THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Is It Science Or Engineering? A Sampling of Recent Research

Michael Reiter

Lawrence M. Slifkin Distinguished Professor
Department of Computer Science
UNC Chapel Hill

Engineering and Science

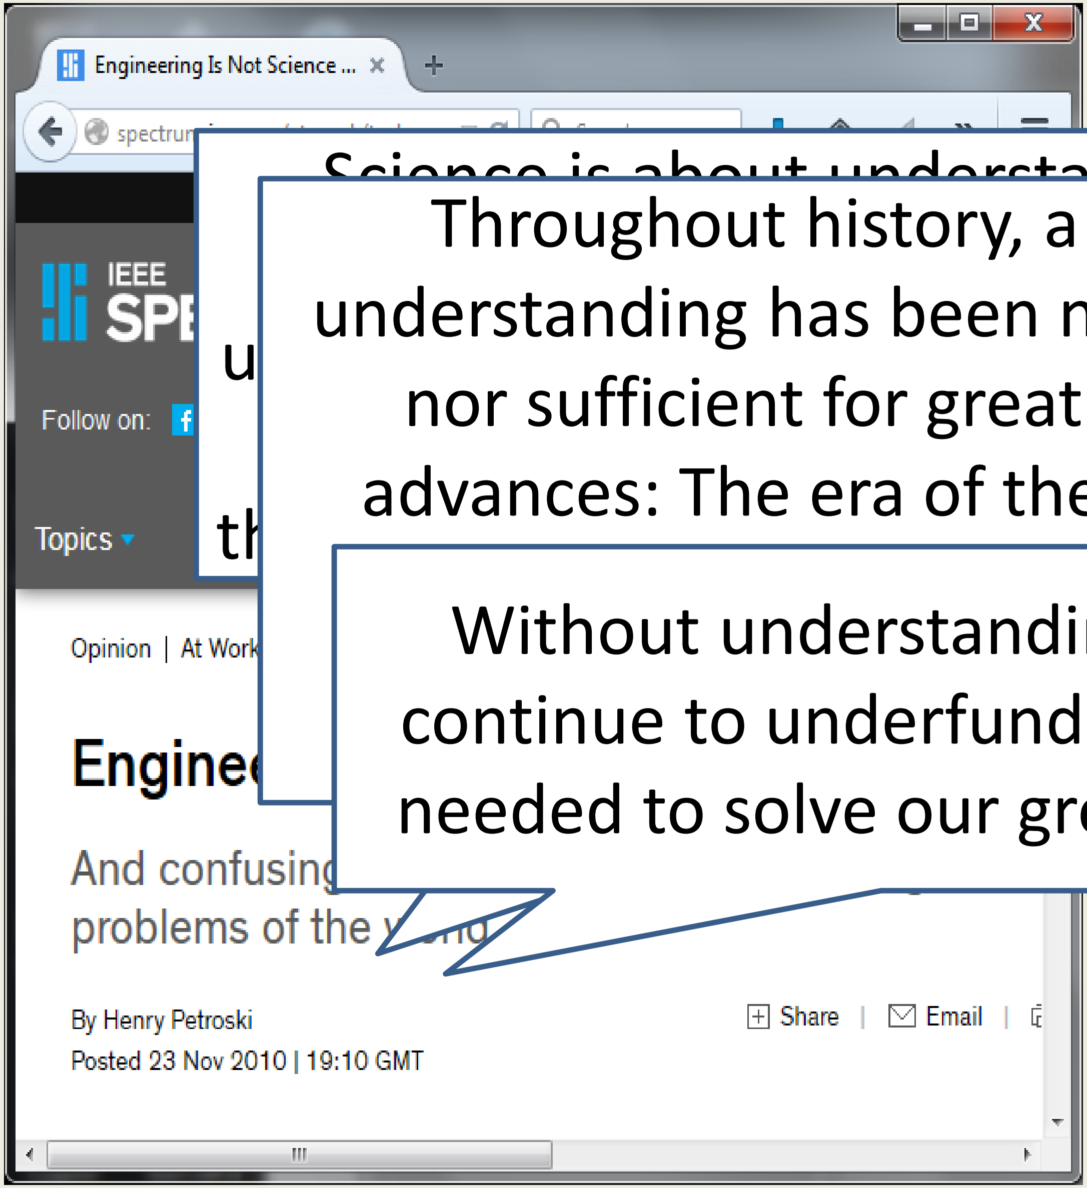

Engineering Is Not Science ...

spectrum

IEEE SPE

Follow on:

Topics

Opinion | At Work

Enginee

And confusing

problems of the w

By Henry Petroski

Posted 23 Nov 2010 | 19:10 GMT

Science is about understanding the

Throughout history, a full scientific understanding has been neither necessary nor sufficient for great technological advances: The era of the steam engine,

Without understanding this, we will continue to underfund the engineering needed to solve our greatest problems.

Engineering and Science

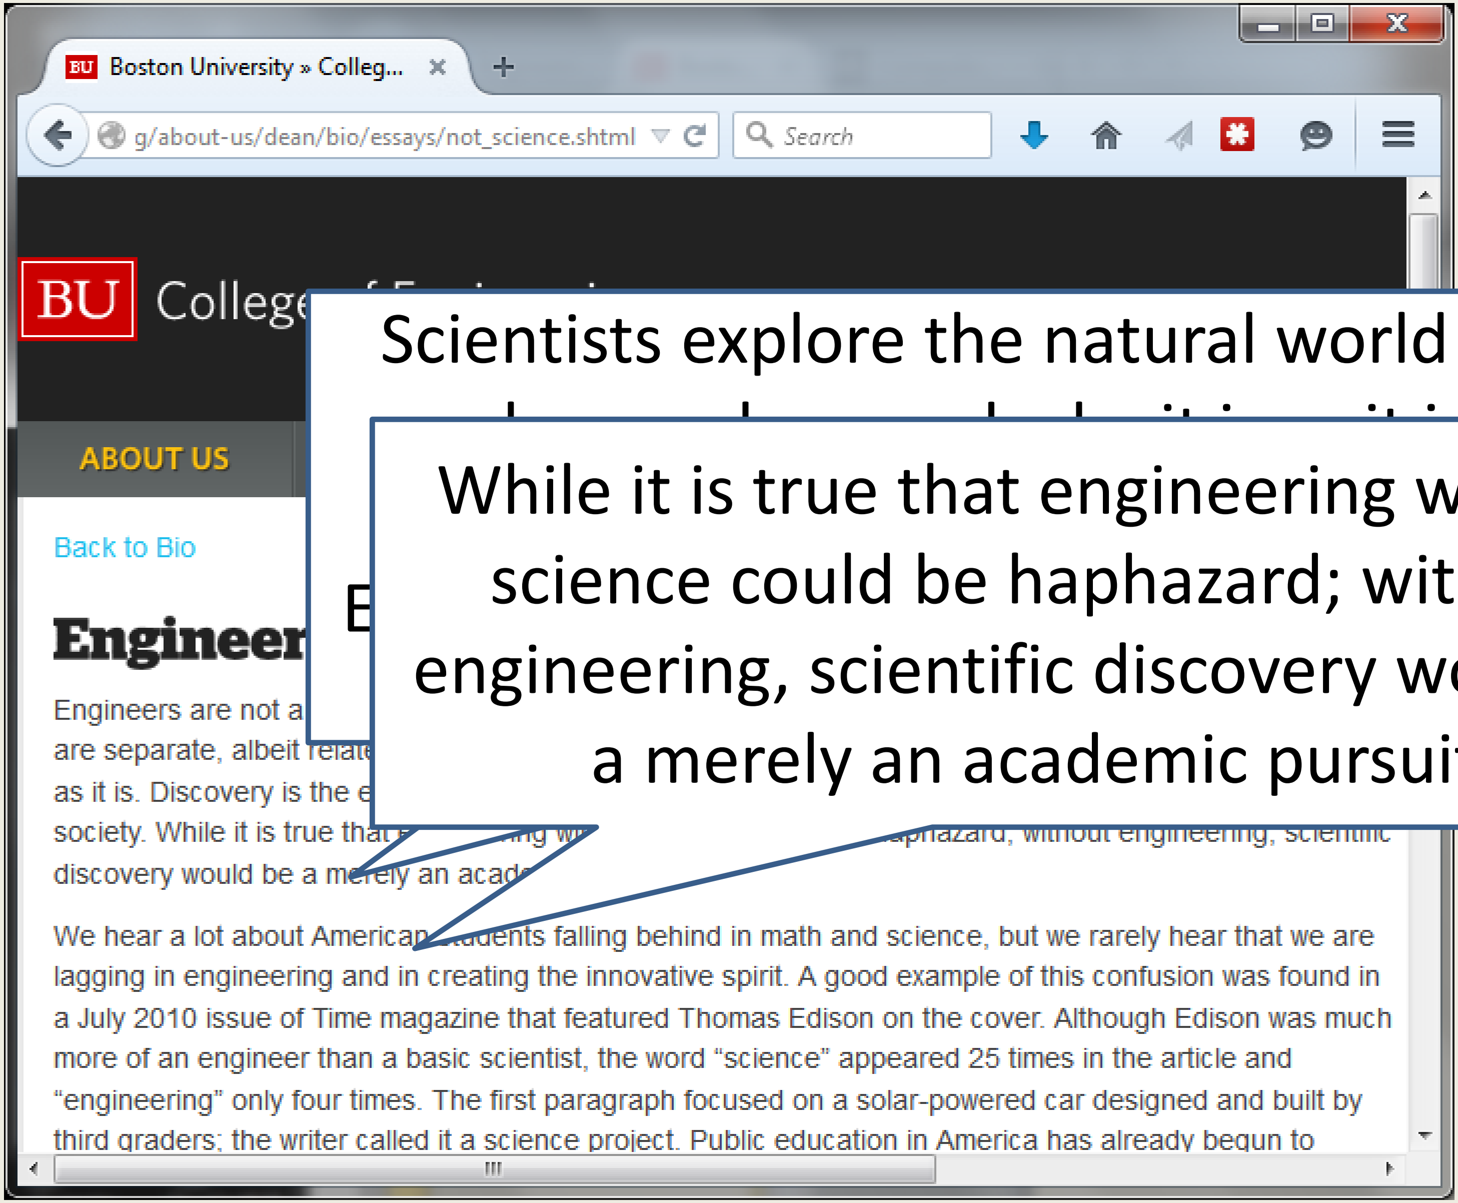
Boston University » Colleg...
g/about-us/dean/bio/essays/not_science.shtml
Search
BU
College
ABOUT US
Back to Bio

Engineer

Engineers are not a

are separate, albeit relate

as it is. Discovery is the e

society. While it is true that e ... ing wi ... aphazard, without engineering, scientific

discovery would be a merely an acad

We hear a lot about American students falling behind in math and science, but we rarely hear that we are lagging in engineering and in creating the innovative spirit. A good example of this confusion was found in a July 2010 issue of Time magazine that featured Thomas Edison on the cover. Although Edison was much more of an engineer than a basic scientist, the word “science” appeared 25 times in the article and “engineering” only four times. The first paragraph focused on a solar-powered car designed and built by third graders; the writer called it a science project. Public education in America has already begun to

Scientists explore the natural world and

E

While it is true that engineering without science could be haphazard; without engineering, scientific discovery would be a merely an academic pursuit.

Science of Security?

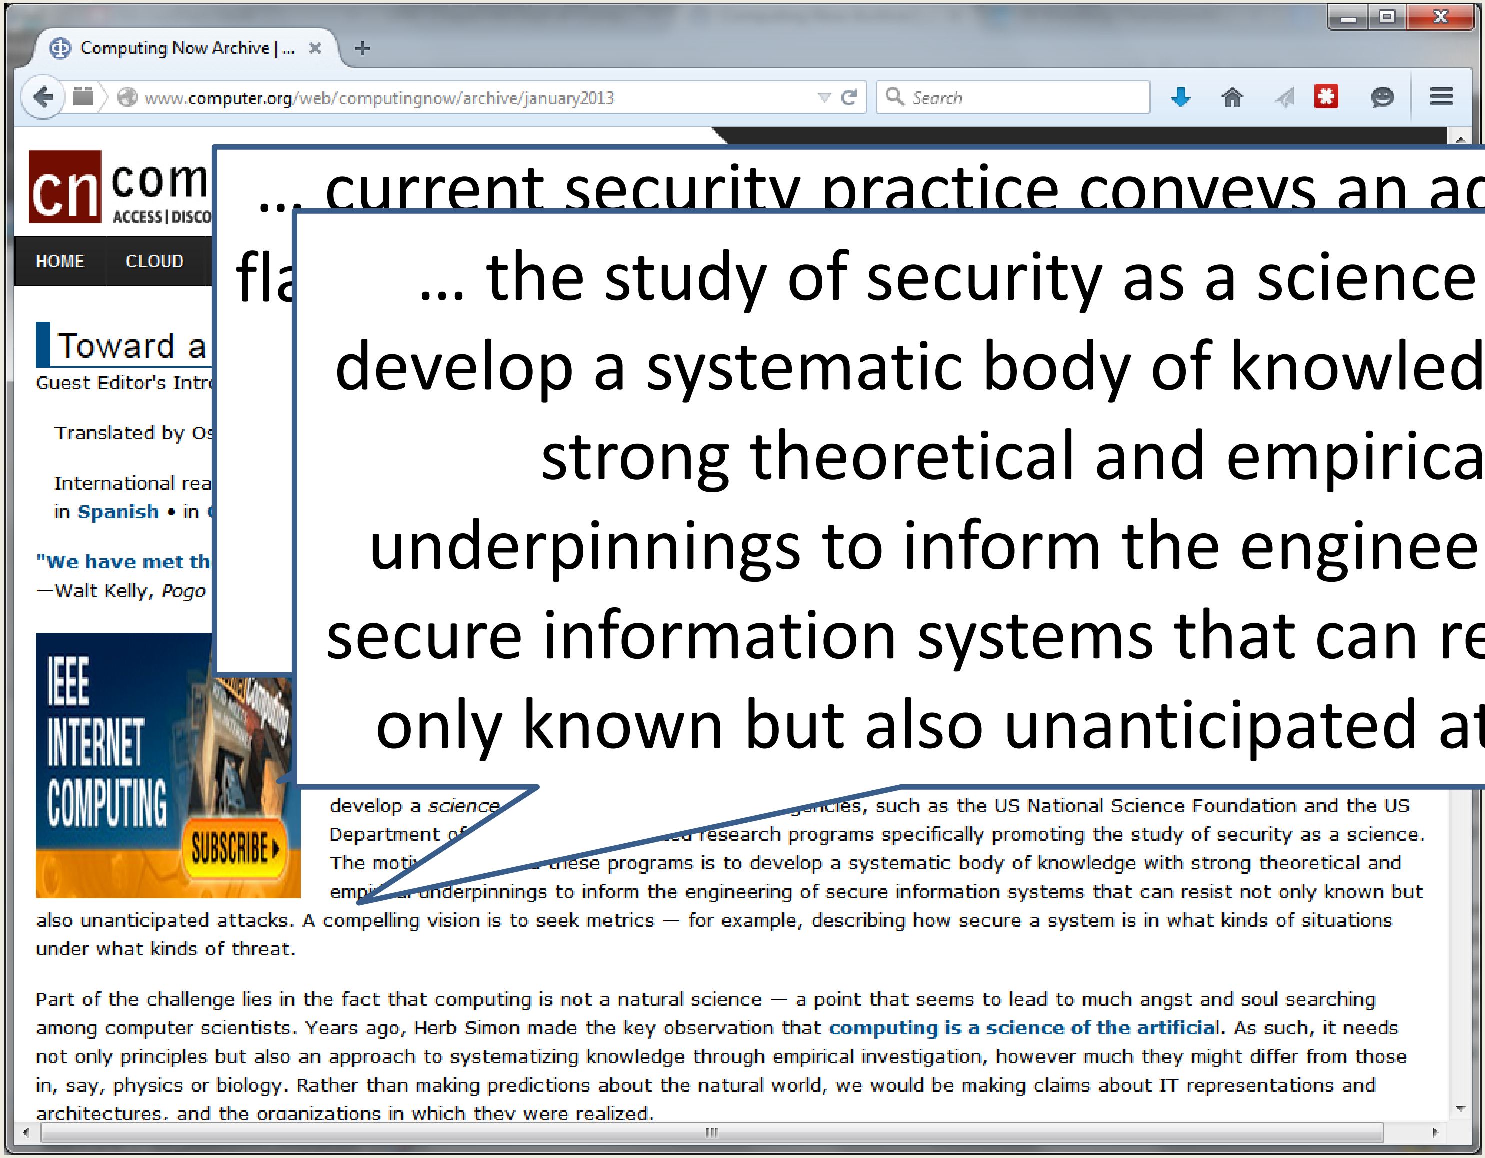

Computing Now Archive | ...
www.computer.org/web/computingnow/archive/january2013
Search
cn com
ACCESS | DISCO
HOME
CLOUD
IEEE INTERNET COMPUTING
SUBSCRIBE

Toward a

Guest Editor's Intr

Translated by Os

International rea
in **Spanish** • in (

"We have met th
—Walt Kelly, *Pogo*

develop a *science* … ncies, such as the US National Science Foundation and the US Department of … research programs specifically promoting the study of security as a science. The moti… these programs is to develop a systematic body of knowledge with strong theoretical and empi… underpinnings to inform the engineering of secure information systems that can resist not only known but also unanticipated attacks. A compelling vision is to seek metrics — for example, describing how secure a system is in what kinds of situations under what kinds of threat.

Part of the challenge lies in the fact that computing is not a natural science — a point that seems to lead to much angst and soul searching among computer scientists. Years ago, Herb Simon made the key observation that **computing is a science of the artificial**. As such, it needs not only principles but also an approach to systematizing knowledge through empirical investigation, however much they might differ from those in, say, physics or biology. Rather than making predictions about the natural world, we would be making claims about IT representations and architectures, and the organizations in which they were realized.

… current security practice conveys an ad hoc
fla

… the study of security as a science … to develop a systematic body of knowledge with strong theoretical and empirical underpinnings to inform the engineering of secure information systems that can resist not only known but also unanticipated attacks.

Science or Engineering?

- Can we recognize science if/when we see it?
- On the one hand, rarely is a paper accepted at a conference that simply patches a bug
- On the other hand, we’re told we need more science
- So, what is all the work we’re doing?

This Talk

- A sampling of research over the past few years in which I’ve been involved, including
  - Large-scale measurement study
  - Design of new cryptographic protocols
  - Attack design
  - Crowdsourcing security configuration
- My goal: to incite discussion about where on the “science vs. engineering” spectrum they lie

An Epidemiological Study of Malware Encounters

[ACM CCS 2014; w/ Yen, Heorhiadi, Oprea, & Juels]

- Goal: To understand malware encounters in a large enterprise network
- Research questions:
  - How did the malware infiltrate network perimeter?
  - Where did encounters occur?
  - How does user behavior affect encounter rate?
  - Can we predict encounters?

Malware Encounters
Data Sets

- Enterprise with 85,000+ hosts instrumented with McAfee anti-virus
- Monitored over four months
  - Jul 10 – Nov 10, 2013
  - ~600,000 AV detections
- Each record contains hostname, virus name, file path, detection time, reporting time

Malware Encounters
Data Sets

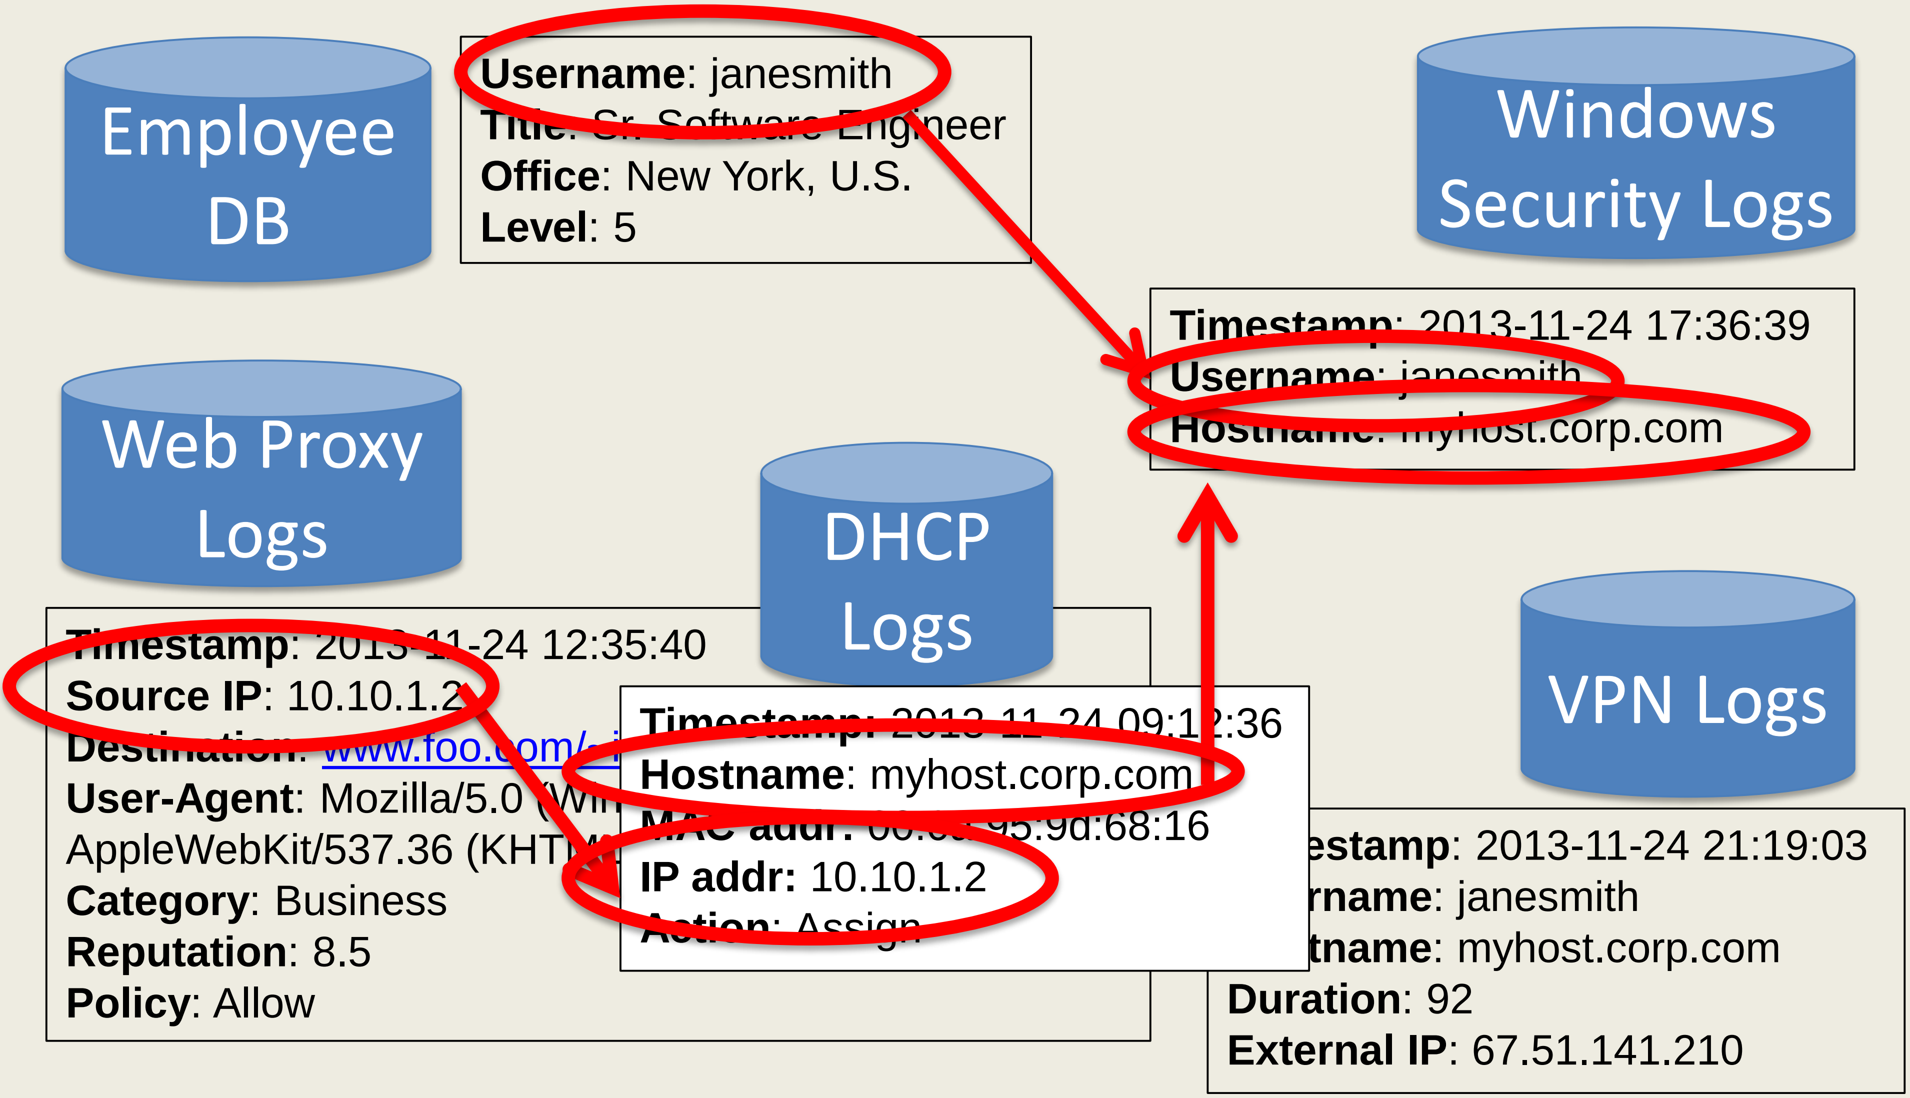

Username: janesmith
Title: Sr. Software Engineer
Office: New York, U.S.
Level: 5

Timestamp: 2013-11-24 17:36:39
Username: janesmith
Hostname: myhost.corp.com

Timestamp: 2013-11-24 12:35:40
Source IP: 10.10.1.2
Destination: www.foo.com/-i
User-Agent: Mozilla/5.0 (Win
AppleWebKit/537.36 (KHTM
Category: Business
Reputation: 8.5
Policy: Allow

Timestamp: 2013-11-24 09:12:36
Hostname: myhost.corp.com
MAC addr: 00:0c:95:9d:68:16
IP addr: 10.10.1.2
Action: Assign

estamp: 2013-11-24 21:19:03
rname: janesmith
tname: myhost.corp.com
Duration: 92
External IP: 67.51.141.210

Malware Encounters

Summary of Findings

- Encounter rate varies widely across countries
- External drive is main malware location
- One-third of web-based encounters originate from websites in “business” category
- Encounters *outside* enterprise network 3 × more common than inside
- Lowest encounter rate among upper management, highest among technical jobs

Malware Encounters

Prioritizing Hosts Based on Risk

- Randomly select half of hosts for training, half for testing; average over 10 runs
- Order by score, compute encounter rate for top N

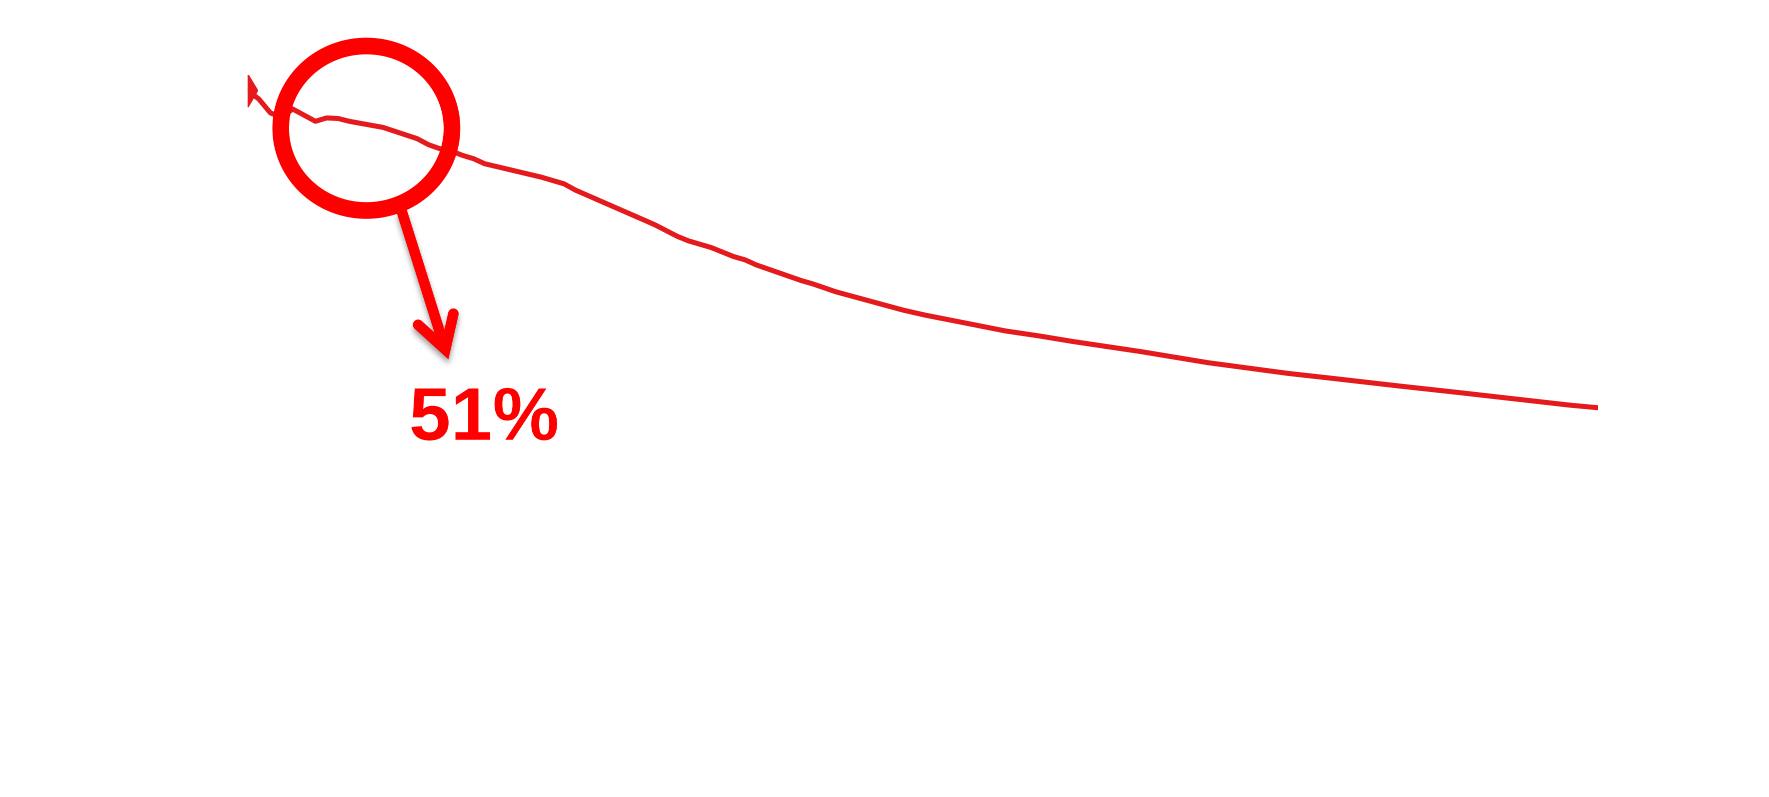

Science or Engineering?

Does this contribute to “a systematic body of knowledge … to inform the engineering of secure information systems that can resist unanticipated attacks?”

- ? Classifier for prioritizing responses to infection indicators … that’s an artifact!
- ? It didn’t have a “hypothesis” and controlled experiment, and so it can’t be science (?)
- ✓ We’re trying to find patterns in malware encounters … that’s science, right?

Science or Engineering?

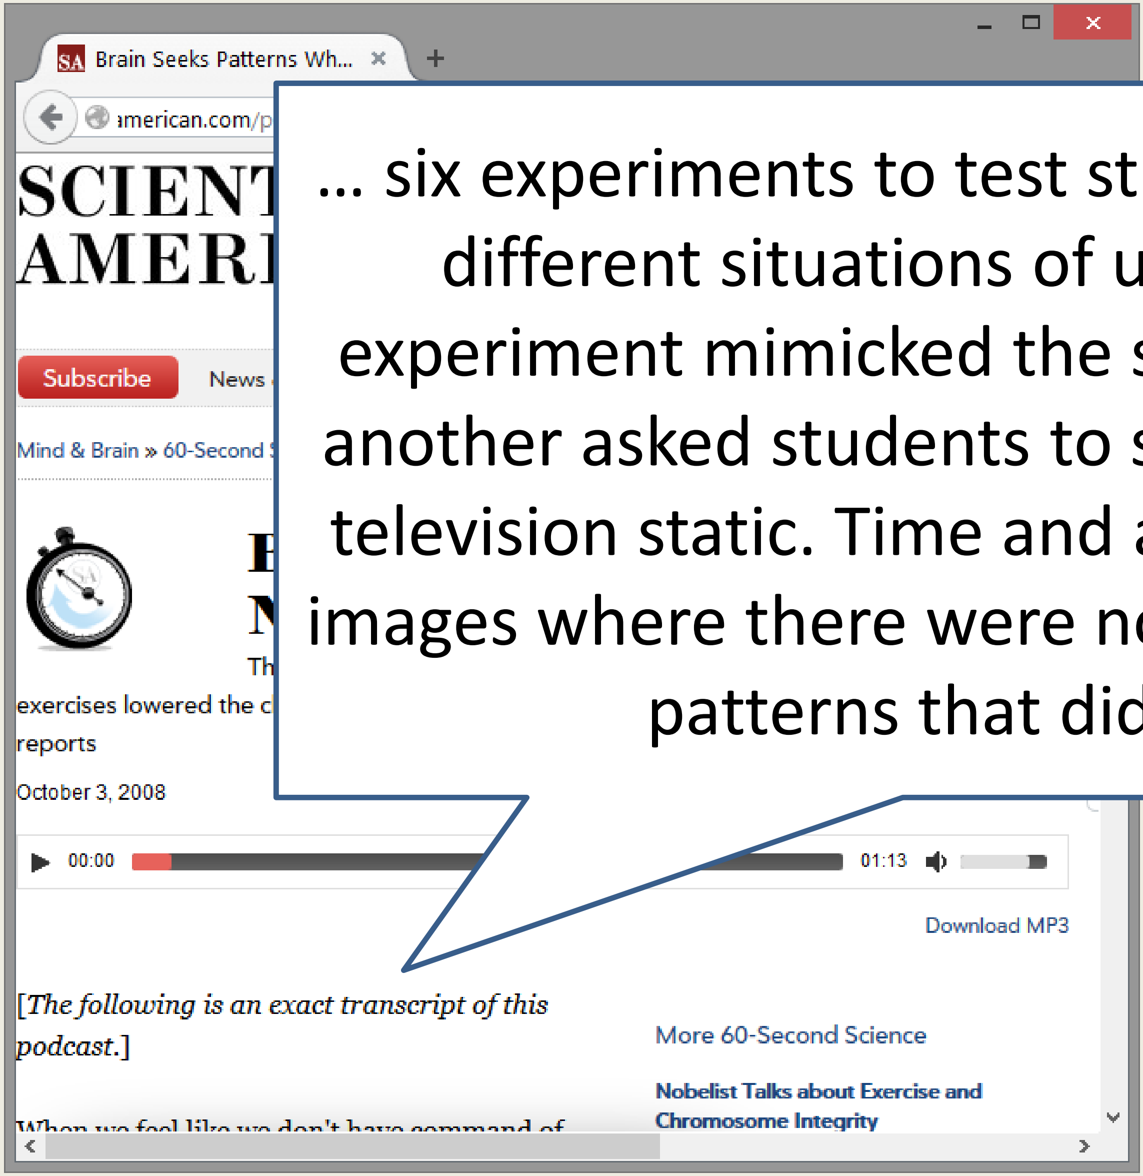

Th

exercises lowered the c

reports

October 3, 2008

[*The following is an exact transcript of this podcast.*]

Nobelist Talks about Exercise and Chromosome Integrity

When we feel like we don't have command of

… six experiments to test students' reactions to different situations of uncertainty. One experiment mimicked the stock market, while another asked students to search for images in television static. Time and again, students saw images where there were none and found stock patterns that didn't exist.

