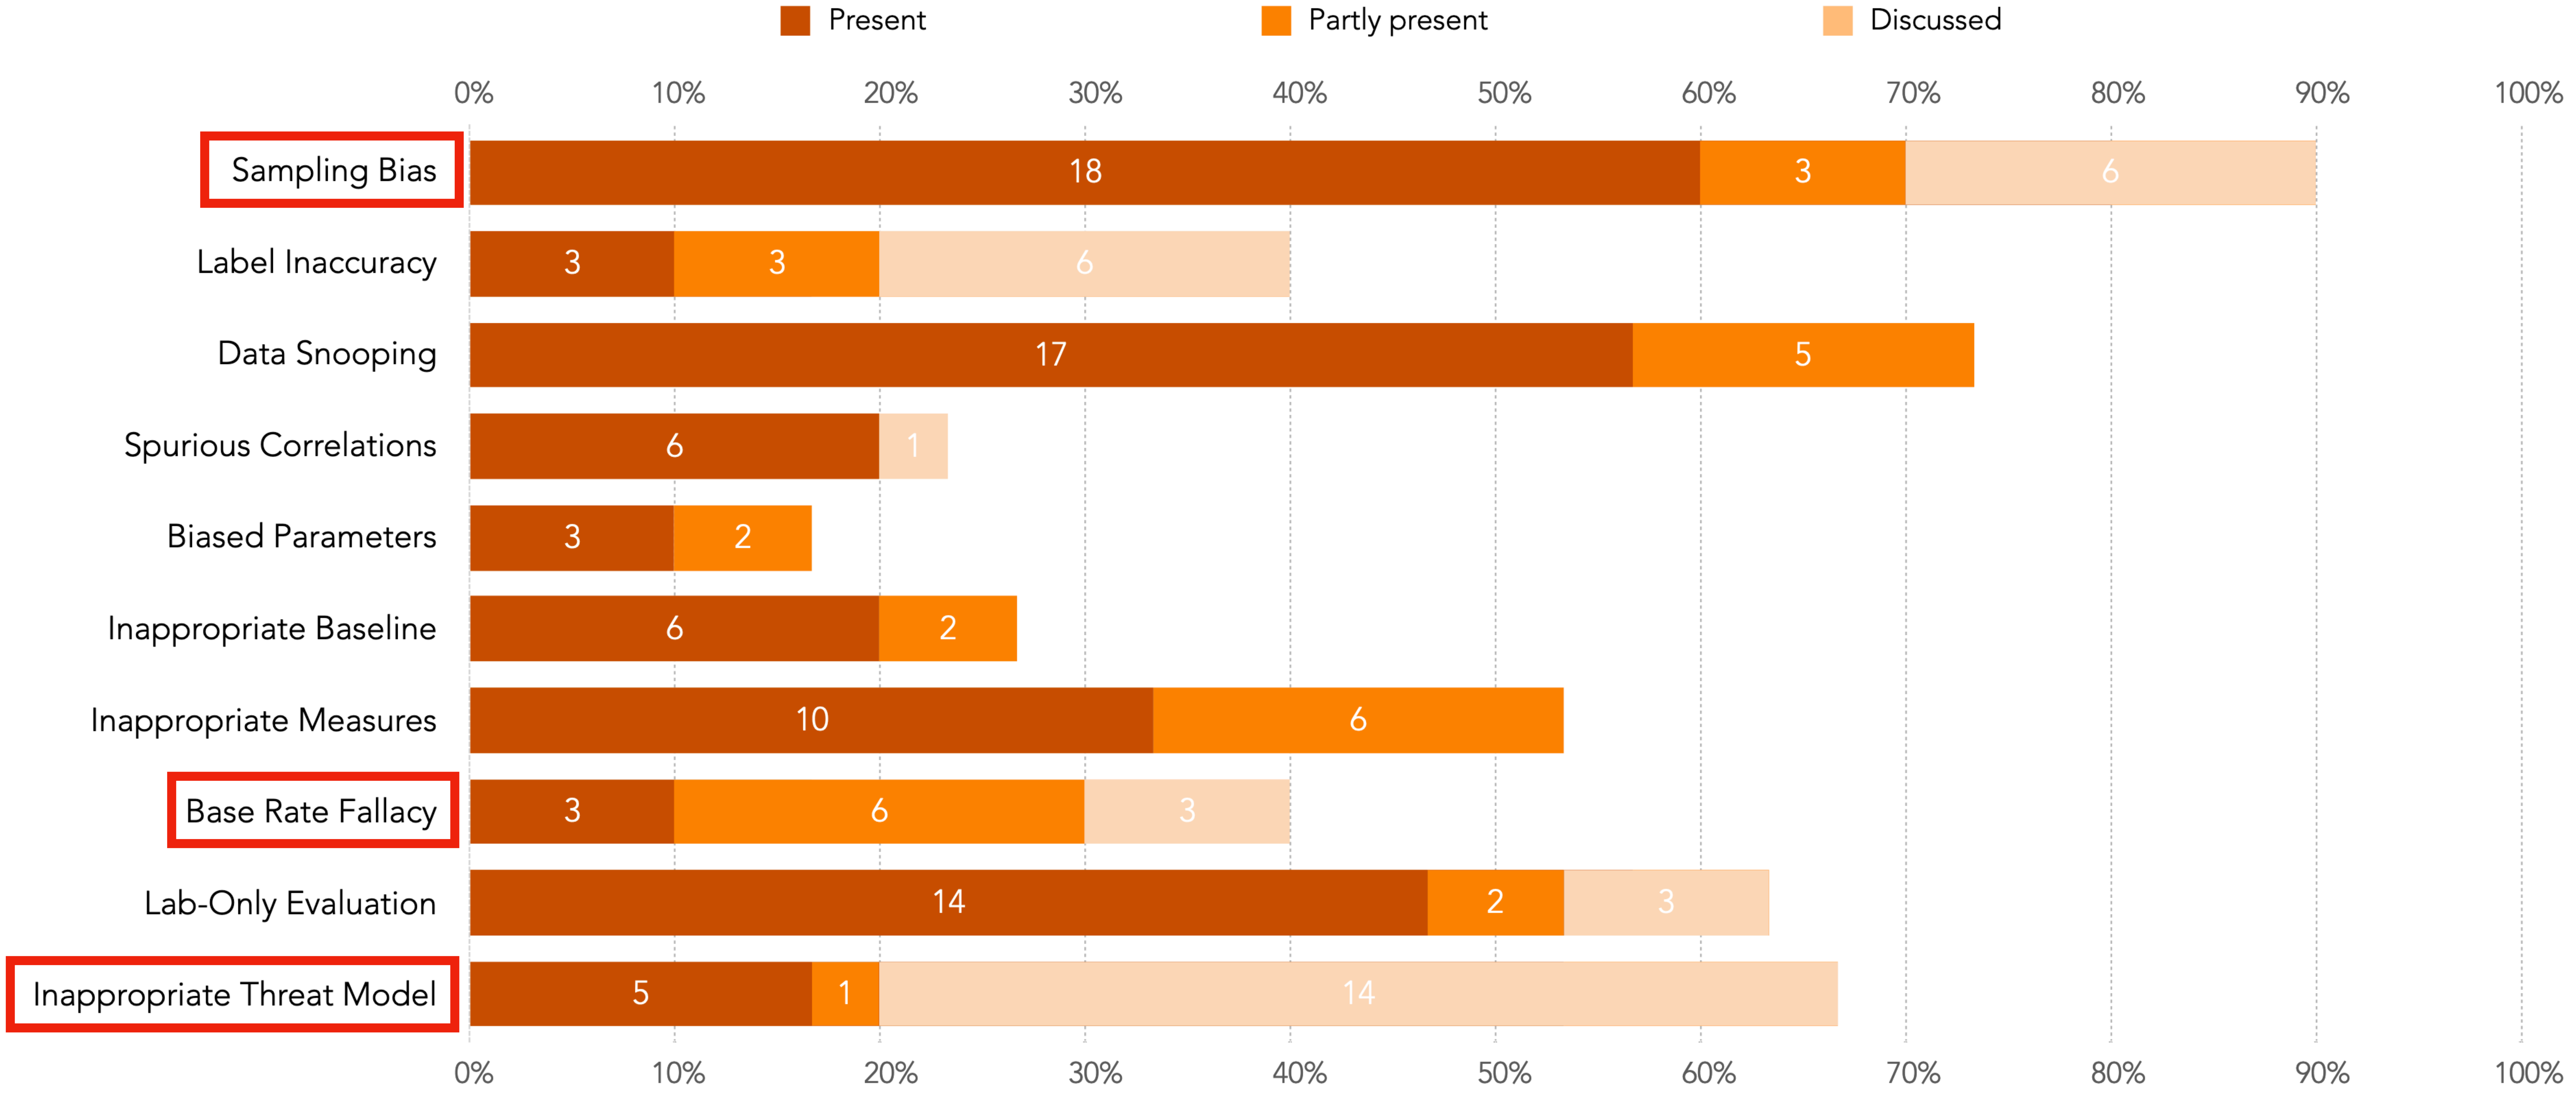

Sampling Bias

"The collected data does not sufficiently represent the true data distribution of the underlying security problem"

