Where's Wally? Precise User Discovery Attacks in Location Proximity Services

Jason Polakis, George Argyros, Theofilos Petsios,
Suphannee Sivakorn, Angelos D. Keromytis

Columbia University

CCS '15, October 14th, Denver, Colorado

Outline

- Introduction
- Location Proximity
- User Discovery Attacks
- Experimental Evaluation
- Countermeasure
- Future Work & Conclusions

polakis@cs.columbia.edu

Introduction

Introduction

- Social networks / LBS have massive user base
  - Facebook > 1.4 billion users

Introduction

- Social networks / LBS have massive user base
  - Facebook > 1.4 billion users
- Smartphones & Internet on-the-go
  - Real-time location information readily accessible

Introduction

- Social networks / LBS have massive user base
  - Facebook > 1.4 billion users
- Smartphones & Internet on-the-go
  - Real-time location information readily accessible
- Locational Privacy is crucial
  - Inference attacks (e.g., medical issues, religion, sexual preference)
  - Surveillance (government, law enforcement)
  - Physical threats (e.g., stalking)

Introduction

- Social networks / LBS have massive user base
  - Facebook > 1.4 billion users
- Smartphones & Internet on-the-go
  - Real-time location information readily accessible
- Locational Privacy is crucial
  - Inference attacks (e.g., medical issues, religion, sexual preference)
  - Surveillance (government, law enforcement)
  - Physical threats (e.g., stalking)
- *Your contacts are not always friendly!*
  - Fake or compromised accounts, law enforcement agents

Why location proximity?

Why location proximity?

- Exact location too pervasive

Why location proximity?

- Exact location too pervasive
- Proximity is good enough for many application scenarios
  - Impromptu meeting with friends
  - Social discovery or dating apps

Why location proximity?

- Exact location too pervasive
- Proximity is good enough for many application scenarios
  - Impromptu meeting with friends
  - Social discovery or dating apps
- Distance-based proximity
  - Who is nearby? How far away?