Malware Encounters

Broader Context

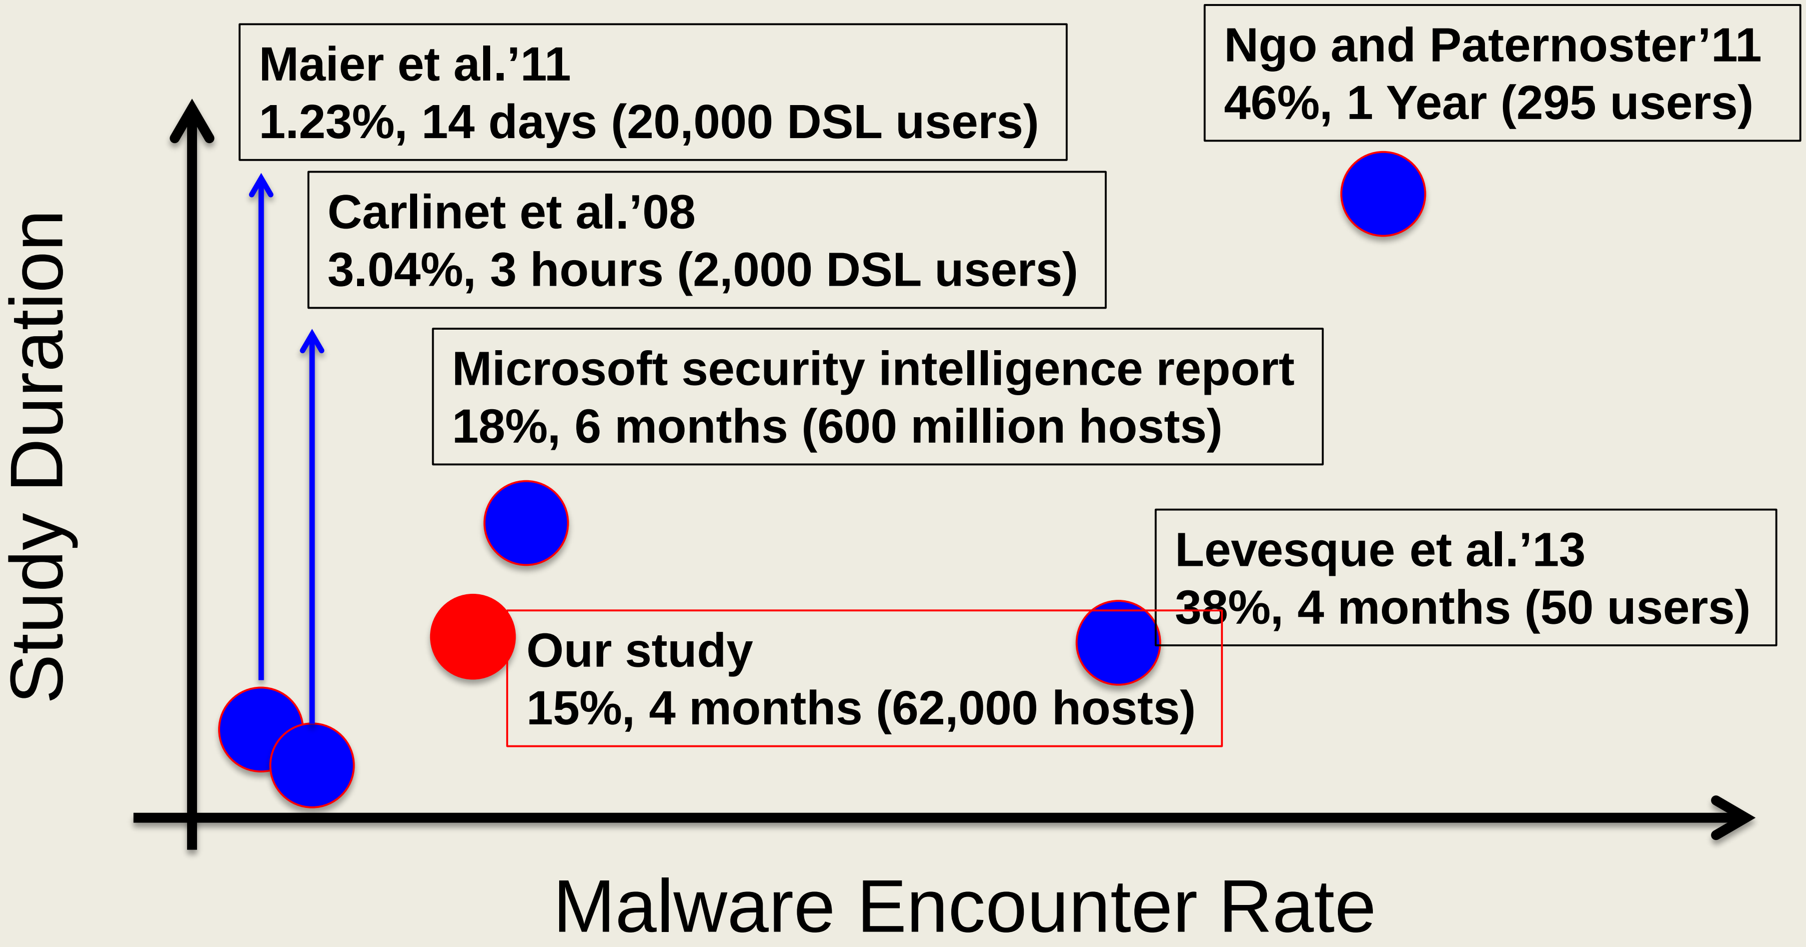

Maier et al.’11
1.23%, 14 days (20,000 DSL users)

Carlinet et al.’08
3.04%, 3 hours (2,000 DSL users)

Microsoft security intelligence report
18%, 6 months (600 million hosts)

Ngo and Paternoster’11
46%, 1 Year (295 users)

Levesque et al.’13
38%, 4 months (50 users)

Our study
15%, 4 months (62,000 hosts)

Cross-tenant Side Channels in PaaS Clouds

[ACM CCS 2014; w/ Zhang, Juels, Ristenpart]

PaaS

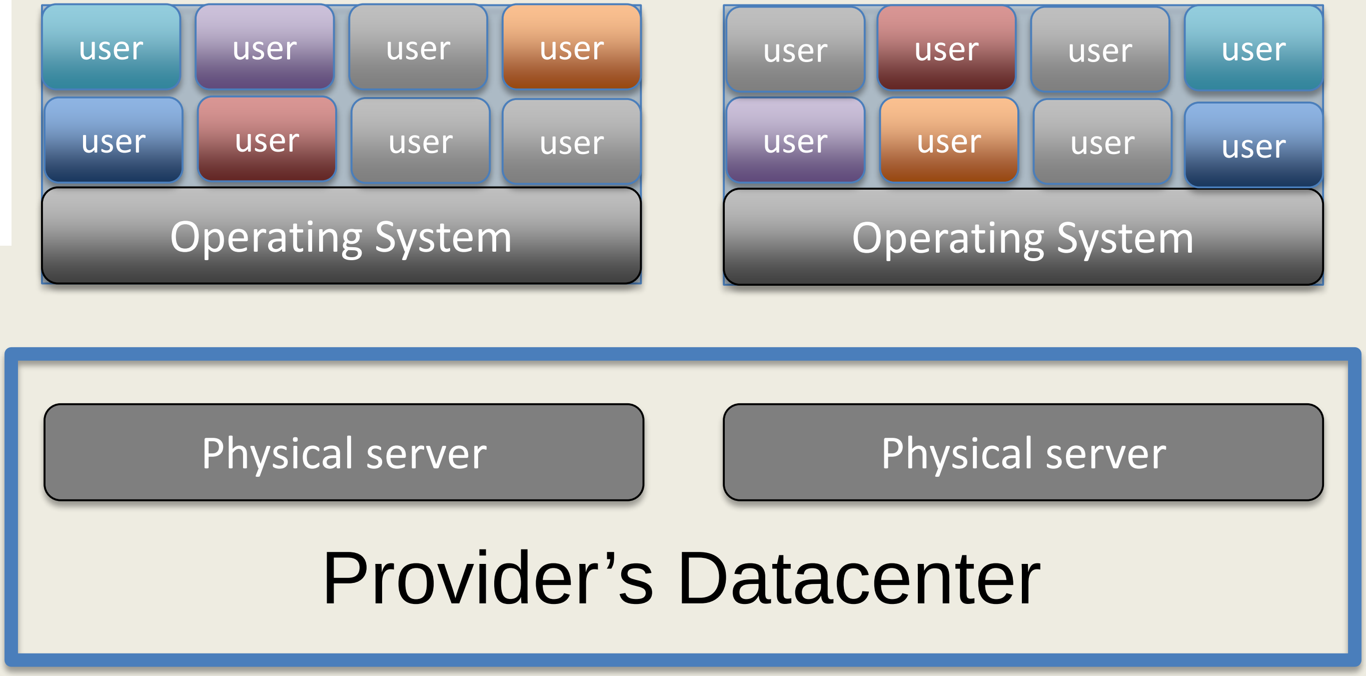

user user user user

user user user user

Operating System

user user user user

user user user user

Operating System

Physical server

Physical server

Provider's Datacenter

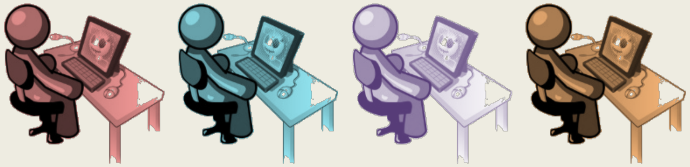

Cross-tenant Side Channels in PaaS Clouds

[ACM CCS 2014; w/ Zhang, Juels, Ristenpart]

PaaS

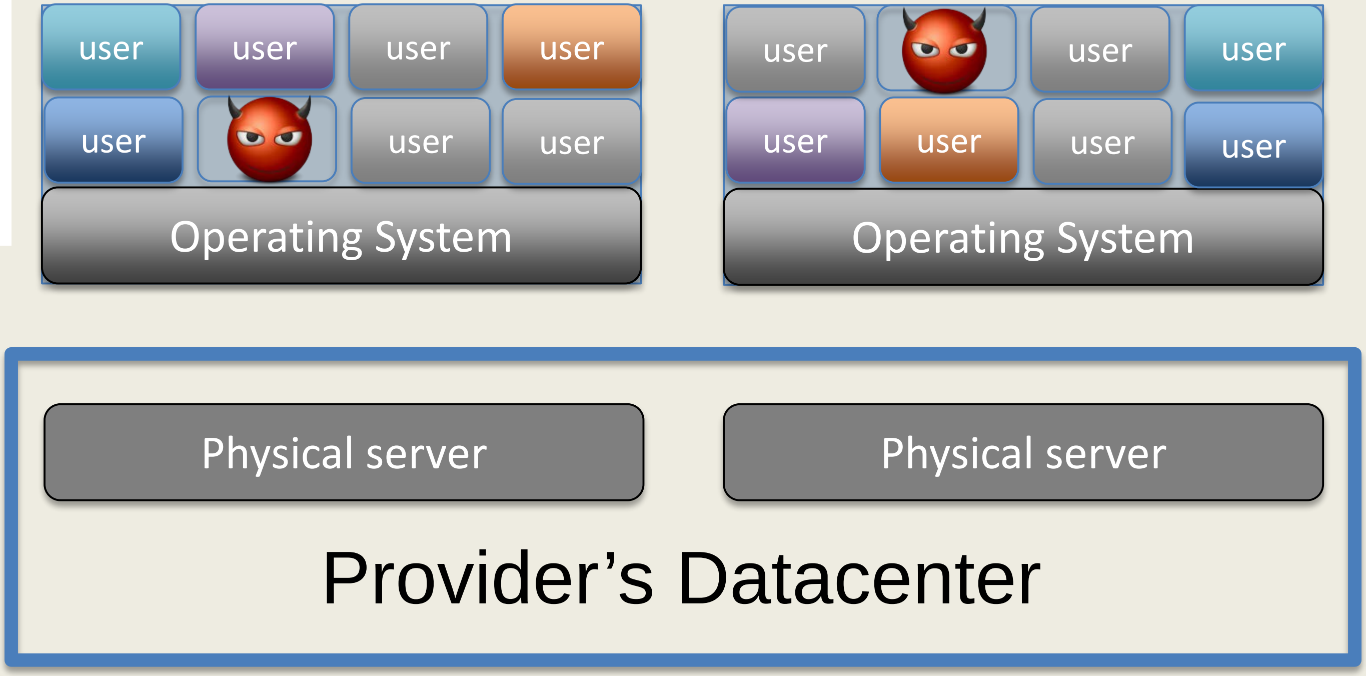

user user user user

user user user

Operating System

user user user

user user user user

Operating System

Physical server

Physical server

Provider’s Datacenter

Flush-Reload Attacks

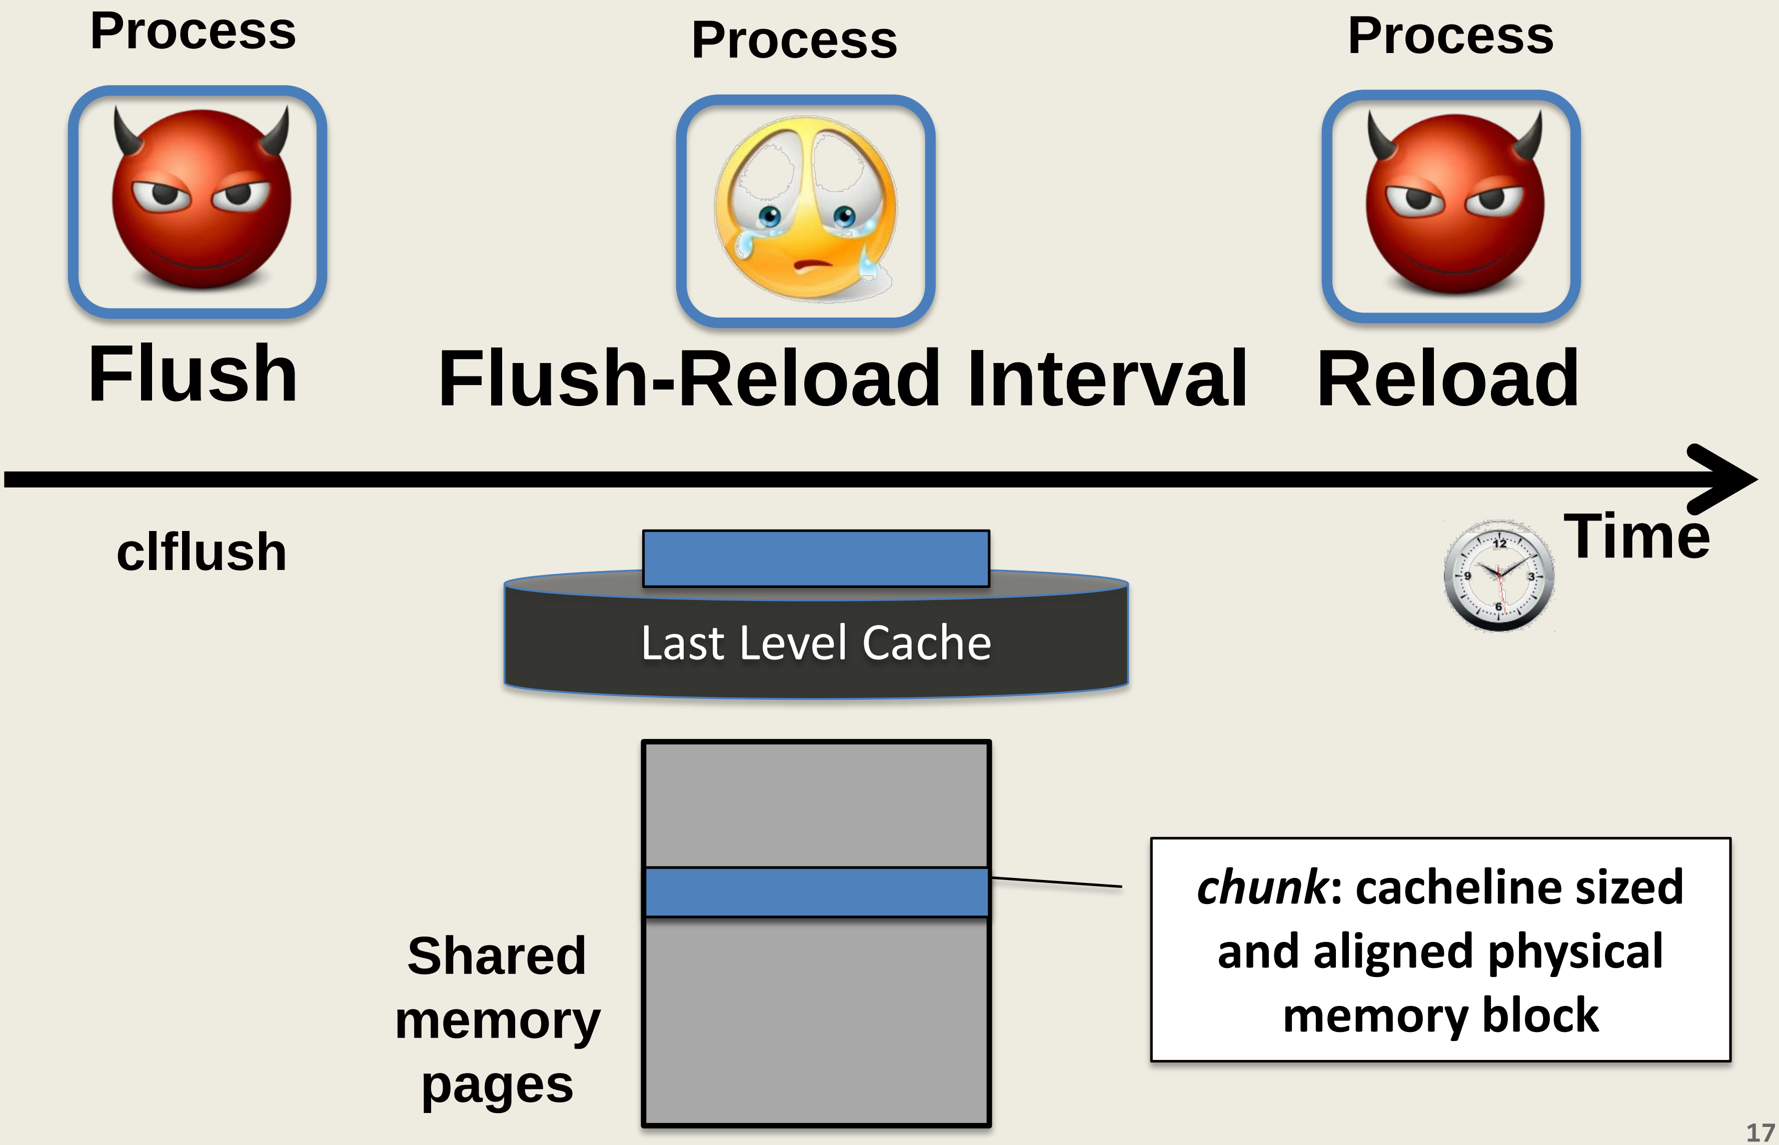

Process

Process

Process

Flush

Flush-Reload Interval

Reload

Time

clflush

Last Level Cache

Shared memory pages

chunk: cacheline sized and aligned physical memory block

A Simple Example

```
#include "stdio.h”
int b;
int inc(int number) {
        return number + 1;
}
int main() {
        int a = 9;
        if (a % 2 == 1)
                a = inc(a);
        b = a;
        return 0;
}
```

Control-Flow Graph

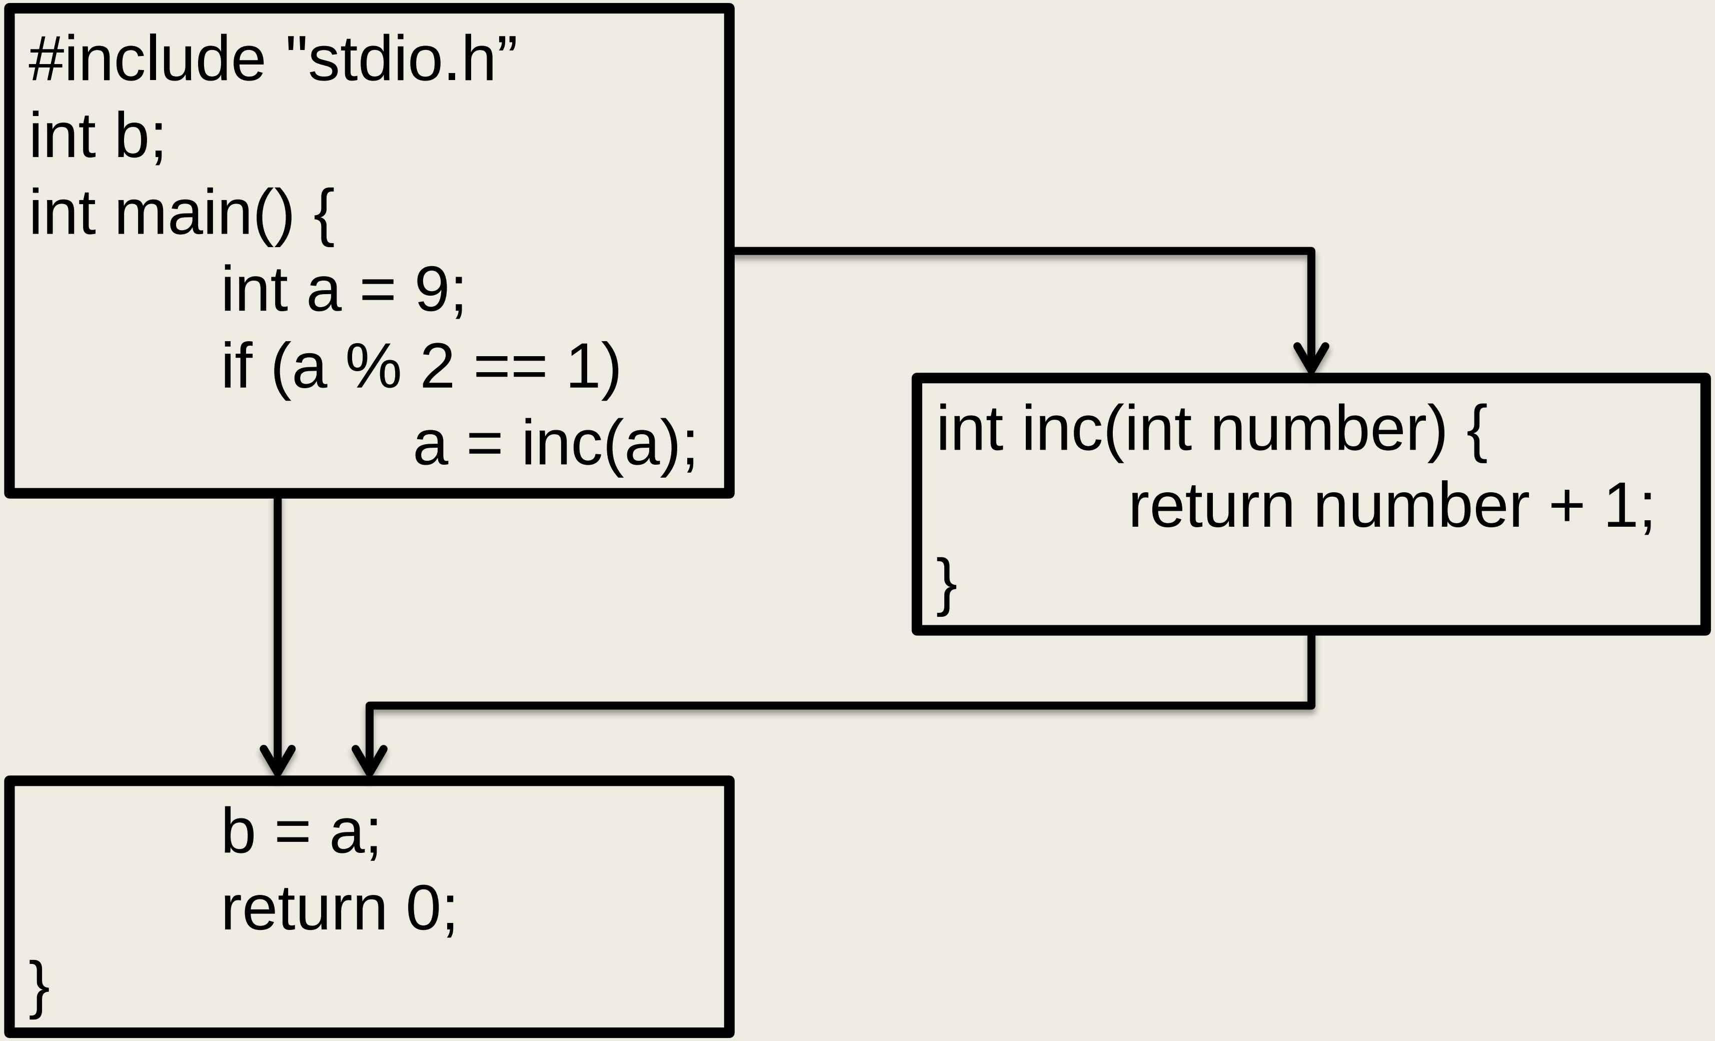

```
#include "stdio.h”
int b;
int main() {
        int a = 9;
        if (a % 2 == 1)
                a = inc(a);
```

```
int inc(int number) {
        return number + 1;
}
```

```
        b = a;
        return 0;
}
```

Control-Flow Graph

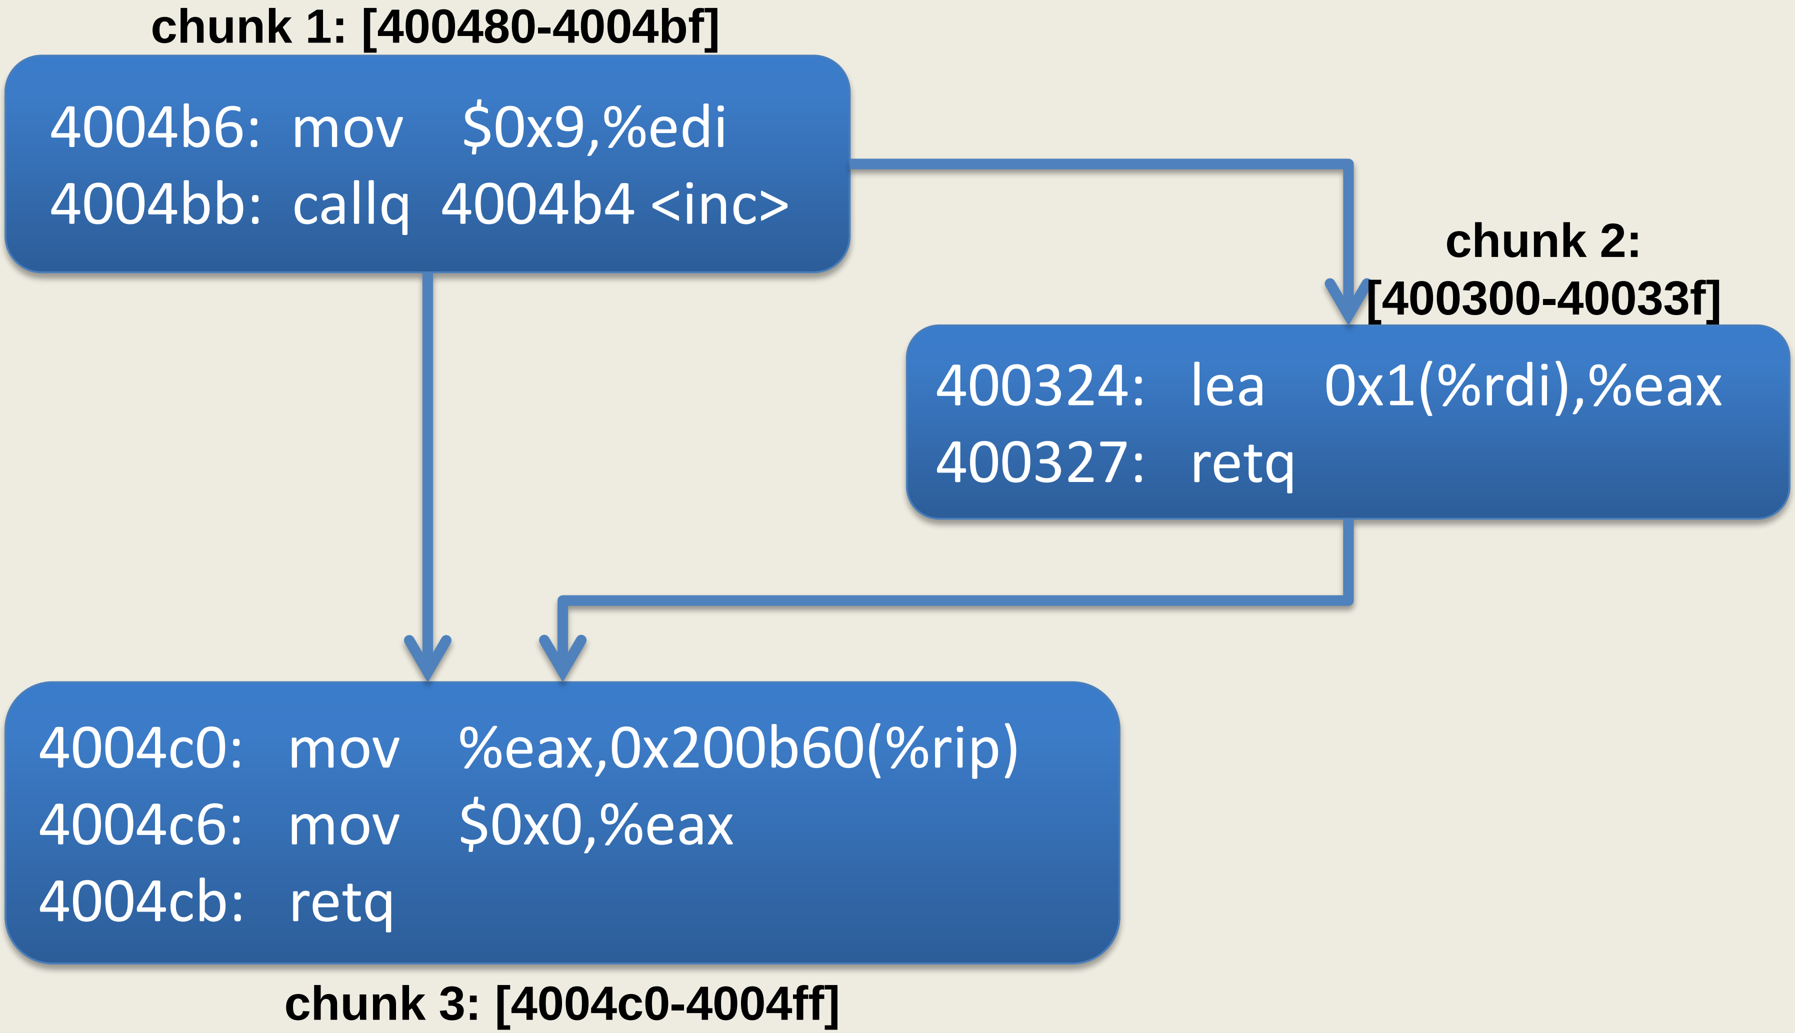

chunk 1: [400480-4004bf]

```
4004b6:  mov    $0x9,%edi
4004bb:  callq  4004b4 <inc>
```

chunk 2: [400300-40033f]

```
400324:   lea    0x1(%rdi),%eax
400327:   retq
```

```
4004c0:   mov    %eax,0x200b60(%rip)
4004c6:   mov    $0x0,%eax
4004cb:   retq
```

chunk 3: [4004c0-4004ff]

Control-Flow Graph

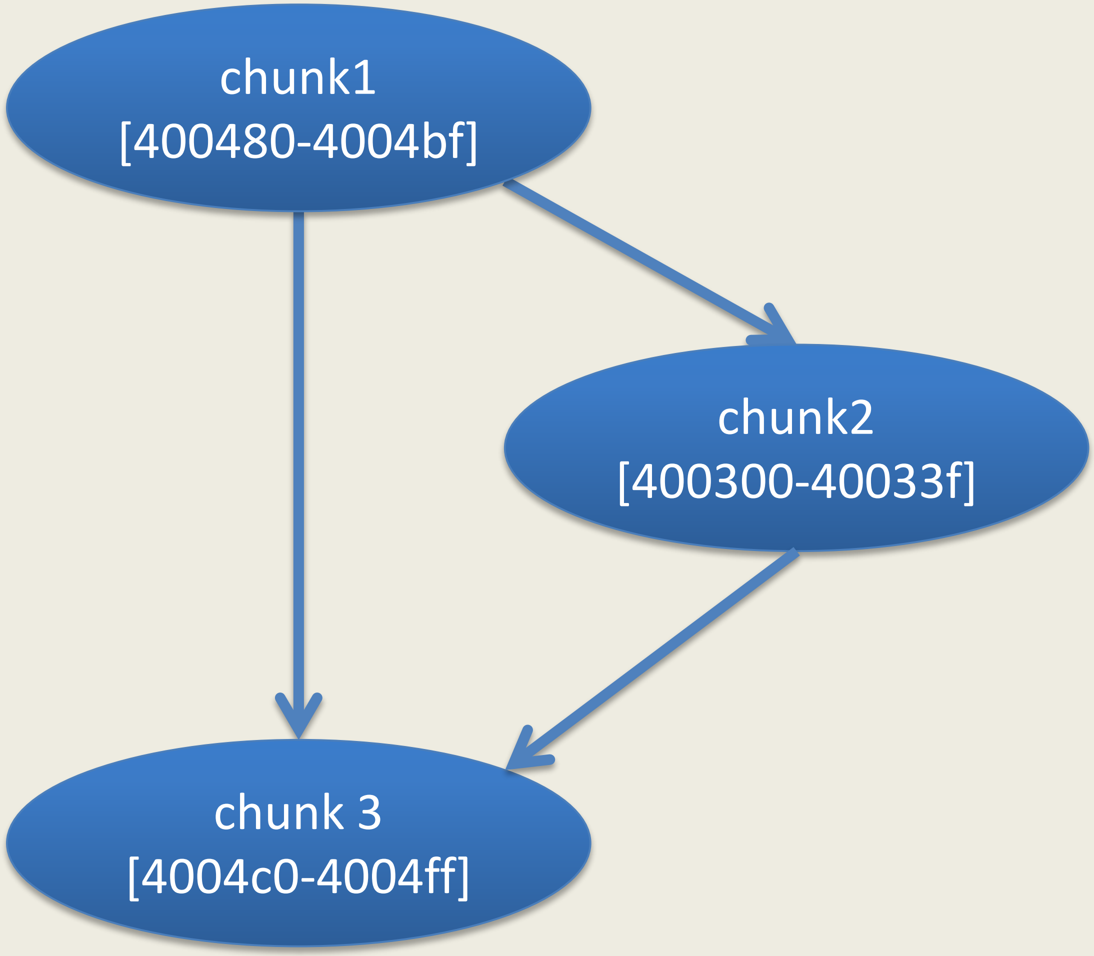

Side Channels in PaaS Clouds
An Attack NFA

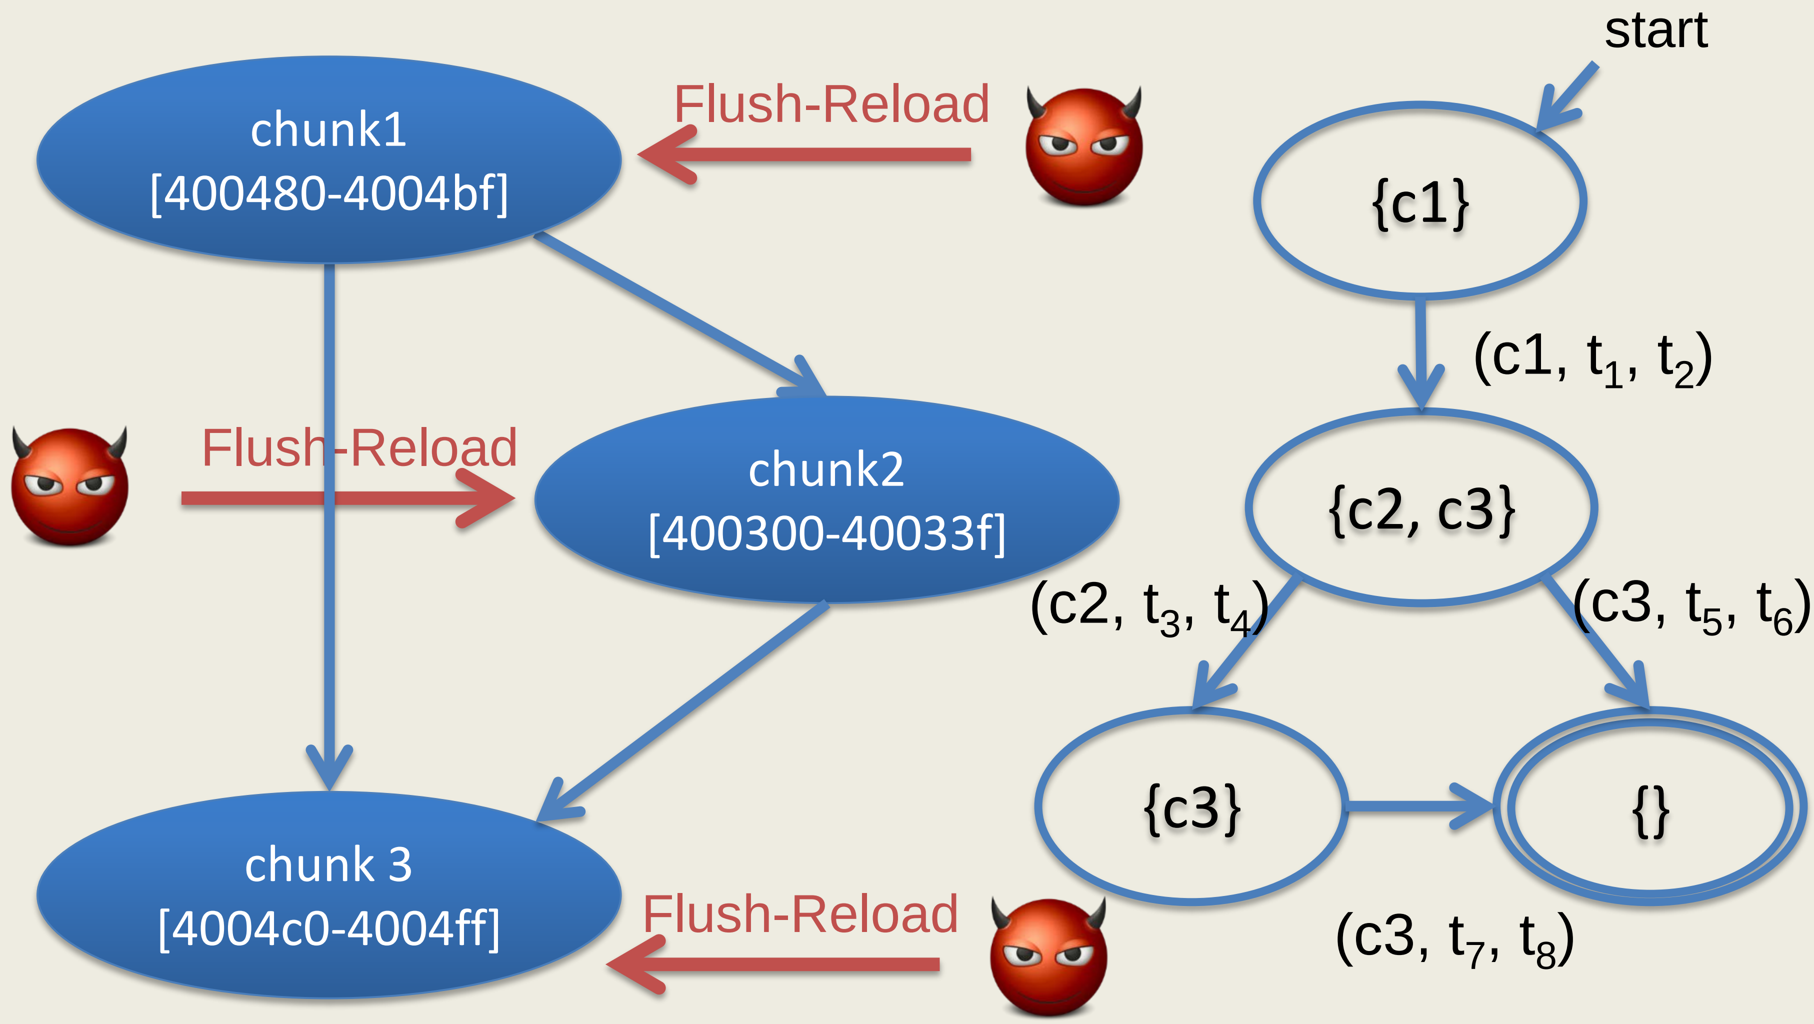

Three Example Attacks

- Inferring sensitive user data
- SAML-based single sign-on attacks
- Password-reset attack

