WEB 2.0
SECURITY&
PRIVACY 2014

I Know Where You've Been: Geo-Inference Attacks via the Browser Cache

Yaoqi Jia*, Xinshu Dong†, Zhenkai Liang *, Prateek Saxena*
*School of Computing, National University of Singapore
†Advanced Digital Sciences Center

Geo-location in Browsers

Benefits

craigslist sg

Threats

NUS National University of Singapore

May I Access Your Geo-location?

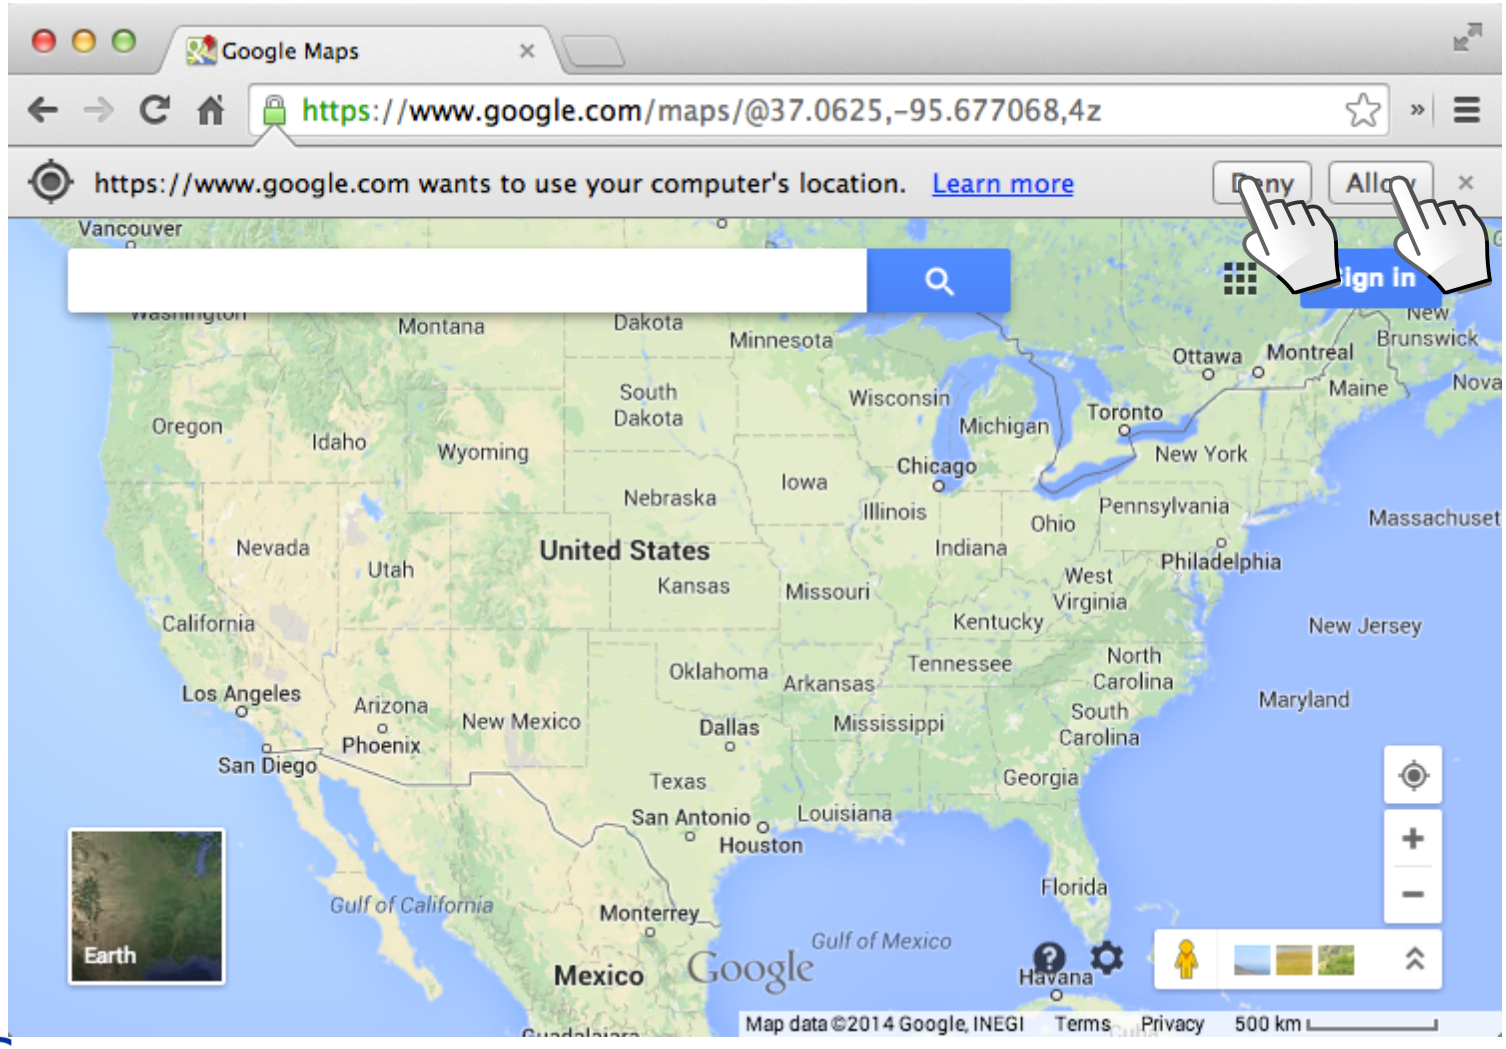

Sources of Users' Geo-locations

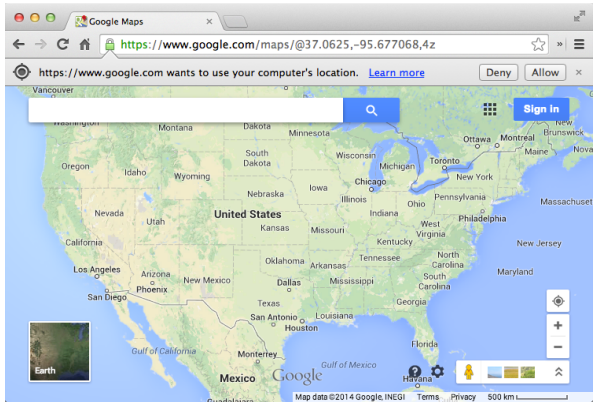

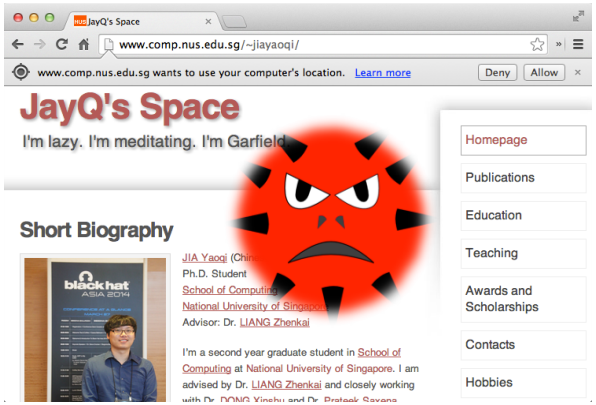

Browser

NUS National University of Singapore

Problem Statement

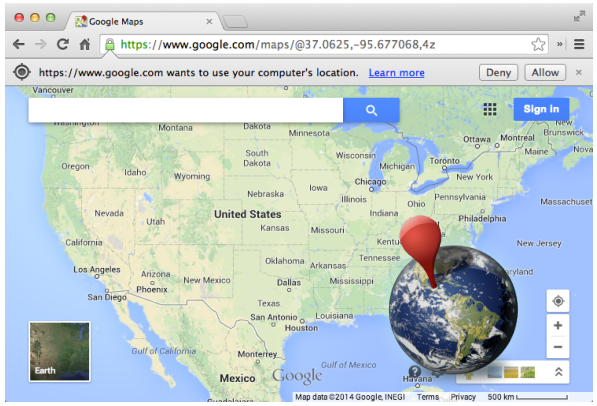

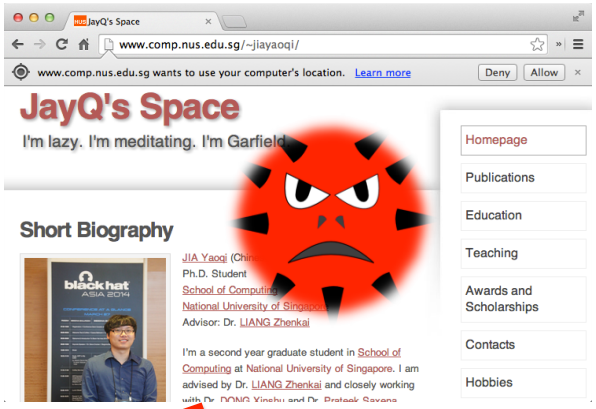

?

Can we infer the user's geo-location from his browser?