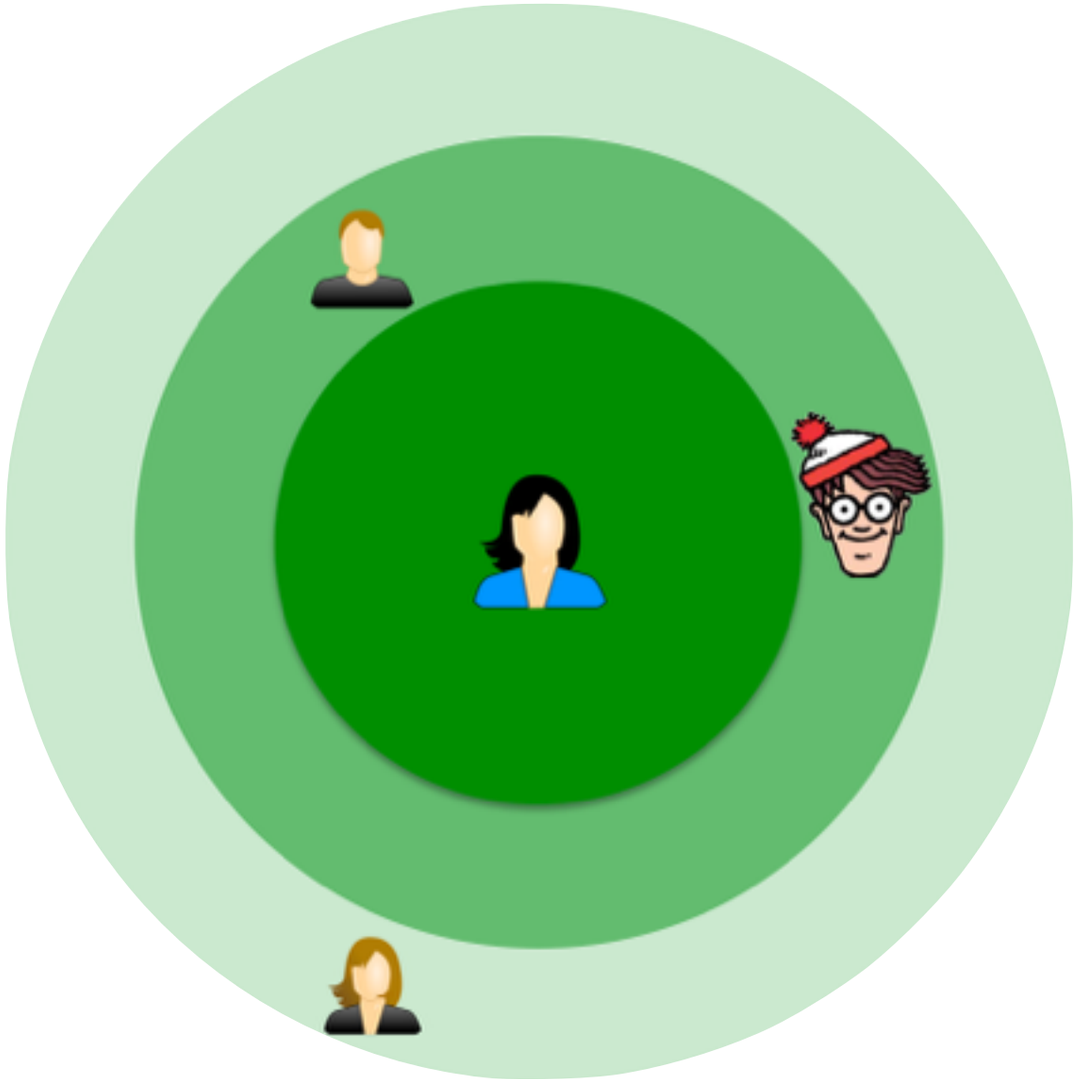

User Discovery Attacks

User Discovery Attacks

User Discovery Attacks

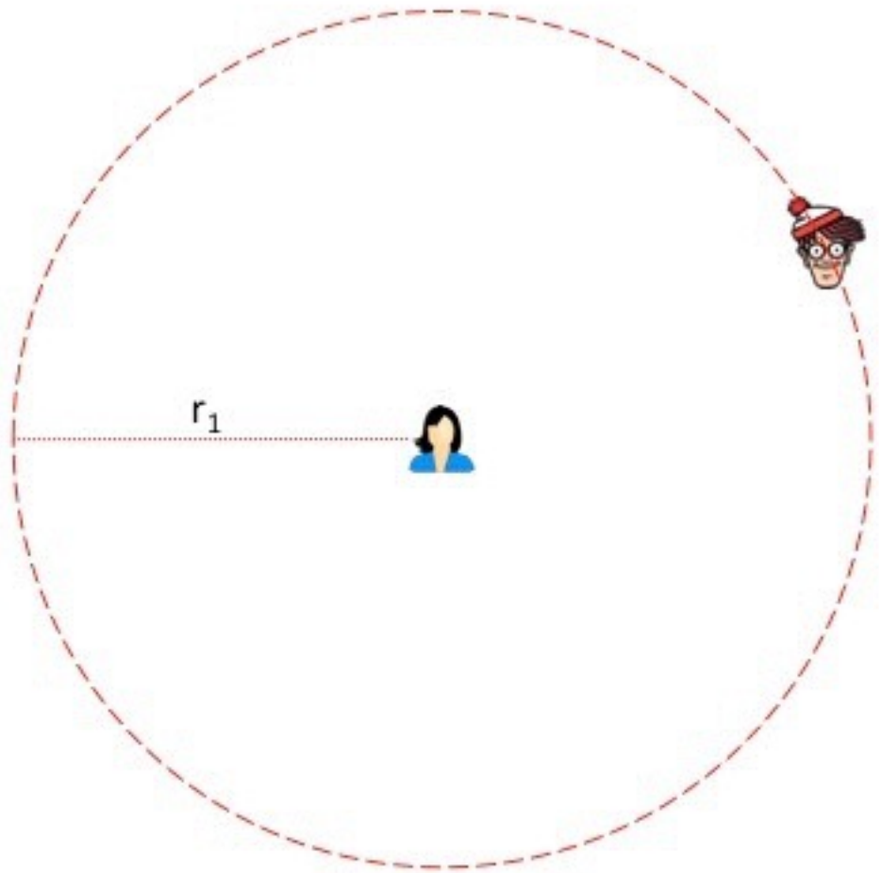

User Discovery Attacks

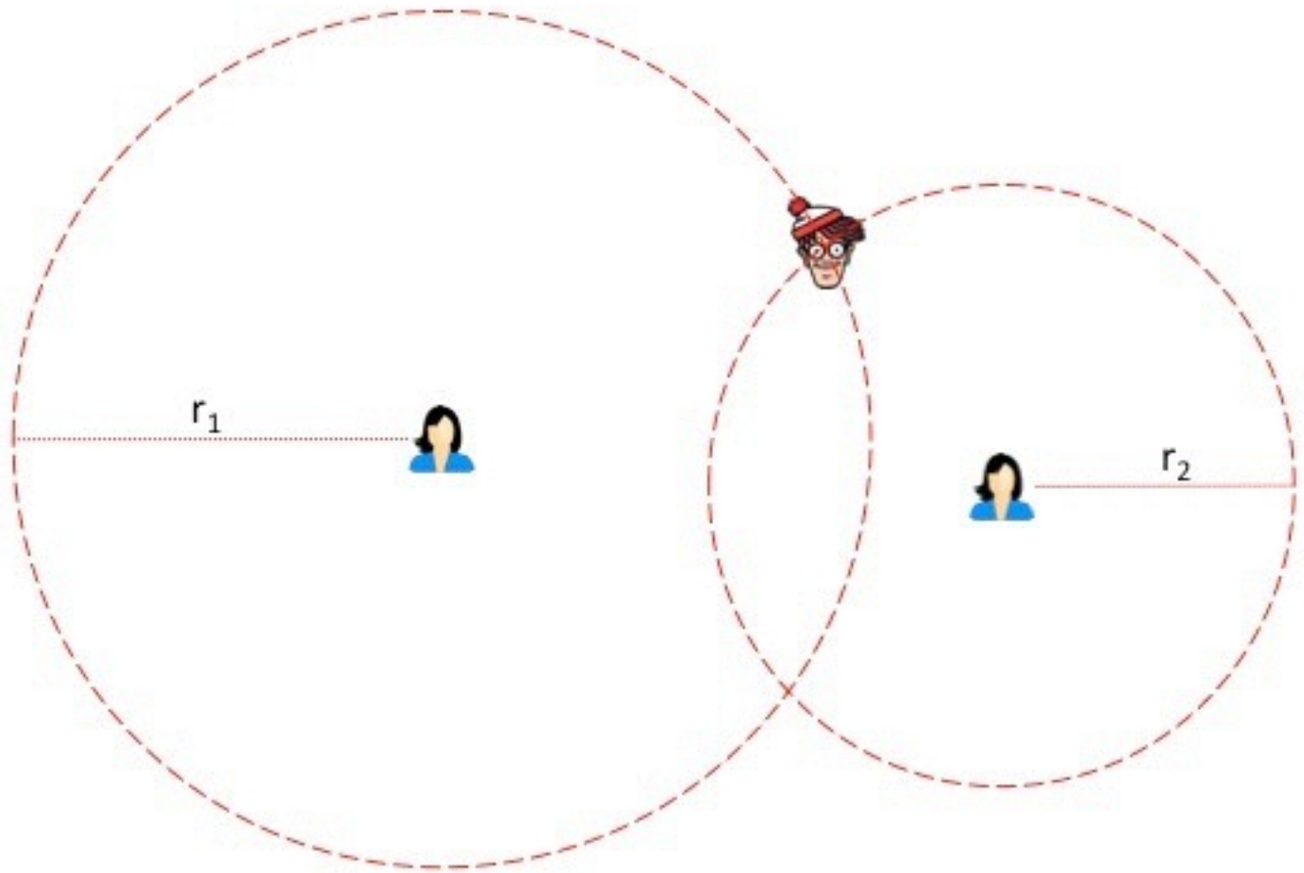

User Discovery Attacks

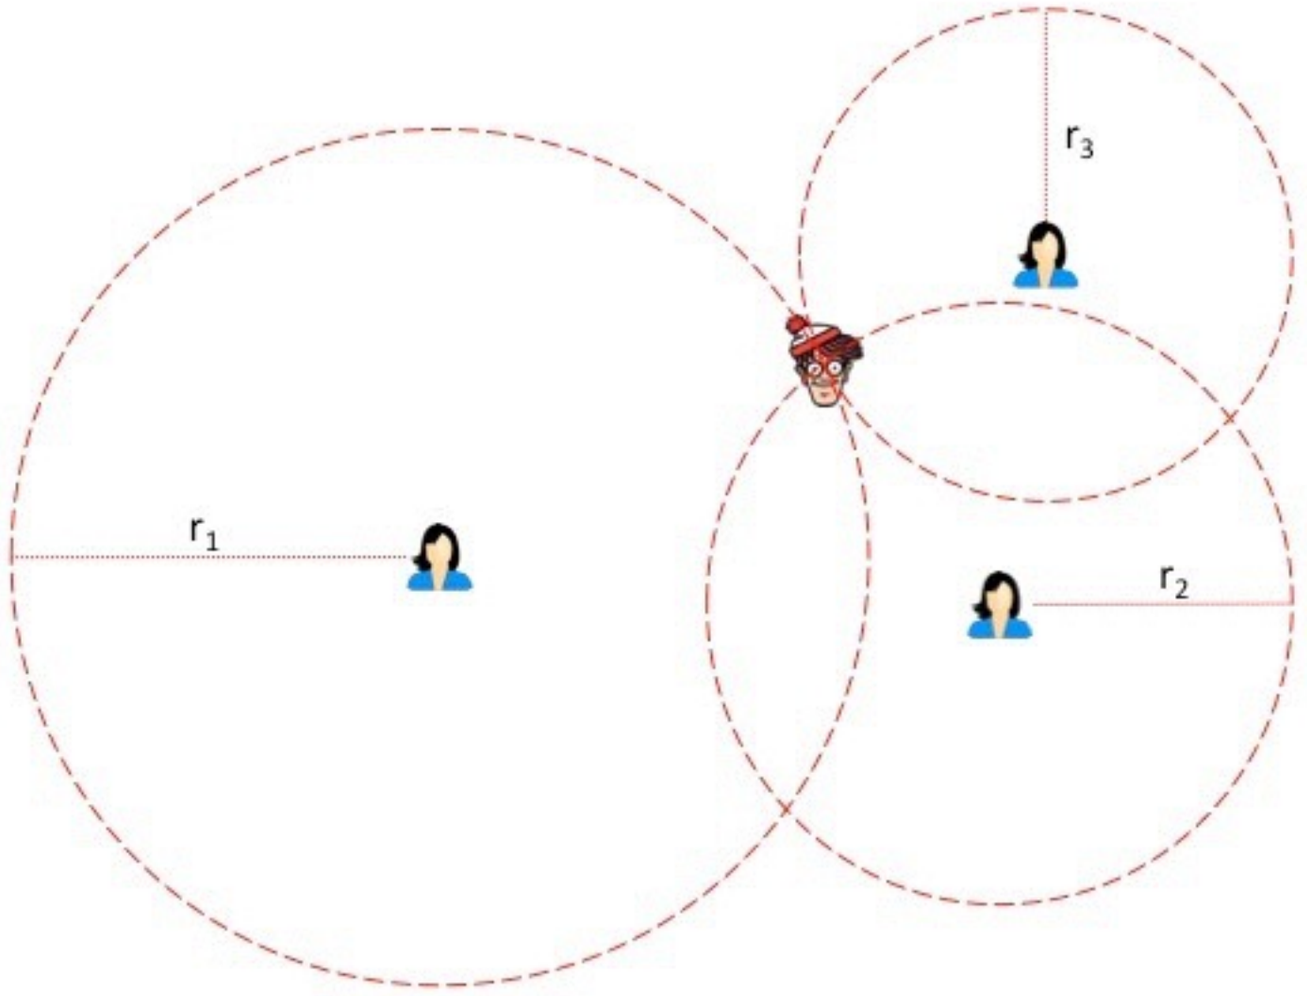

User Discovery Attacks

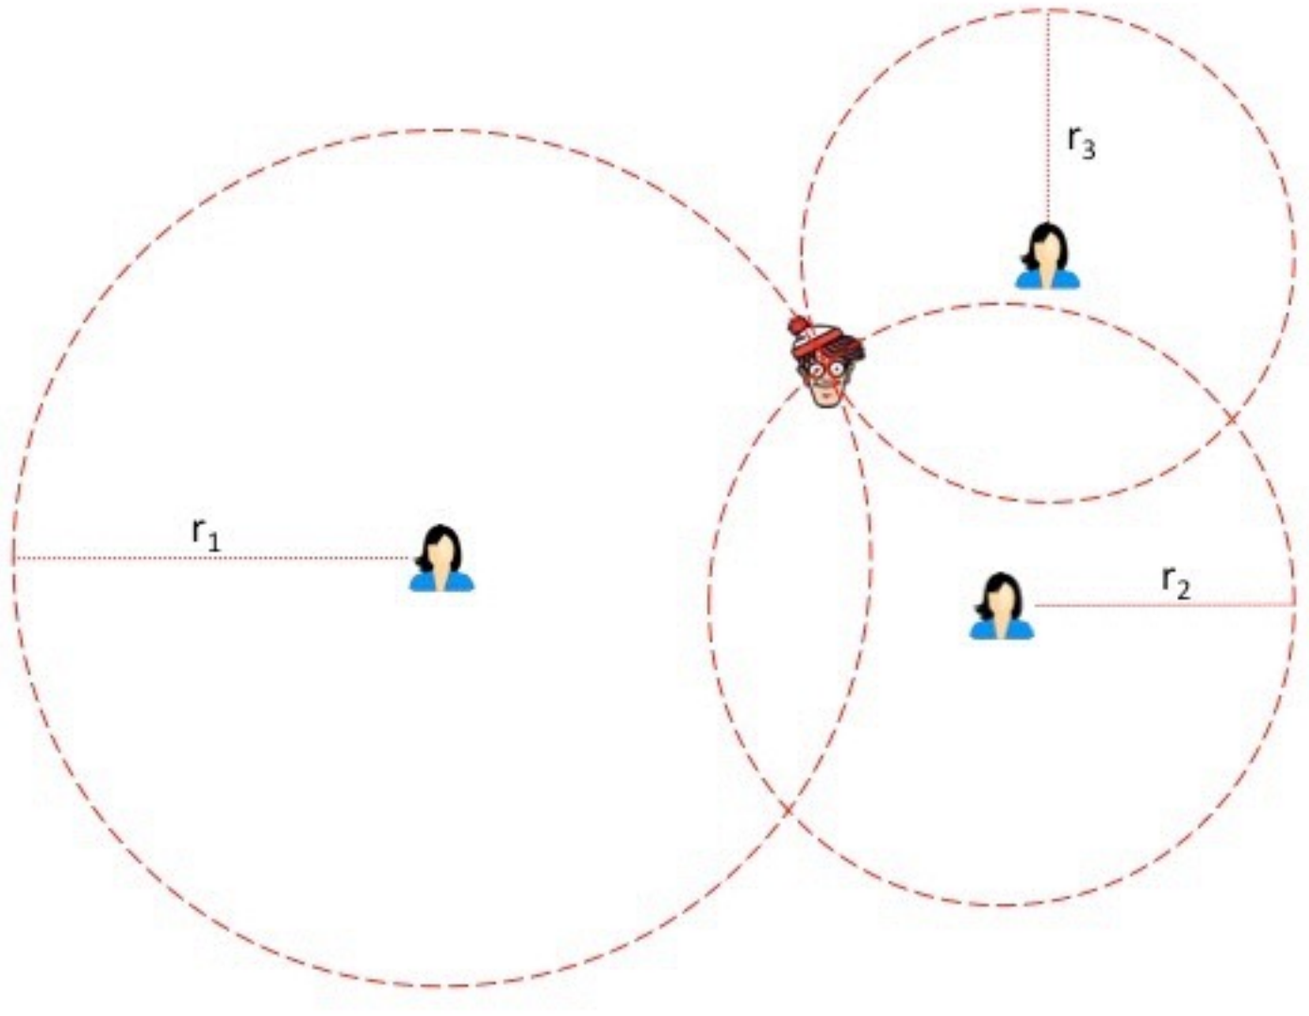

Exact distances enable trilateration

How do popular services prevent attacks?

Can we neutralize those defenses?

Location Proximity Models

Location Proximity Models

"less than 1km away"

Location Proximity Models

“less than 1km away”

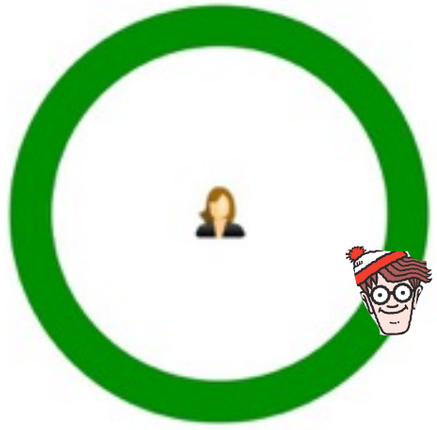

“1.5 km away”

polakis@cs.columbia.edu

Location Proximity Models

“less than 1km away”

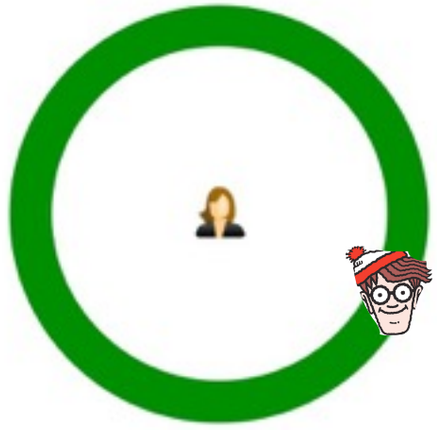

“1.5 km away”

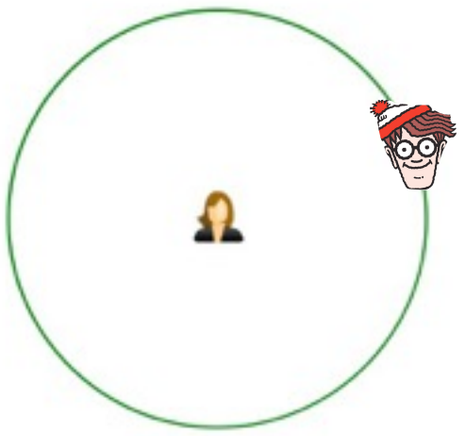

“167 meters away”

Location Proximity Models

“less than 1km away”

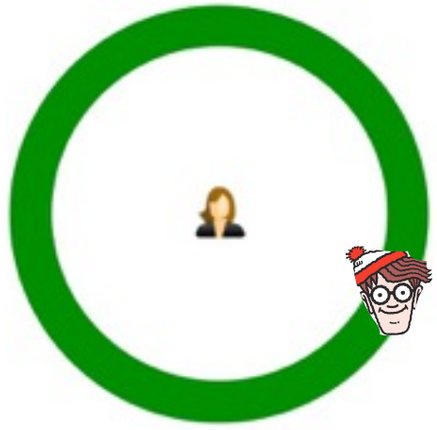

“1.5 km away”

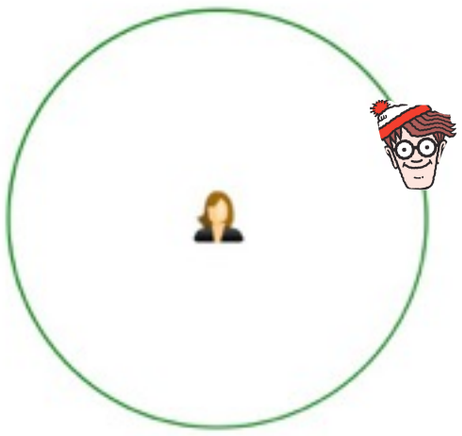