When the training data for a model does not represent the intended use case

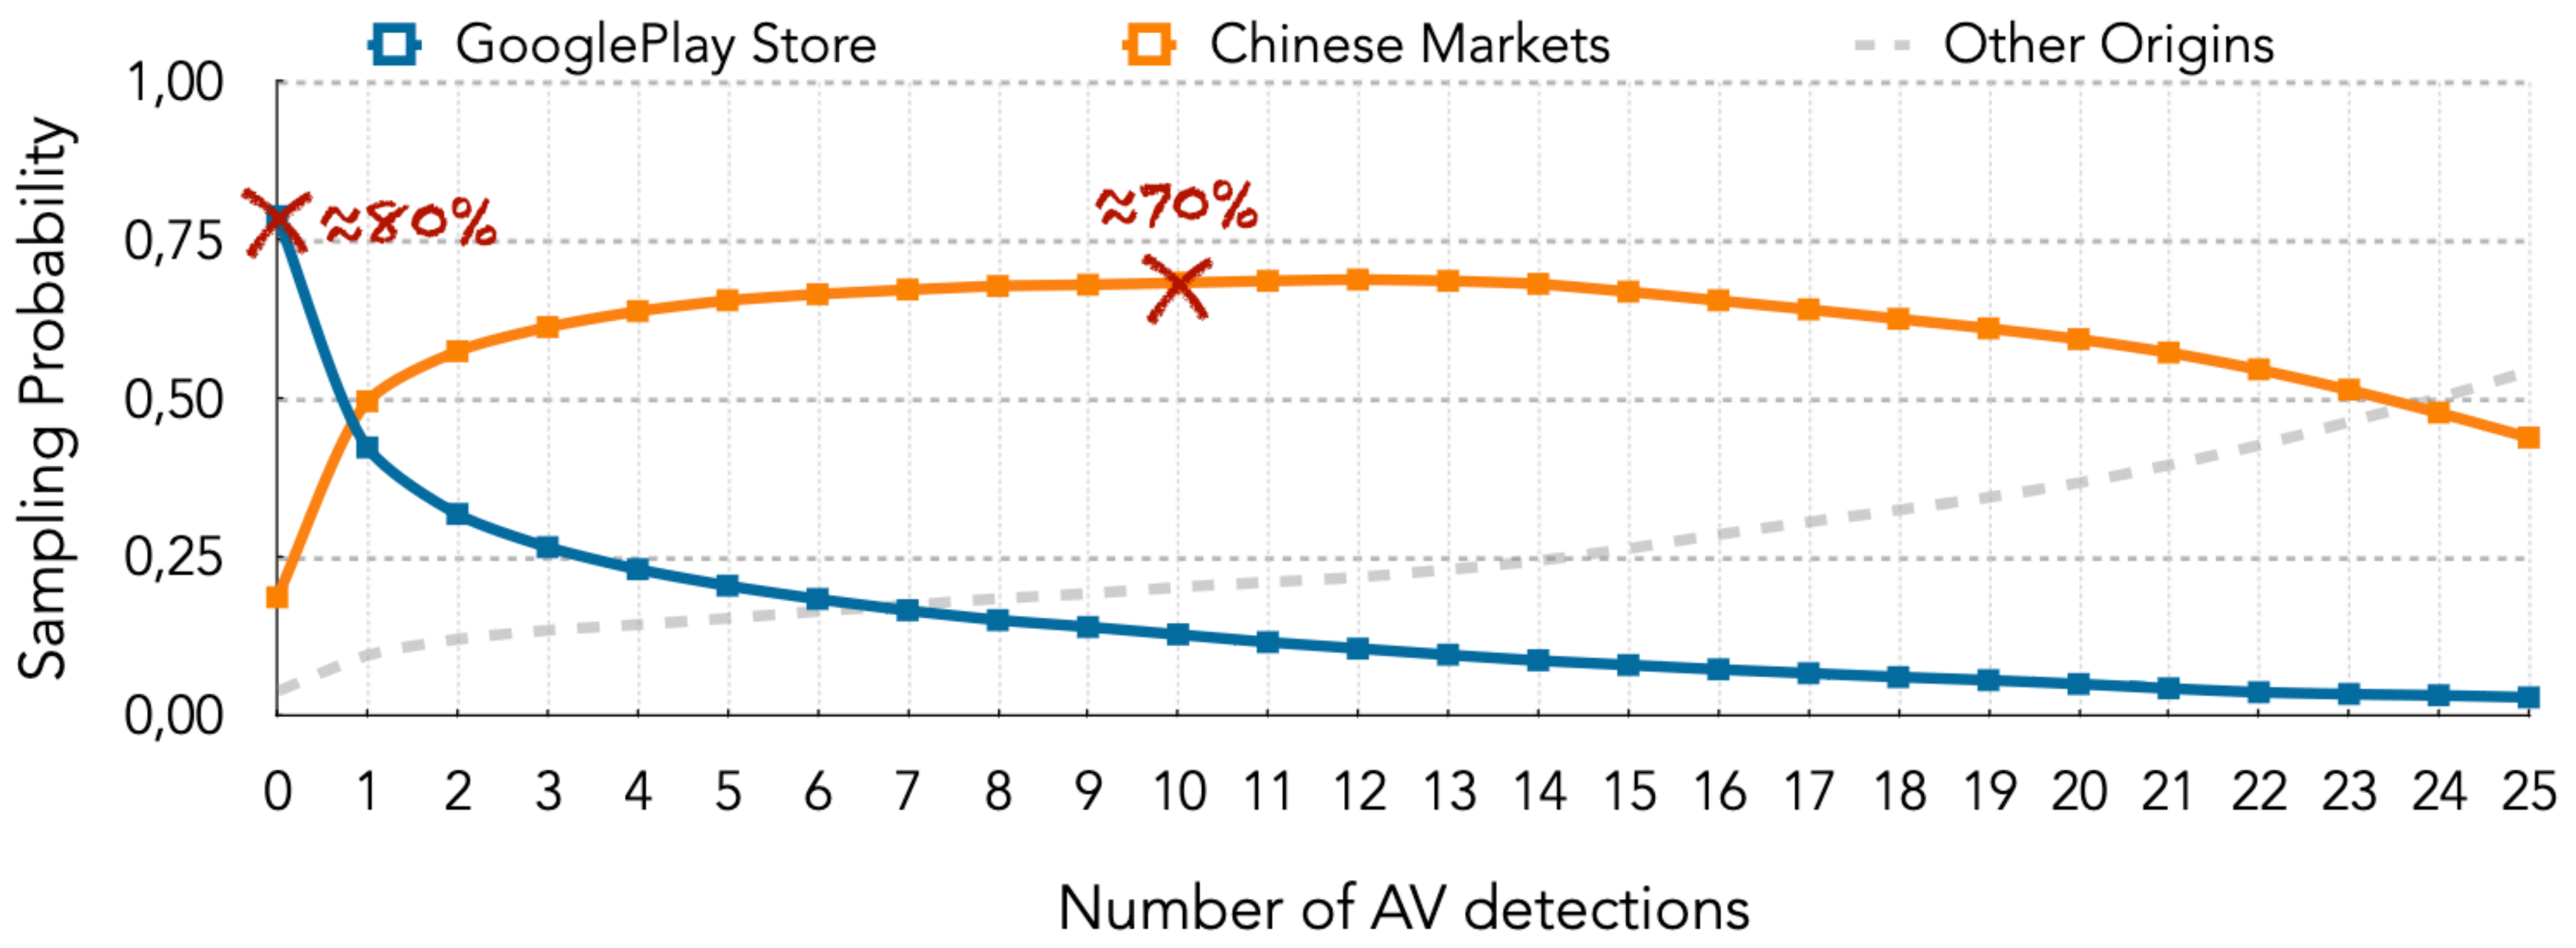

How should we collect benign (0 AV detection) and malicious (10+ AV detections) datasets?

Sampling Bias

What prior study did: randomly sample from all benign apps and all malicious apps to generate training / test data

Outcome: the URL "play.google.com" is one of the top distinguishing features for malware detection (Problem #4: Spurious correlations)