NUS National University of Singapore

Site-Related States in Browser

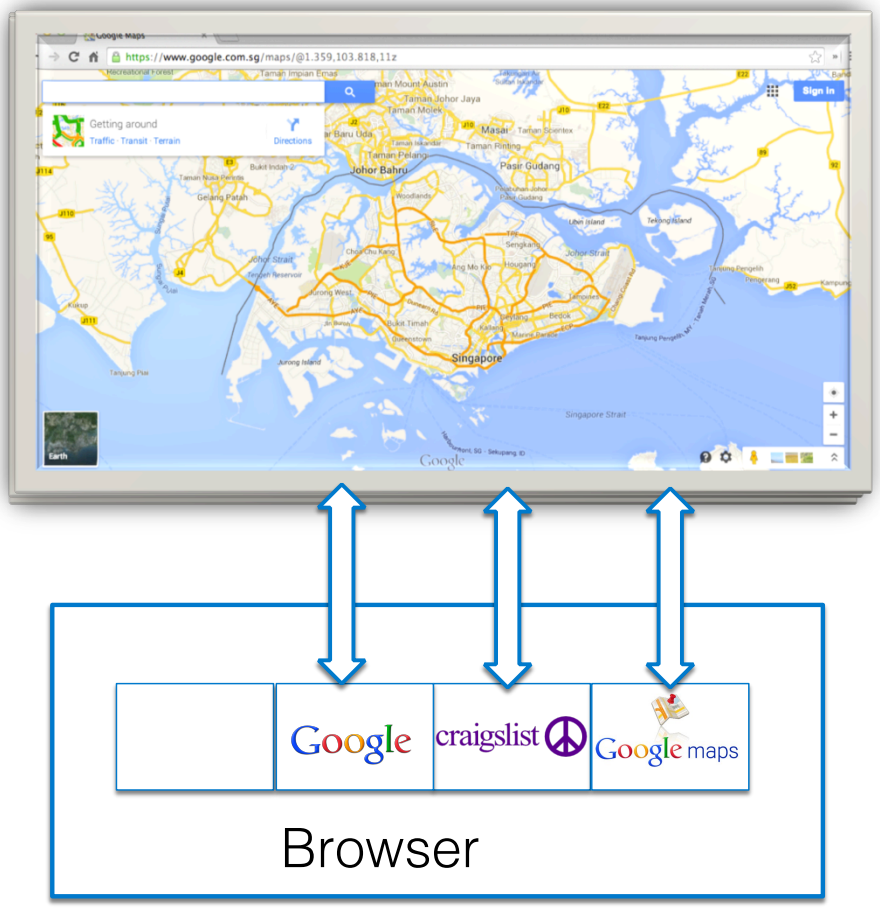

Browser Cache Saves Loading Time

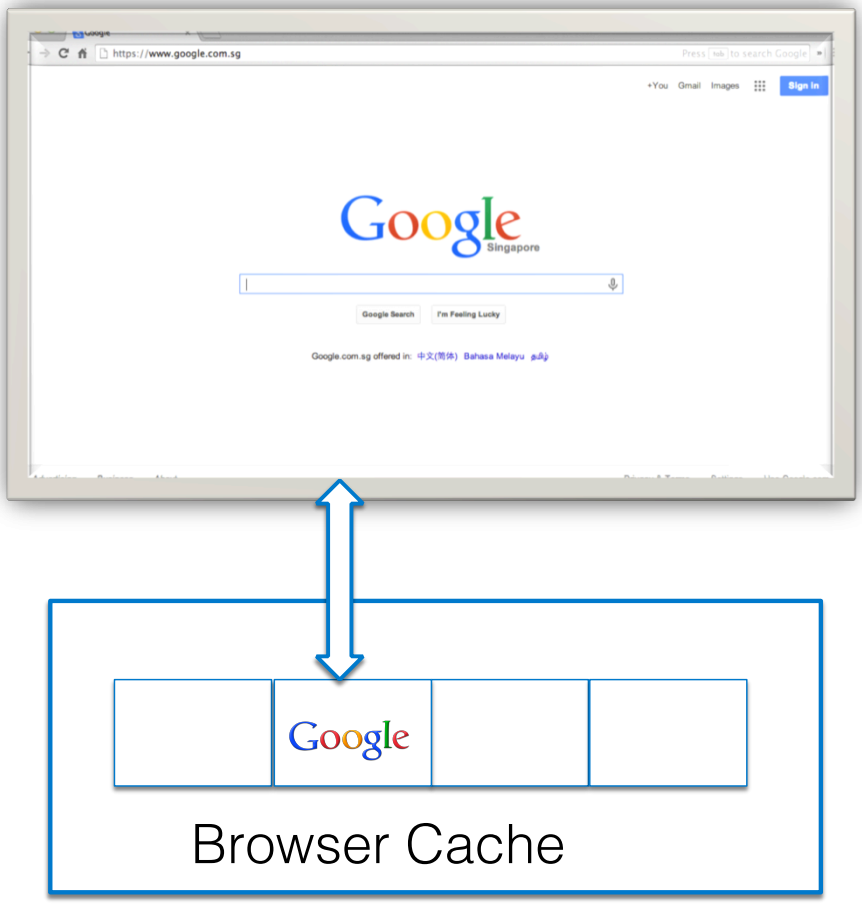

1st: 1360ms

2nd: 320ms

3rd: 350ms

Browser Cache Abused: Timing Channels of Leakage

Felten and Shneider, CCS'00

Browser cache is shared across all sites

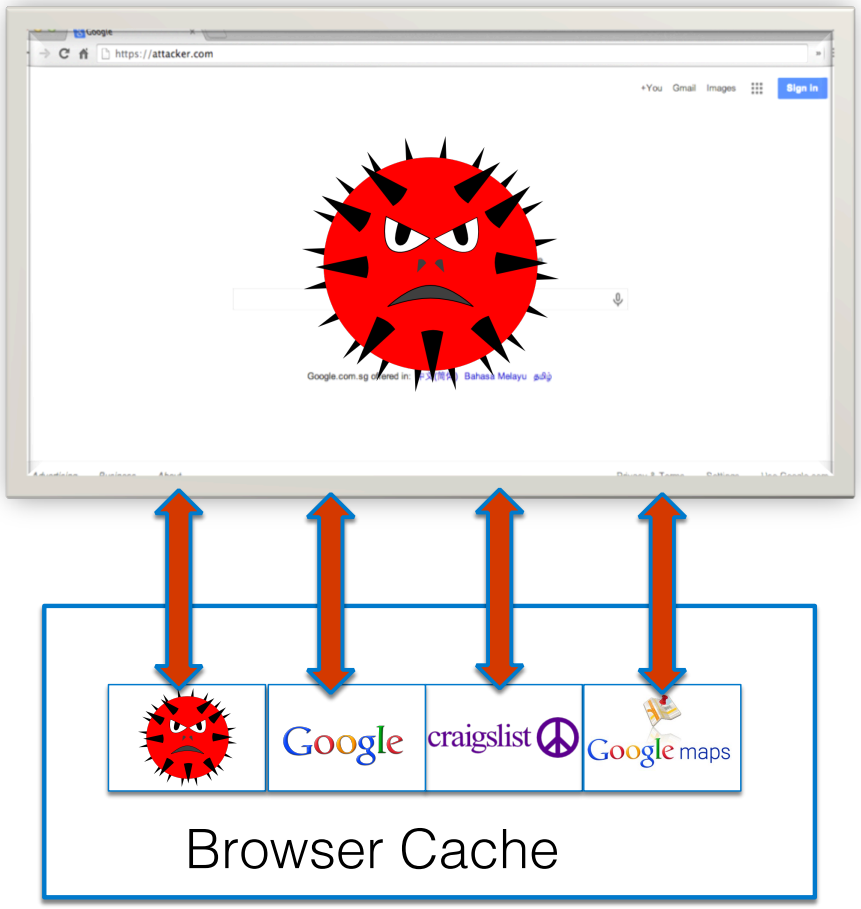

Our Contributions

- Geo-inference attacks via the browser cache
  - Infer a user's country, city or even neighborhood
- Prevalence of geo-inference attacks
  - Five mainstream browsers and TorBrowser
  - Top 55 Alexa and 11 map websites
- Pros & cons of potential solutions

Outline

- Problem Statement
- Case Studies
- Evaluation
- Discussion

Case Studies

- Can we infer a user's country?
- Can we infer a user's city?
- Can we infer a user's neighborhood?

How to Infer a User's Country?

- Google has 191 regional sites, and one site represents one country or region.
- Measure image load time of Google's logo from Google's 191 regional sites

Measuring Image Load Time

Before Loading → **img.onload Fires**

```
var image = document.createElement(`img');

image.setAttribute(`startTime', (new
Date().getTime()));

image.onload = function()

{

    var endTime = new Date().getTime();

    var loadTime = endTime -
parseInt(this.getAttribute(`startTime'));

    ......

}
```

How to Infer a User's City?