“167 meters away”

distance-based ordering

polakis@cs.columbia.edu

Modeling Discovery Attacks

- Formalized problem of discovering user's location under different proximity models

- Treat service as an oracle answering queries about distance to other users

- We **do not** use info about user's prior locations

[Check paper for analysis and more details…]

Disk Proximity Oracle

polakis@cs.columbia.edu

Disk Proximity Oracle

- Is Wally within distance X from Mallory?
  (*binary*)

Disk Proximity Oracle

- Is Wally within distance X from Mallory?
  (*binary*)

- Service can have disks of varying radius

Disk Proximity Oracle

- Is Wally within distance X from Mallory?
  (*binary*)

- Service can have disks of varying radius

Disk Proximity Oracle

- Is Wally within distance X from Mallory?
  (*binary*)
- Service can have disks of varying radius
- Attack algorithm

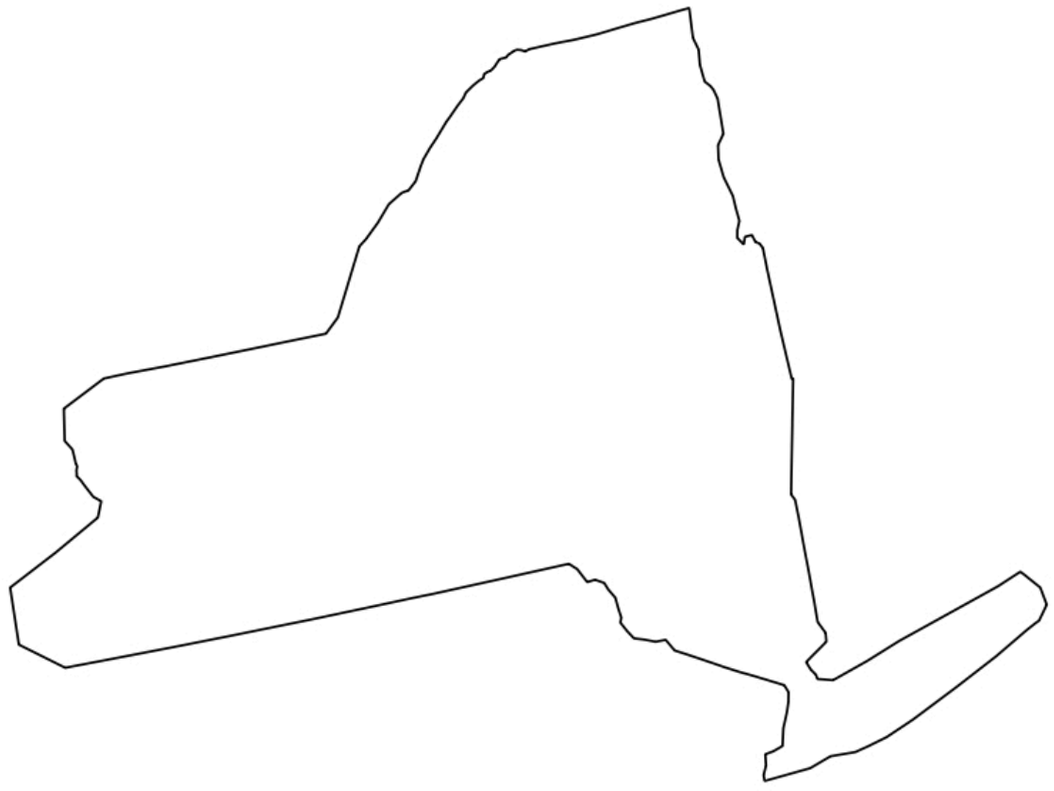

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)