Three Example Attacks

- Inferring sensitive user data
- SAML-based single sign-on attacks
- Password-reset attack

Password Reset

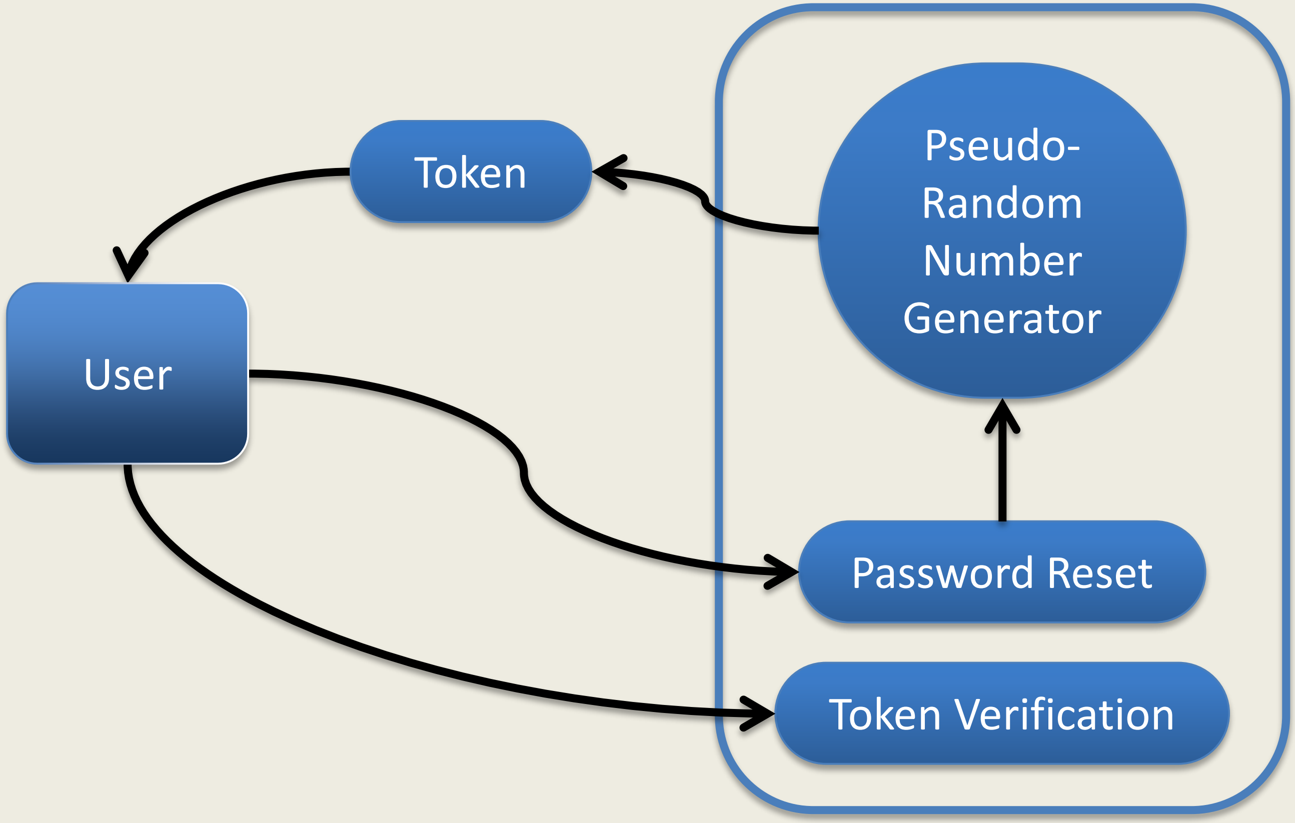

Side Channels in PaaS Clouds

Password Reset Attack

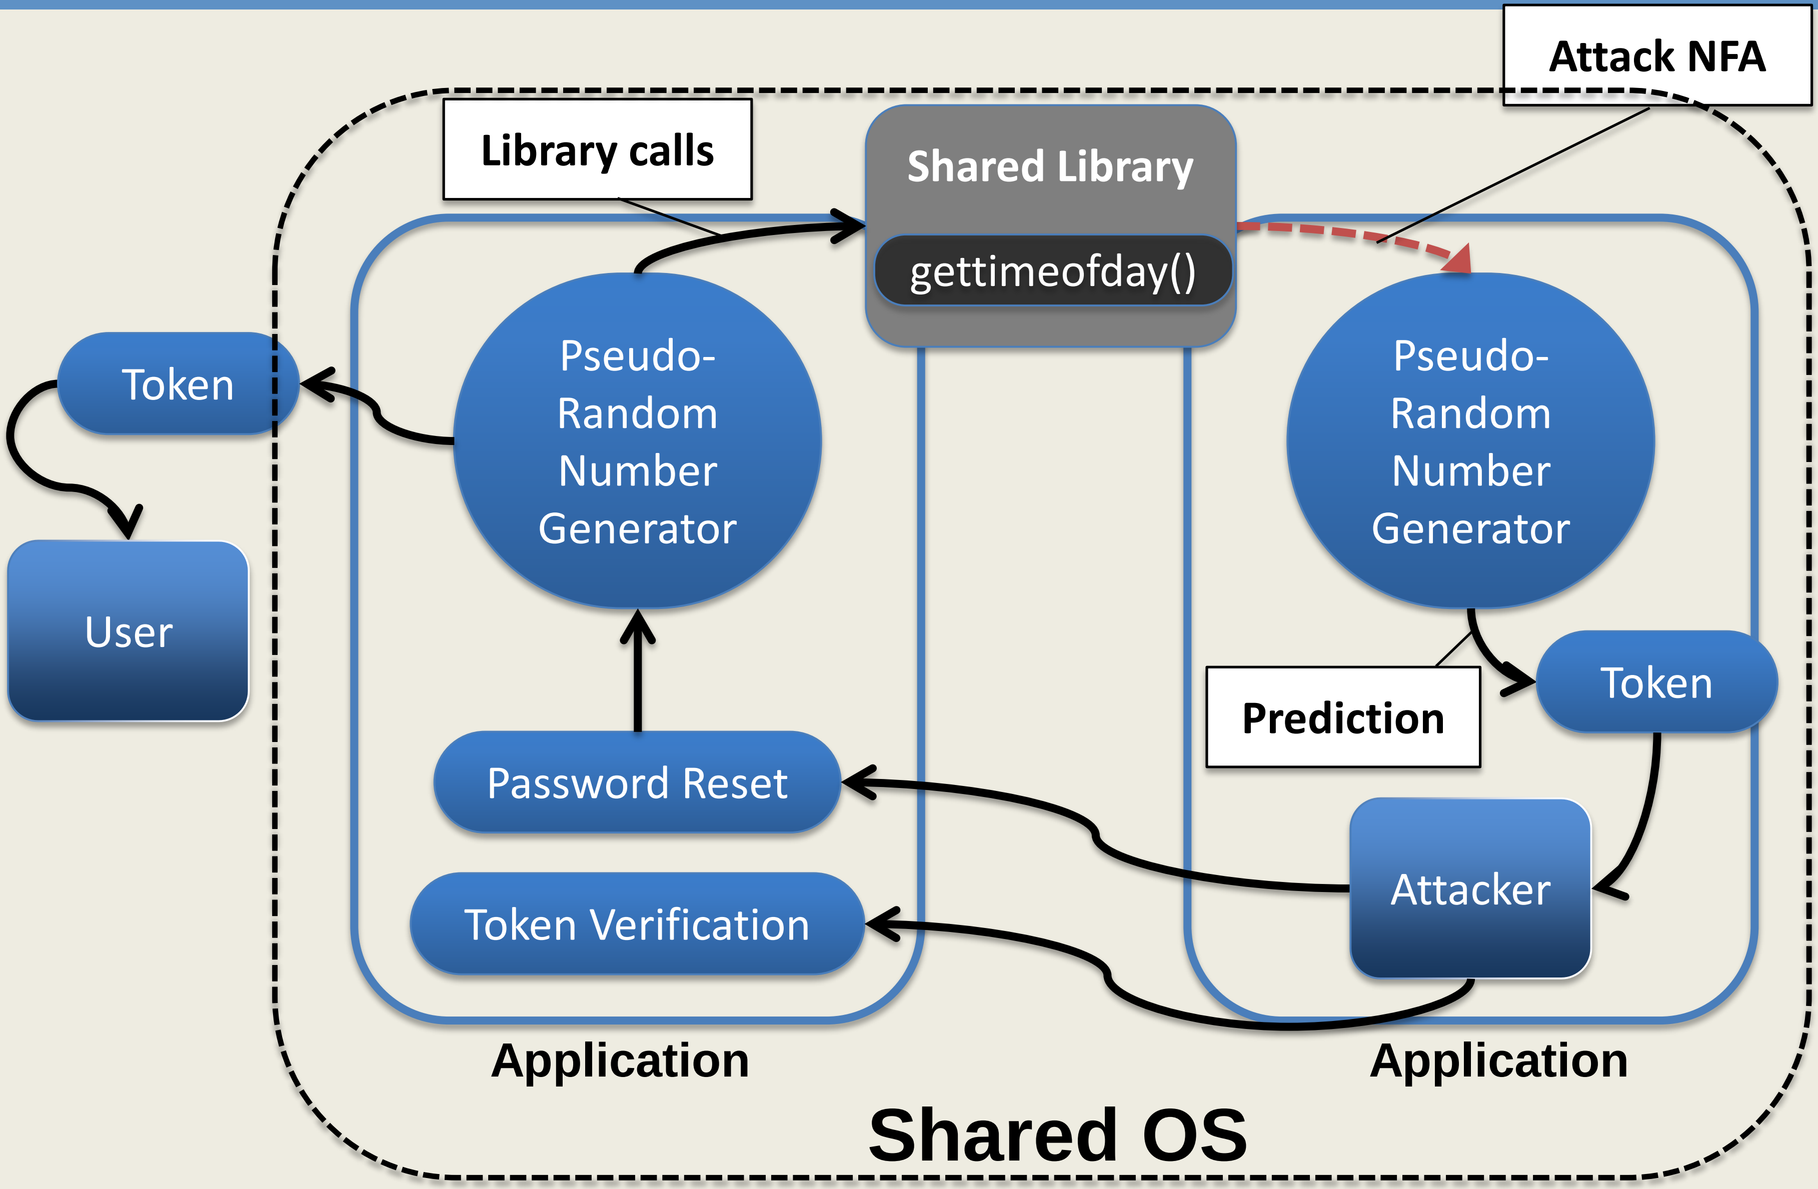

Side Channels in PaaS Clouds
The Attack NFA

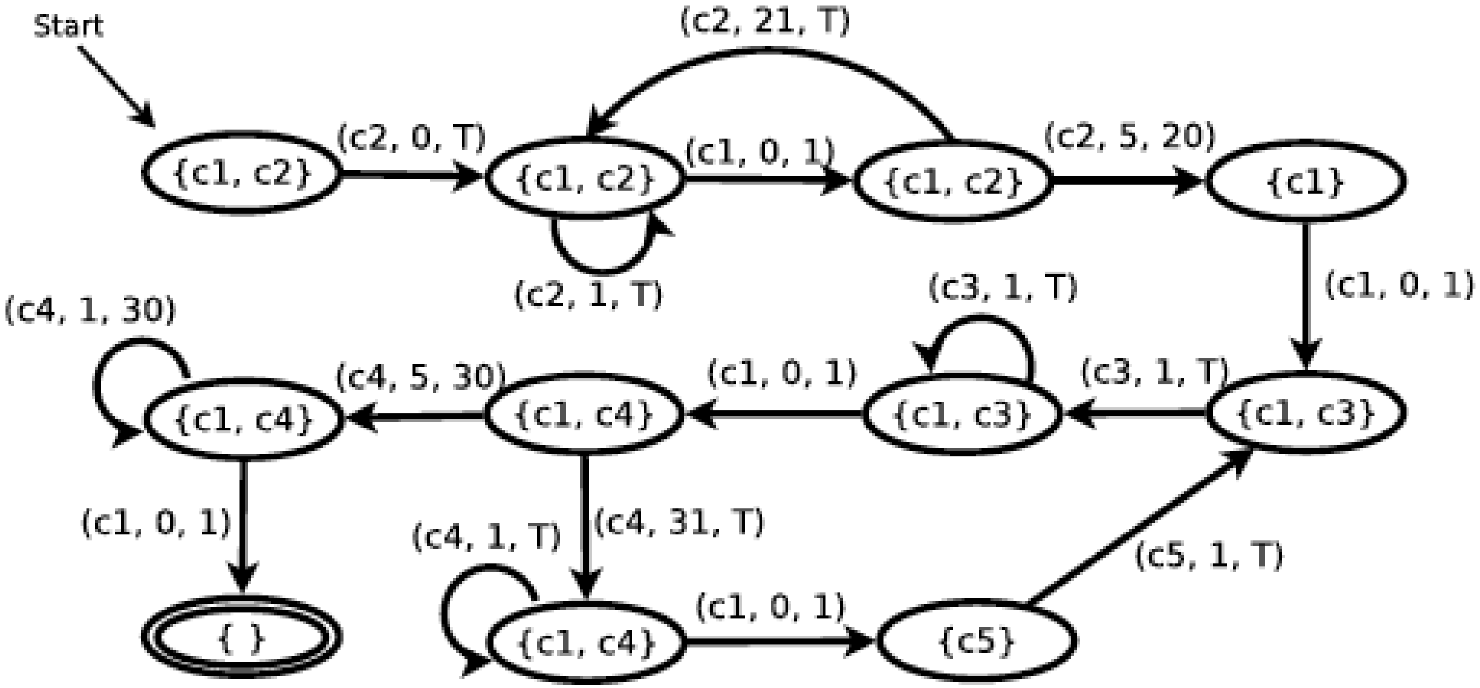

c1. gettimeofday@plt	php5-fpm [0x42ee40 - 0x42ee7f]
c2. lcg_seed	php5-fpm [0x5eab00 - 0x5eab3f]
c3. php_gettimeofday	php5-fpm [0x5f0380 - 0x5f03bf]
c4. uniqid	php5-fpm [0x6028c0 - 0x6028ff]
c5. php_combined_lcg	php5-fpm [0x5eab40 - 0x5eab7f]

Evaluation

- Demonstrated successful attacks against Magento (controlled by ourselves) in a public PaaS cloud.
- After 2^{20} offline computation, the attacker can narrow down the password reset token to 2^{2} possible values---easy to brute-force online.

Side Channels in PaaS Clouds

Science or Engineering?

Does this contribute to “a systematic body of knowledge … to inform the engineering of secure information systems that can resist unanticipated attacks?”

✓ Attack NFAs are maybe “systematic”

? Showed the problem, but not how to resist it

? Can attack papers be part of a “science”?

3rd-Party Private DFA Eval on Encrypted Data

[ESORICS 2012, 2013; IJIS 2014; w/ Wei]

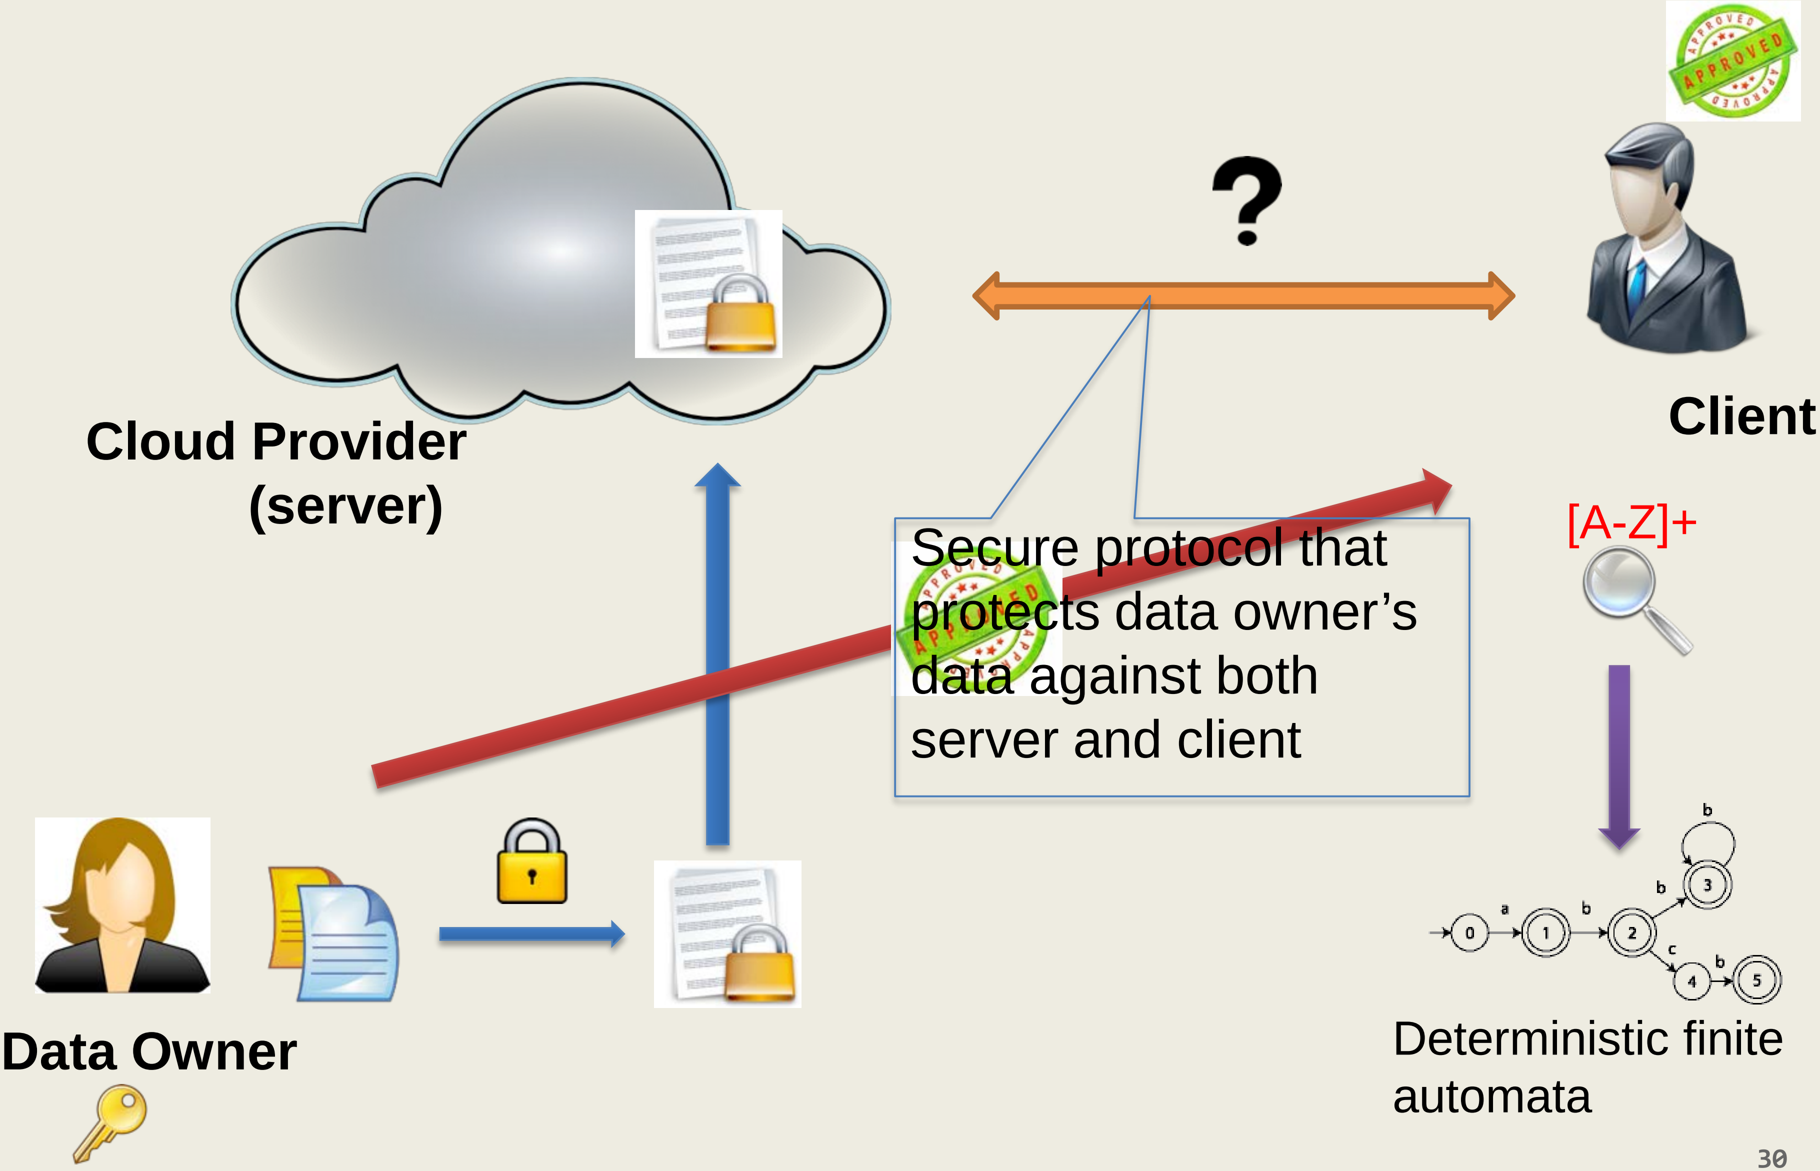

Cloud Provider (server)

Client

Data Owner

[A-Z]+

Secure protocol that protects data owner's data against both server and client

Deterministic finite automata

Private DFA Eval
Protocol Idea

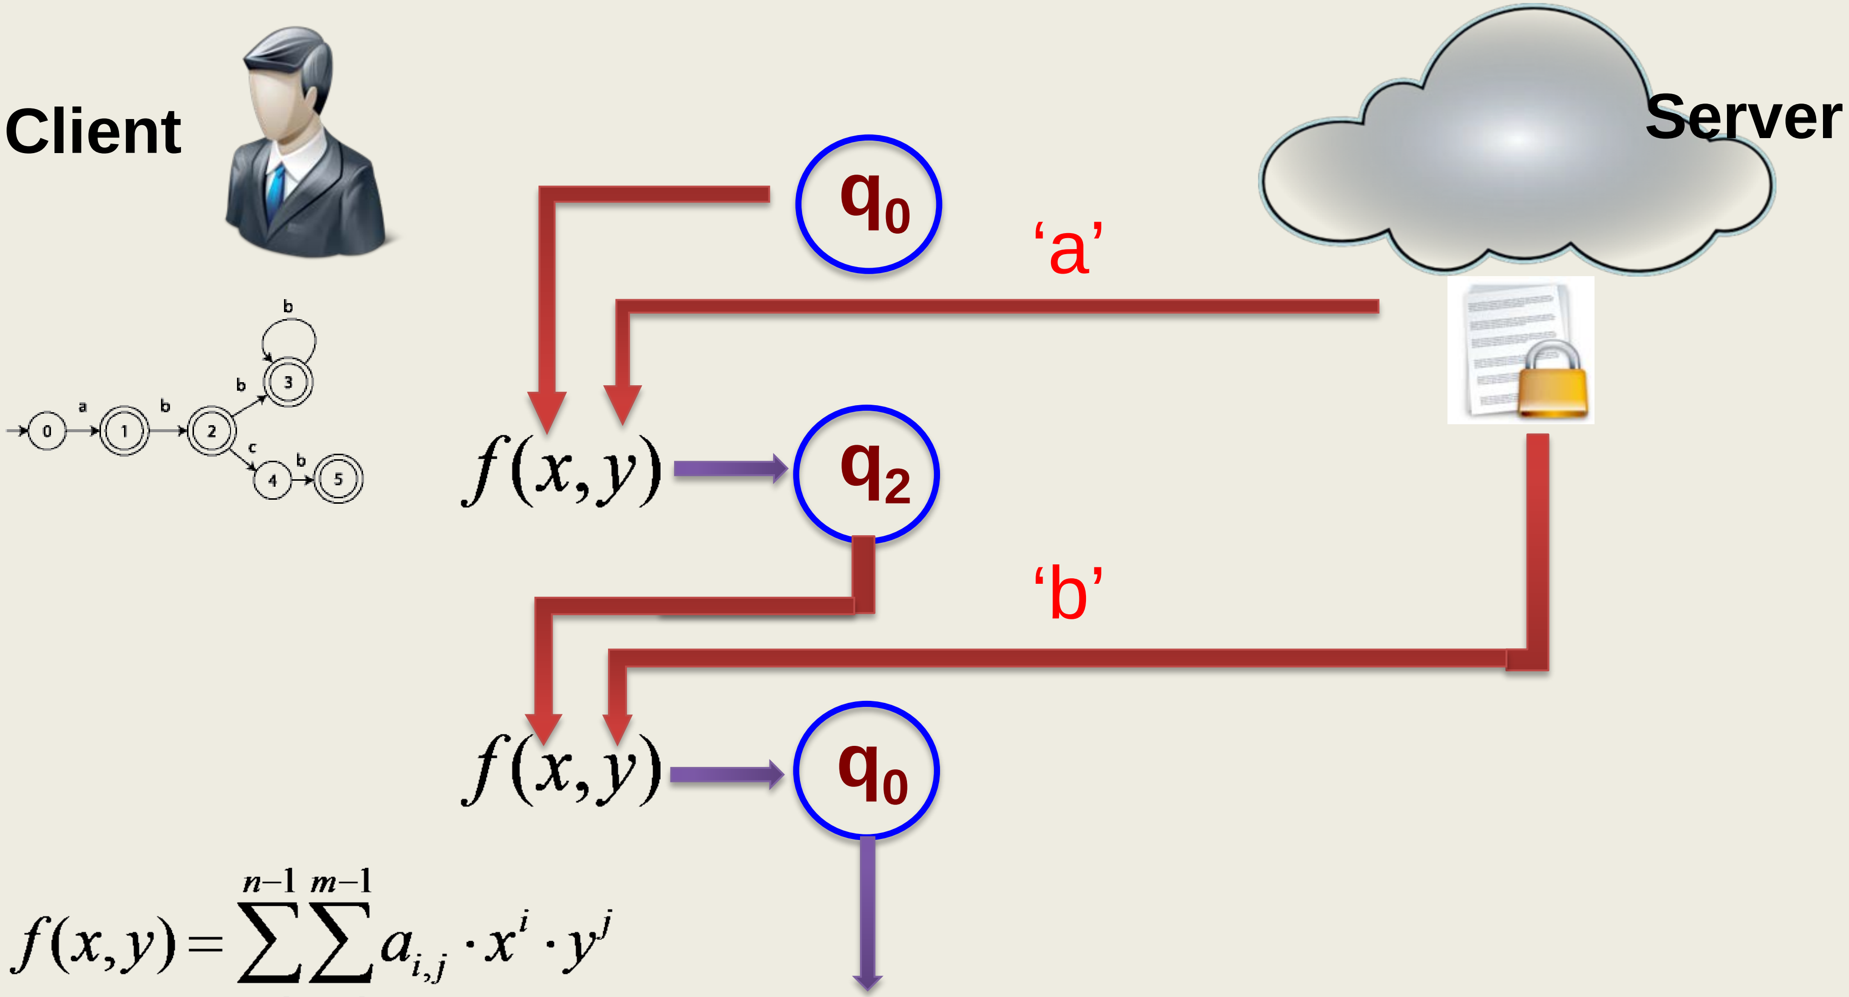

$$f(x,y)=\sum_{i=0}^{n-1}\sum_{j=0}^{m-1}a_{i,j}\cdot x^{i}\cdot y^{j}$$

1. Model the DFA transition function as a bivariate polynomial
2. Simulate the state transition by evaluating the polynomial

Polynomial Evaluation on Ciphertext

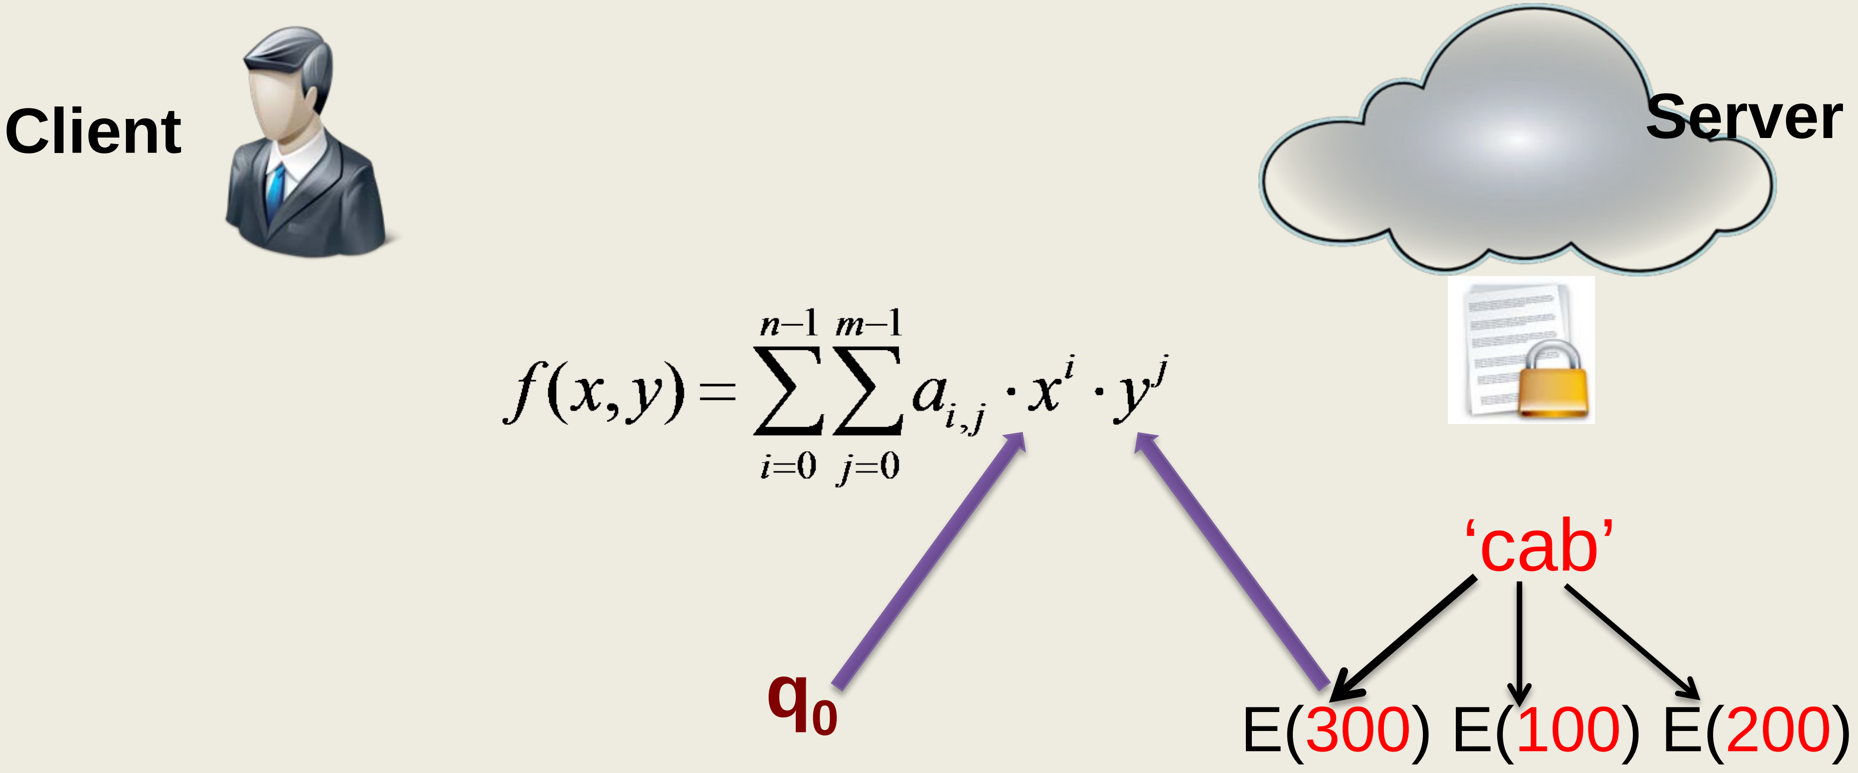

Client

Server

$$f(x,y)=\sum_{i=0}^{n-1}\sum_{j=0}^{m-1}a_{i,j}\cdot x^{i}\cdot y^{j}$$

‘cab’

q_0

E(300) E(100) E(200)

- Need to evaluate the polynomial using ciphertext as input

Additively Homomorphic Encryption

- Additively homomorphic encryption scheme, e.g., Paillier cryptosystem
  - Additive homomorphism:

$$E(m_1 + m_2) = E(m_1) \boxed{\oplus} E(m_2)$$

  - Given E(m) and a constant c, multiply the constant into the ciphertext:

$$E(m \cdot c) = c \boxed{\otimes} E(m)$$

Evaluate Polynomial on Encrypted Data

$$f(x,y)=\sum_{i=0}^{n-1}\sum_{j=0}^{m-1}a_{i,j}\cdot x^{i}\cdot y^{j}$$

- Input: q_0 and $\boxed{E(\sigma^{0}),E(\sigma^{1}),\ldots,E(\sigma^{m-1})}$
- Compute:

$$q_0^{i}\otimes E(\sigma^{j}) \Longrightarrow E(q_0^{i}\cdot\sigma^{j}),i\in[n],j\in[m]$$

- Then:

$$a_{i,j}\otimes E(q_0^{i}\cdot\sigma^{j}) \Longrightarrow E(a_{i,j}\cdot q_0^{i}\cdot\sigma^{j}),i\in[n],j\in[m]$$

- Finally:

$$\bigoplus_{i}\bigoplus_{j}E(a_{i,j}\cdot q_0^{i}\cdot\sigma^{j})=E(\sum_{i}\sum_{j}a_{i,j}\cdot q_0^{i}\cdot\sigma^{j})$$

$$=E(f(q_0,\sigma))=E(q_{next})$$

Private DFA Eval

Protocol Outline

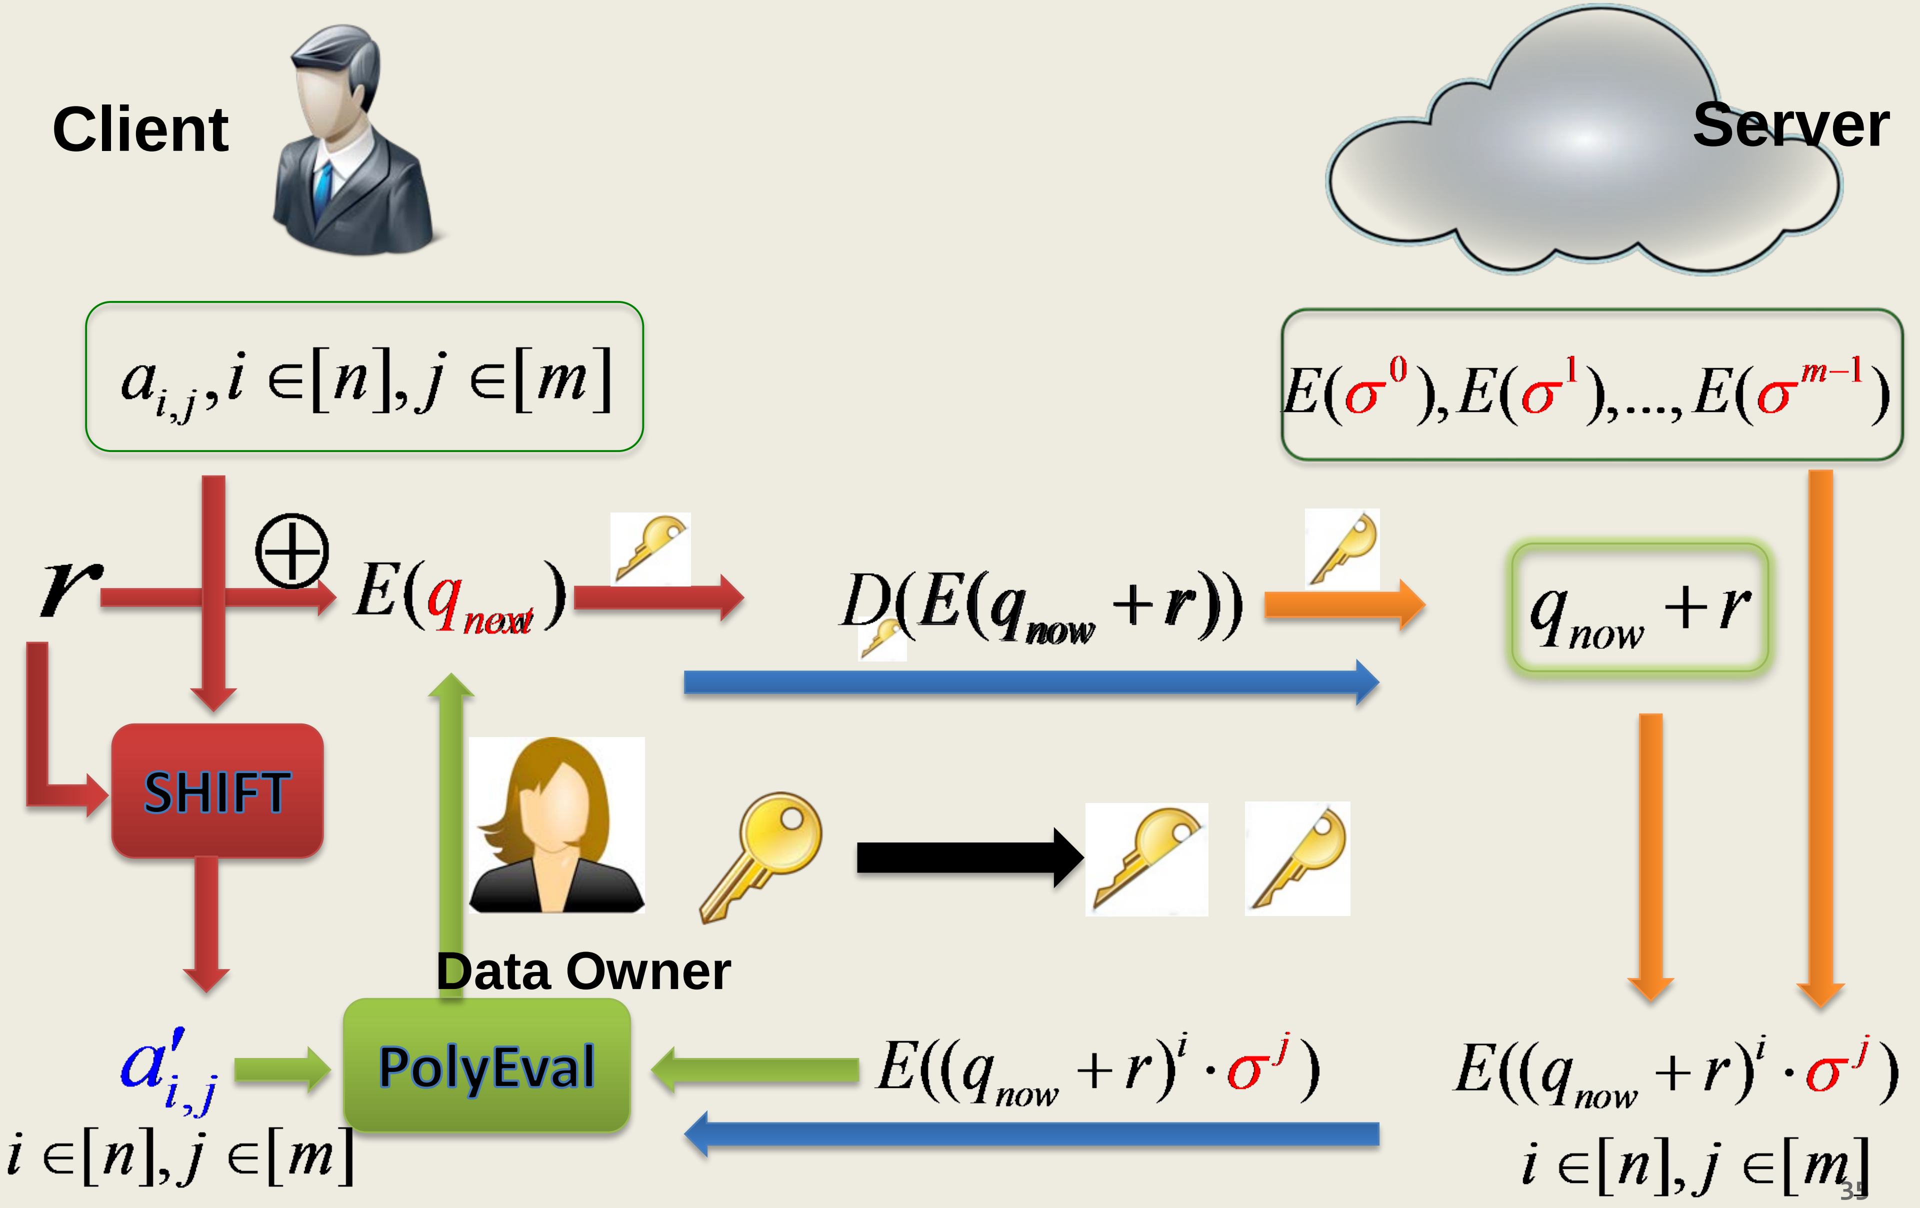

Client

Server

$a_{i,j}, i \in [n], j \in [m]$

$E(\sigma^0), E(\sigma^1), \ldots, E(\sigma^{m-1})$

r

$\oplus$

$E(q_{next})$

$D(E(q_{now} + r))$

$q_{now} + r$

SHIFT

Data Owner

$a'_{i,j}$

$i \in [n], j \in [m]$

PolyEval

$E((q_{now} + r)^i \cdot \sigma^j)$

$E((q_{now} + r)^i \cdot \sigma^j)$

$i \in [n], j \in [m]$

Security Properties

- DFA query privacy:
  - Provably protected against malicious server
- File content privacy:
  - Provably protected against malicious server and honest-but-curious client (heuristically against a malicious client) except for the evaluation result