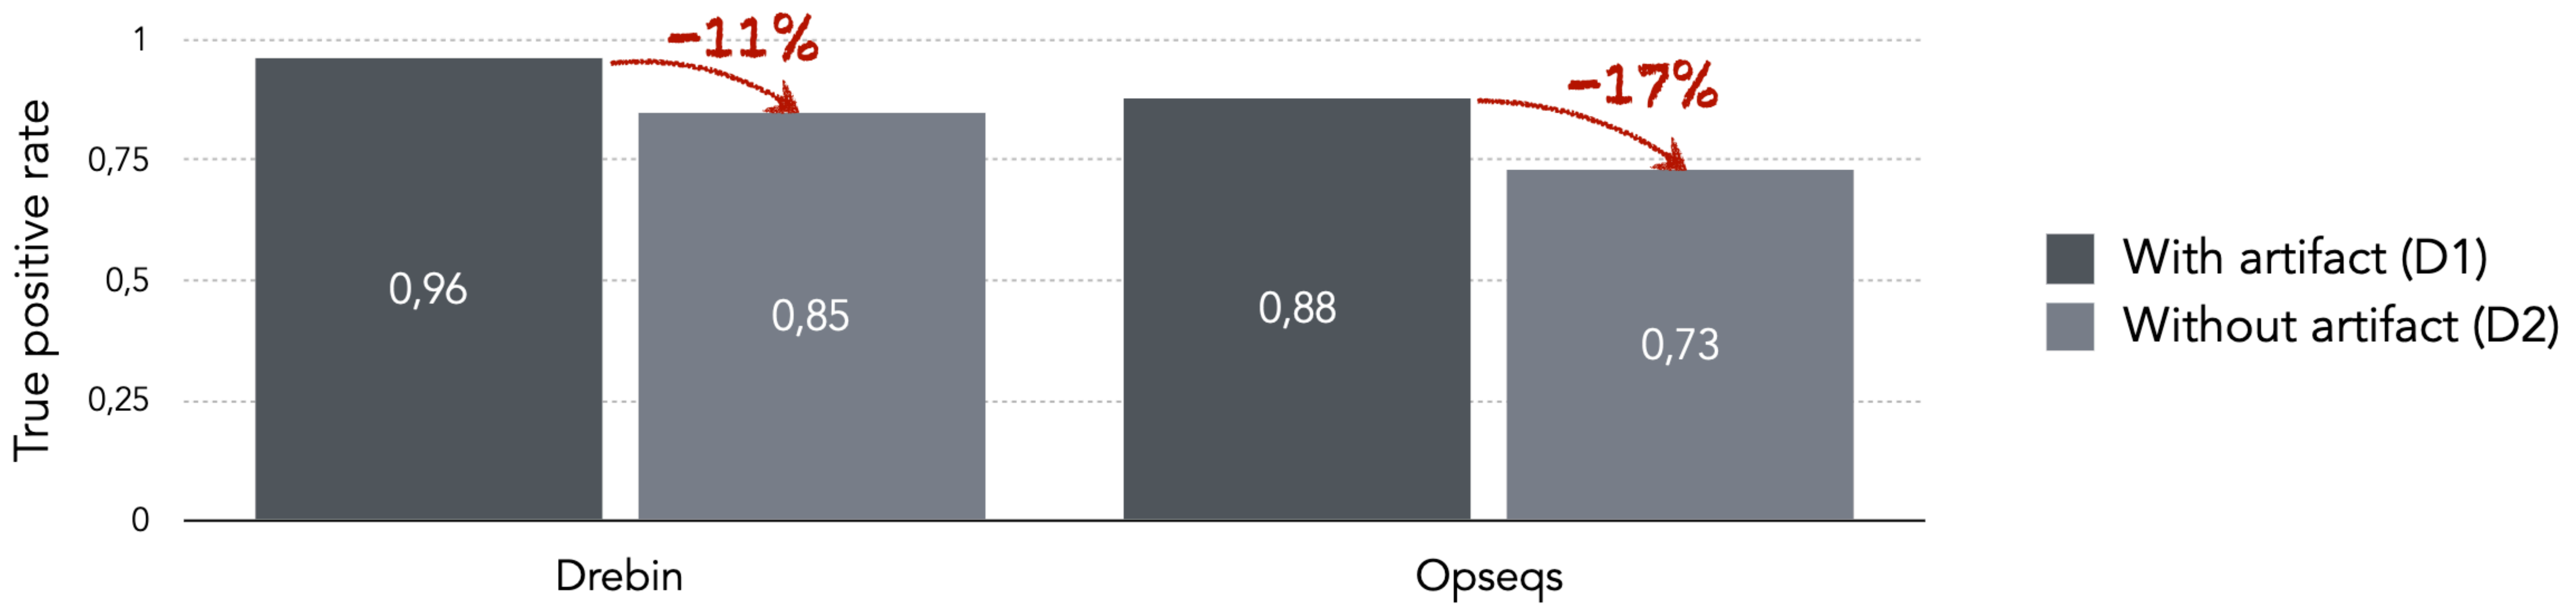

Base rate fallacy

Assume: medical test with 5% false positive rate and no false negative rate

How good is this test when the base rate of infection in the population is 40%?

400 infected / 430 positive = 93% confident

Number of people	Infected	Uninfected	Total
Test positive	*400* (true positive)	*30* (false positive)	430
Test negative	0 (false negative)	570 (true negative)	570
Total	400	600	**1000**

How good is this test when the base rate of infection in the population is 2%?

20 infected / 69 positive = 29% confident

Number of people	Infected	Uninfected	Total
Test positive	*20* (true positive)	*49* (false positive)	69
Test negative	0 (false negative)	931 (true negative)	931
Total	20	980	**1000**

https://en.wikipedia.org/wiki/Base_rate_fallacy

Base rate fallacy

A tendency to ignore the base rate (across a full population) in favor of the accuracy of an individual test

Takeaway: Low positive rate (FPR) is super critical for security systems that handle large amounts of data, especially when base rate is relatively low (e.g., malicious network packets, APT detection)

Also when cost of false positive is high! For example, blocking a legitimate email, or requiring manual analysis of a (not-actually) malicious network signal

https://en.wikipedia.org/wiki/Base_rate_fallacy

Improper threat model

Building a ML model is not enough to counter a threat - it's possible, often trivial, to break machine learning models.

Example: model for code authorship, 95% accuracy - can reveal relationships between malware, potential cheating / copying for assignments

Attack: removing unused code decreased code attribution accuracy by 48%

How to mitigate? Think like an attacker! Take Prof. Sanghyun Hong's class, CS499/579, AI539 :: Trustworthy Machine Learning

Recap

Measurement + empirical research is tricky to do correctly!

1. Methodology can require careful design and evaluation - ZMap

2. Ethical considerations are essential, but sometimes subjective - frameworks can facilitate discussion

3. Analyzing data with Machine Learning is fraught with many pitfalls - important to follow best practices, when possible

TODOs for you

Specify presentation preferences by **9PM tonight**. Sign-up link on the syllabus at https://empirical-security.net/syllabus

> I will send out presentation + reading (which 1 of the 2 papers to read for each class) assignments tomorrow morning on Canvas

First paper reading + questions will be due by 6PM **Tuesday, October 8th**.

Create a project team by **Friday, October 4th**. Reach out if you need help