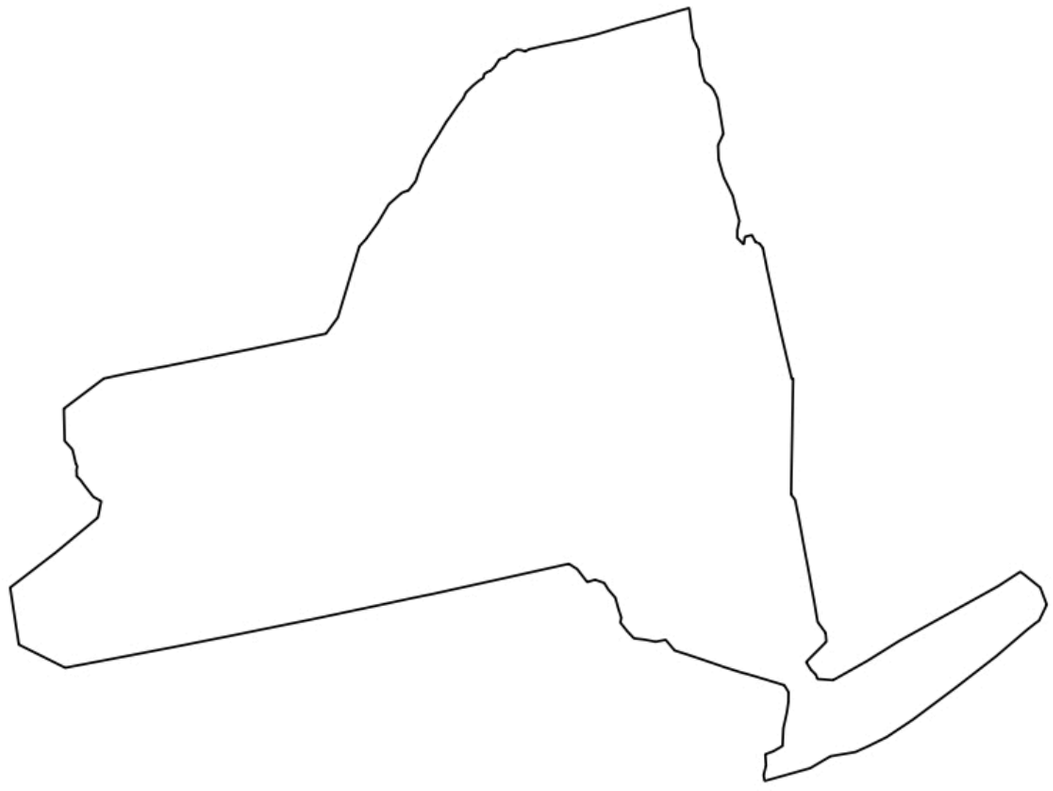

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)

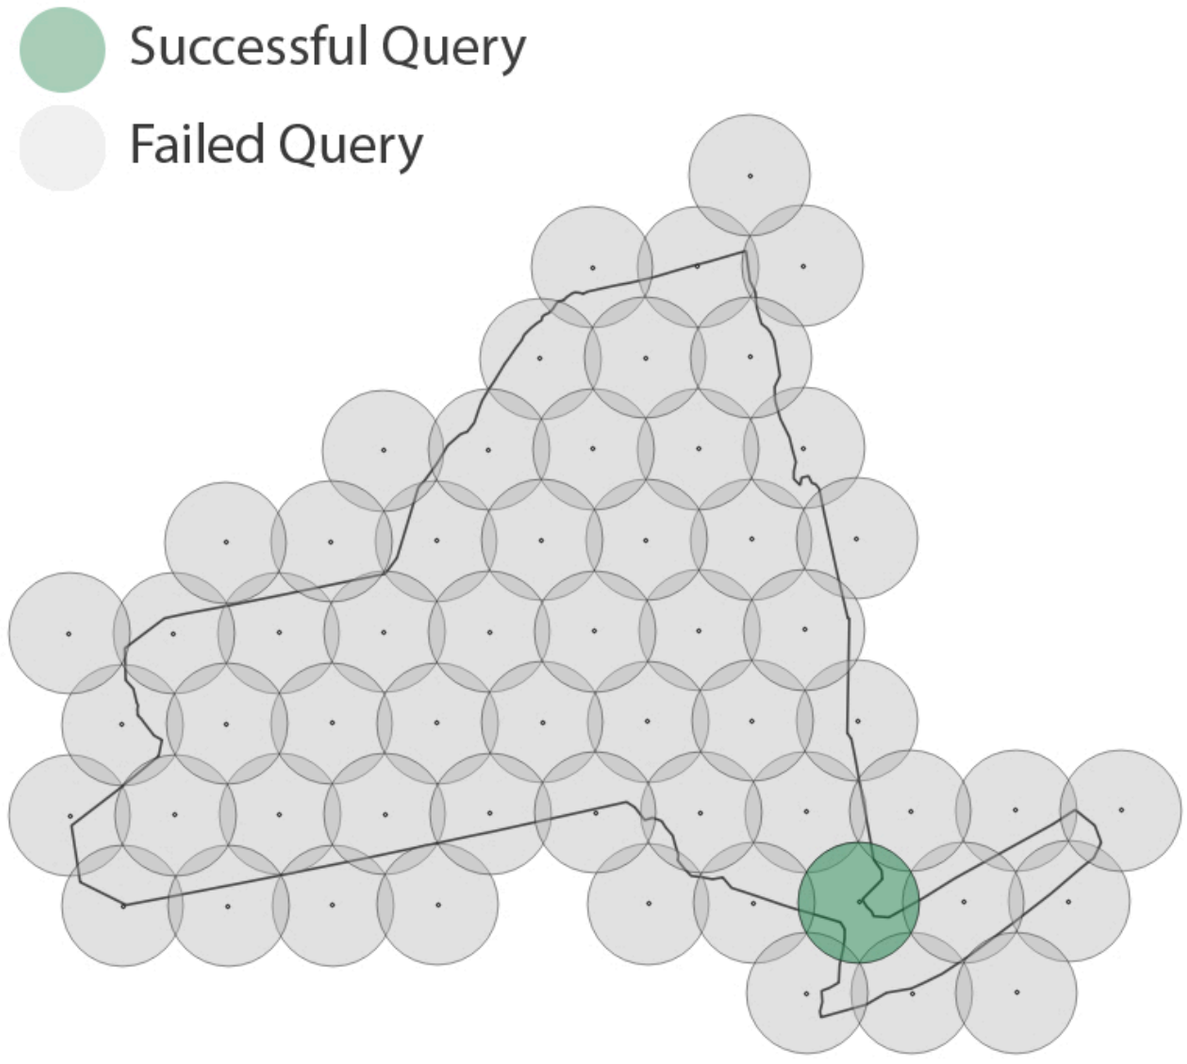

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)
    2. Reduce search area to a single point

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)
    2. Reduce search area to a single point
        - Recursively cut in half with disks (**DUDP**)

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
    1. Restrict user's location to a single disk (Disk Coverage)
    2. Reduce search area to a single point
        - Recursively cut in half with disks (**DUDP**)

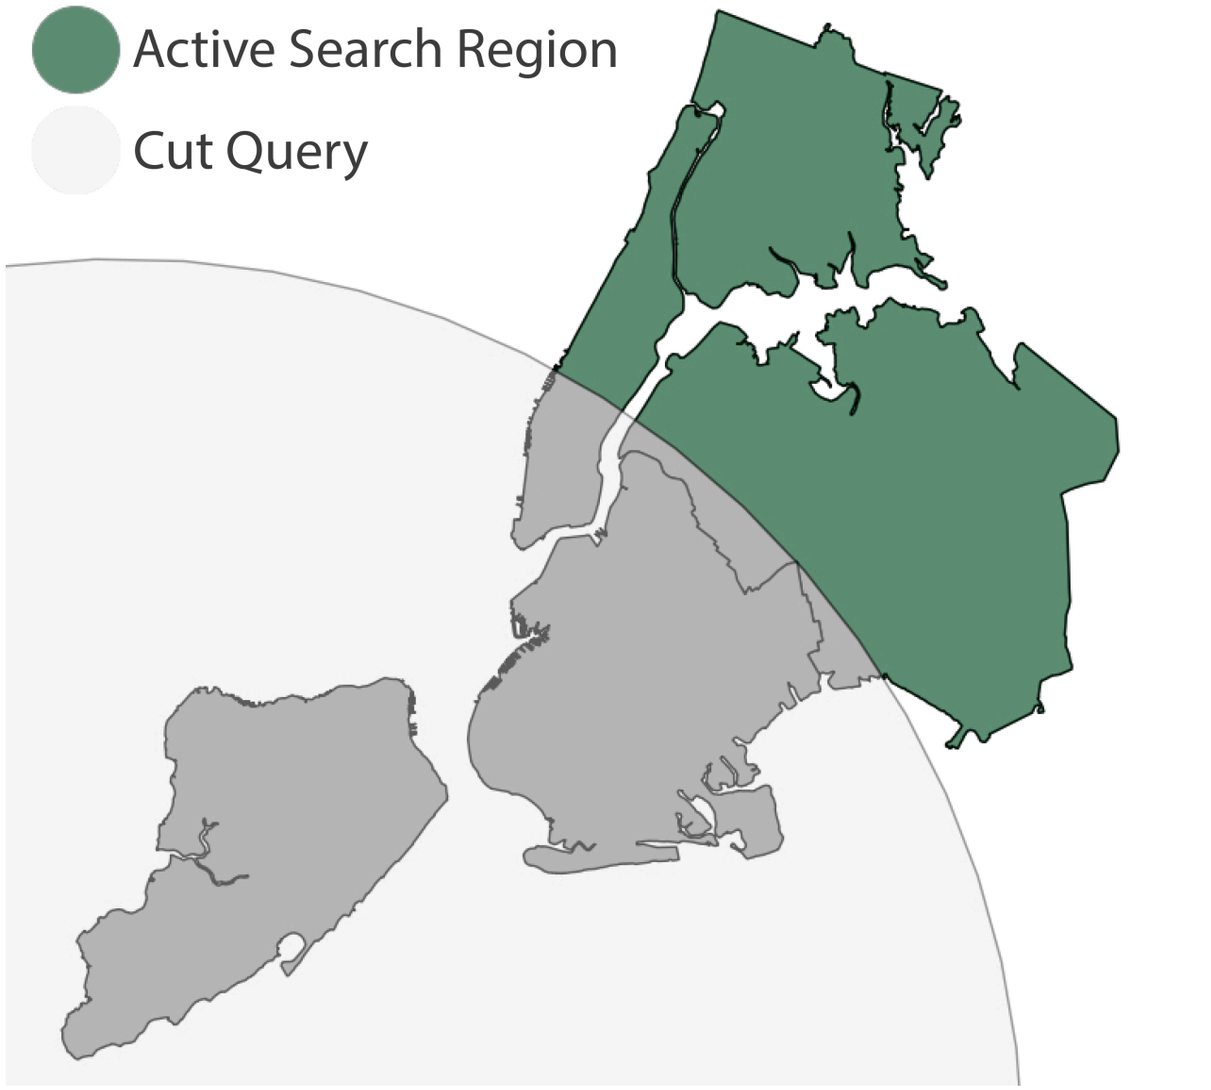

polakis@cs.columbia.edu

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm

1. Restrict user's location to a single disk (Disk Coverage)
2. Reduce search area to a single point
   - Recursively cut in half with disks (**DUDP**)