Private DFA Eval

RegExp Searching on Encrypted Email

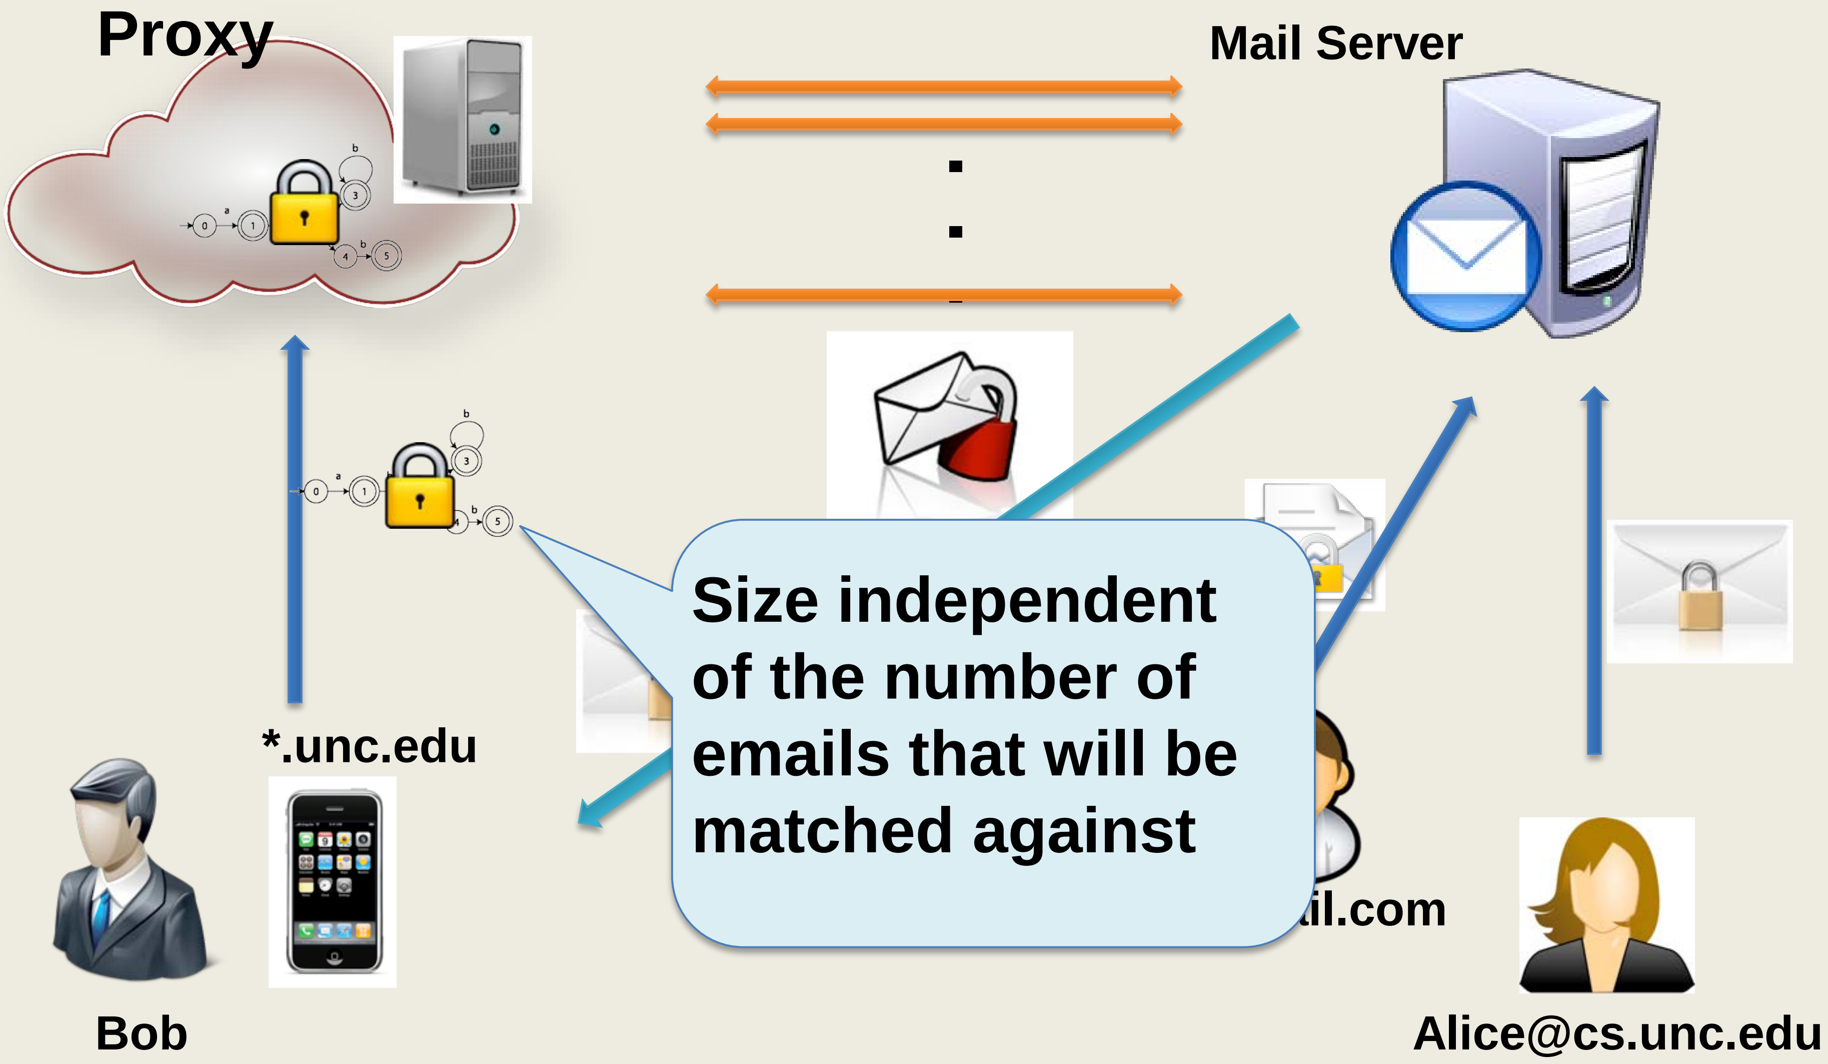

Size independent of the number of emails that will be matched against

Example: Query on Date Field

- **Range query**: 2001/09/10-2002/04/20
- Corresponding regular expression:

(0109 (10|11|...|31)) **|** (01(10|11|12)(01|...|31)) **|**
(02|(01|02|03)(01|...|31)) **|** (0204(01|02|...|20))

Time spent per email in seconds

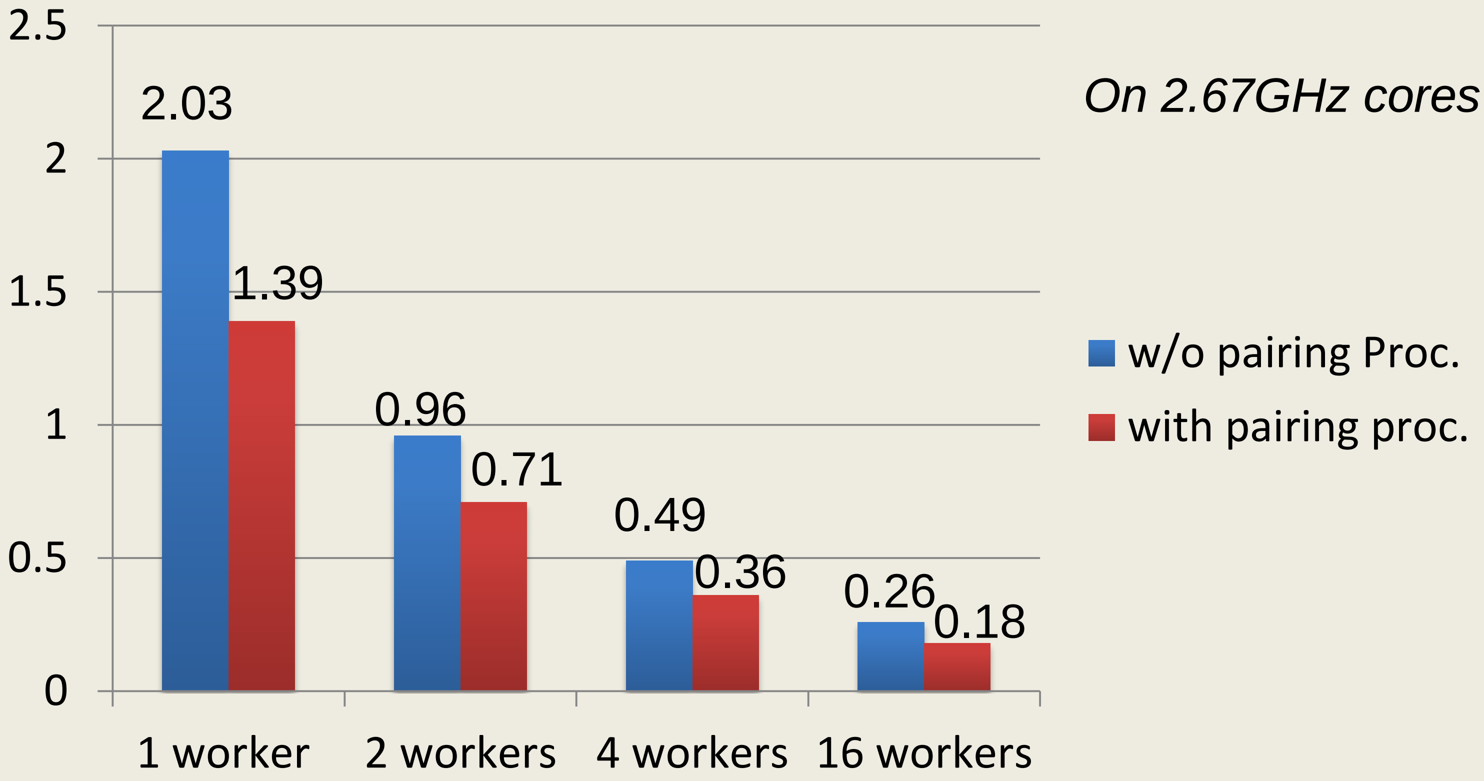

On 2.67GHz cores

Science or Engineering?

Does this contribute to “a systematic body of knowledge … to inform the engineering of secure information systems that can resist unanticipated attacks?”

✓ It has theorems!

? But of course the theorems apply only to the attacks we’ve considered in our threat model

? Produced an artifact, and so maybe it’s just principled engineering?

Crowdsourced Exploration of Security Configs

[CHI 2015; w/ Ismael, Ahmed, and Kapadia]

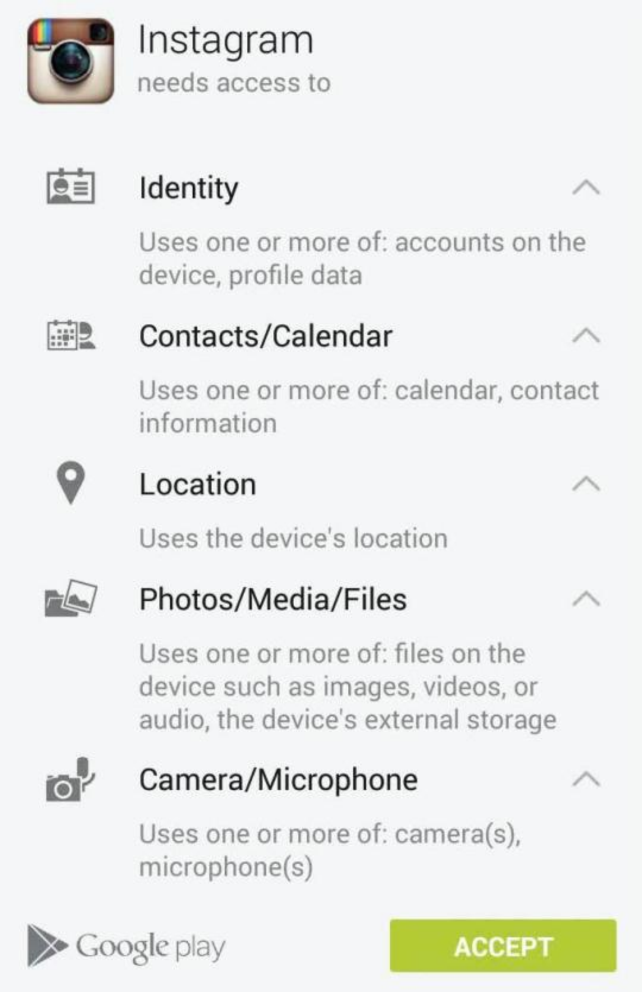

Instagram
needs access to

Identity

Uses one or more of: accounts on the device, profile data

Contacts/Calendar

Uses one or more of: calendar, contact information

Location

Uses the device's location

Photos/Media/Files

Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera/Microphone

Uses one or more of: camera(s), microphone(s)

Android

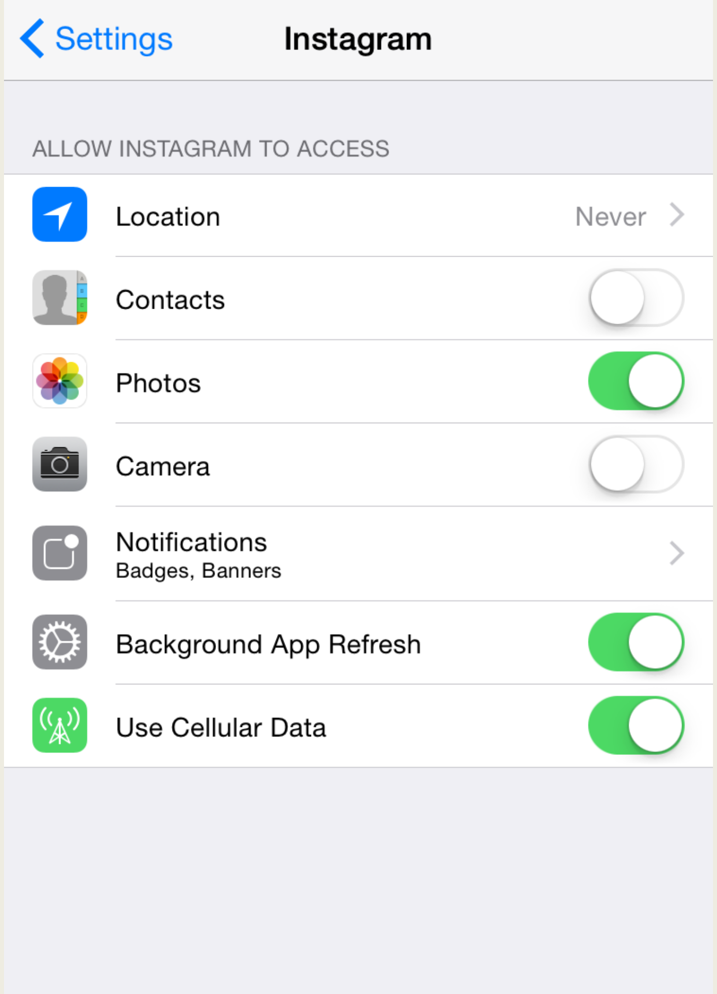

< Settings **Instagram**

ALLOW INSTAGRAM TO ACCESS

- Location — Never >
- Contacts ○
- Photos ●
- Camera ○
- Notifications — Badges, Banners >
- Background App Refresh ●
- Use Cellular Data ●

iOS

Are All Those Permissions Really Necessary?

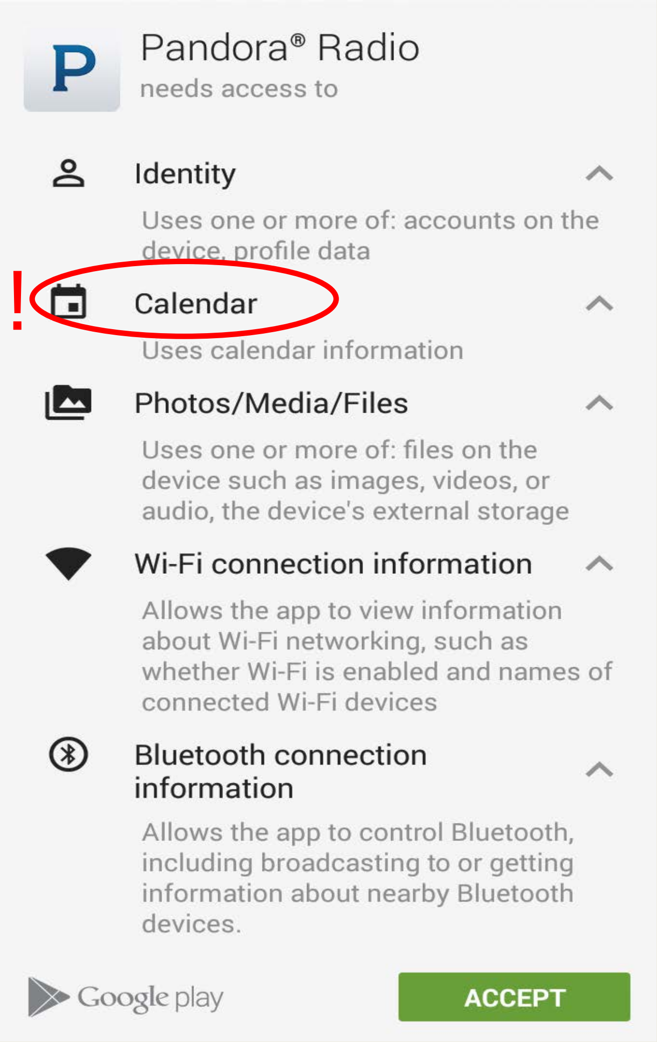

Pandora® Radio
needs access to

Identity
Uses one or more of: accounts on the device, profile data

Calendar
Uses calendar information

Photos/Media/Files
Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Wi-Fi connection information
Allows the app to view information about Wi-Fi networking, such as whether Wi-Fi is enabled and names of connected Wi-Fi devices

Bluetooth connection information
Allows the app to control Bluetooth, including broadcasting to or getting information about nearby Bluetooth devices.

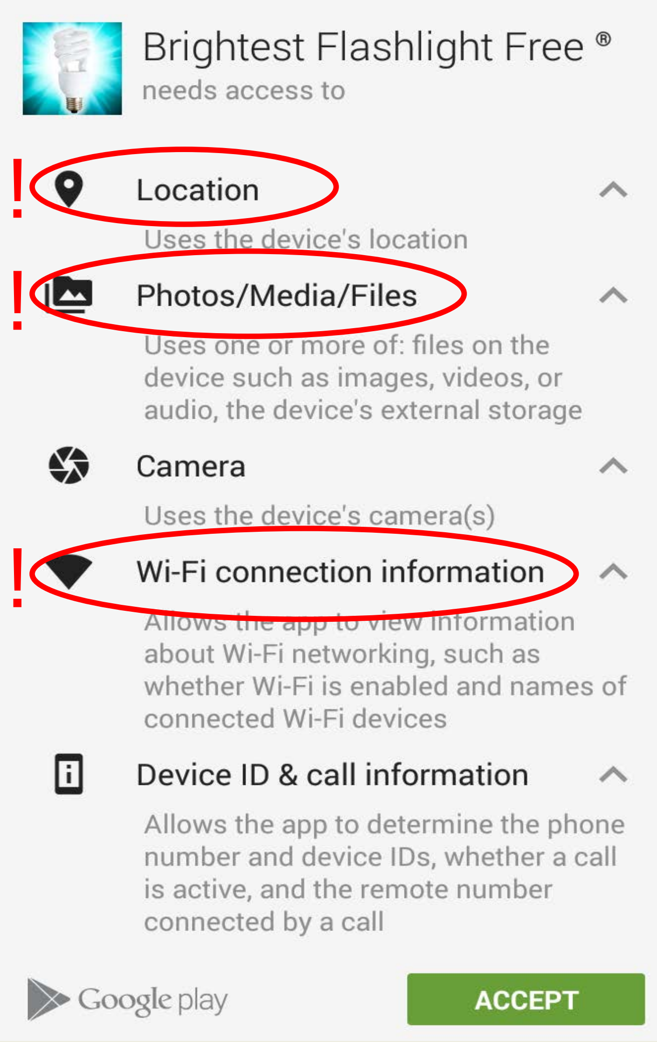

Brightest Flashlight Free ®
needs access to

Location
Uses the device's location

Photos/Media/Files
Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera
Uses the device's camera(s)

Wi-Fi connection information
Allows the app to view information about Wi-Fi networking, such as whether Wi-Fi is enabled and names of connected Wi-Fi devices

Device ID & call information
Allows the app to determine the phone number and device IDs, whether a call is active, and the remote number connected by a call

What is “Necessary” Depends on the User

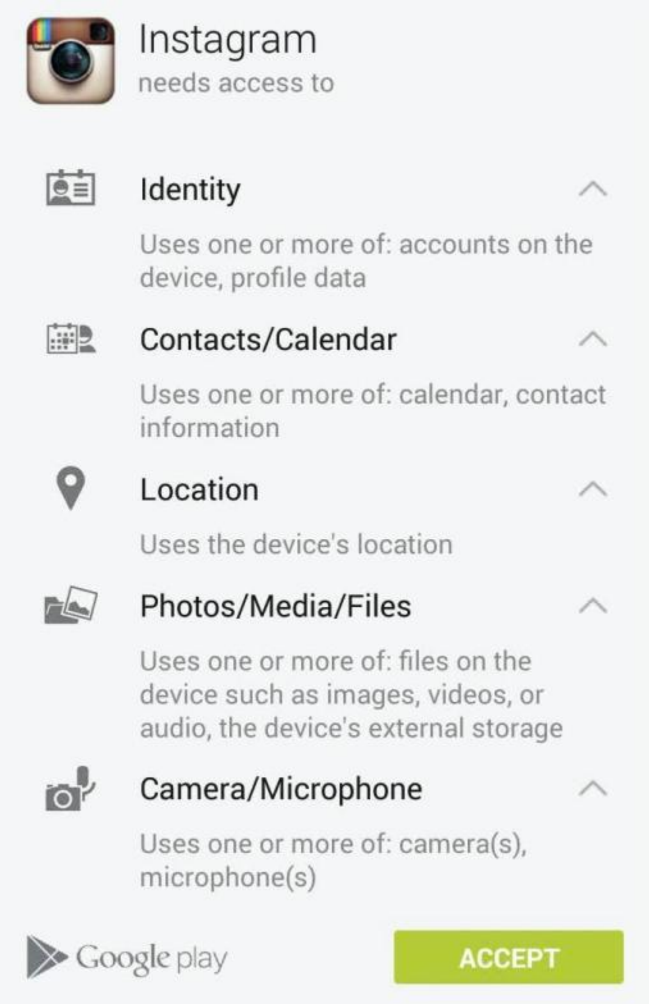

Instagram
needs access to

Identity

Uses one or more of: accounts on the device, profile data

Contacts/Calendar

Uses one or more of: calendar, contact information

Location

Uses the device's location

Photos/Media/Files

Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera/Microphone

Uses one or more of: camera(s), microphone(s)

What is “Necessary” Depends on the User

I never tag my **location** in Instagram

I rarely access my **contacts** to look for friends

I rarely **post** photos and videos

Which Permissions Can Be Disabled?

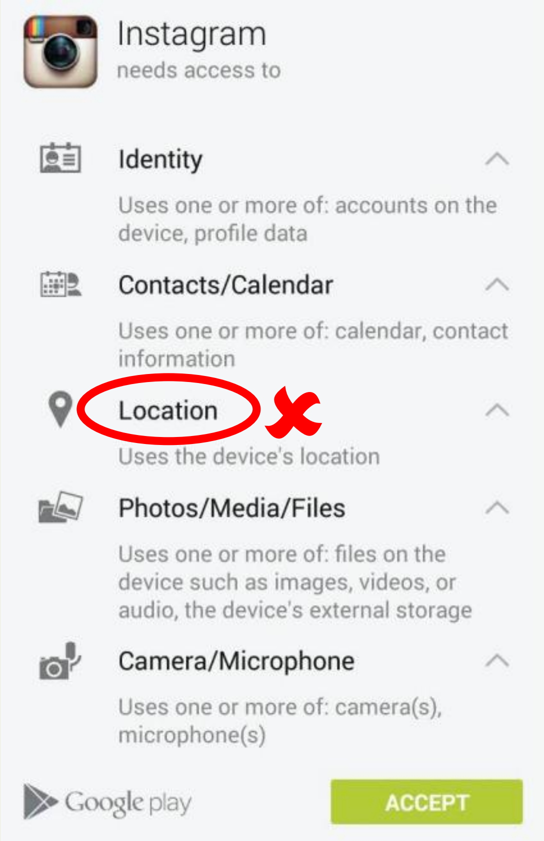

Instagram
needs access to

Identity
Uses one or more of: accounts on the device, profile data

Contacts/Calendar
Uses one or more of: calendar, contact information

Location
Uses the device's location

Photos/Media/Files
Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera/Microphone
Uses one or more of: camera(s), microphone(s)

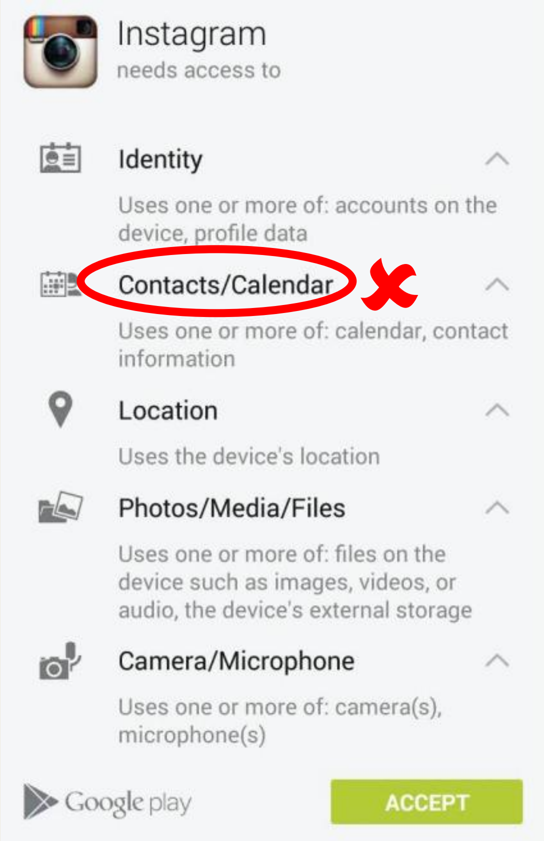

Instagram
needs access to

Identity
Uses one or more of: accounts on the device, profile data

Contacts/Calendar
Uses one or more of: calendar, contact information

Location
Uses the device's location

Photos/Media/Files
Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera/Microphone
Uses one or more of: camera(s), microphone(s)

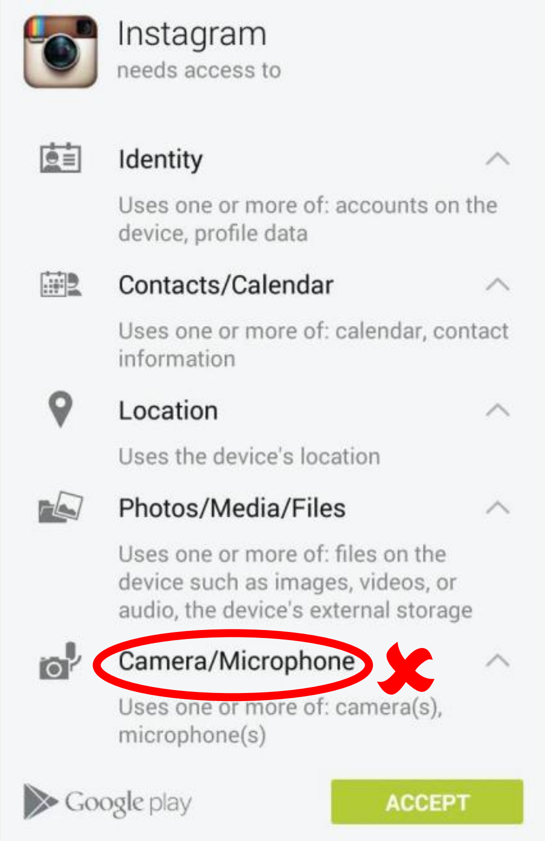

Instagram
needs access to

Identity
Uses one or more of: accounts on the device, profile data

Contacts/Calendar
Uses one or more of: calendar, contact information

Location
Uses the device's location

Photos/Media/Files
Uses one or more of: files on the device such as images, videos, or audio, the device's external storage

Camera/Microphone
Uses one or more of: camera(s), microphone(s)