Measure page load time of Craigslist's 712 city sites, determine which page is cached

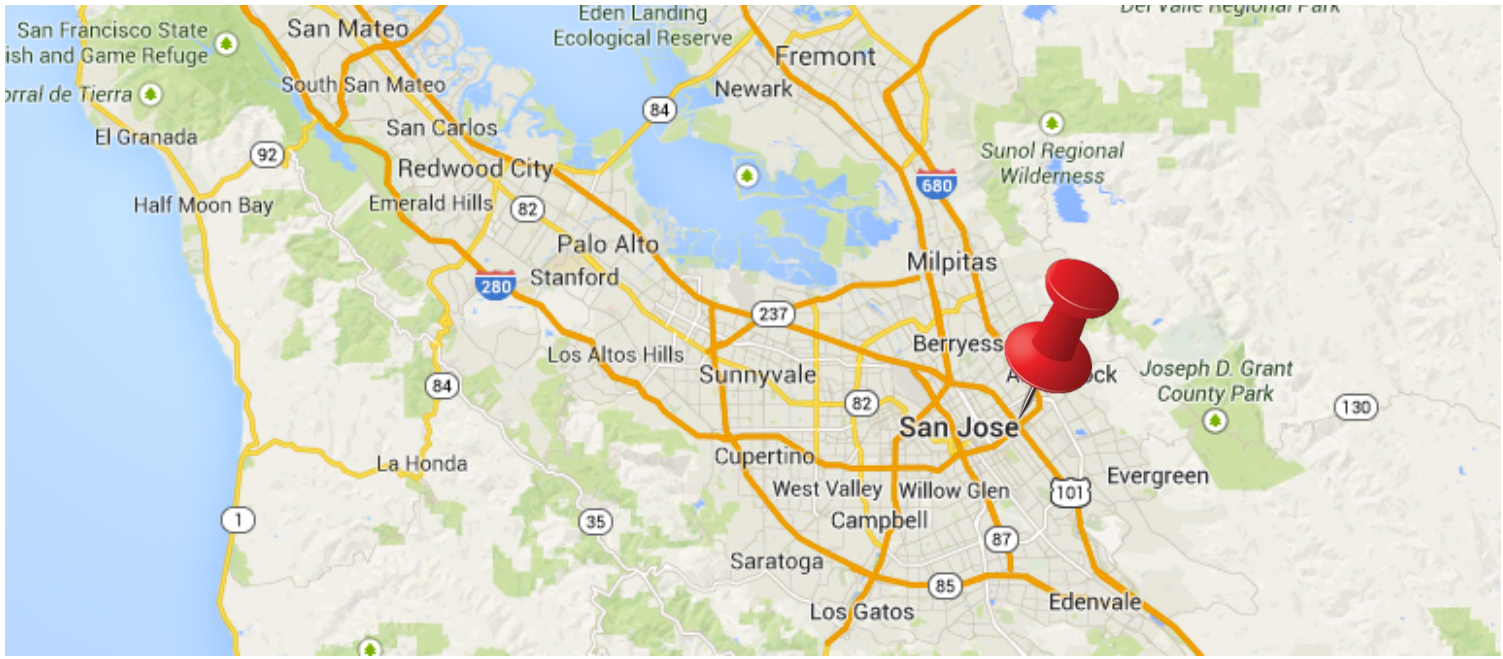

Measuring Page Load Time

Before Loading → **iframe.onload Fires**

```
var page = document.createElement(`iframe');

page.setAttribute(`startTime', (new
Date()).getTime());

page.onload = function ()

{

    var endTime = (new Date()).getTime();

    var loadTime = ( endTime -
parseInt(this.getAttribute(`startTime')));

    ......

}
```

How to Infer a User's Neighborhood?

Measure the image load time of map tiles of the user's city from Google Maps, determine which tiles are cached

Evaluation

Questions to be answered:

- (Prevalence) How many browsers and websites are susceptible to geo-inference attacks?
- (Reliability) How big is the time difference between resources load time without cache and that with cache?

Evaluation Setup

- Websites: 191 Google's regional sites, 100 Craigslist's city sites, and 4,646 map tiles of New York City from Google Maps.

- Browsers: Five mainstream browsers, i.e., Chrome, Firefox, Safari, Opera and IE, as well as TorBrowser (version 3.5.2.1) on both desktop and available mobile platforms.

- Locations: US, UK, Australia, Singapore, and Japan, via VPN service Hotspot Shield.