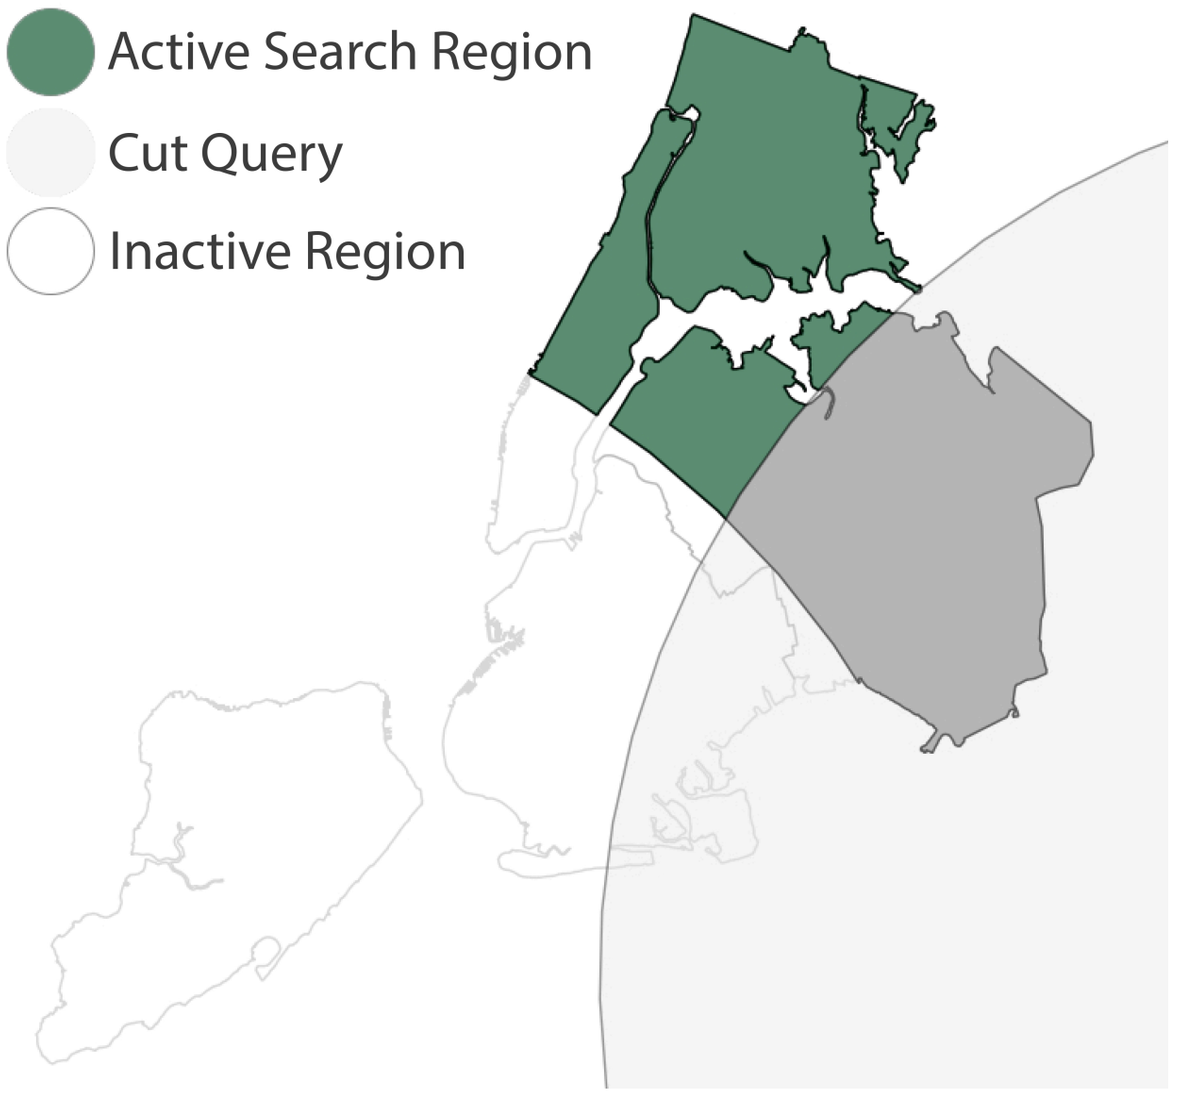

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
  1. Restrict user's location to a single disk (Disk Coverage)
  2. Reduce search area to a single point
     - Recursively cut in half with disks (**DUDP**)

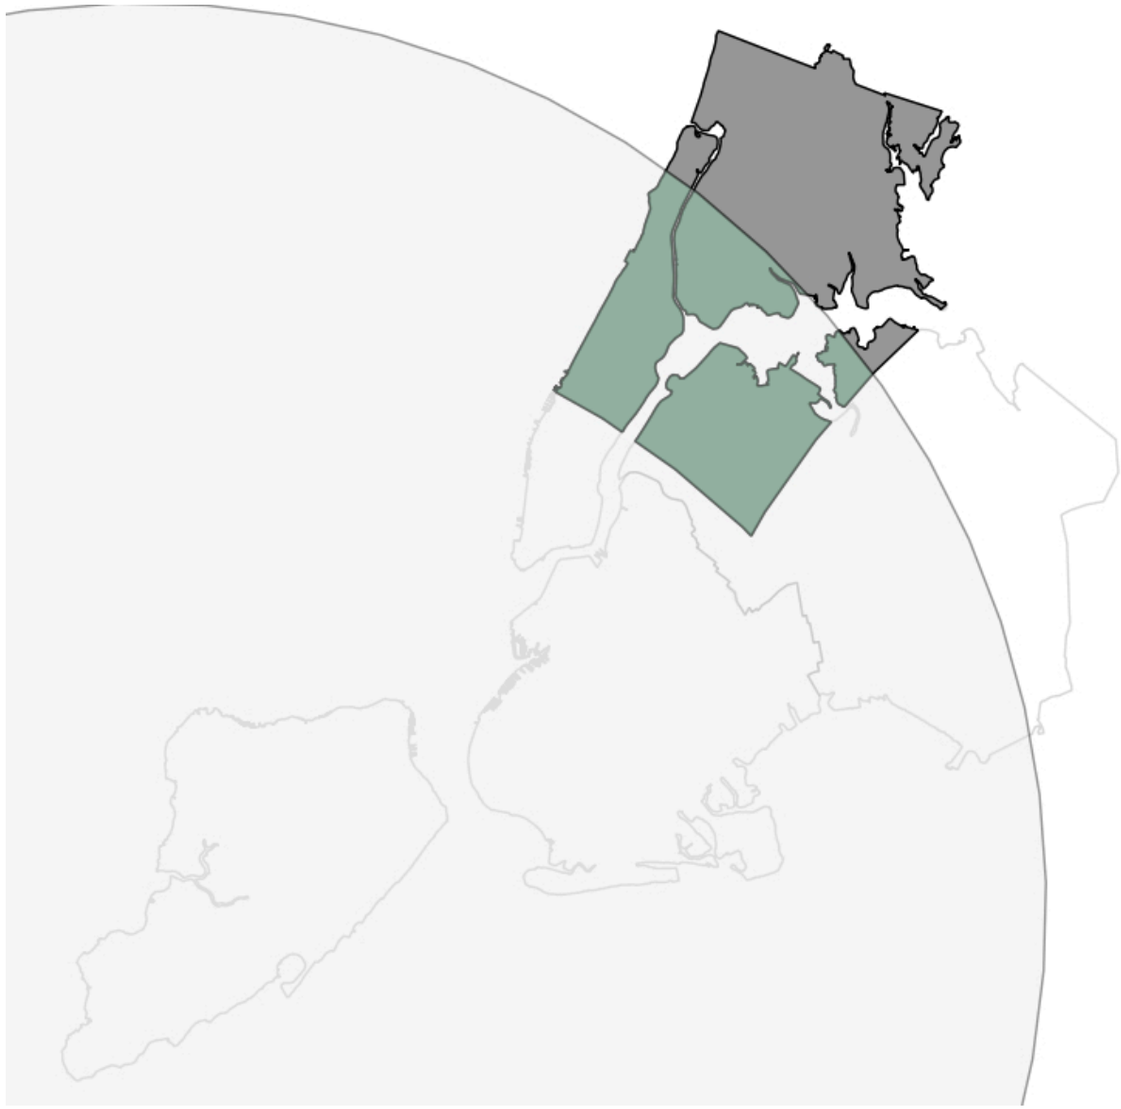

Disk Proximity Oracle

- Is Wally within distance X from Mallory? (*binary*)
- Service can have disks of varying radius
- Attack algorithm
  1. Restrict user's location to a single disk (Disk Coverage)
  2. Reduce search area to a single point
     - Recursively cut in half with disks (**DUDP**)