Crowdsourcing Security Configs

Have Crowd Identify Usable Configs

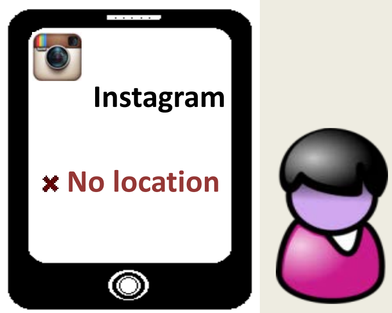

Instagram

✖ No location

Instagram

✖ No contacts

Instagram

✖ No camera

Crowdsourcing Security Configs

Have Crowd Identify Usable Configs

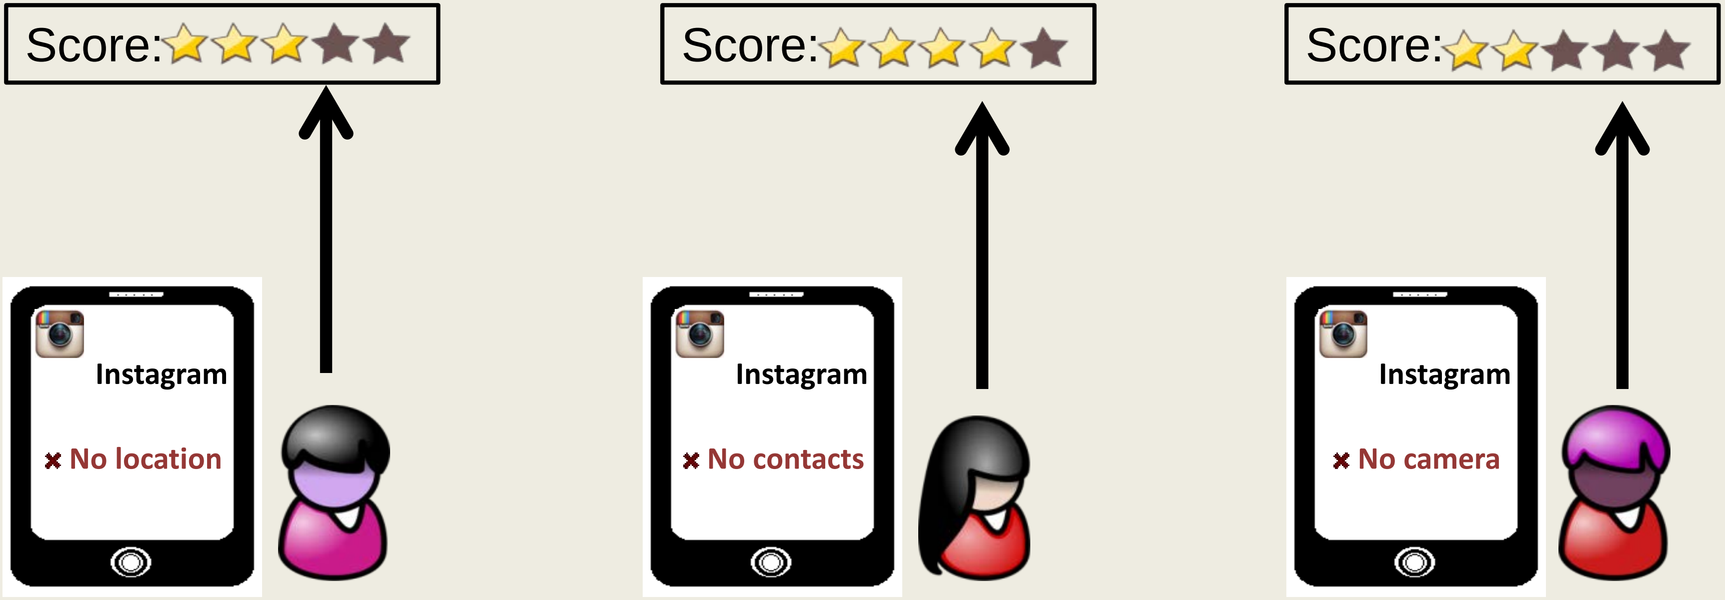

Score: ★★★☆☆

Instagram

✖ No location

Score: ★★★★☆

Instagram

✖ No contacts

Score: ★★☆☆☆

Instagram

✖ No camera

Crowdsourcing Security Configs

Have Crowd Identify Usable Configs

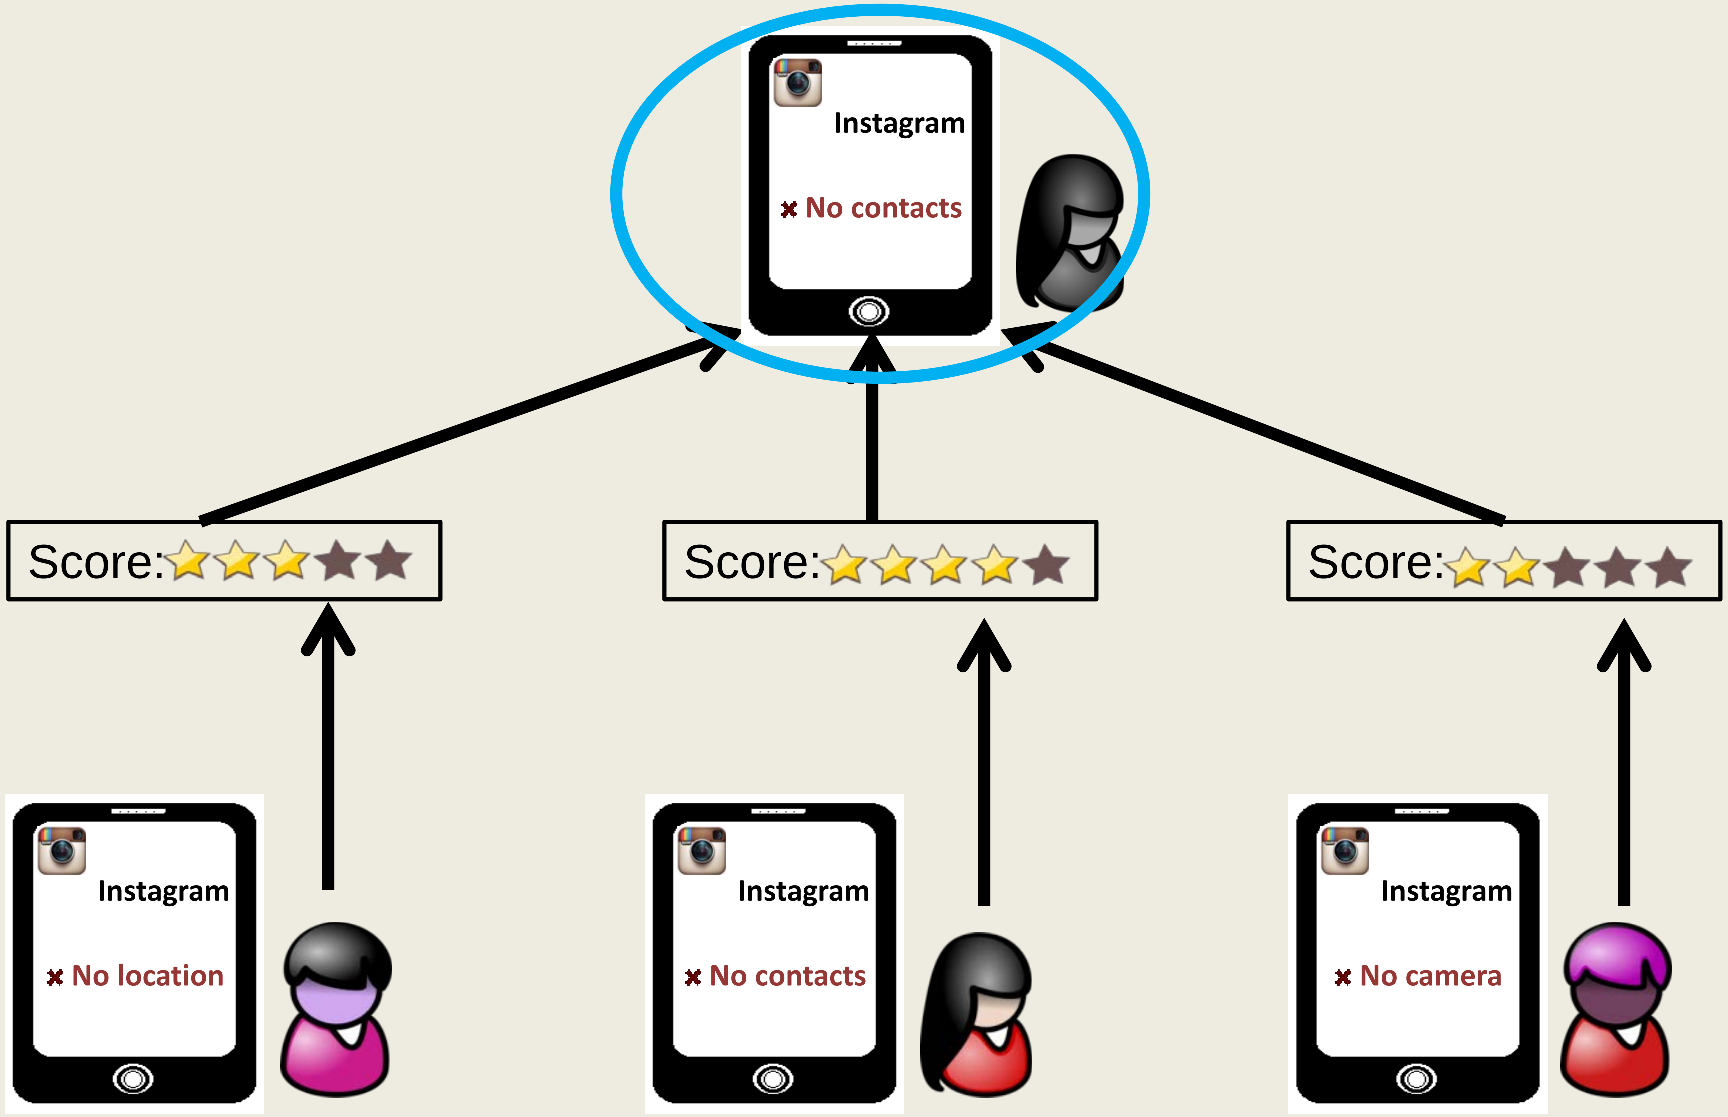

Crowdsourcing Security Configs

Have Crowd Identify Usable Configs

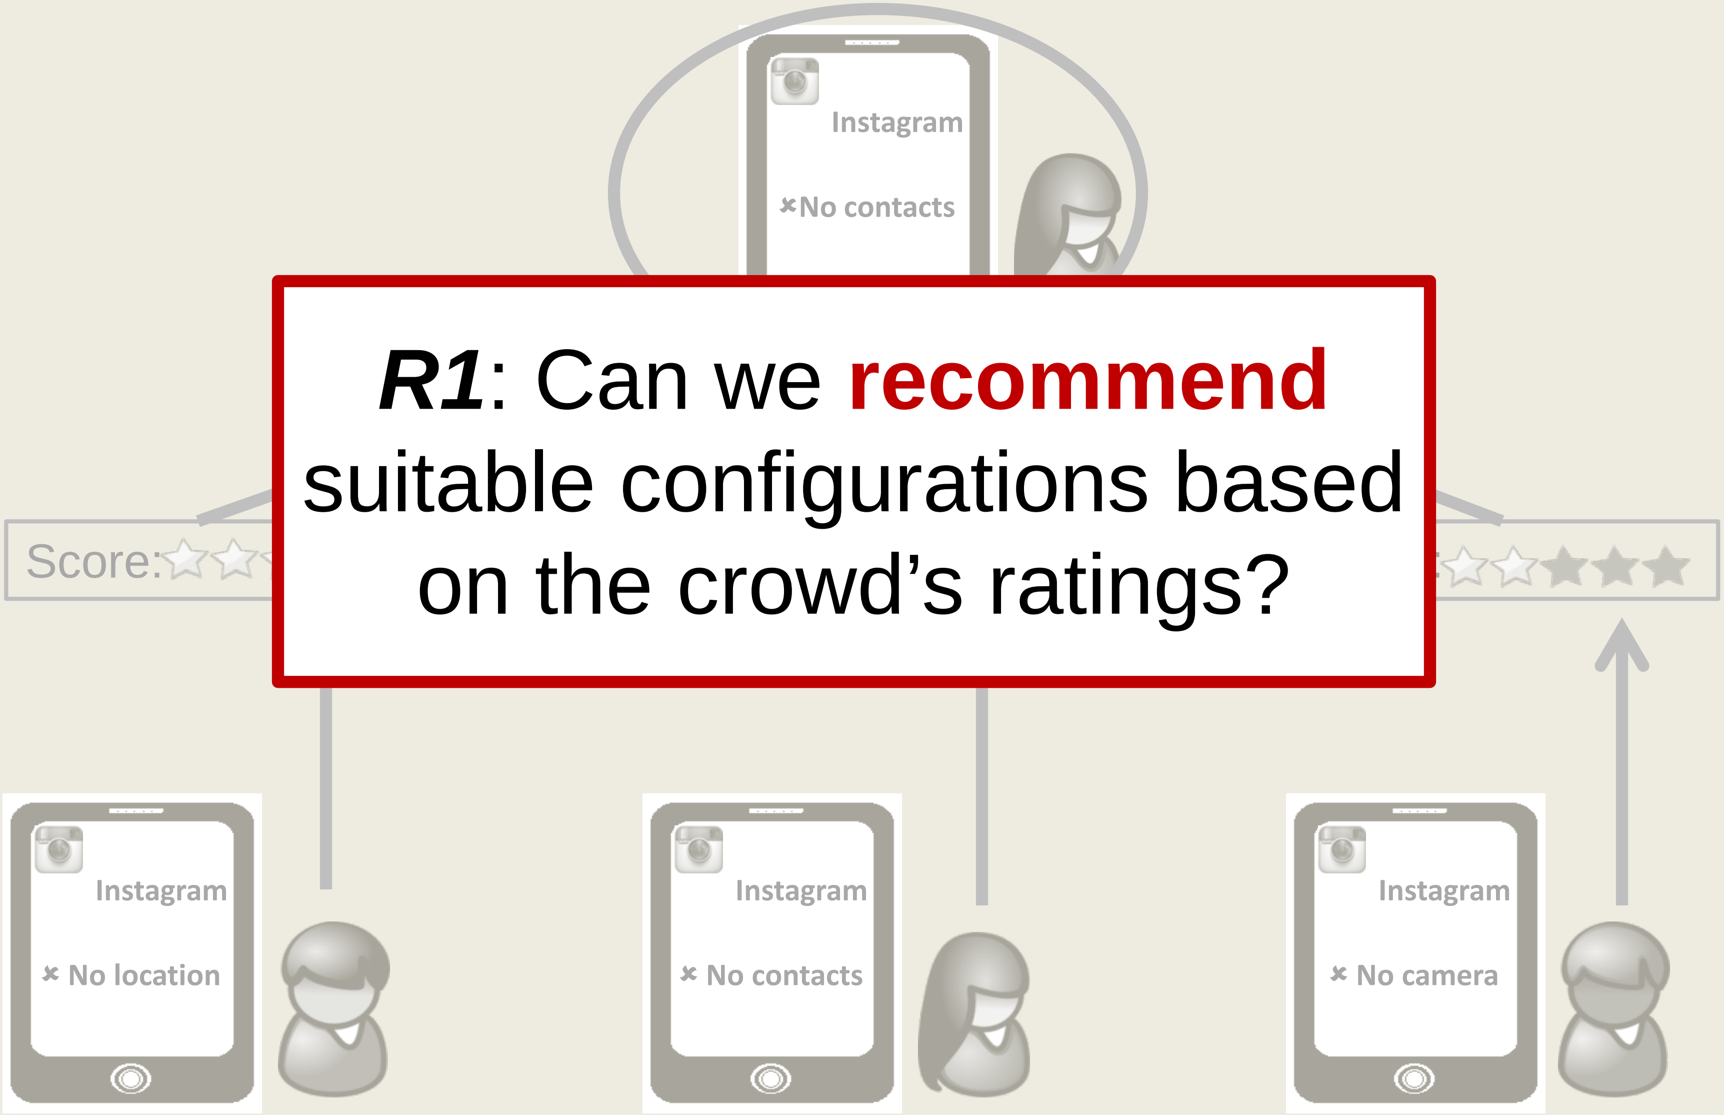

R1: Can we **recommend** suitable configurations based on the crowd’s ratings?

Crowdsourcing Security Configs

User-Based Collaborative Filtering

	✖Contacts	✖ Location	✖Camera	✖Mic
1	👍	👎	👍	👍
2		👍	👎	👎
3	👍	👍	👎	
4	👎		👍	
5	👍	👍		👎

Crowdsourcing Security Configs

User-Based Collaborative Filtering

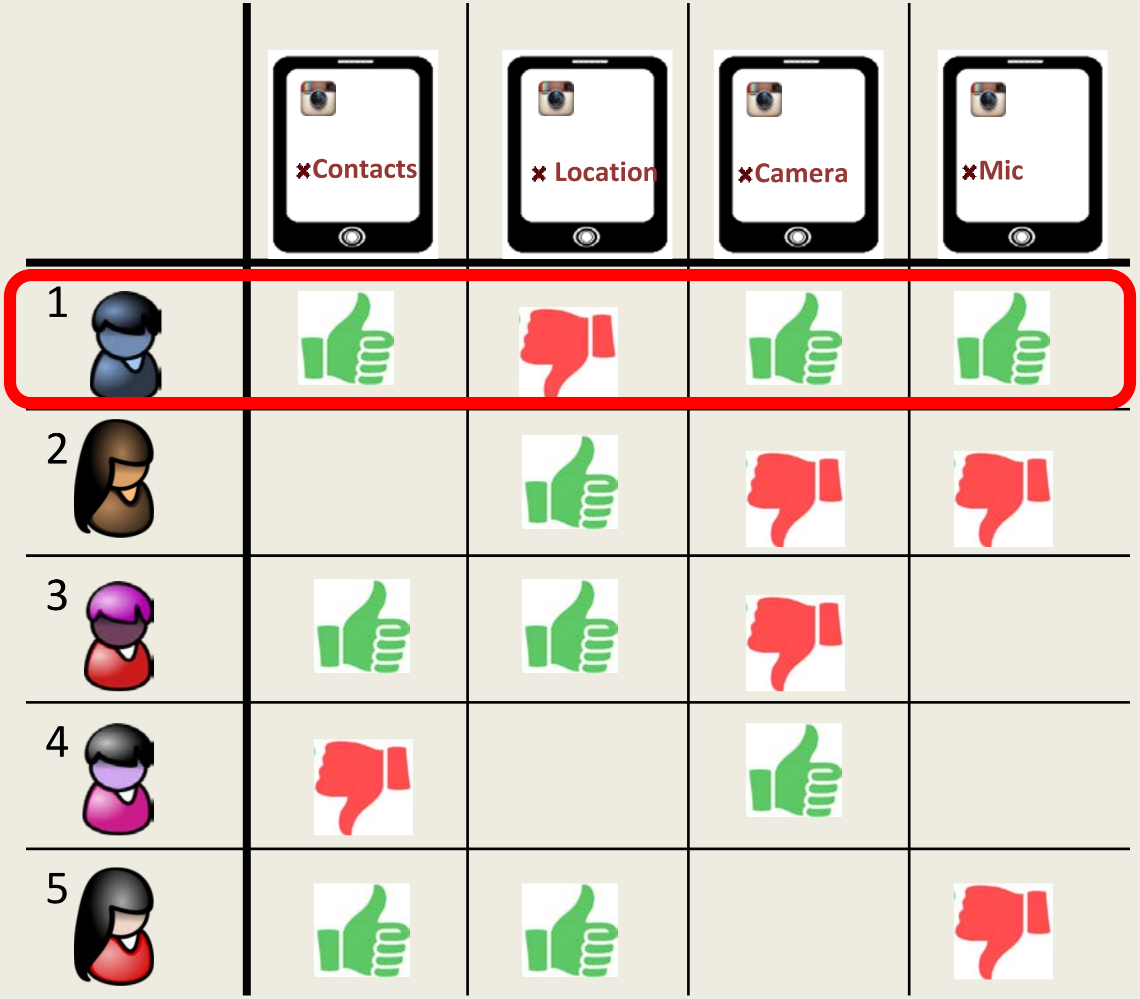

	✖Contacts	✖ Location	✖Camera	✖Mic
1	👍	👎	👍	👍
2		👍	👎	👎
3	👍	👍	👎	
4	👎		👍	
5	👍	👍		👎

Will 👤 Like the Config with No Camera?

	✖Contacts	✖ Location	✖Camera	✖Mic
1	👍	👎	👍	👍
2		👍	👎	👎
3	👍	👍	👎	
4	👎		👍	
5	👍	👍	?	👎

Will 👤 Like the Config with No Camera?

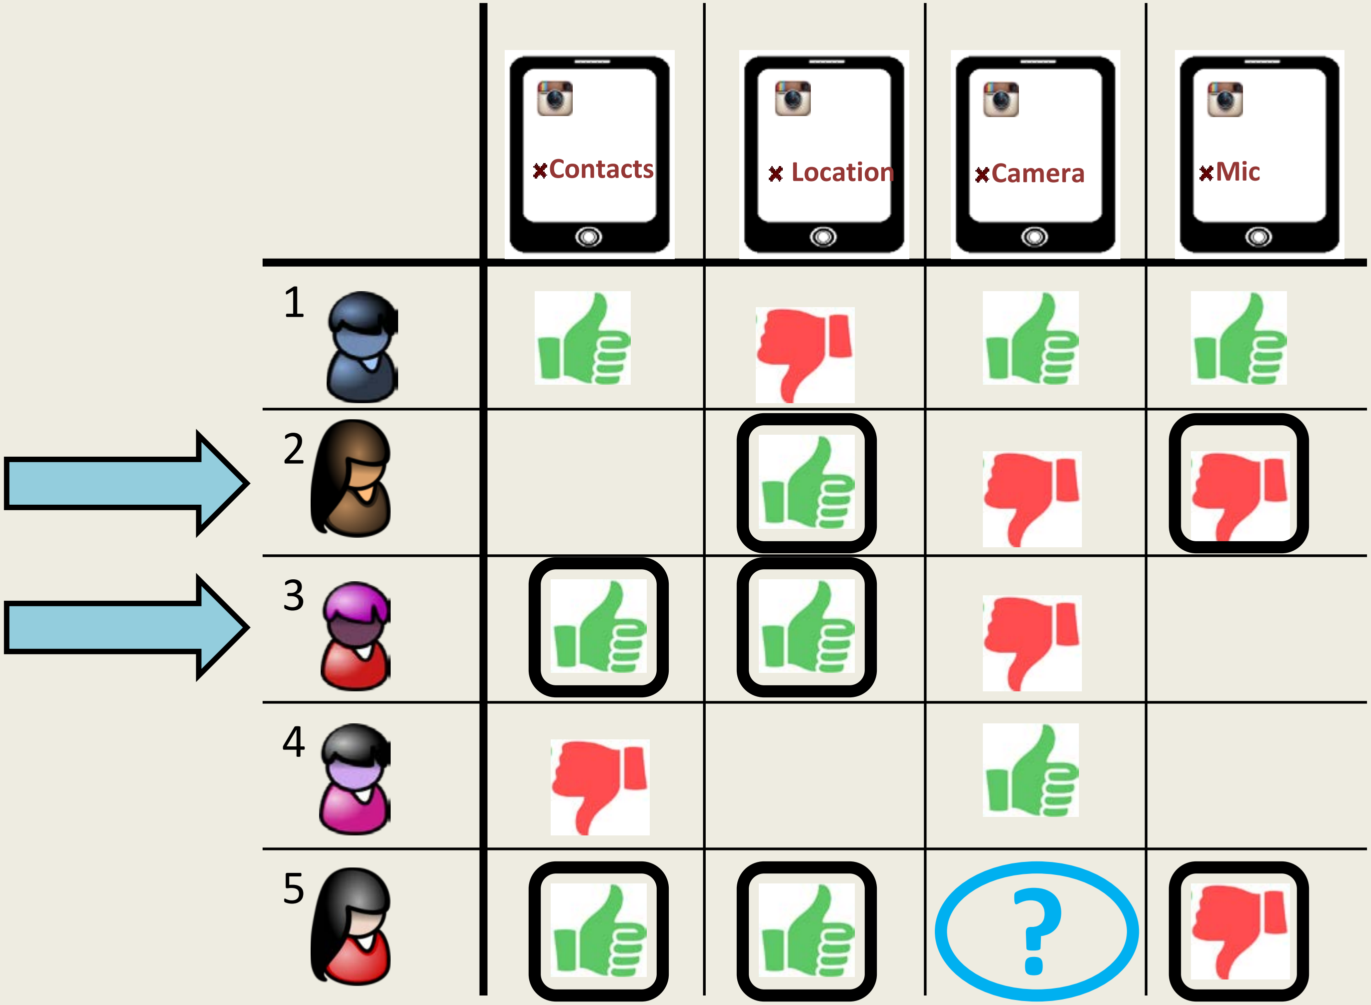

	✗Contacts	✗ Location	✗Camera	✗Mic
1 👤	👍	👎	👍	👍
2 👤		👍	👎	👎
3 👤	👍	👍	👎	
4 👤	👎		👍	
5 👤	👍	👍	?	👎

Will 👤 Like the Config with No Camera?

		✖Contacts	✖ Location	✖Camera	✖Mic
	1 👤	👍	👎	👍	👍
➡	2 👤		👍	👎	👎
➡	3 👤	👍	👍	👎	
	4 👤	👎		👍	
	5 👤	👍	👍	👎	👎

Crowdsourcing Security Configs

How Many Configurations to Test?

Crowdsourcing Security Configs

How Many Configurations to Test?

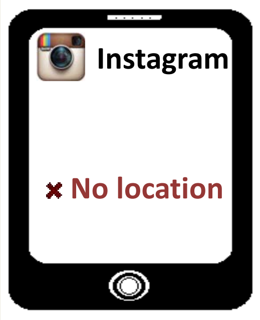

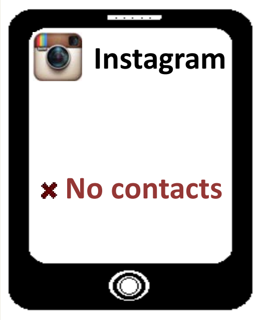

How Many Configurations to Test?

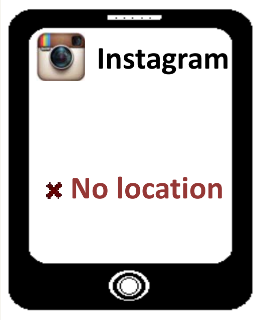

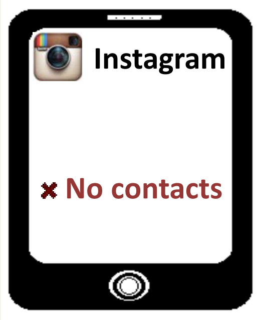

Instagram
✖ No microphone

Instagram
✖ No camera

Crowdsourcing Security Configs

How Many Configurations to Test?

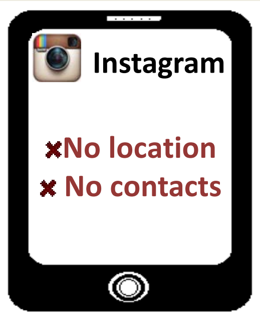
Instagram

✖No location
✖ No contacts

Instagram

✖ No location
✖ No microphone

Instagram

✖ No location
✖ No camera

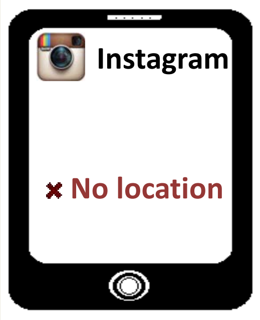
Instagram

✖ No location

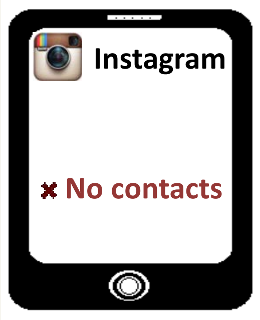
Instagram

✖ No contacts

Instagram

✖ No microphone

Instagram

✖ No camera

How Many Configurations to Test?

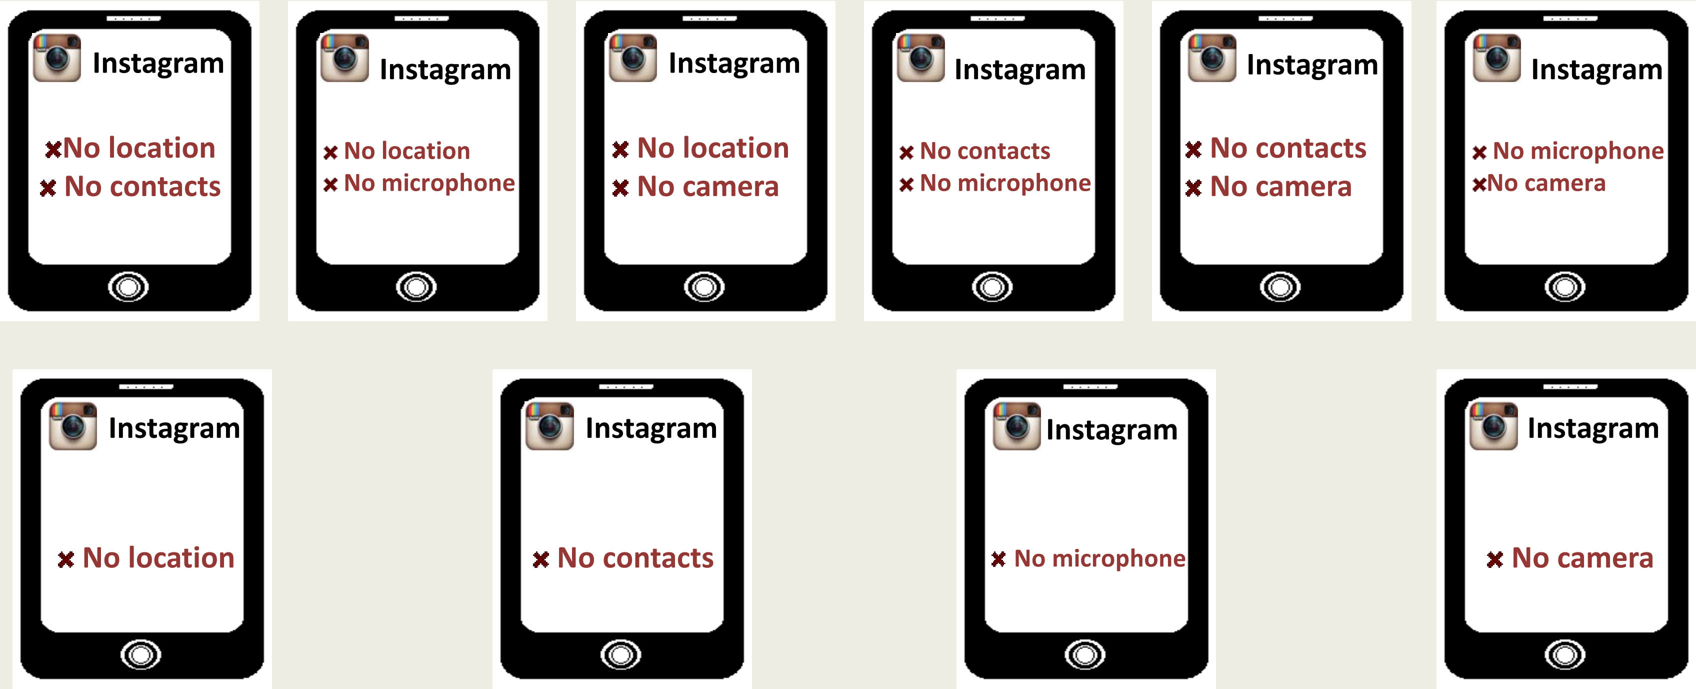

Instagram
✖No location
✖ No contacts

Instagram
✖ No location
✖ No microphone

Instagram
✖ No location
✖ No camera

Instagram
✖ No contacts
✖ No microphone

Instagram
✖ No contacts
✖ No camera

Instagram
✖ No microphone
✖No camera

Instagram
✖ No location

Instagram
✖ No contacts

Instagram
✖ No microphone

Instagram
✖ No camera

Crowdsourcing Security Configs

How Many Configurations to Test?

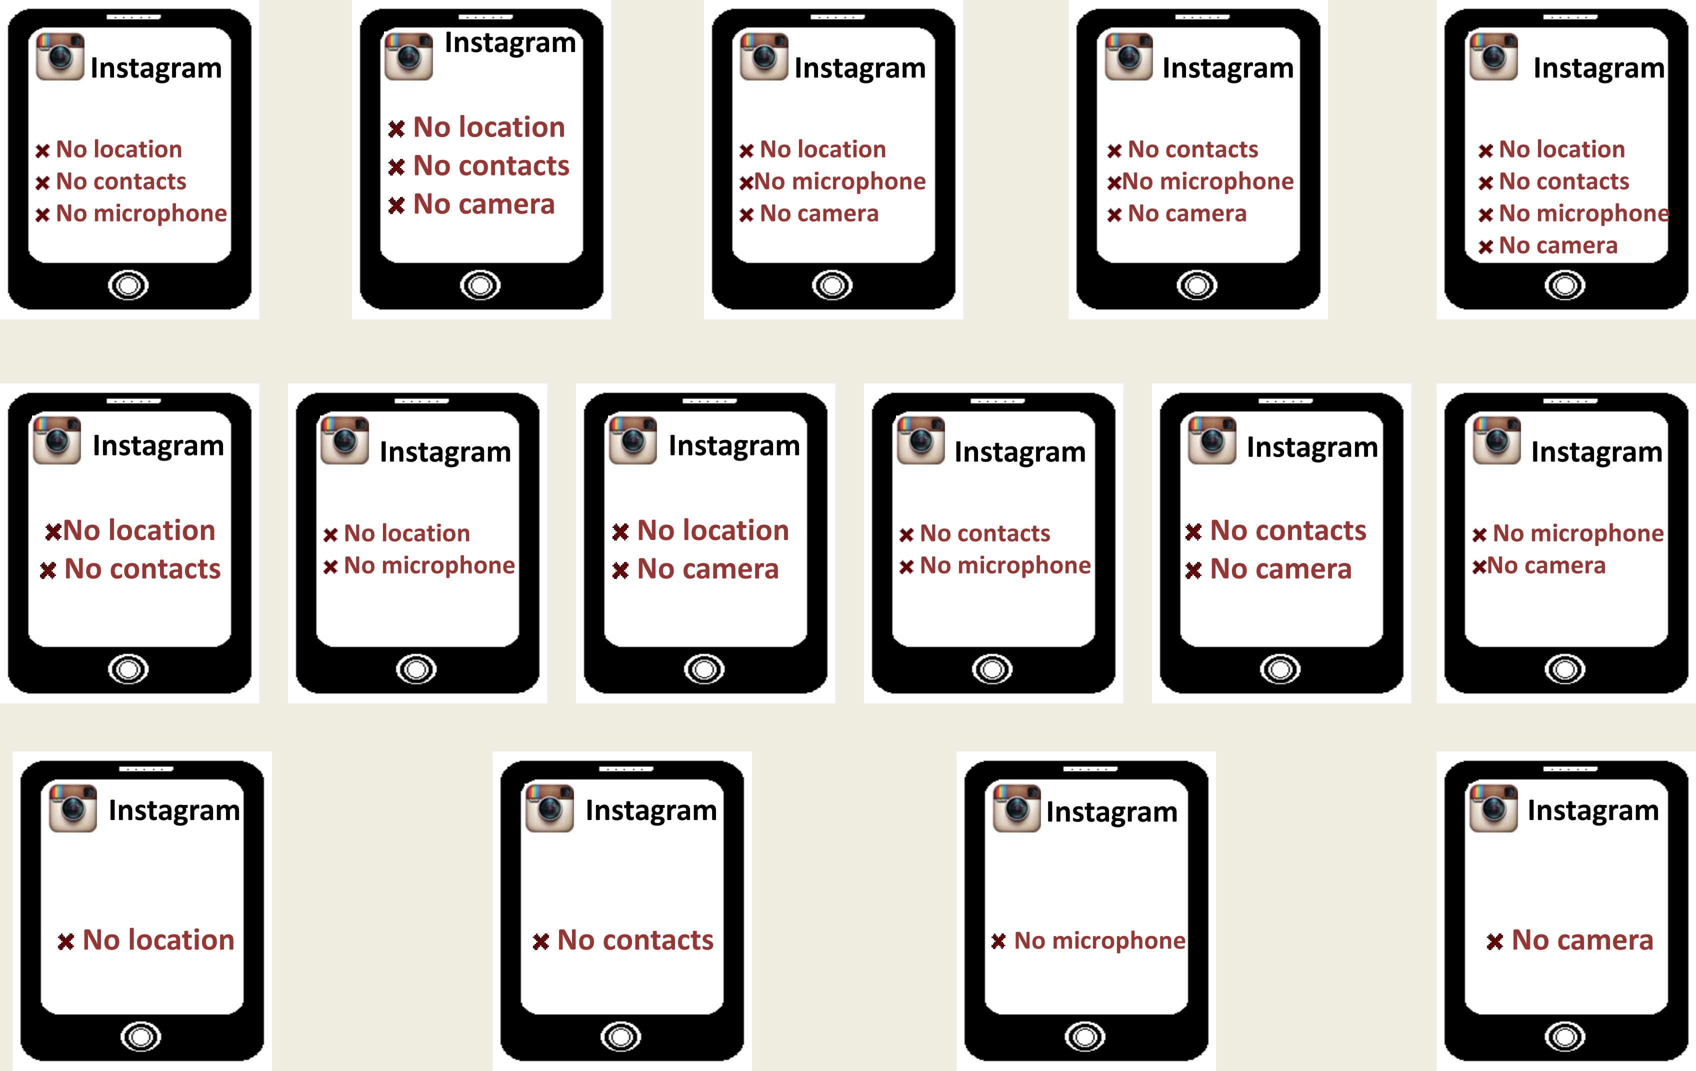

Instagram
✖ No location
✖ No contacts
✖ No microphone

Instagram
✖ No location
✖ No contacts
✖ No camera

Instagram
✖ No location
✖No microphone
✖ No camera

Instagram
✖ No contacts
✖No microphone
✖ No camera

Instagram
✖ No location
✖ No contacts
✖ No microphone
✖ No camera

Instagram
✖No location
✖ No contacts

Instagram
✖ No location
✖ No microphone

Instagram
✖ No location
✖ No camera

Instagram
✖ No contacts
✖ No microphone

Instagram
✖ No contacts
✖ No camera

Instagram
✖ No microphone
✖No camera

Instagram
✖ No location

Instagram
✖ No contacts

Instagram
✖ No microphone

Instagram
✖ No camera

Crowdsourcing Security Configs

How Many Configurations to Test?

How Many Configurations to Test?

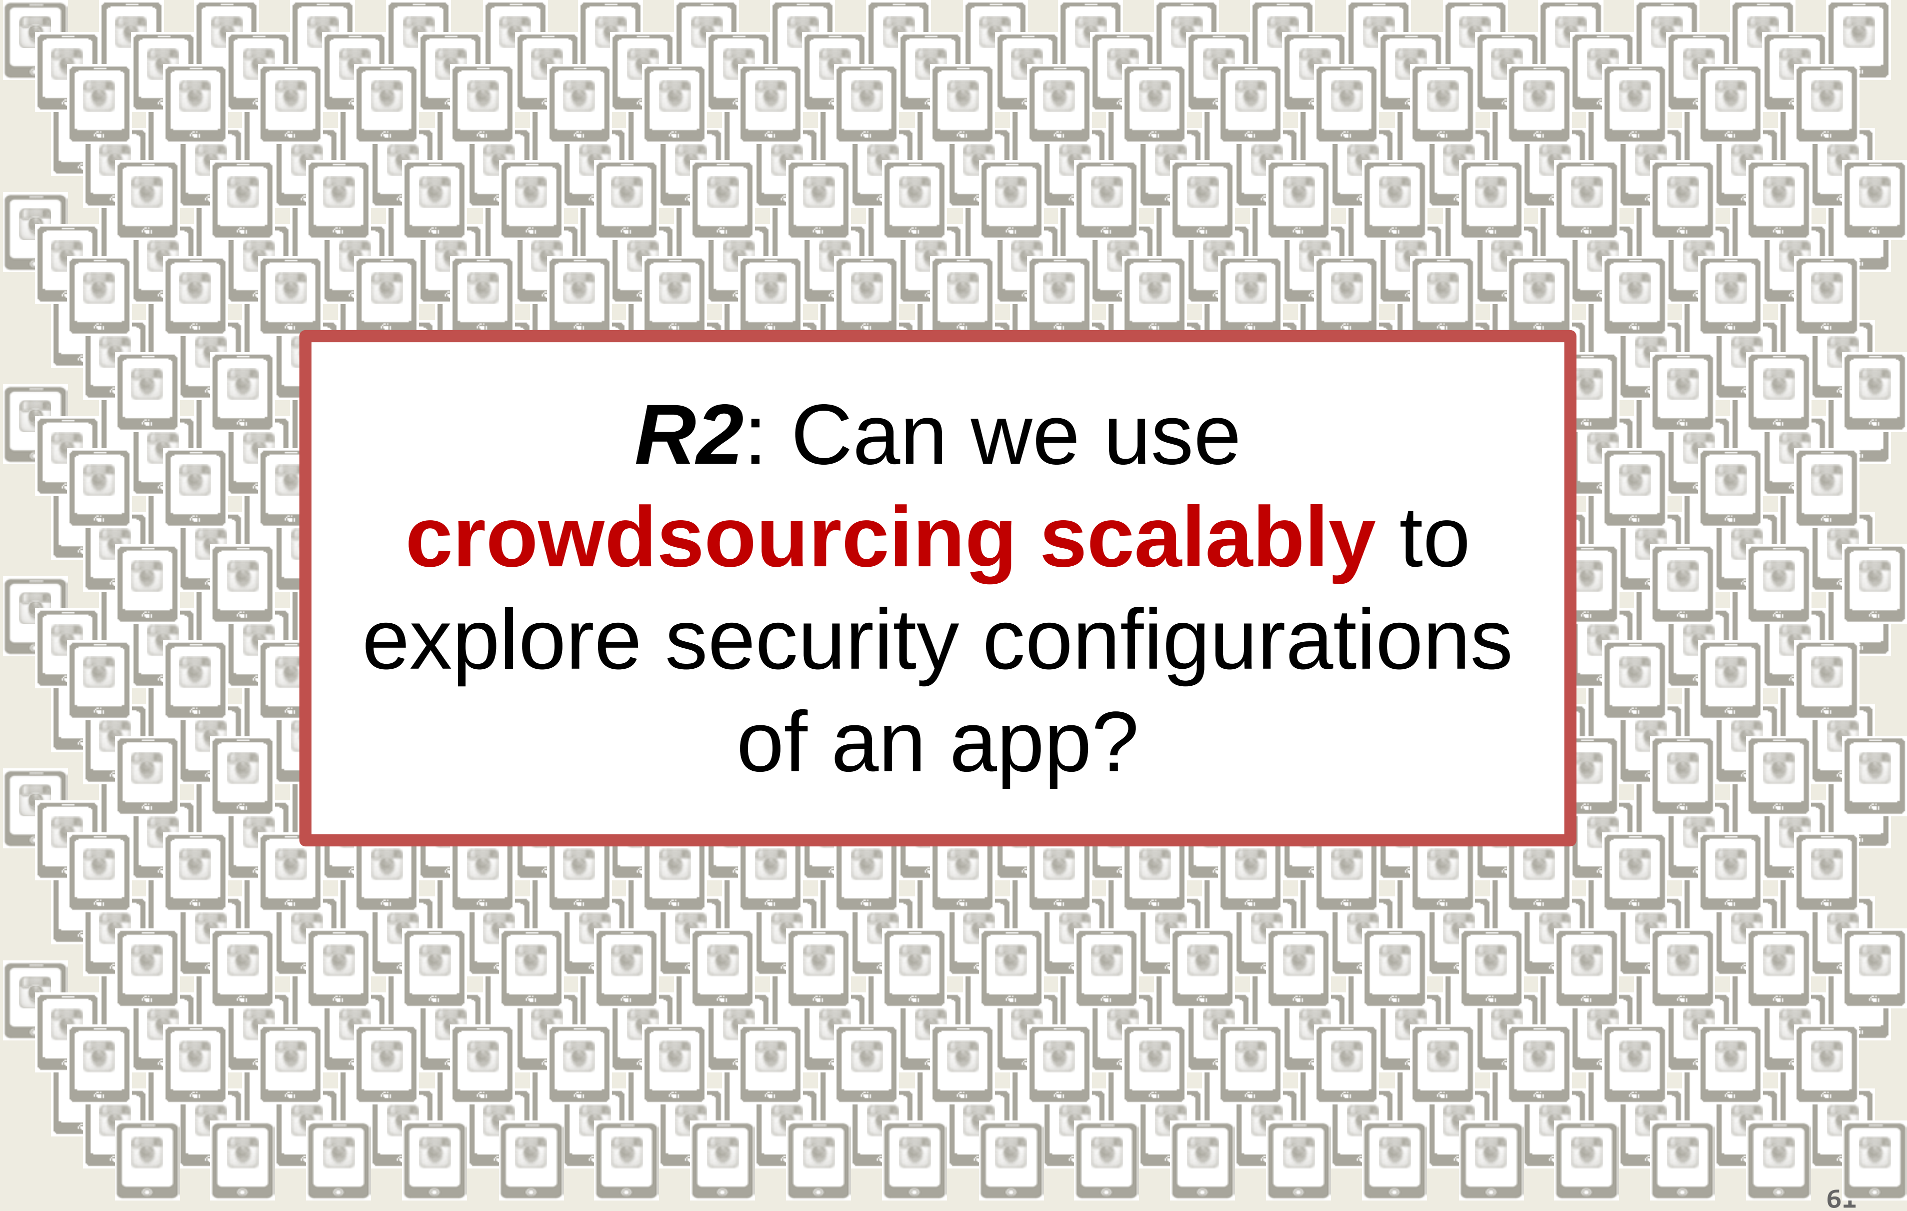

R2: Can we use **crowdsourcing scalably** to explore security configurations of an app?