Websites with Location-Related Resources in Browser Cache

Total 11 map service sites

62% of 55 top Alexa global sites

Browsers Susceptible to Geo-Inference Attacks

Mainstream Browsers

Desktop Platforms

Mobile Platforms

Reliability (Time Difference)

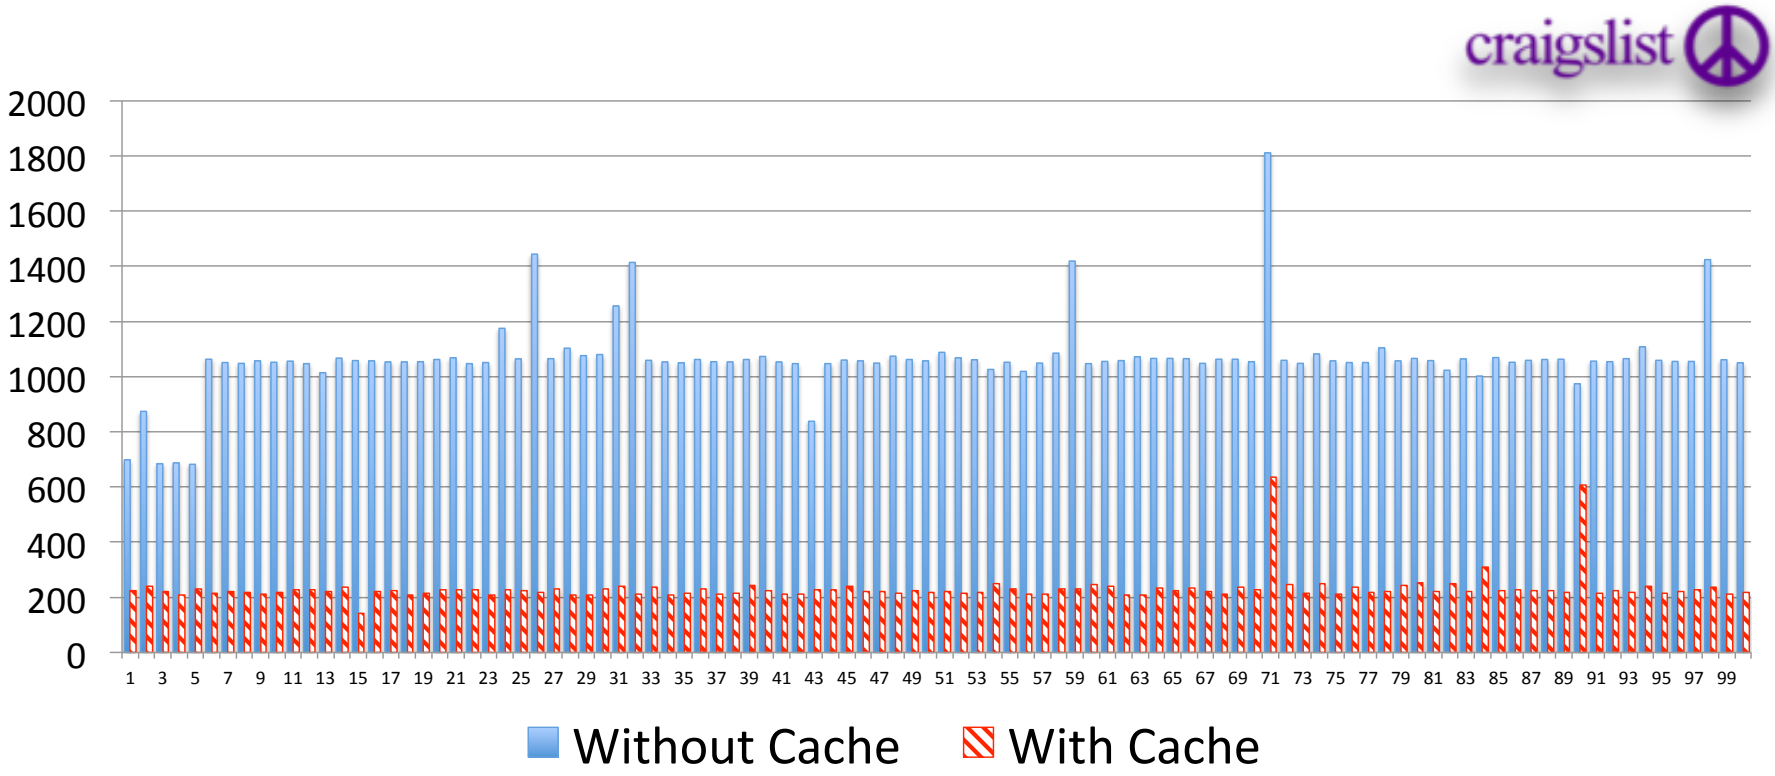

The huge difference between the page load time (in millisecond) of 100 Craigslist sites without cache (> 1000 ms) and with cache (≈ 220 ms) indicates geo-inference attacks with Craigslist

NUS National University of Singapore

Discussion of Defense Solutions

- Private Browsing Mode and TorBrowser
- Randomizing timing measurements
- Segregating browser cache