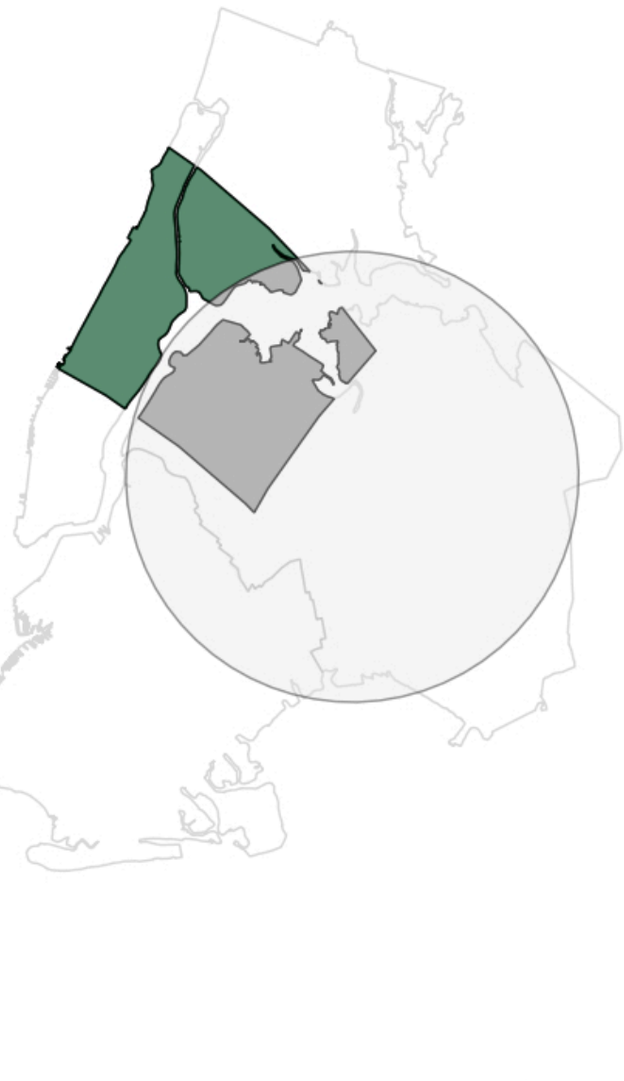

Other Oracles

Other Oracles

- Rounding Proximity Oracle (**RUDP** attack)
  - Returns rounded distance

polakis@cs.columbia.edu

Other Oracles

- Rounding Proximity Oracle (**RUDP** attack)
  - Returns rounded distance

- Randomized Proximity Oracle (**RANDUDP** attack)
  - Oracle lies with certain probability depending on distance

Practical Aspects

Practical Aspects

- “Connection” to user
  - Social networks more restricting, require connection
  - Dating apps give proximity info to anyone

Practical Aspects

- "Connection" to user
  - Social networks more restricting, require connection
  - Dating apps give proximity info to anyone
- Location spoofing
  - Service may run detection heuristics (e.g., user's moving speed)
  - Attacker can identify heuristic thresholds and remain undetected [Polakis et al. ACSAC '13]
  - May impact attack's performance

Auditing Popular Apps

	#	Dst	GPS	Grid	Query	Speed	Rand
Facebook	1-5B	◎	✗	✗	✗	✓	✗
Swarm	5-10M	◯	✓	✗	✗	✓	✓
Grindr	5-10M	⊙ | ∅	✗	✗	✓	✗	✗
Skout	10-50M	◎	✗	✗	✓	✓	✓
MeetMe	10-50M	◯ | ◎	✗	✓	✗	✓	✗
Lovoo	10-50M	◎	✗	✗	✗	✗	✗
Tinder	10-50M	◎	✗	✗	✗	✓	✗
SayHi	10-50M	⊙	✓	✗	✗	✗	✗
Jaumo	5-10M	◎	✓	✗	✗	✗	✗
HelloWorld	1K-5K	◎	✗	✗	✗	✗	✗

Distance: | exact ⊙ | rings ◎ | disks ◯ | none ∅

None of the audited services sufficiently protects users

Auditing Popular Apps

	#	Dst.	GPS	Grid	Query	Speed	Rand
Facebook	1-5B	◎	✗	✗	✗	✓	✗
Swarm	5-10M	◯	✓	✗	✗	✓	✓
Grindr	5-10M	⊙ | ∅	✗	✗	✓	✗	✗
Skout	10-50M	◎	✗	✗	✓	✓	✓
MeetMe	10-50M	◯ | ◎	✗	✓	✗	✓	✗
Lovoo	10-50M	◎	✗	✗	✗	✗	✗
Tinder	10-50M	◎	✗	✗	✗	✓	✗
SayHi	10-50M	⊙	✓	✗	✗	✗	✗
Jaumo	5-10M	◎	✓	✗	✗	✗	✗
HelloWorld	1K-5K	◎	✗	✗	✗	✗	✗

Distance: | exact ⊙ | rings ◎ | disks ◯ | none ∅

None of the audited services sufficiently protects users

Foursquare's Swarm

- Presents users in proximity disks
  - Server sends rounded distances to app
  - Distance calculations contain noise
- Reverse engineered private API
  - 2 API calls needed
  - We can arbitrarily spoof location without speed constraints