Crowdsourcing Security Configs

A Lattice-Based Approach

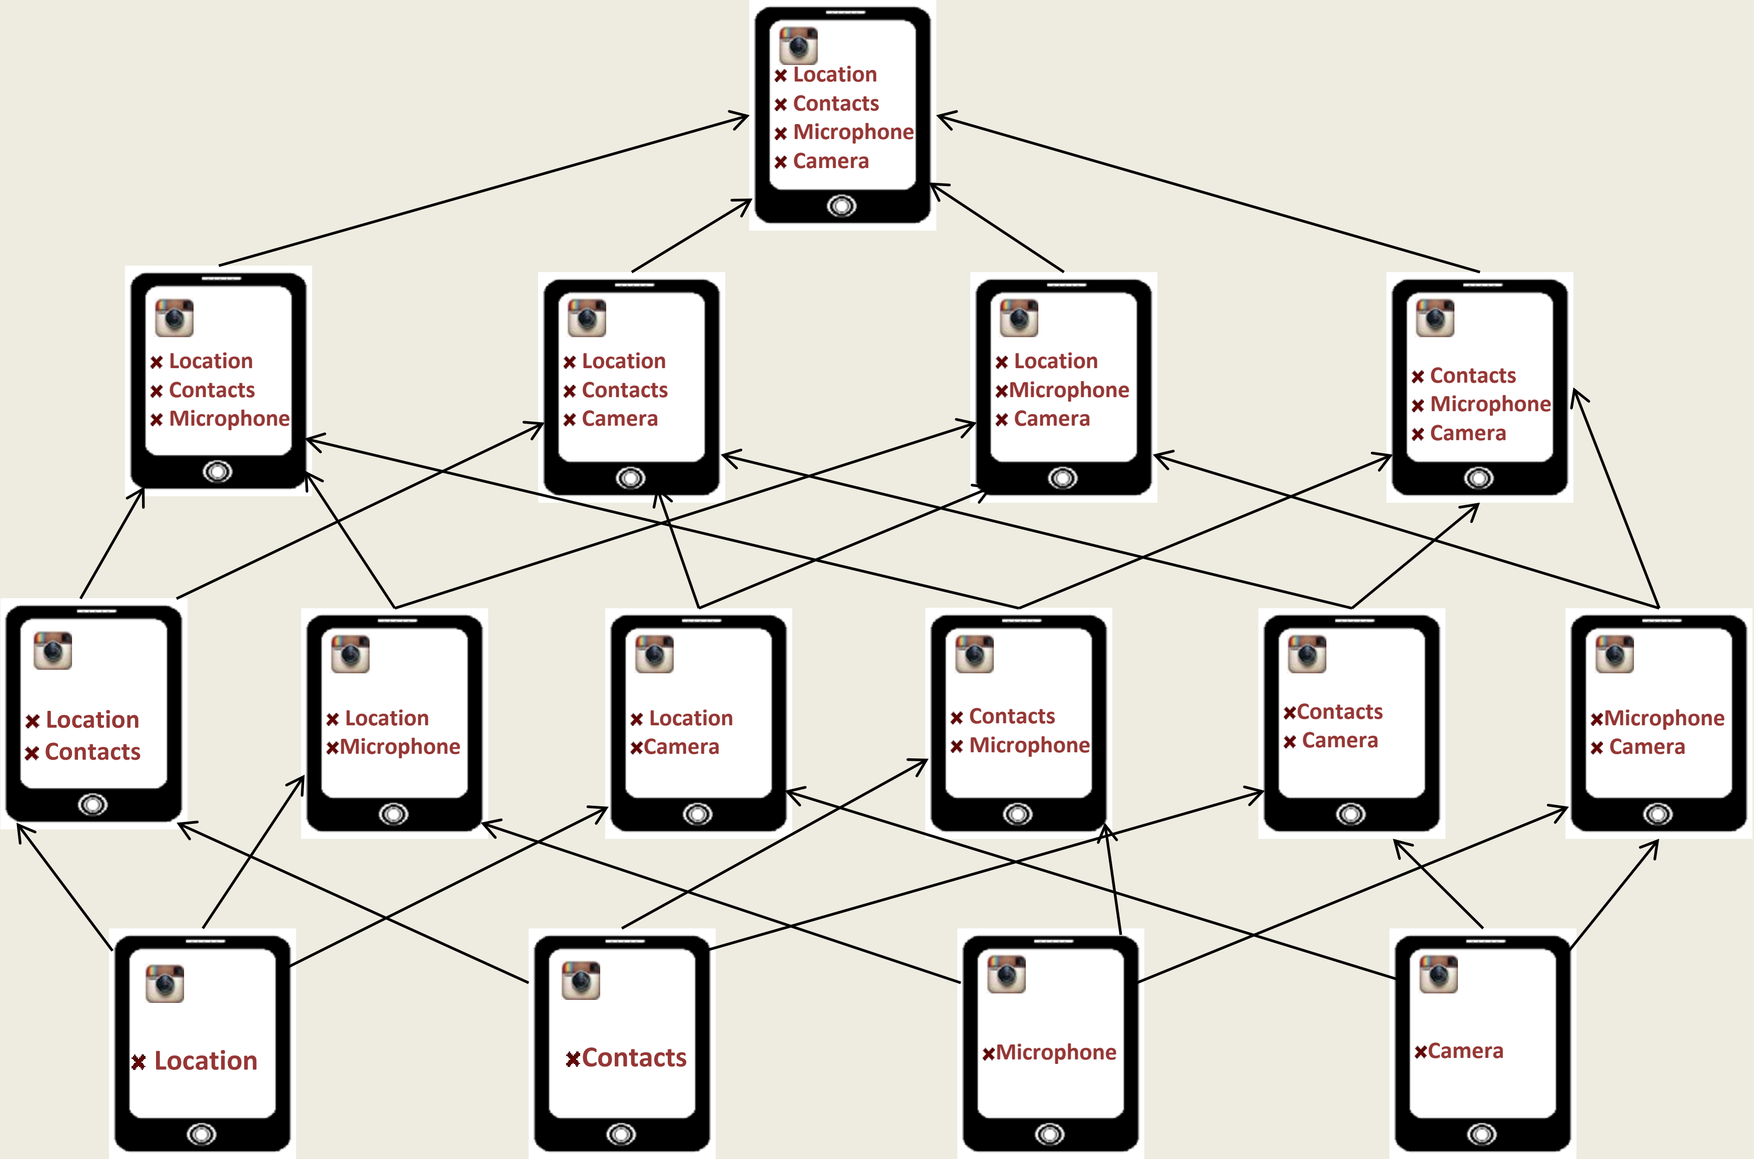

Crowdsourcing Security Configs
A Lattice-Based Approach

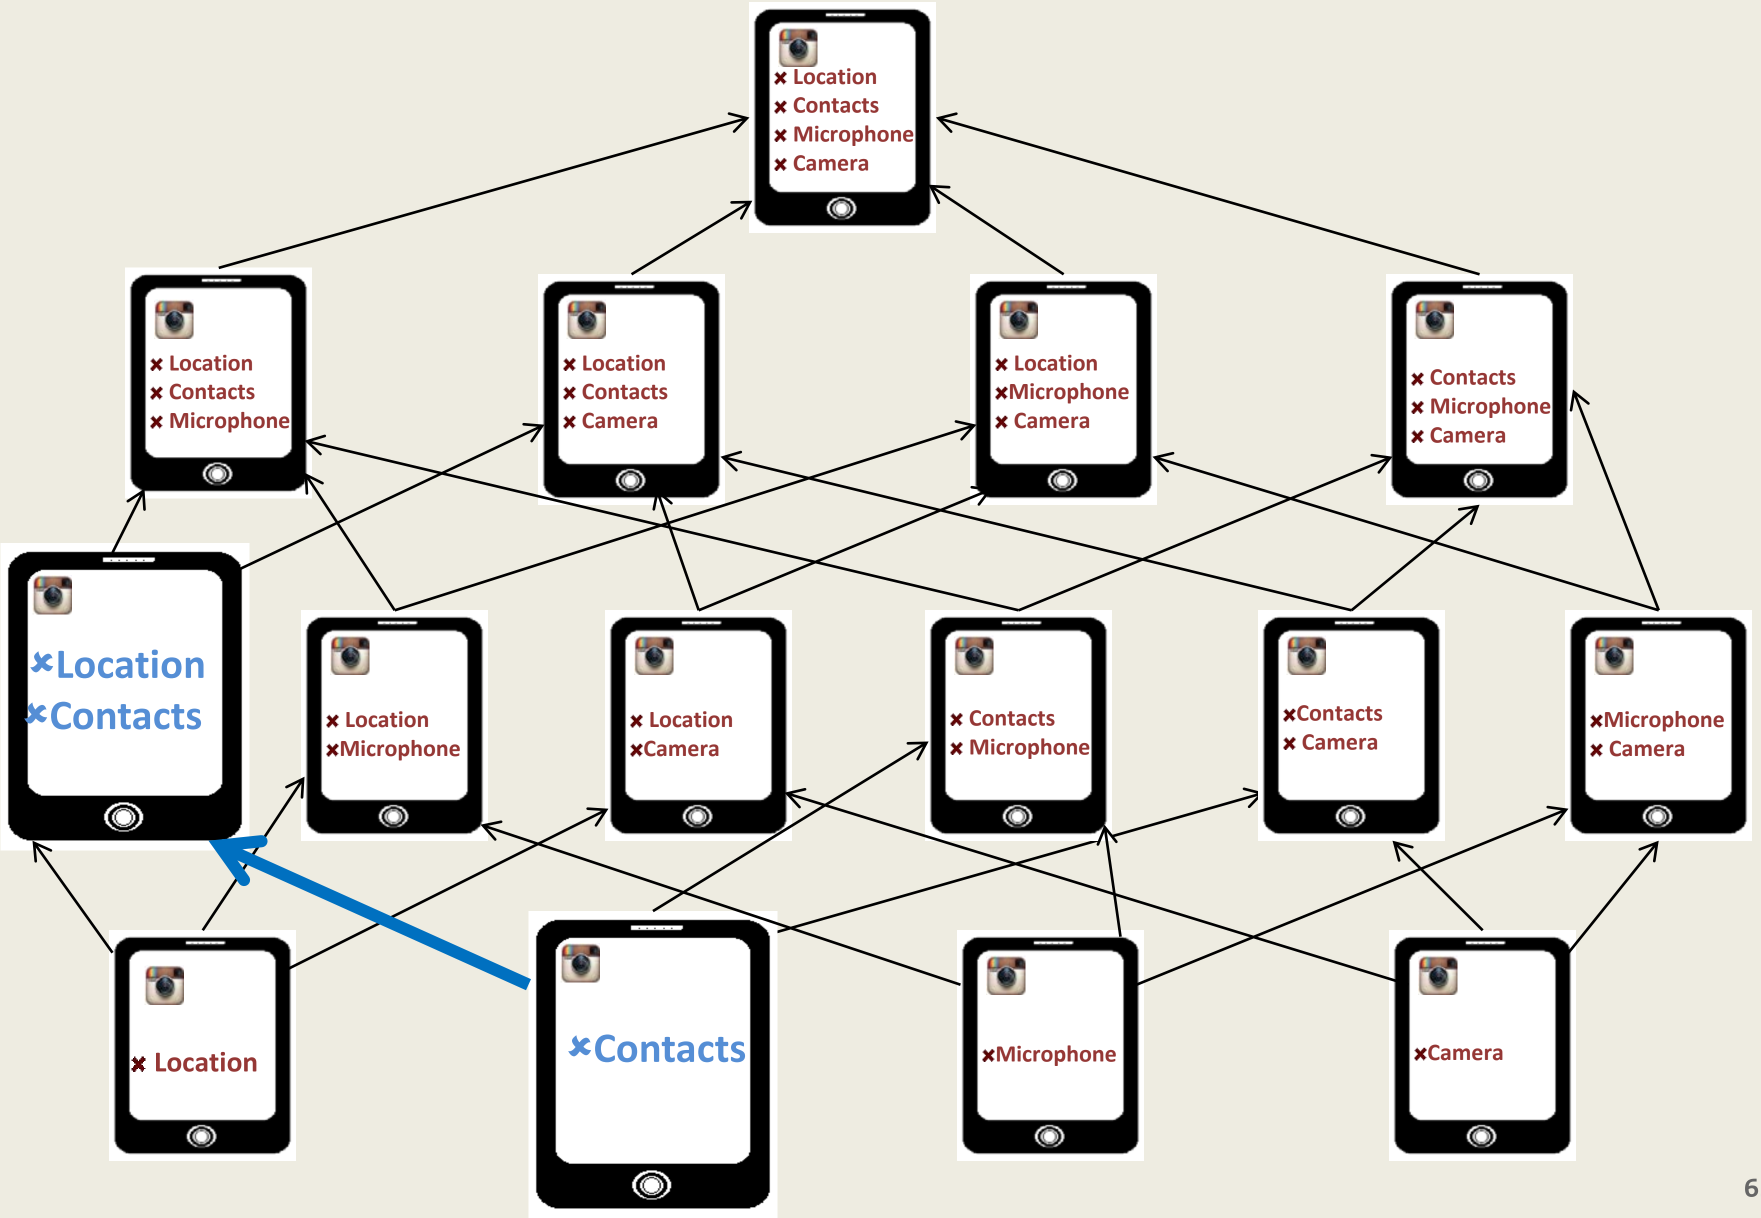

Crowdsourcing Security Configs

A Lattice-Based Approach

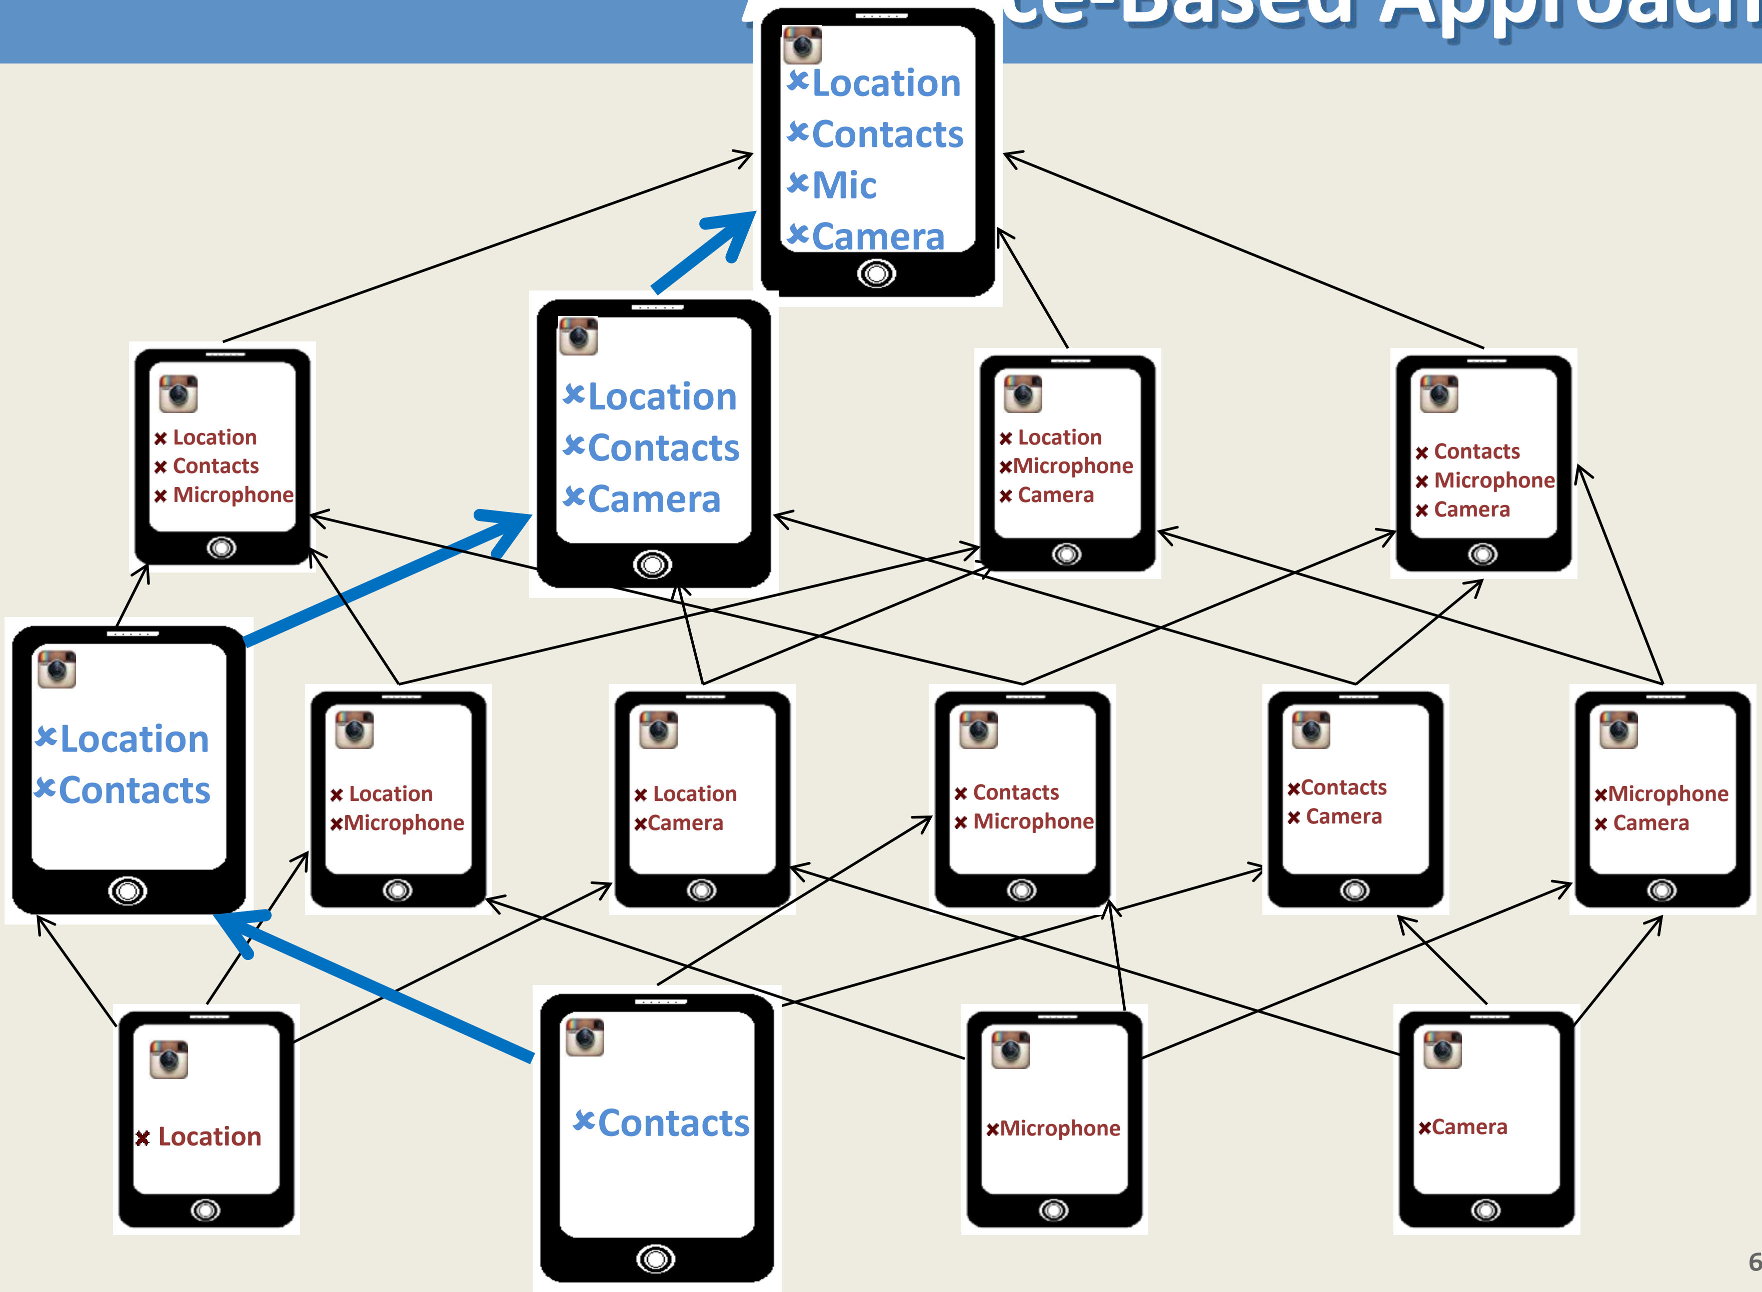

Crowdsourcing Security Configs

Crowd Explores the First Level

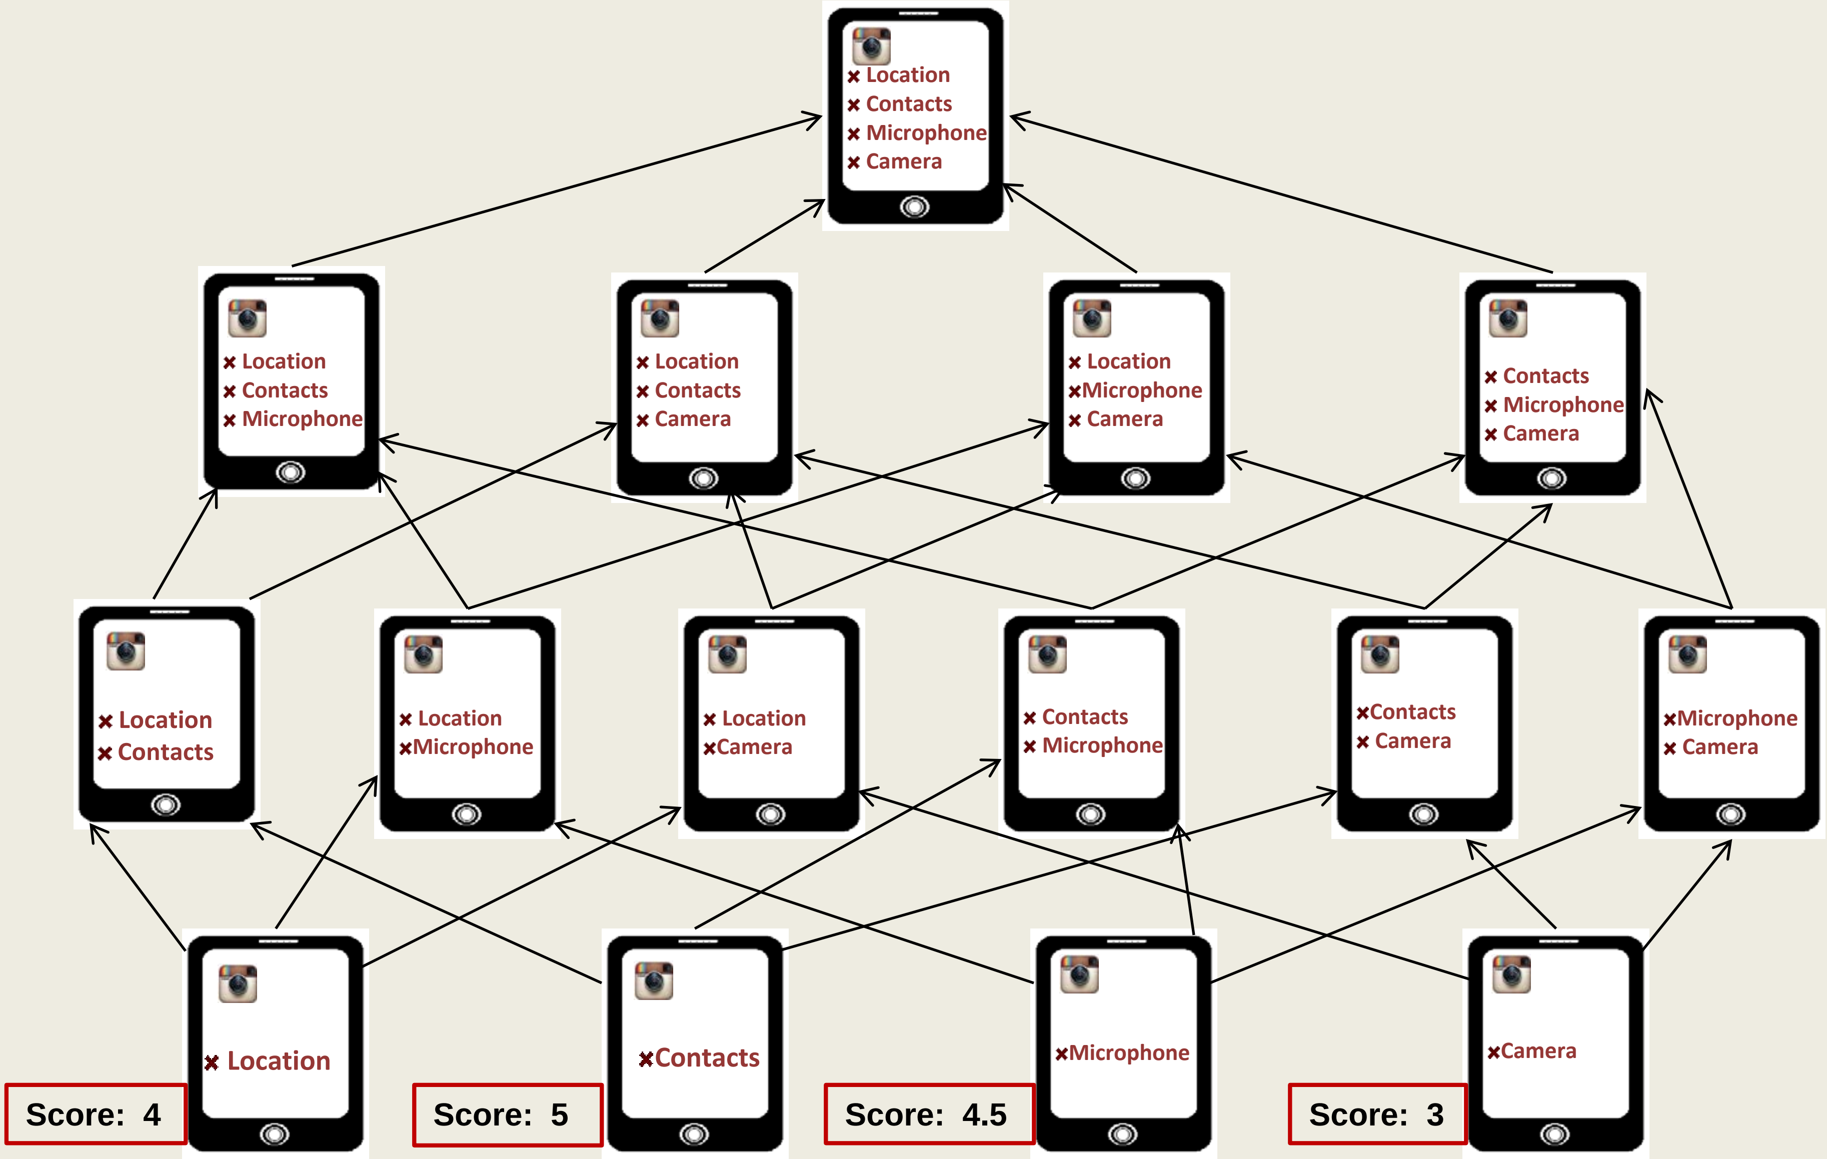

× Location
× Contacts
× Microphone
× Camera

× Location
× Contacts
× Microphone

× Location
× Contacts
× Camera

× Location
×Microphone
× Camera

× Contacts
× Microphone
× Camera

× Location
× Contacts

× Location
×Microphone

× Location
×Camera

× Contacts
× Microphone

×Contacts
× Camera

×Microphone
× Camera

× Location

×Contacts

×Microphone

×Camera

Score: 4

Score: 5

Score: 4.5

Score: 3

Low Score ⇒ Prune Node & Ancestors

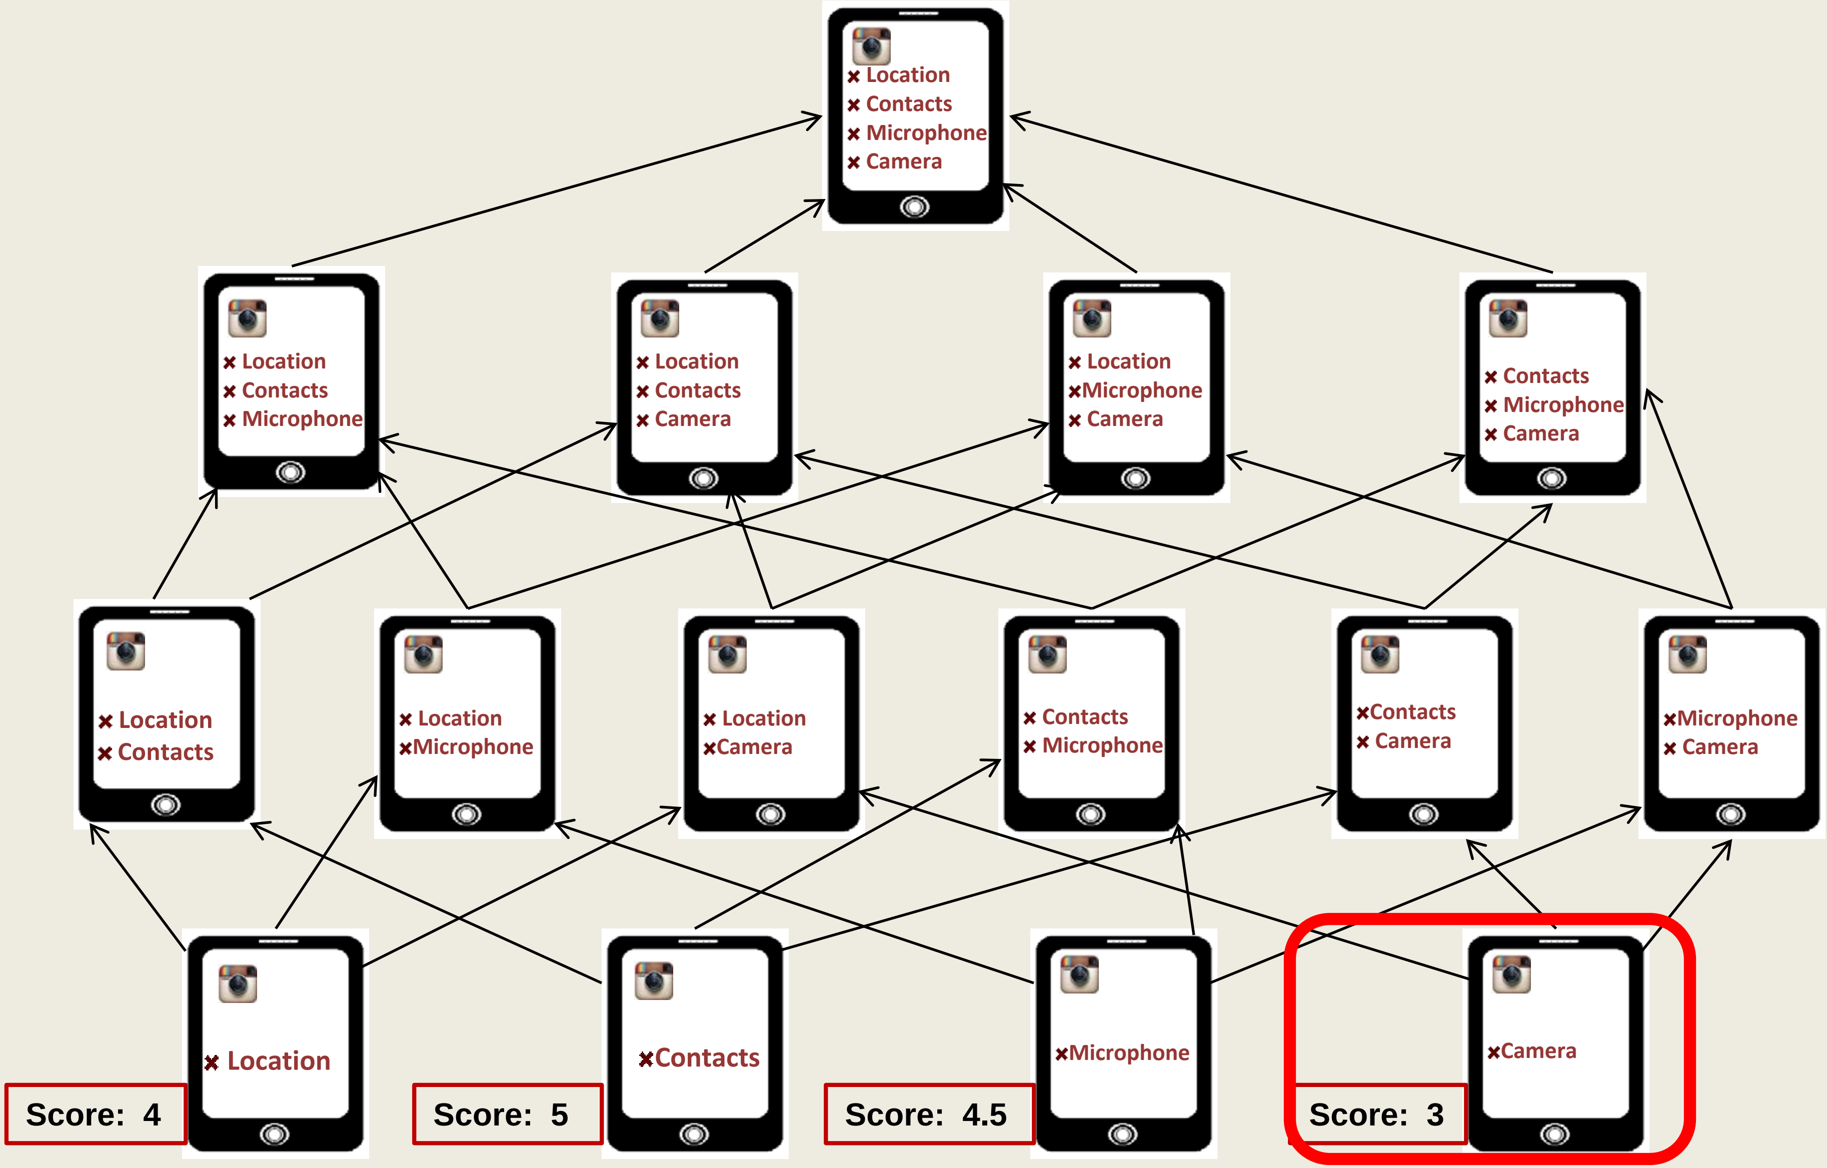

Crowdsourcing Security Configs

Low Score ⇒ Prune Node & Ancestors

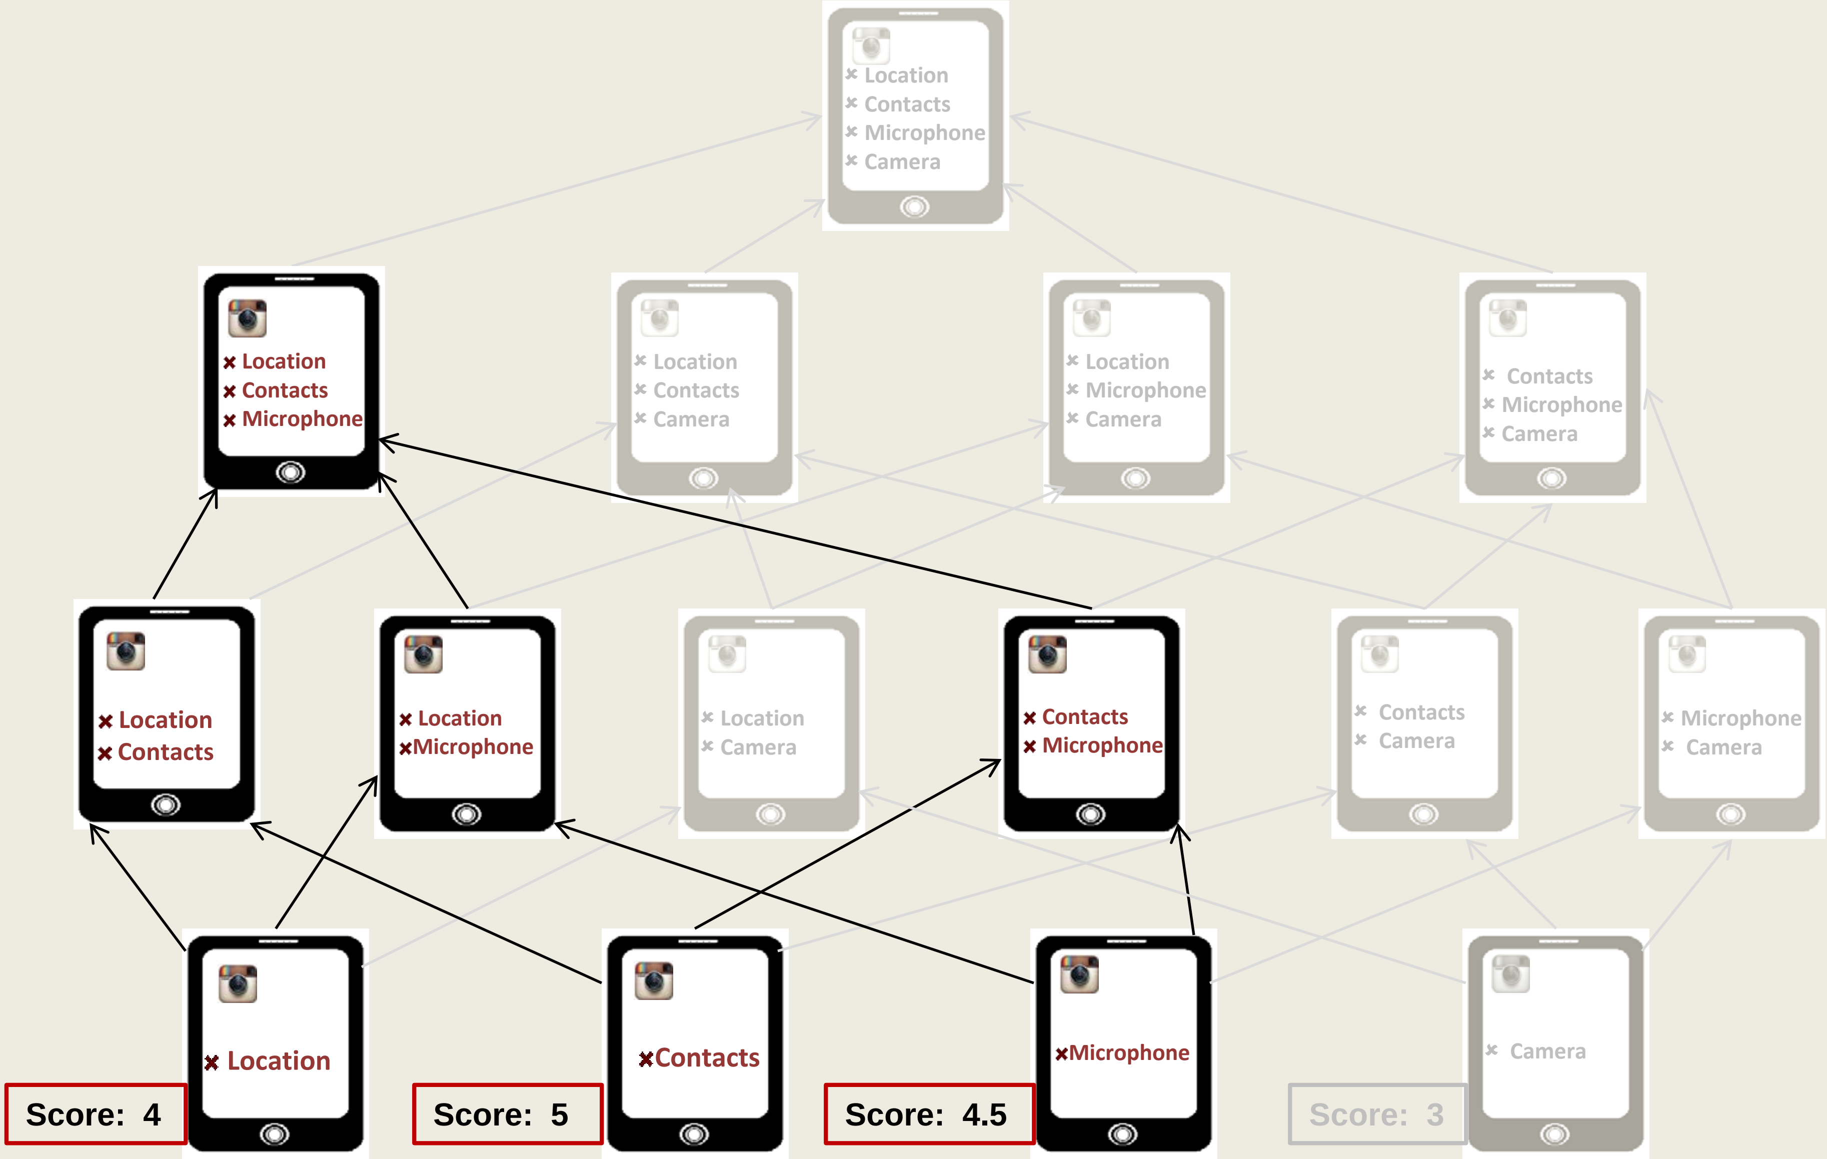

Crowdsourcing Security Configs

Crowd Explores Rest of Second Level

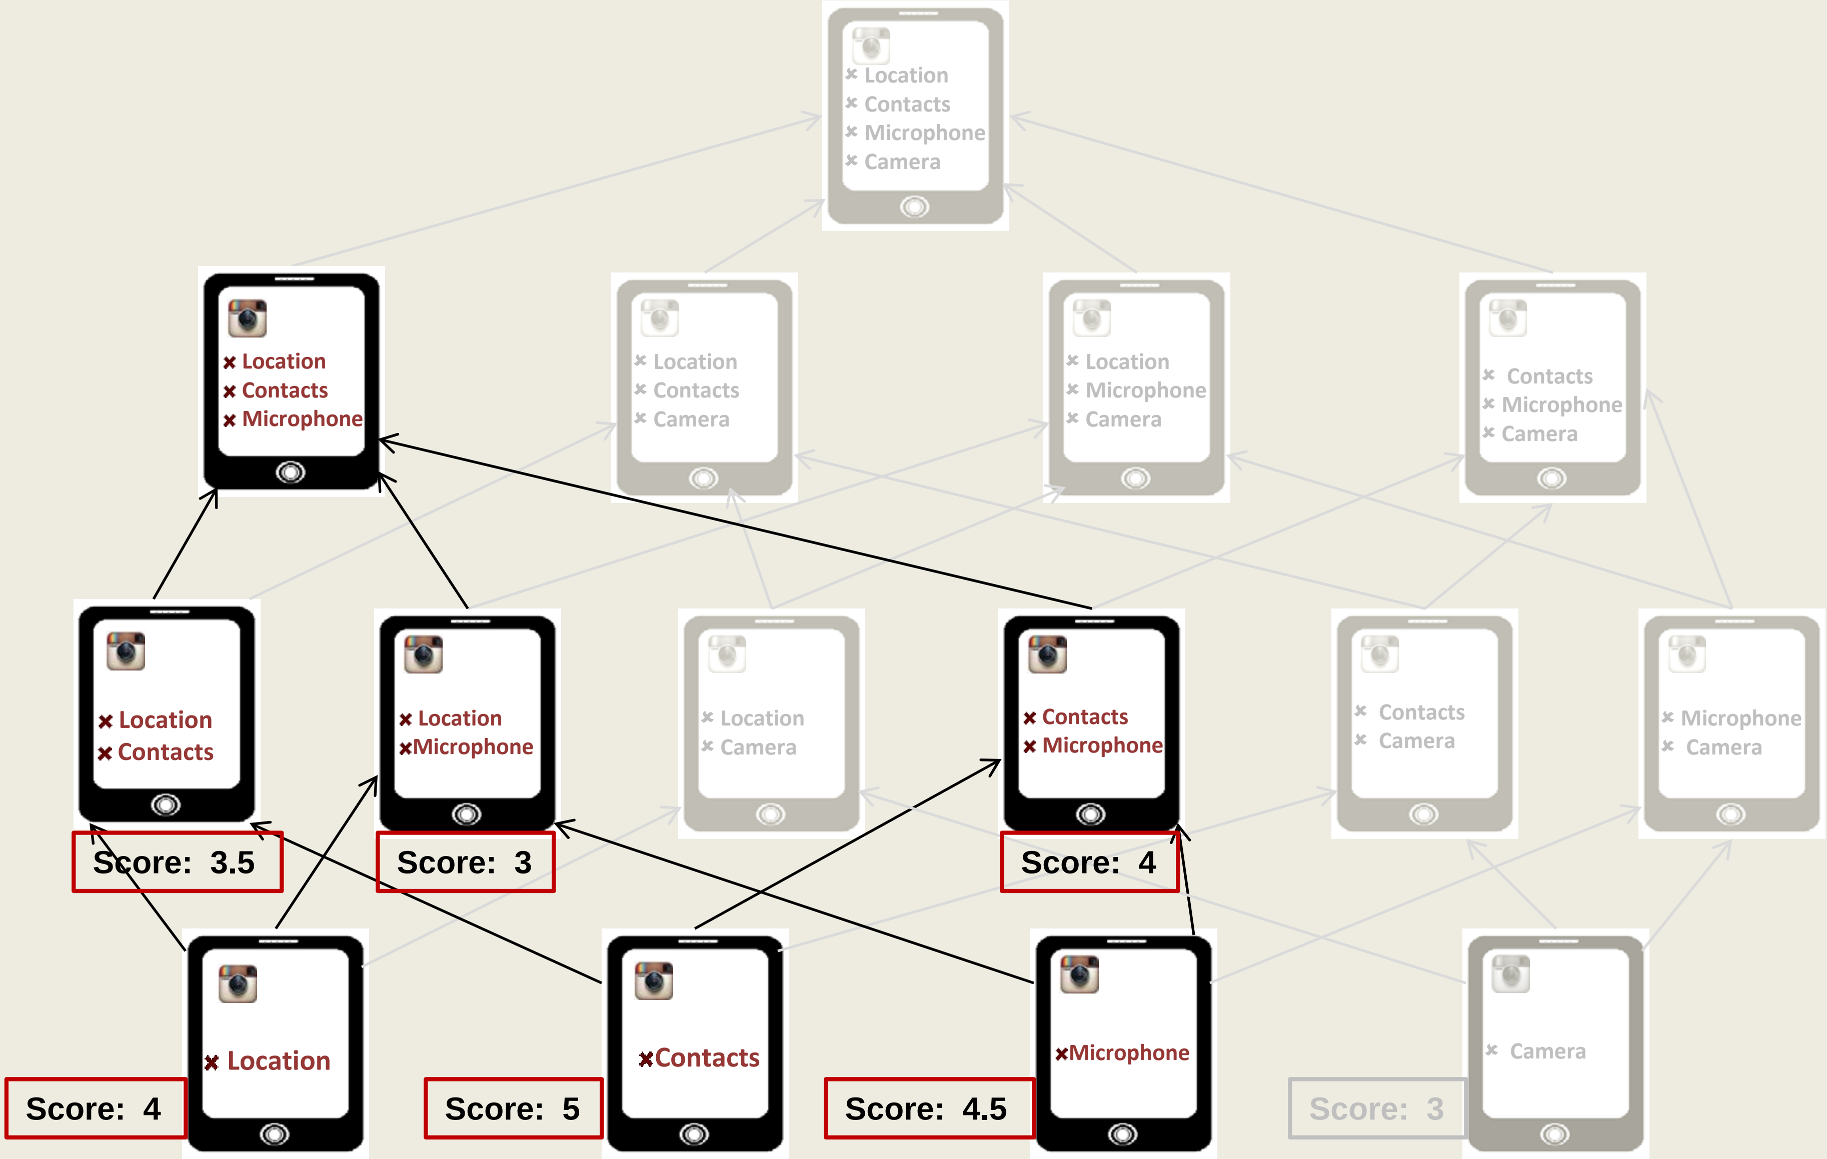

Low Score ⇒ Prune Node & Ancestors

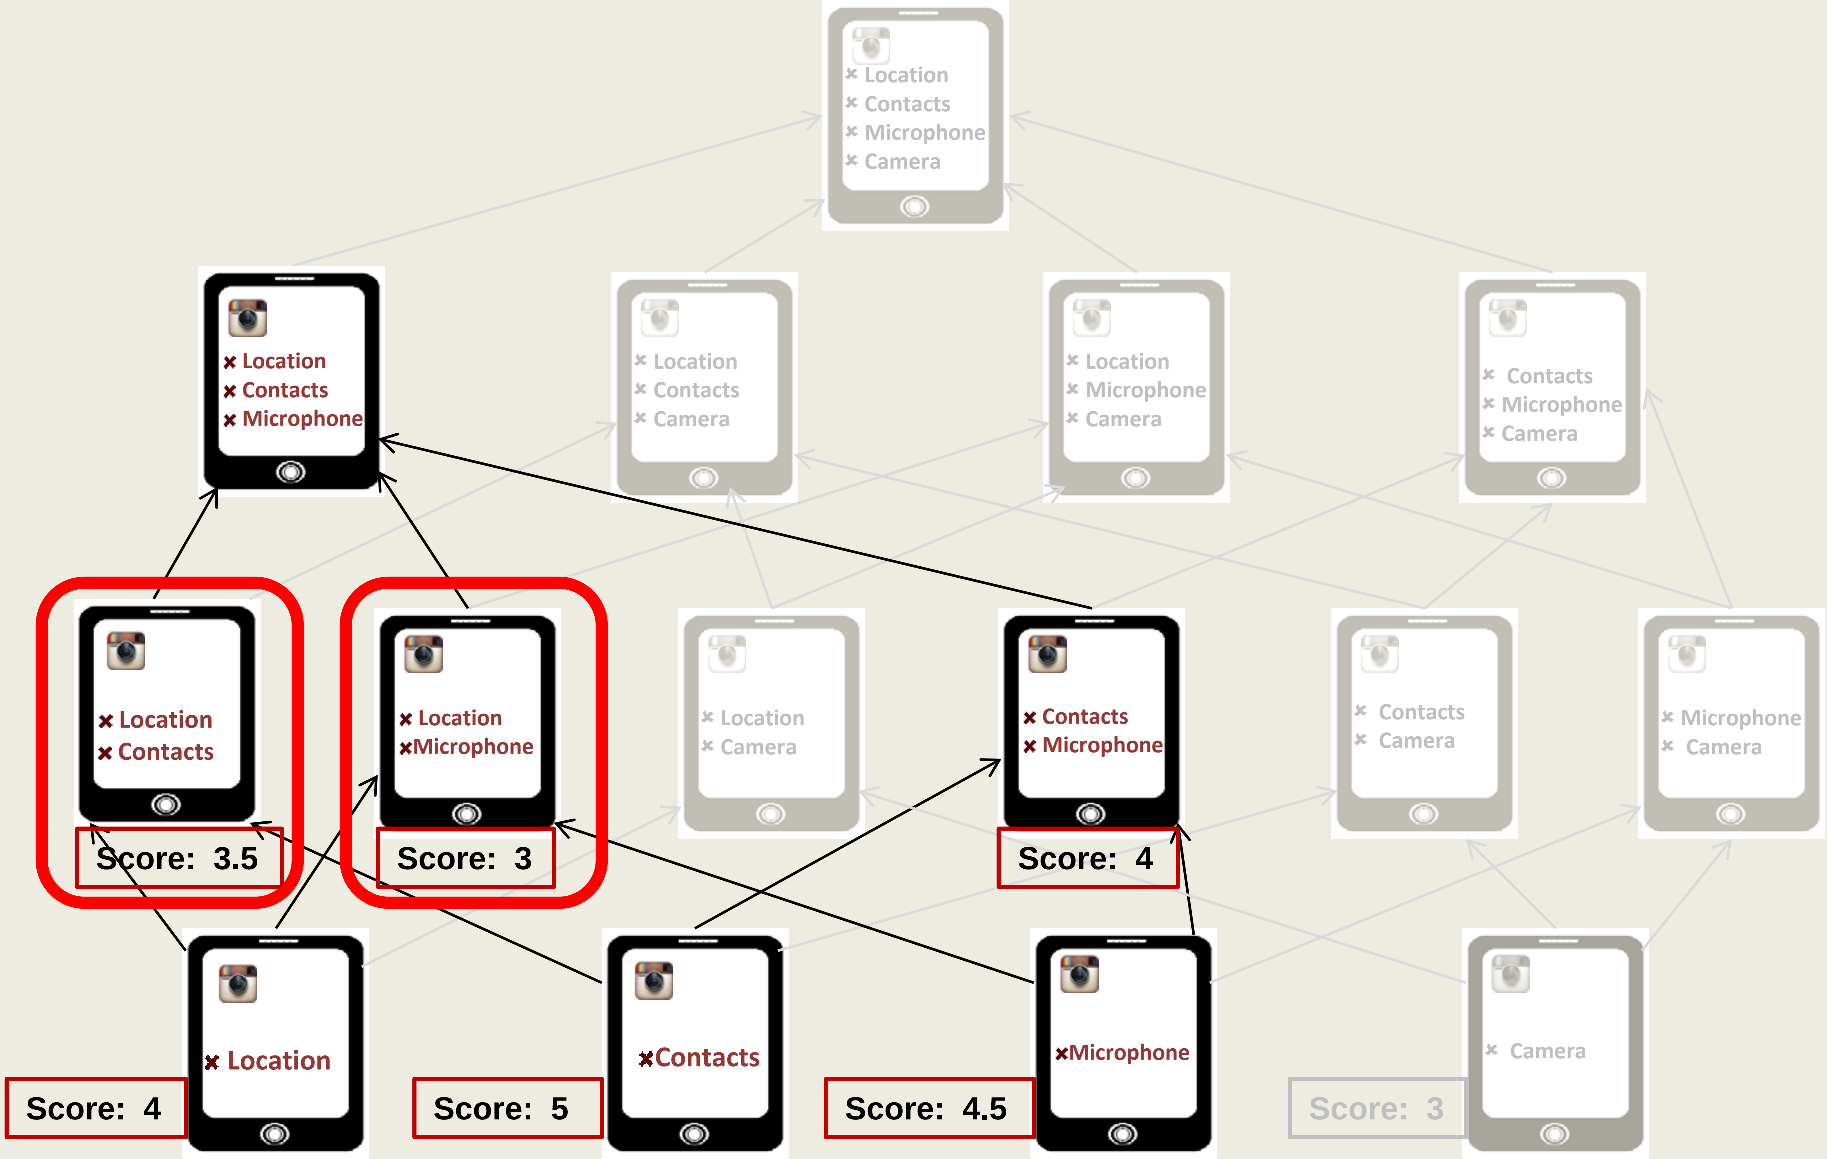

Crowdsourcing Security Configs

Low Score ⇒ Prune Node & Ancestors

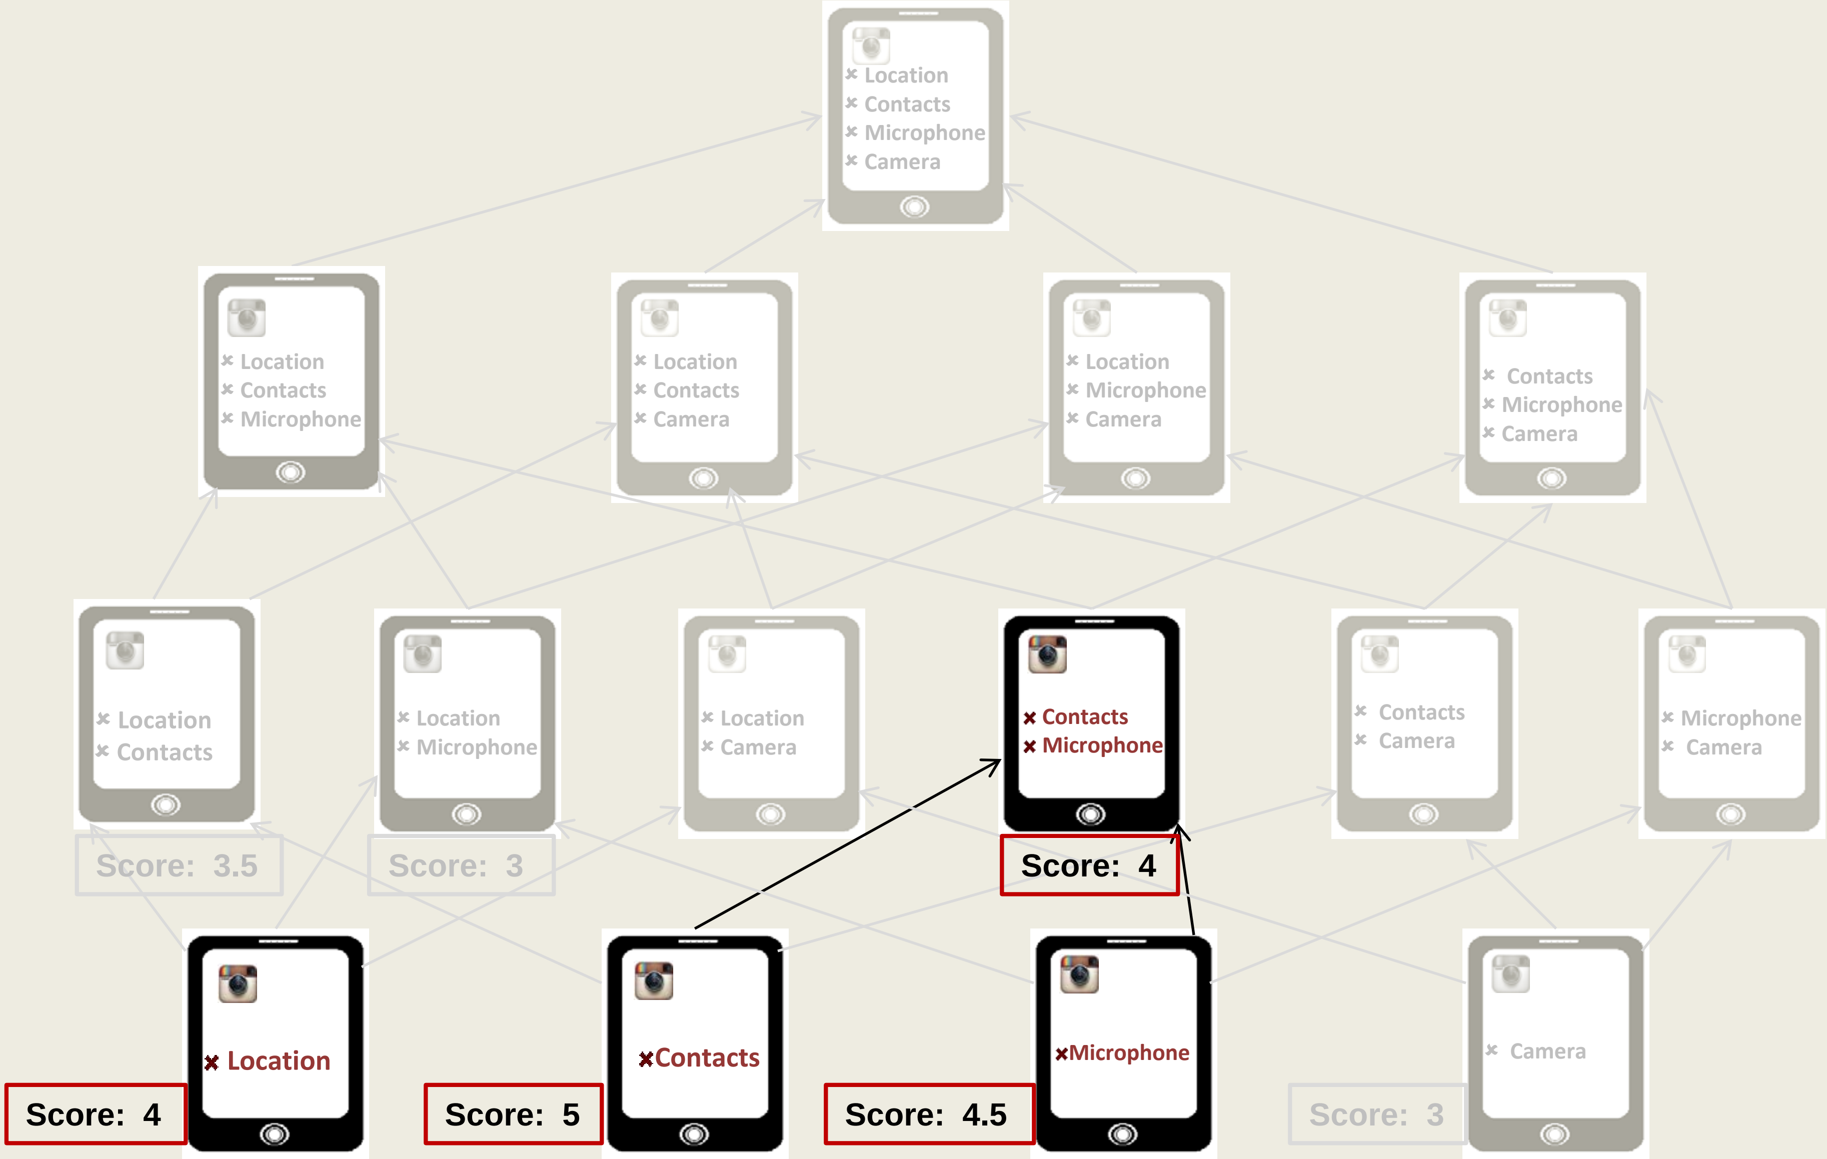

Crowdsourcing Security Configs

Most Acceptable Configurations Remain

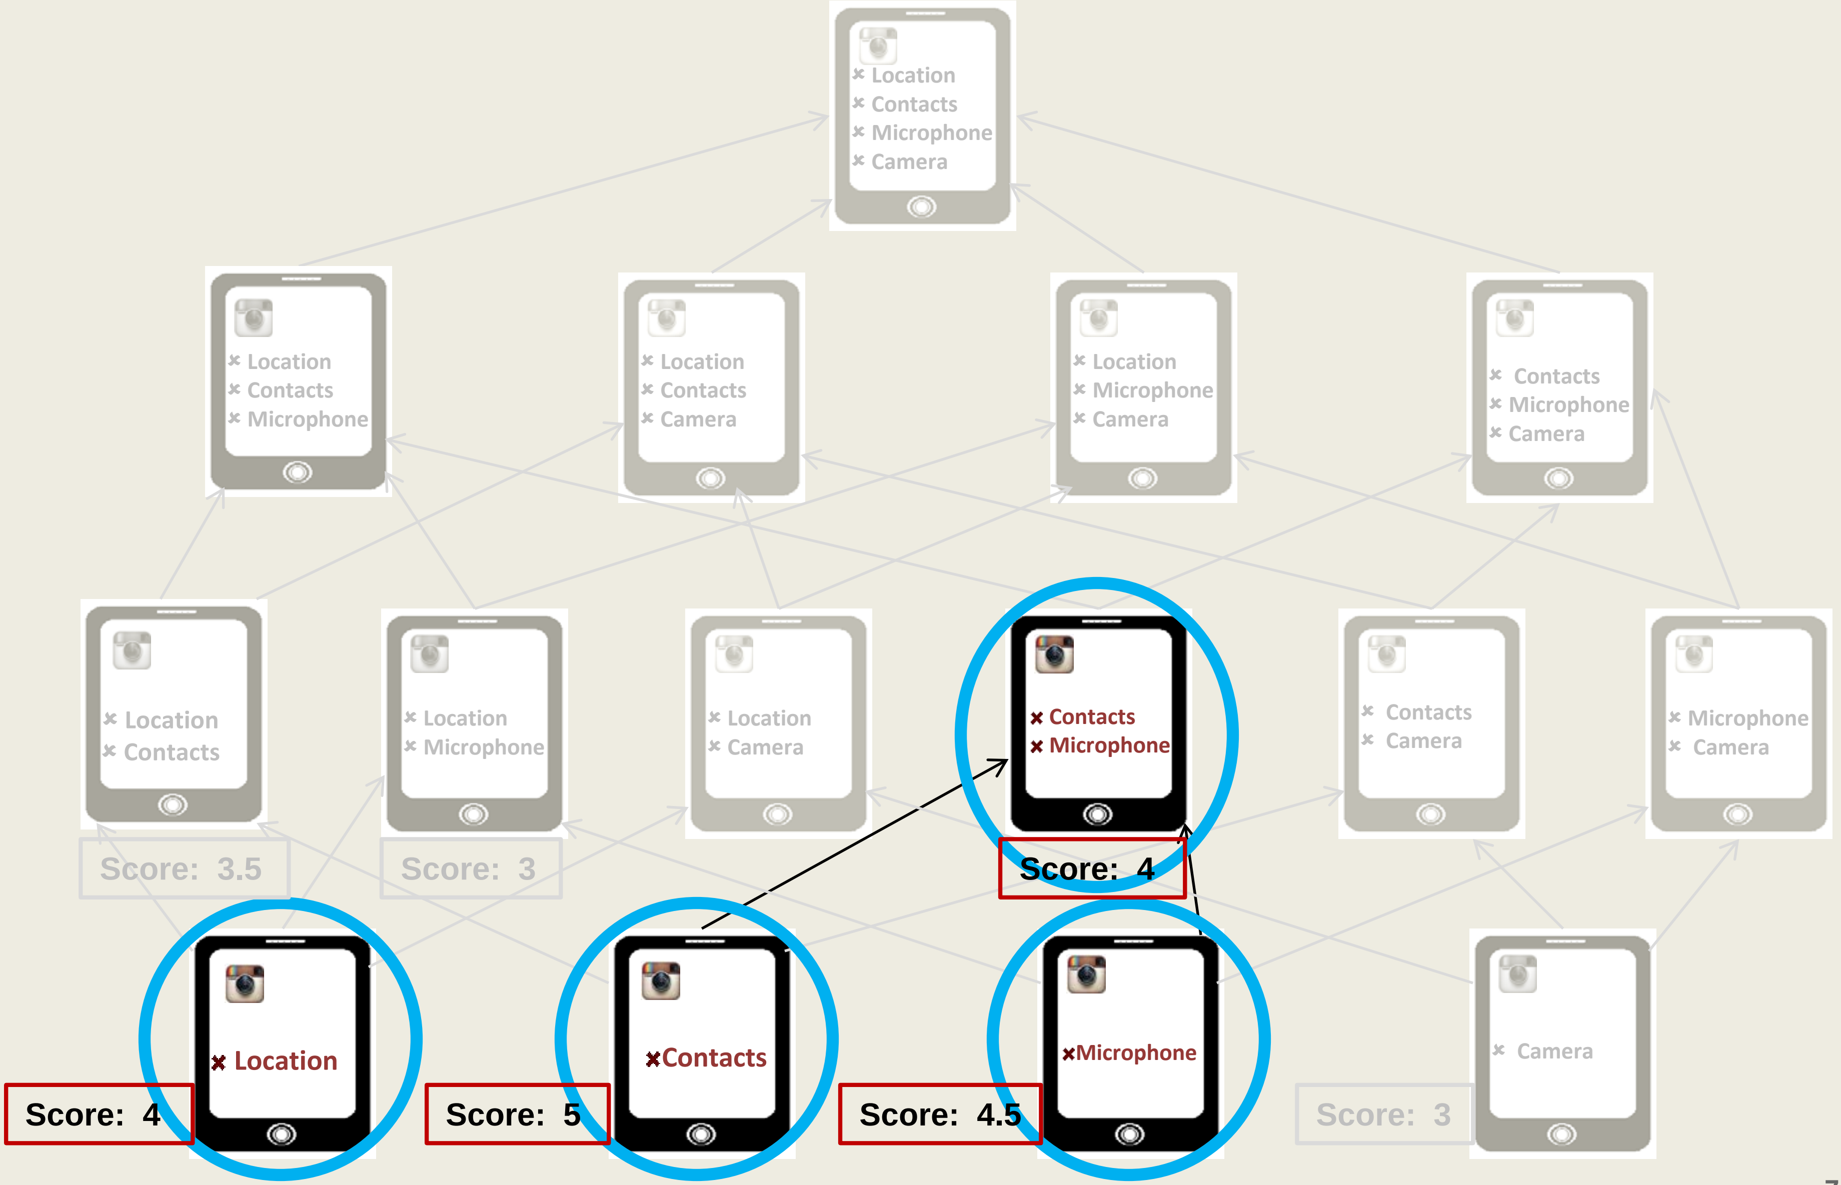

✖ Location
✖ Contacts
✖ Microphone
✖ Camera

✖ Location
✖ Contacts
✖ Microphone

✖ Location
✖ Contacts
✖ Camera

✖ Location
✖ Microphone
✖ Camera

✖ Contacts
✖ Microphone
✖ Camera

✖ Location
✖ Contacts

Score: 3.5

✖ Location
✖ Microphone

Score: 3

✖ Location
✖ Camera

✖ Contacts
✖ Microphone

Score: 4

✖ Contacts
✖ Camera

✖ Microphone
✖ Camera

✖ Location

Score: 4

✖Contacts

Score: 5

✖Microphone

Score: 4.5

✖ Camera

Score: 3

Crowdsourcing Security Configs

Most Acceptable Configurations Remain

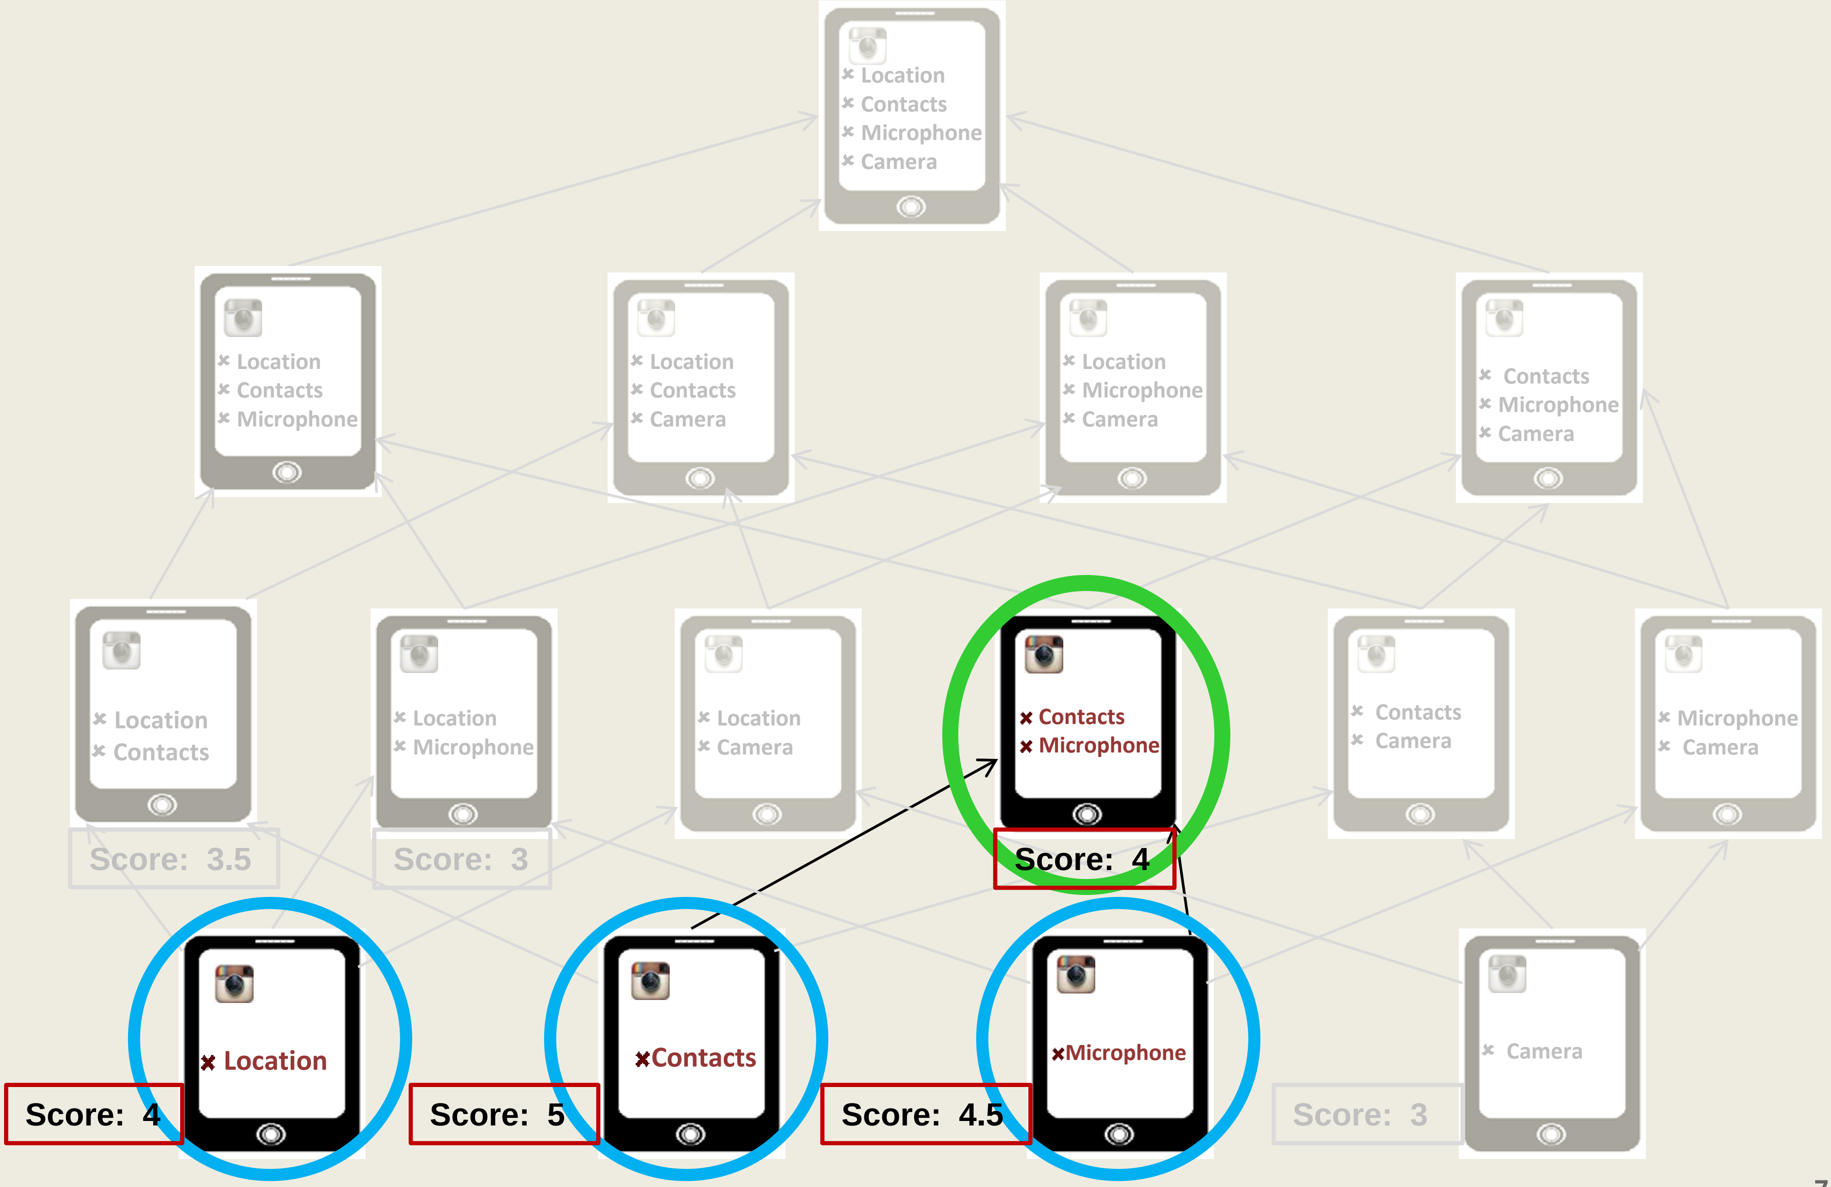

Most Acceptable Configurations Remain

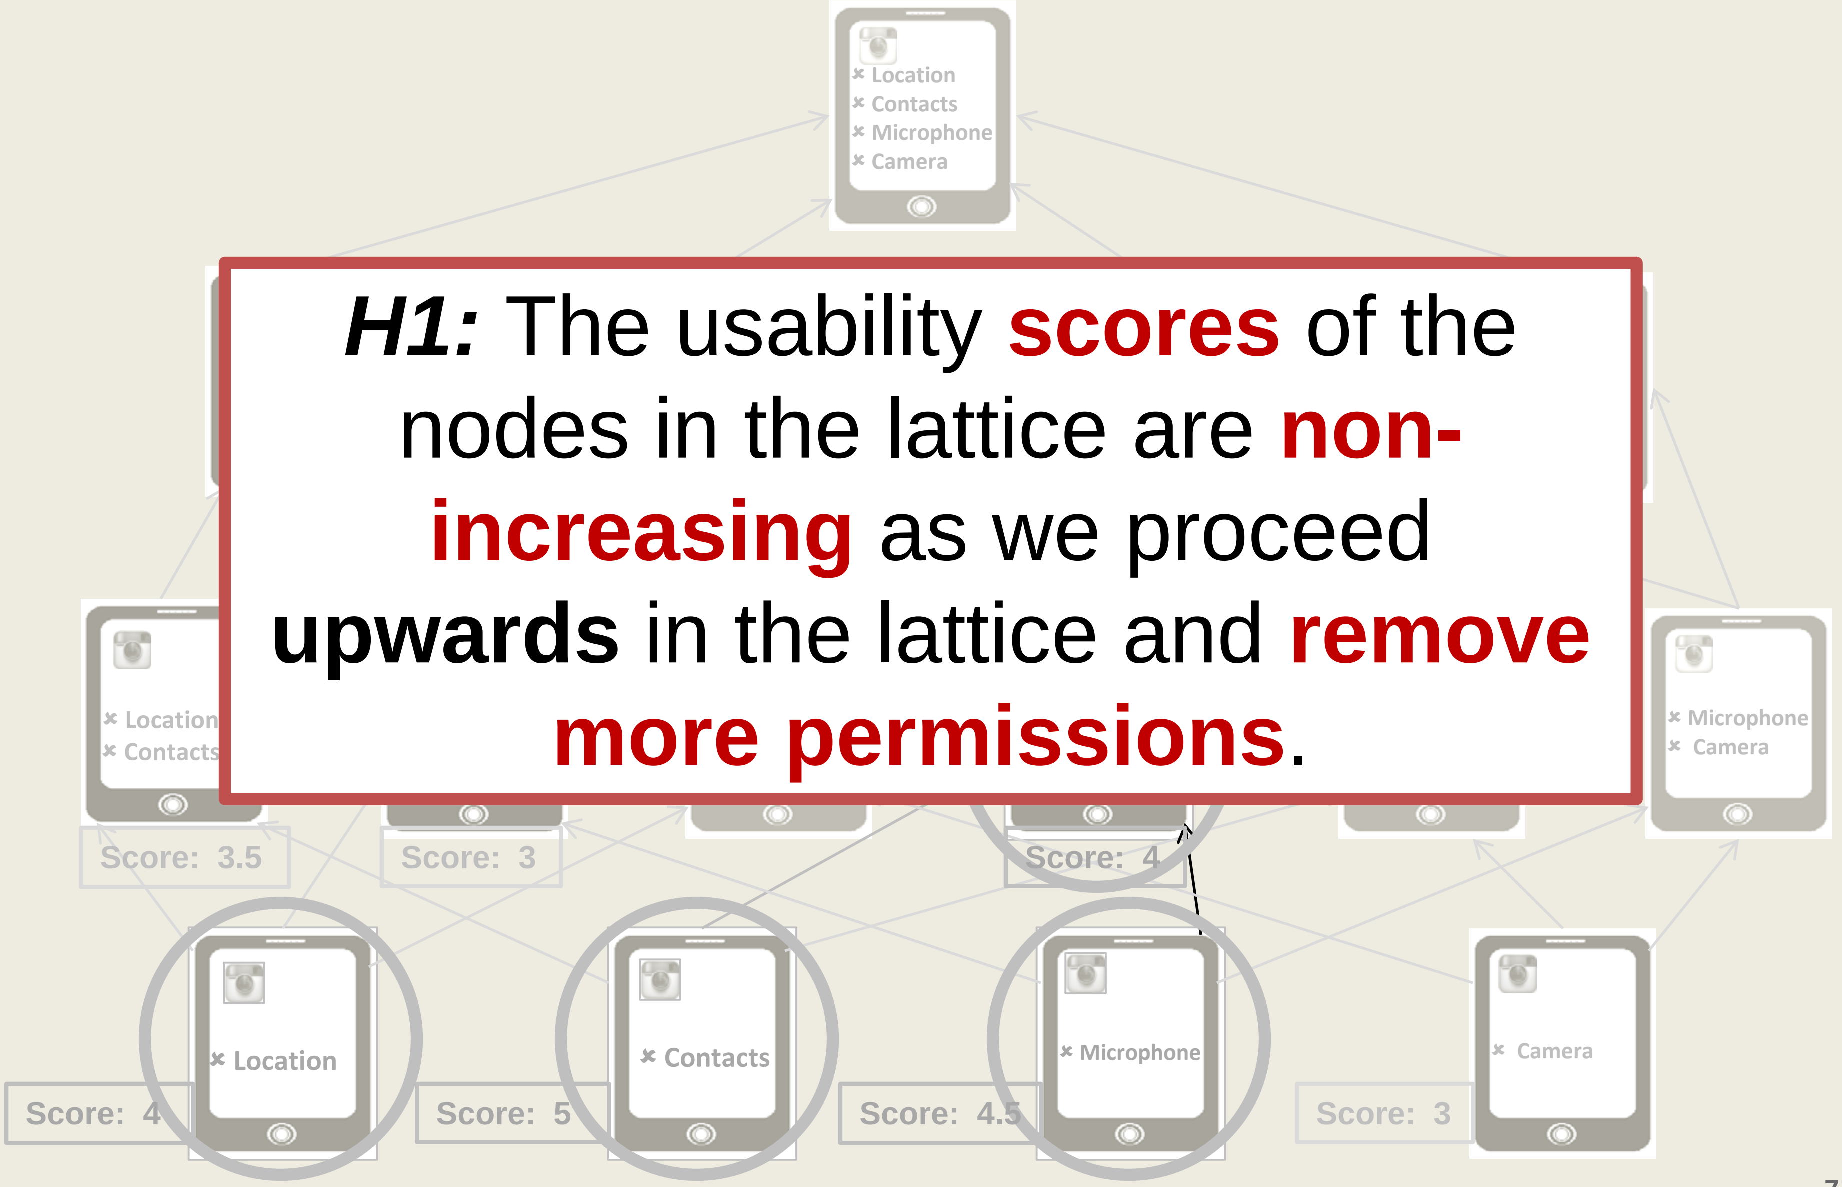

H1: The usability **scores** of the nodes in the lattice are **non-increasing** as we proceed **upwards** in the lattice and **remove more permissions**.

Research Questions and Hypotheses

R1: Can we **recommend** suitable permission sets based on the crowd's ratings?

R2: Can we use **crowdsourcing scalably** to explore security configurations of an app?