Private Browsing Mode is not the Cure

Private Browsing Mode

- Clear browser cache after closing window.
- Disable disk cache, enable memory cache.
- It cannot prevent one site from inferring geo-location of another site
  - Confirmed by experiments.
- TorBrowser is VPN + Private Browsing Mode

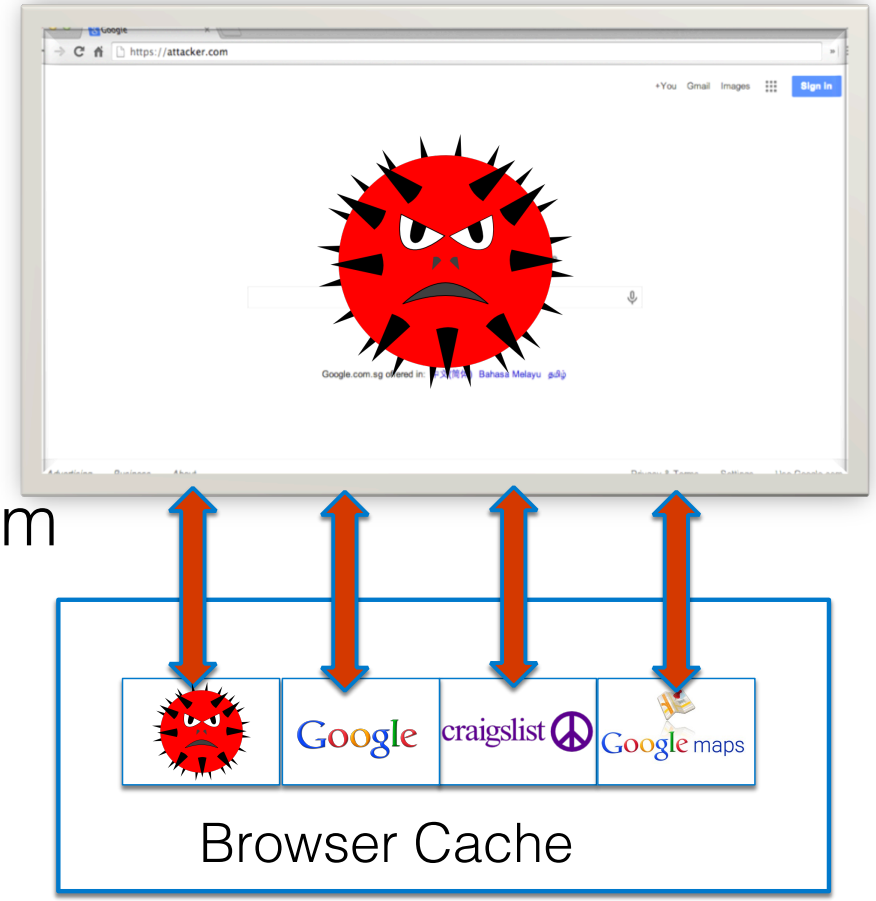

NUS National University of Singapore

Randomizing Timing Measurements

- Add noise into timing measurement mechanisms.
- Intricate engineering effort.