Swarm: Locating Static Users

Swarm: Locating Static Users

Wally is somewhere in NY
Mallory uses one account

Swarm: Locating Static Users

Wally is somewhere in NY
Mallory uses one account

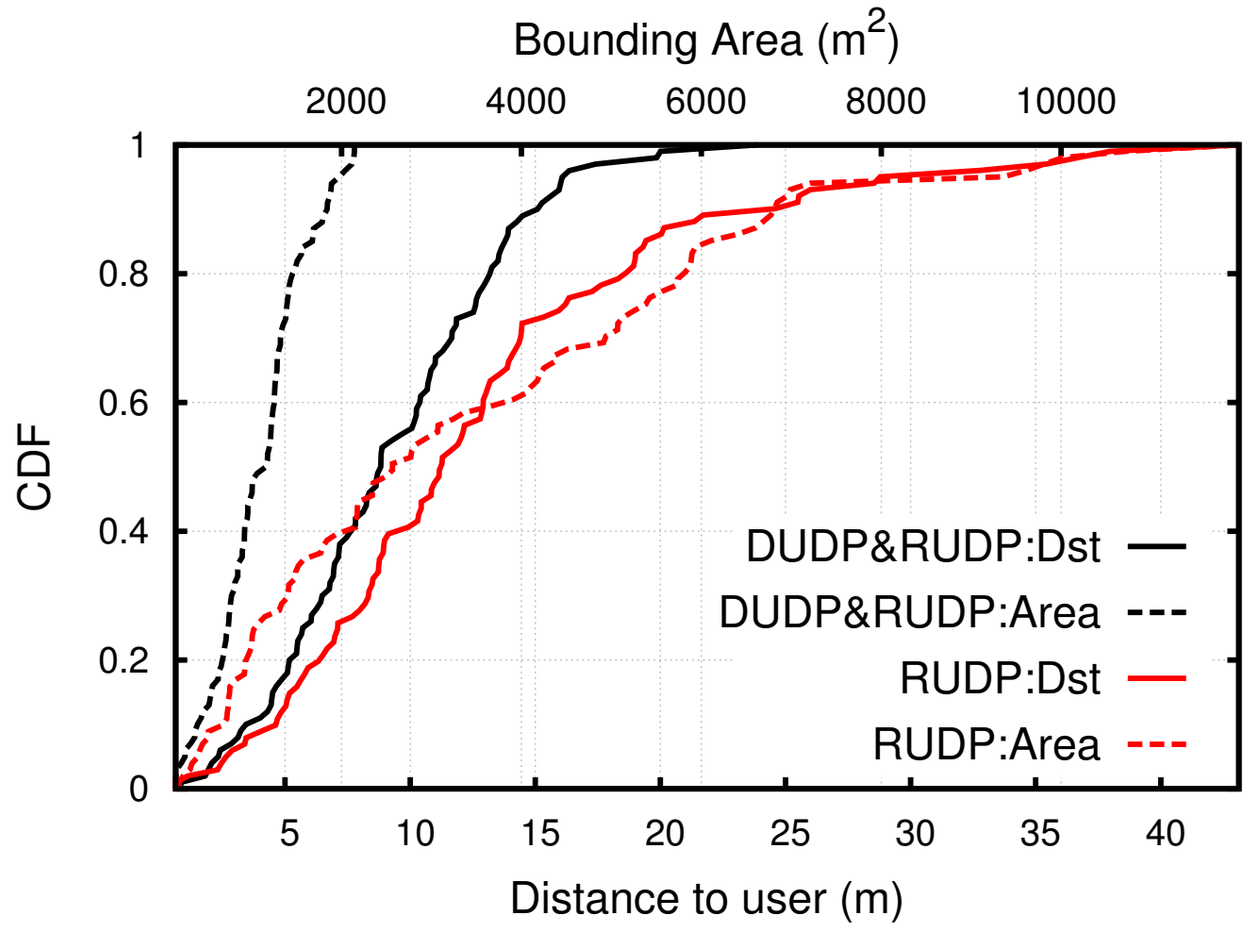

polakis@cs.columbia.edu

Swarm: Locating Static Users

Wally is somewhere in NY
Mallory uses one account

DUDP&RUDP (accurate)
- 56 queries
- 6.9 seconds
- avg 9.5m

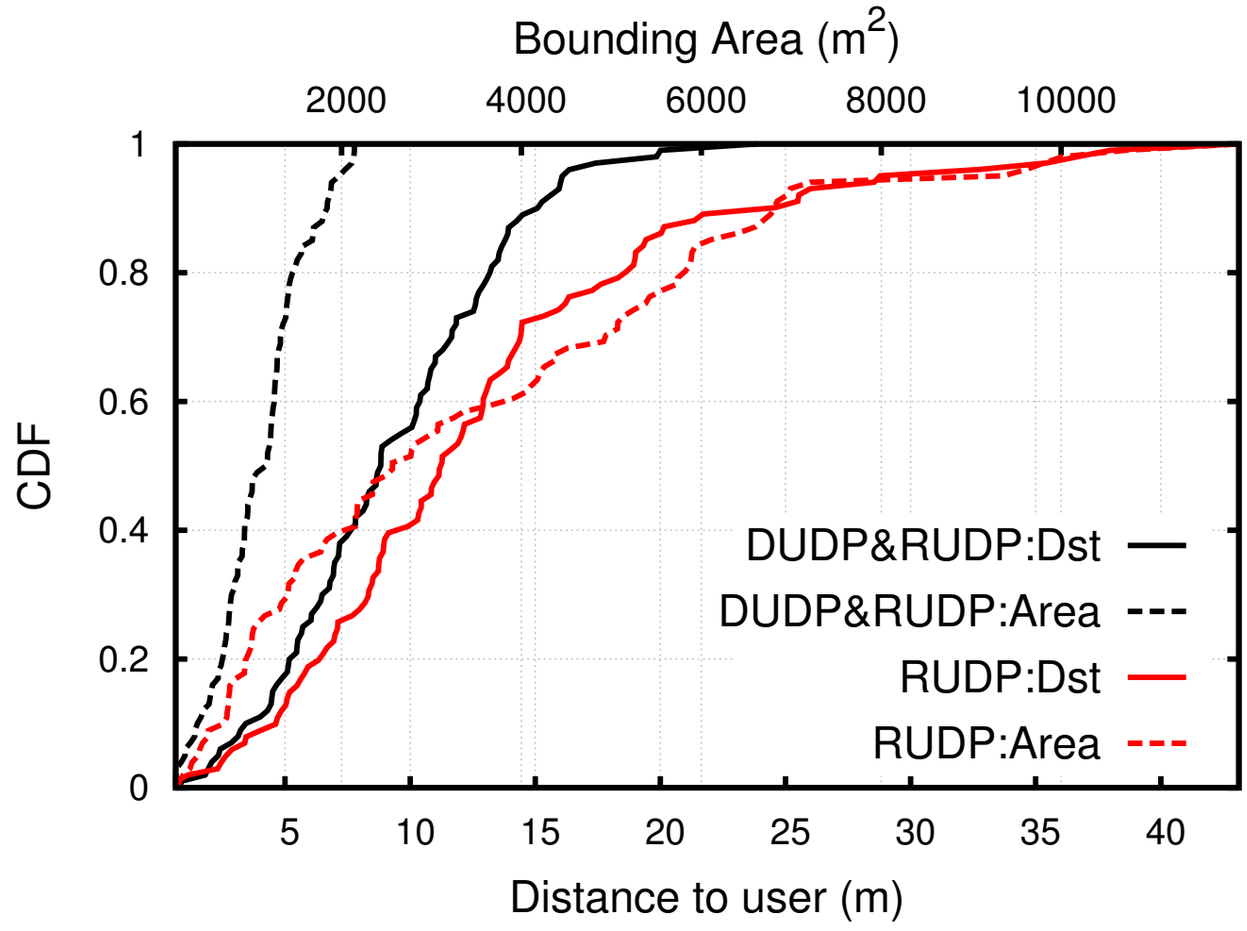

Swarm: Locating Static Users

Wally is somewhere in NY
Mallory uses one account

DUDP&RUDP (accurate)
56 queries
6.9 seconds
avg 9.5m

RUDP (efficient)
18 queries
2.6 seconds
73% within 15m

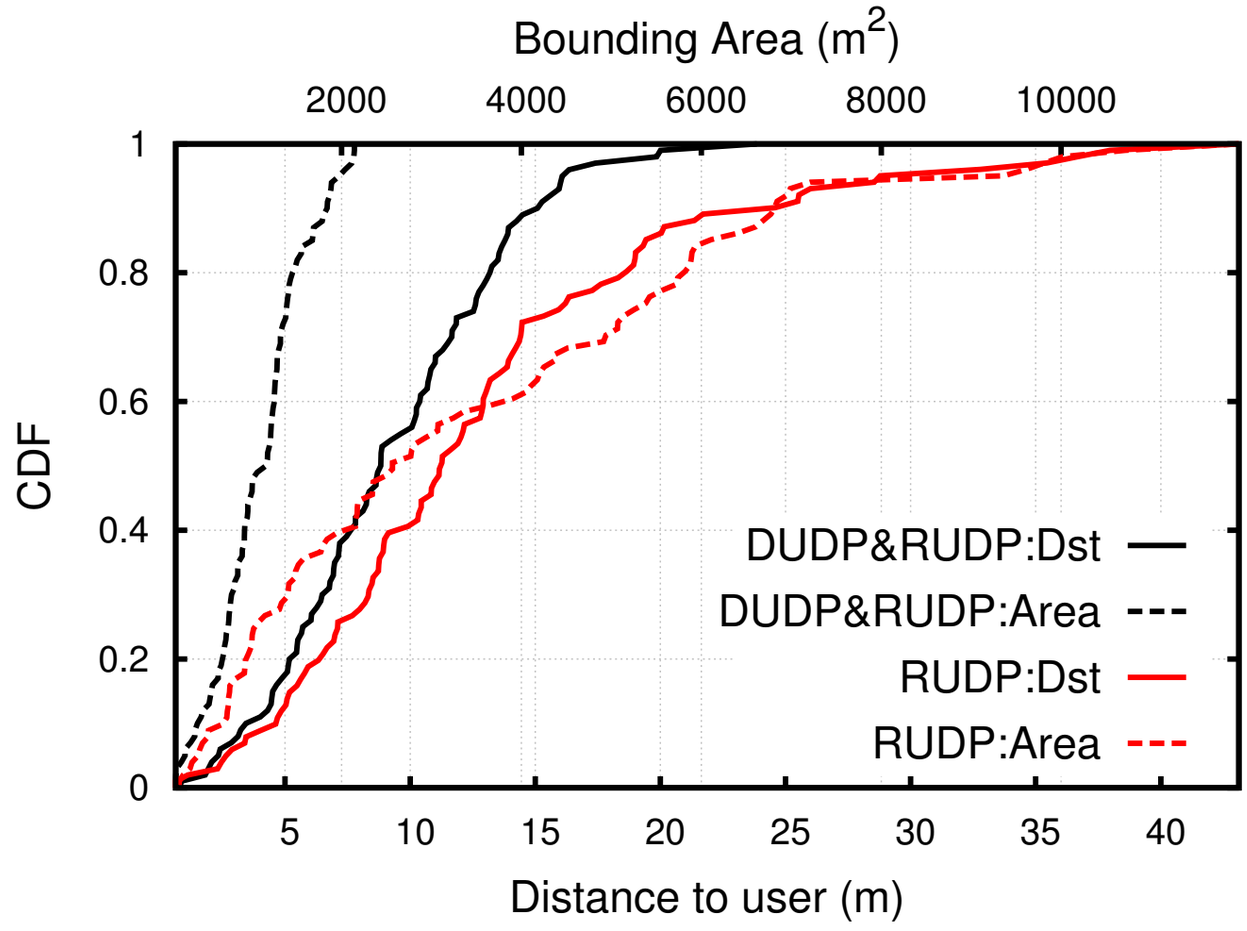

polakis@cs.columbia.edu

Swarm: Locating Moving Users

Swarm: Locating Moving Users

Let's stress-test our attack!
Wally is moving at constant speed
Location updated every 10 sec

Swarm: Locating Moving Users

Let's stress-test our attack!
Wally is moving at constant speed
Location updated every 10 sec

DUDP&RUDP (accurate)
10kmh: 90% within 16m
60kmh: 84% within 16m

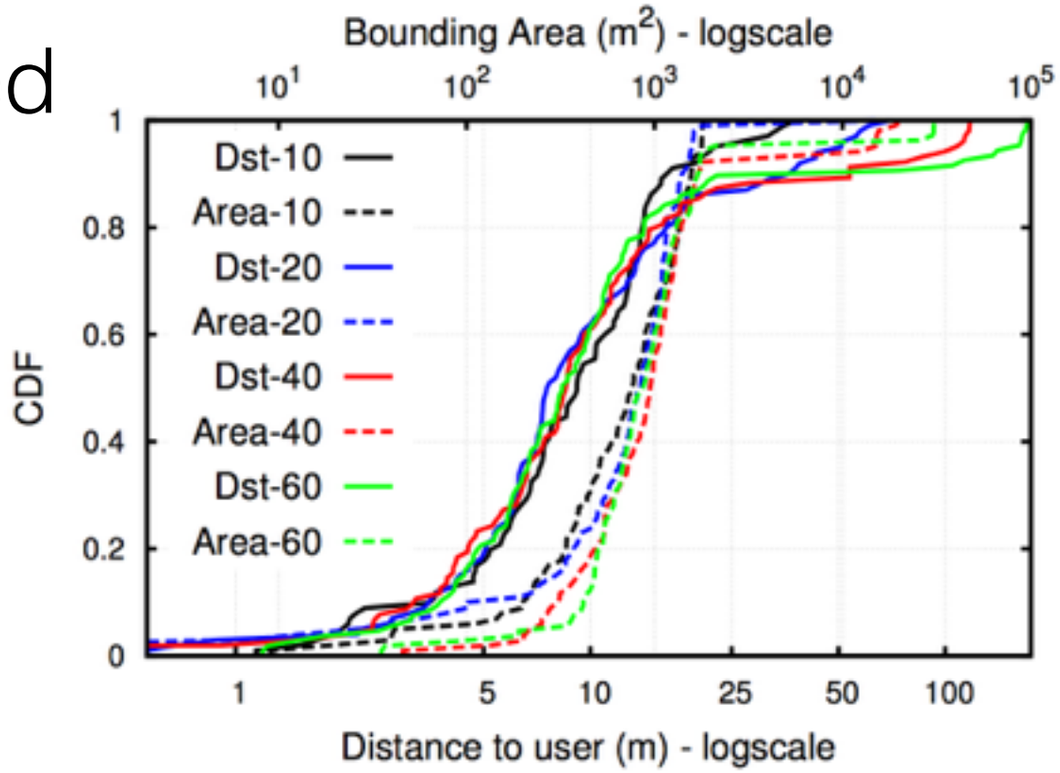

Swarm: Locating Moving Users

Let's stress-test our attack!
Wally is moving at constant speed
Location updated every 10 sec