H1: The usability **scores** of the nodes in the lattice are **non-increasing** as we proceed **upwards** in the lattice and **remove more permissions**.

Science or Engineering?

Does this contribute to “a systematic body of knowledge … to inform the engineering of secure information systems that can resist unanticipated attacks?”

✓ It has hypotheses, and controlled experiments to test them!

✓ The envisioned system strives for “least privilege”, to limit unanticipated attacks

? But is the knowledge gained “systematic”, or is it specific to our envisioned usage?

Parting Thoughts on Science

- I have difficulty distinguishing “science” and “engineering”, even in my own research
  - I hope I’ve confused you a bit, too ☺

- Our ability to recognize good science (or good research?) is itself worthy of skepticism

Some Cautionary Notes

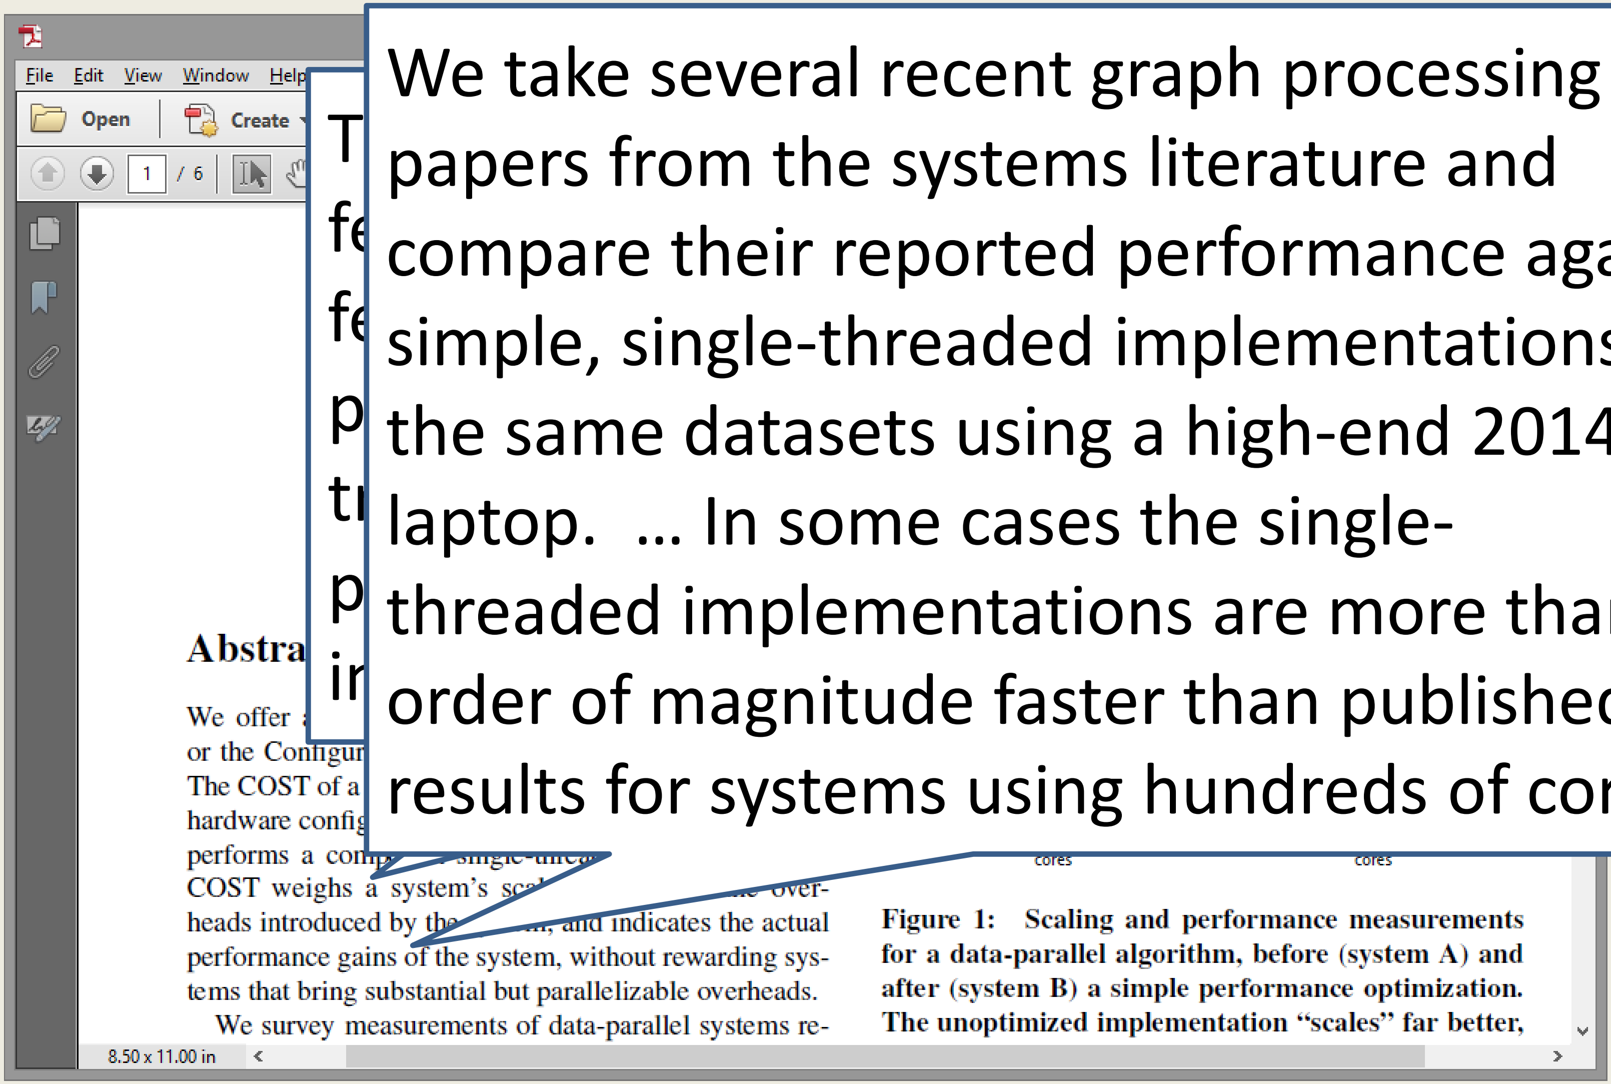

Abstra

We offer …
or the Configur…
The COST of a …
hardware config…
performs a comp… single-threa…
COST weighs a system's sca… over-
heads introduced by the … and indicates the actual
performance gains of the system, without rewarding systems that bring substantial but parallelizable overheads.

We survey measurements of data-parallel systems re-

Figure 1: Scaling and performance measurements for a data-parallel algorithm, before (system A) and after (system B) a simple performance optimization. The unoptimized implementation "scales" far better,

We take several recent graph processing papers from the systems literature and compare their reported performance against simple, single-threaded implementations on the same datasets using a high-end 2014 laptop. … In some cases the single-threaded implementations are more than an order of magnitude faster than published results for systems using hundreds of cores.

To appear in HotOS, May 2015.

Some Cautionary Notes

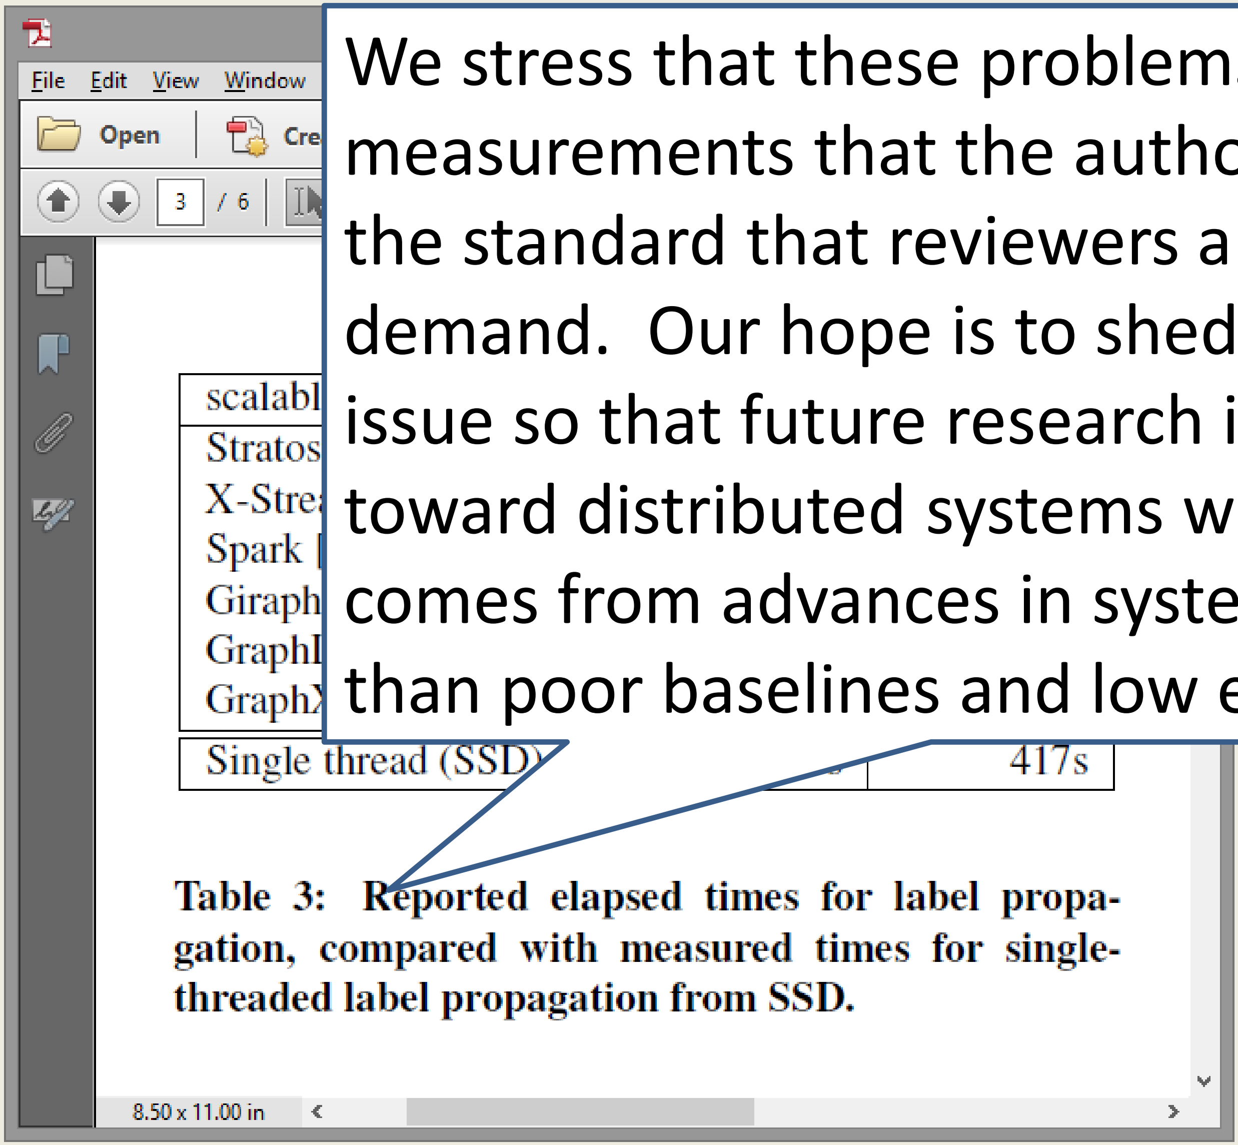

scalabl
Stratos
X-Stre
Spark [
Giraph
GraphI
GraphX
Single thread (SSD)
417s

Table 3: Reported elapsed times for label propagation, compared with measured times for single-threaded label propagation from SSD.

We stress that these problems lie … with the measurements that the authors provide and the standard that reviewers and readers demand. Our hope is to shed light on this issue so that future research is directed toward distributed systems whose scalability comes from advances in system design rather than poor baselines and low expectations.