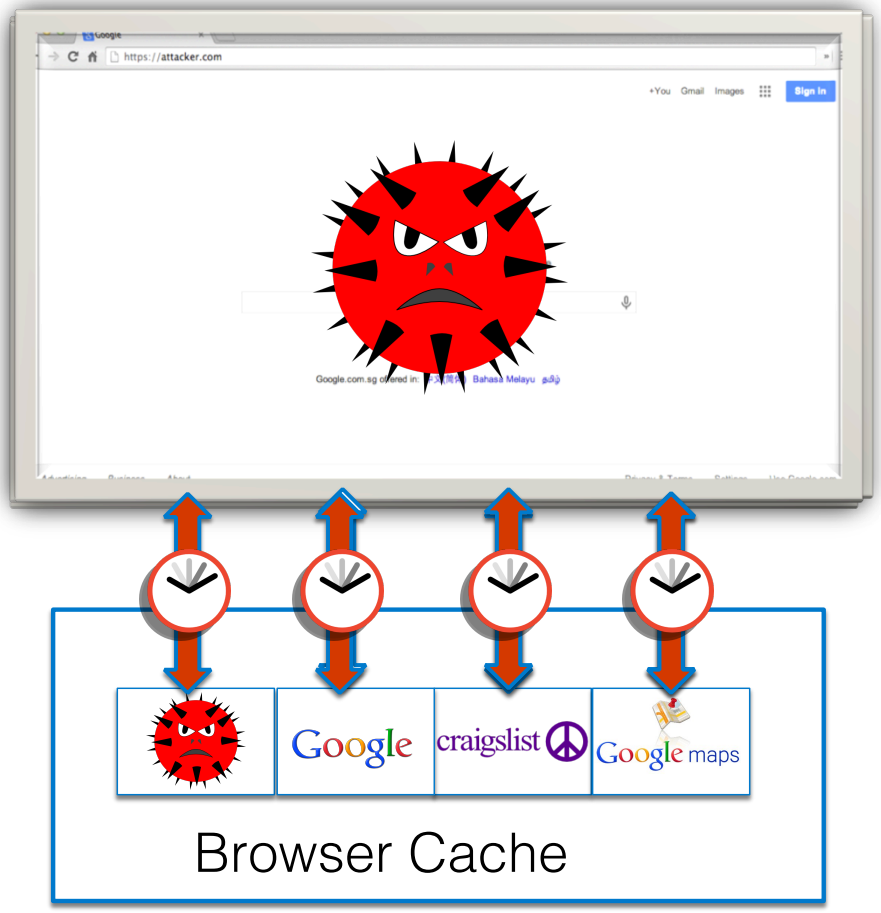

NUS National University of Singapore

Segregating Browser Cache

- Deploy Same-Origin Policy on browser cache. [Jackson et al. WWW'06]
- High performance overhead measured in our experiment

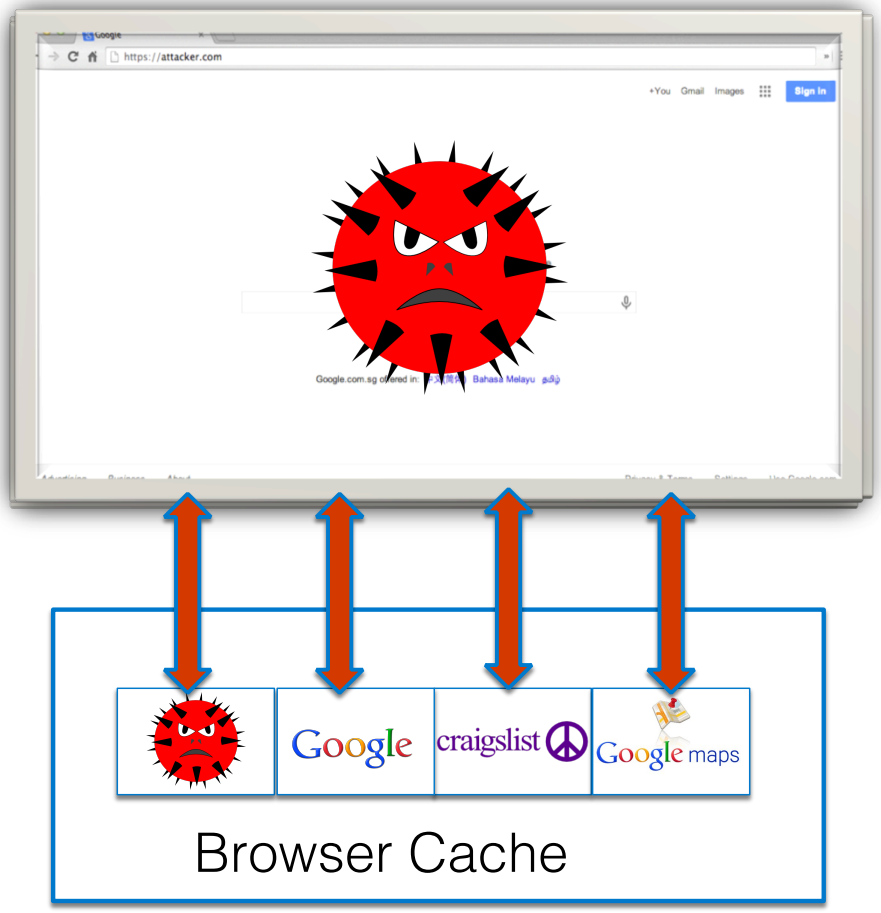

To Cache or Not To Cache?

- No cache for location-sensitive resources.
  - Cache-Control: no-cache HTTP response header
- Identifying location-sensitive resource
  - Developer assistance
  - Automated tool to detect location-sensitive resources

Conclusion

- Geo-inference attacks via the browser cache
- All five mainstream browsers and TorBrowser, as well as 11 map service sites and 62% of Alexa Top 100 websites, are susceptible to such attacks.
- Discussion of existing and potential defenses.
  - Calling for actions

Yaoqi Jia
E-mail: jiayaoqi@comp.nus.edu.sg