From McSherry, et al., “Scalability! But at what COST?” HotOS 2015.

Some Cautionary Notes

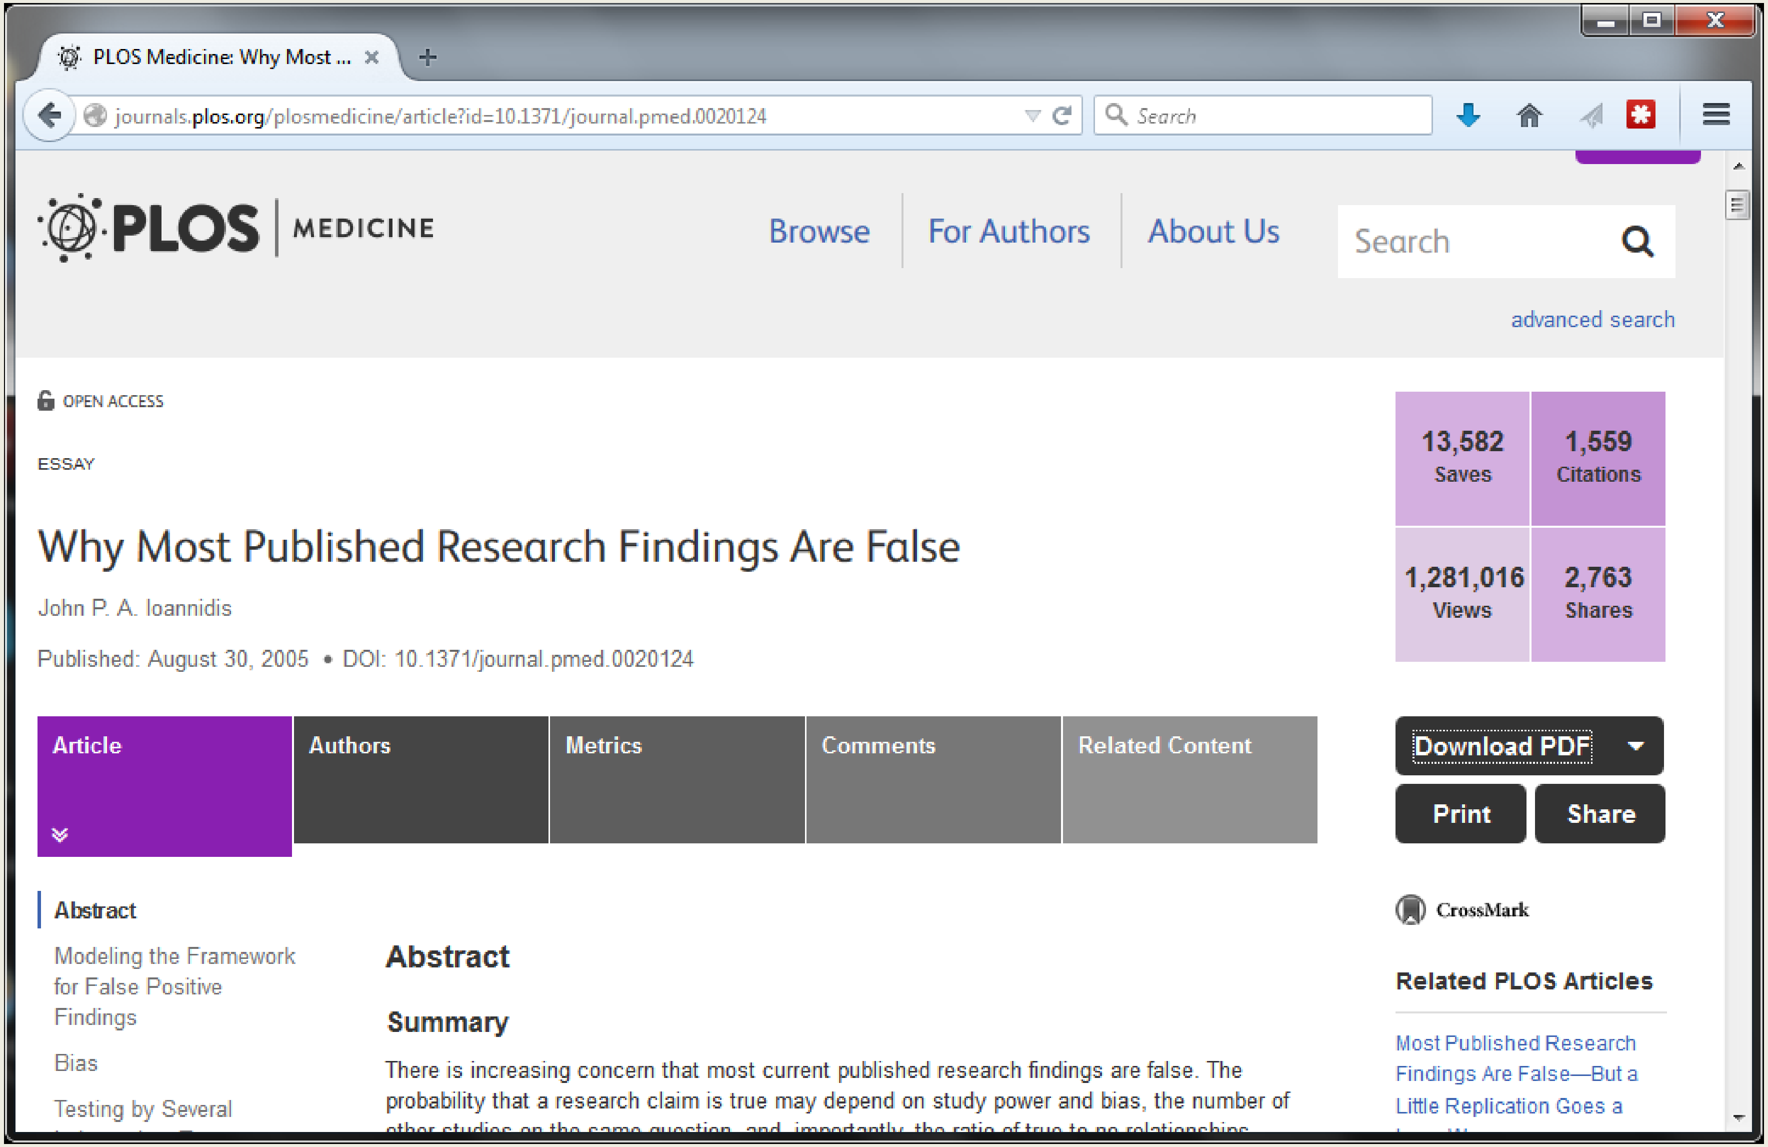

ESSAY

Why Most Published Research Findings Are False

John P. A. Ioannidis

Published: August 30, 2005 • DOI: 10.1371/journal.pmed.0020124

13,582 Saves	1,559 Citations
1,281,016 Views	2,763 Shares

Abstract

Summary

There is increasing concern that most current published research findings are false. The probability that a research claim is true may depend on study power and bias, the number of other studies on the same question, and, importantly, the ratio of true to no relationships

Related PLOS Articles

Most Published Research Findings Are False—But a Little Replication Goes a

Some Cautionary Notes

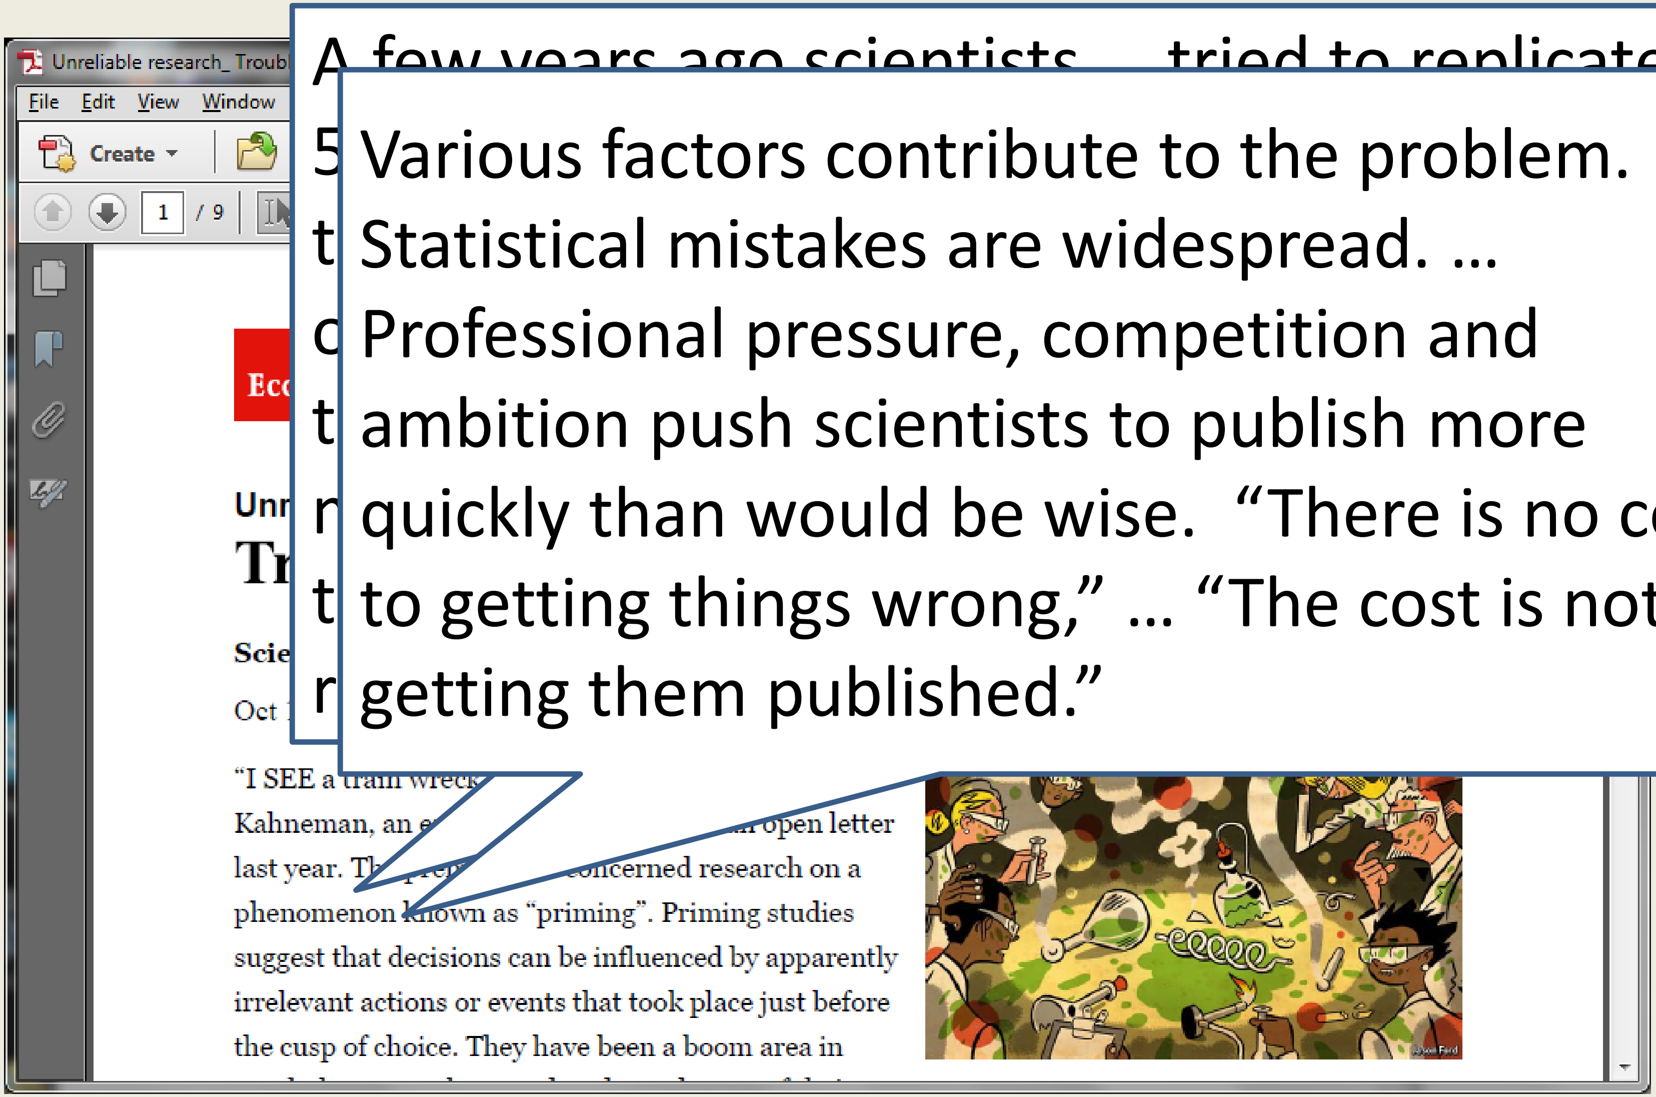

Unr

Tr

Scie

Oct 1

“I SEE a train wreck … Kahneman, an e… open letter last year. The … concerned research on a phenomenon known as “priming”. Priming studies suggest that decisions can be influenced by apparently irrelevant actions or events that took place just before the cusp of choice. They have been a boom area in

A few years ago scientists … tried to replicate

Various factors contribute to the problem. Statistical mistakes are widespread. … Professional pressure, competition and ambition push scientists to publish more quickly than would be wise. “There is no cost to getting things wrong,” … “The cost is not getting them published.”

Some Cautionary Notes

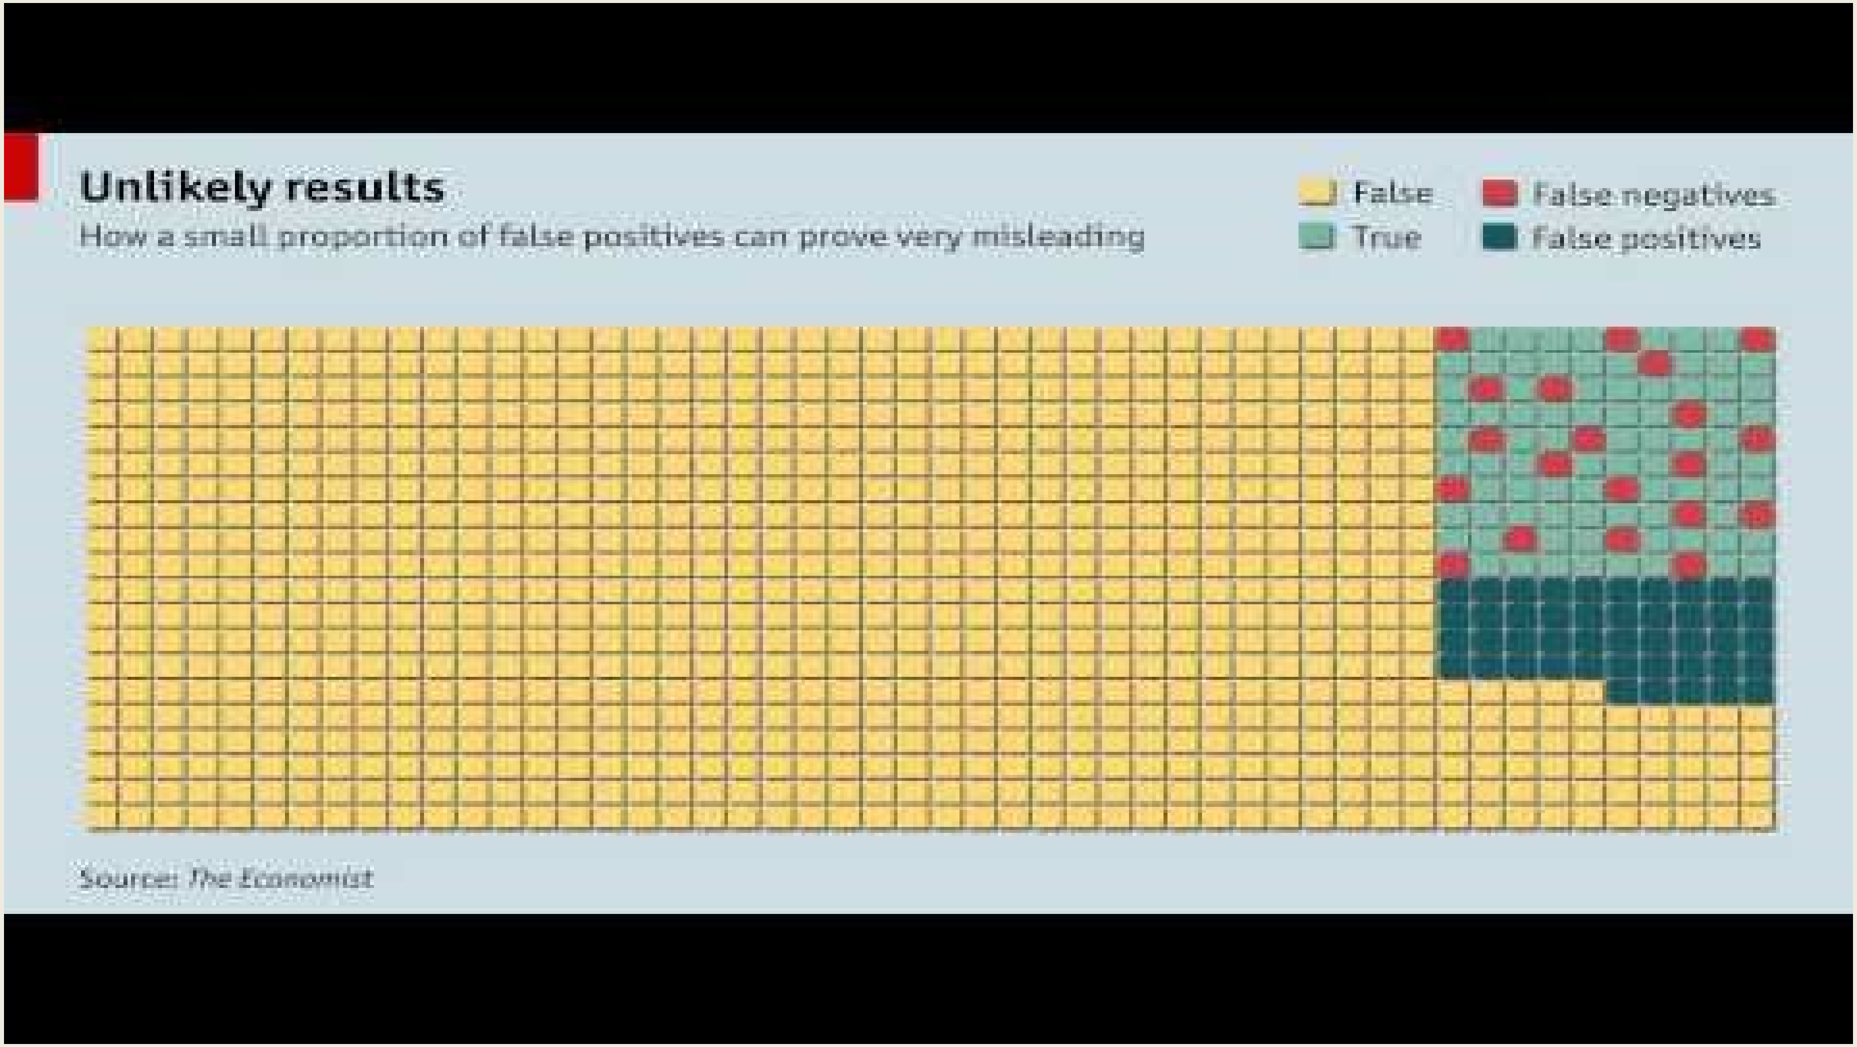

Unlikely results

How a small proportion of false positives can prove very misleading

Source: *The Economist*

Summary

- Distinguishing science from engineering is hard even in specific cases
  - Epidemiological study of malware encounters
  - Cross-tenant side-channel attacks in PaaS clouds
  - Private regular-expression matching on encrypted data
  - Crowdsourced exploration of security configurations
- Social pressures decay even prevailing scientific methods to sometimes an alarming extent

My Opinion on "Fixing" Security Practice

- Yes, more science is important, but we also need *better* science!

- *Accountability* for negligence is also needed to fix how security is done in practice
  - But that is a talk for a different time …
  - … and an agenda for different people

- Both of these are fundamentally cultural problems that won't be easy to fix