DUDP&RUDP (accurate)
10kmh: 90% within 16m
60kmh: 84% within 16m

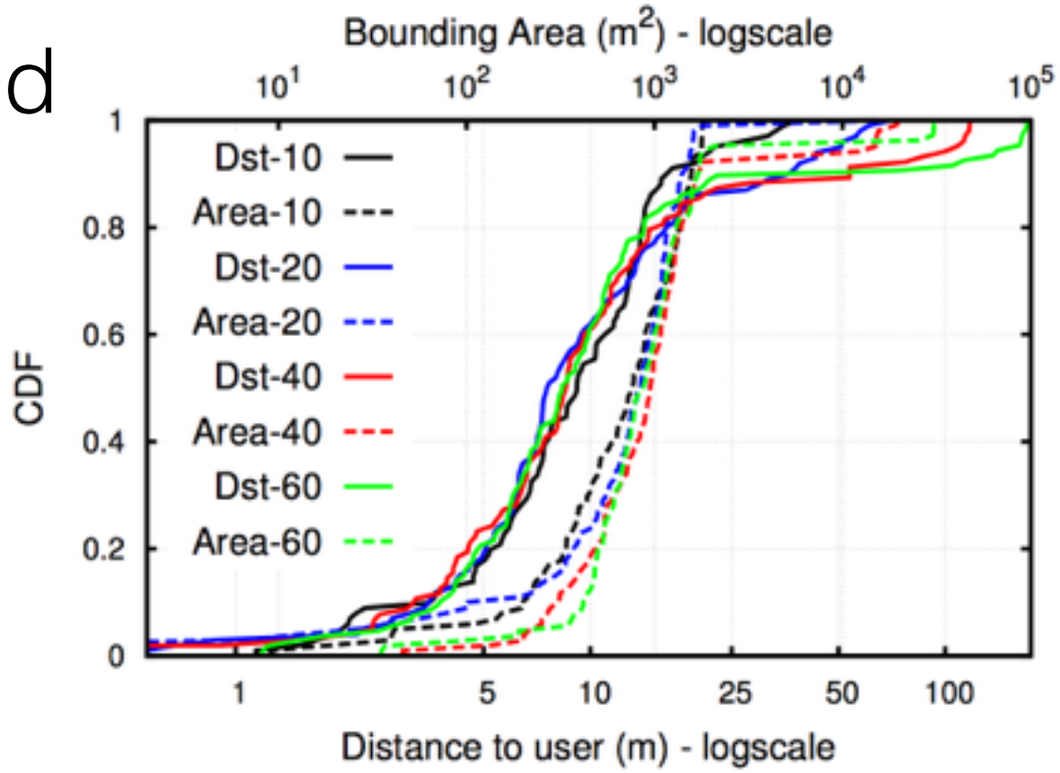

Potential modifications
restrict active search space (e.g., only roads)

polakis@cs.columbia.edu

Facebook (Nearby Friends)

Facebook (Nearby Friends)

- Uses rounded distances

Facebook (Nearby Friends)

- Uses rounded distances
- Reverse engineered private API
  - 2 API calls similar to Swarm
  - Again, we can arbitrarily spoof location without speed constraints

Facebook (Nearby Friends)

- Uses rounded distances
- Reverse engineered private API
  - 2 API calls similar to Swarm
    - Again, we can arbitrarily spoof location without speed constraints
- **RUDP** attack
  - 20.5 queries, 3 sec
  - Always within 5m
  - Bounding area < 100m^2

Facebook (Nearby Friends)

- Uses rounded distances
- Reverse engineered private API
  - 2 API calls similar to Swarm
    - Again, we can arbitrarily spoof location without speed constraints
- **RUDP** attack
  - 20.5 queries, 3 sec

  - Always within 5m
  - Bounding area < 100m^2

polakis@cs.columbia.edu

Grindr

- Social discovery / dating app
  - No connection required (attack with many accounts)
- Articles reported Egyptian government deploying trilateration attacks to locate users (not verified)
  - Homosexuality punishable in many countries
  - Imprisonment (70 countries), Death (5 countries) [BBC, 2014]
- Defense: hide distance for users in oppressive regimes [users still sorted based on distance]

Grindr: Locating Static Users

Grindr: Locating Static Users

- Use pair of colluding accounts and distance-based sorting, to define a disk proximity oracle

Grindr: Locating Static Users

- Use pair of colluding accounts and distance-based sorting, to define a disk proximity oracle

Grindr: Locating Static Users

- Use pair of colluding accounts and distance-based sorting, to define a disk proximity oracle

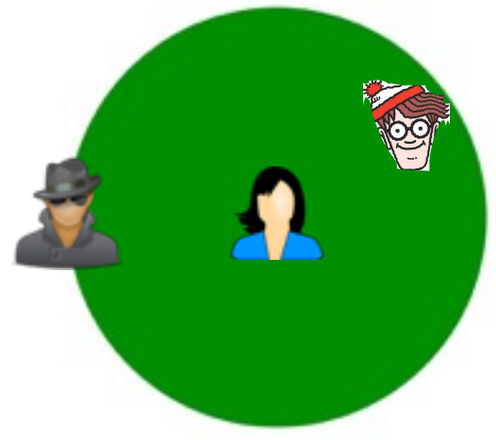

Grindr: Locating Static Users

- Use pair of colluding accounts and distance-based sorting, to define a disk proximity oracle

DUDP attack: avg 8.5m

[even without distance information]

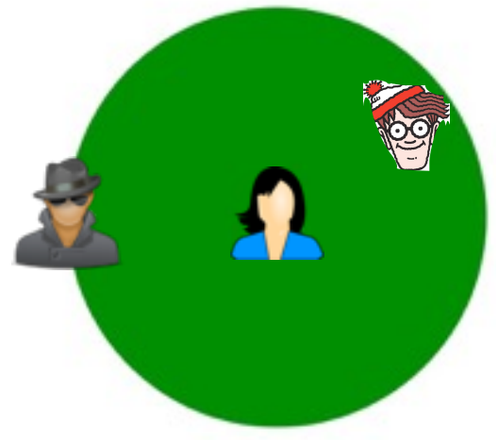

Skout

- Most advanced defense we encountered
  - Randomized proximity oracle
- Small distances rounded by 0.5 miles
- Errors for distances in [0.4, 0.6] miles
  - Error probability follows Gaussian distribution
- Nearby points “clustered” with same label

Skout: Locating Static Users

Skout: Locating Static Users

Wally is somewhere in NY

Skout: Locating Static Users

Wally is somewhere in NY

- 1400 queries per attack

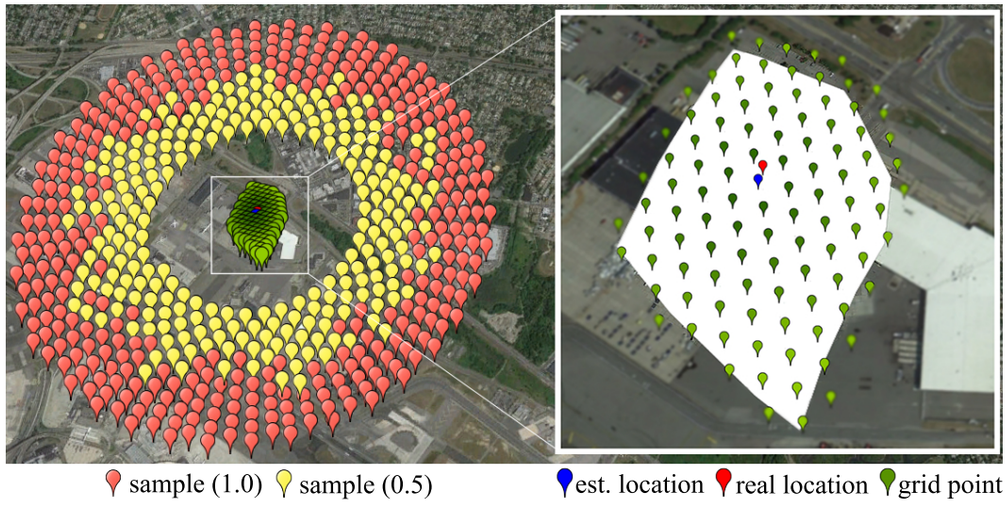

Skout: Locating Static Users

Wally is somewhere in NY

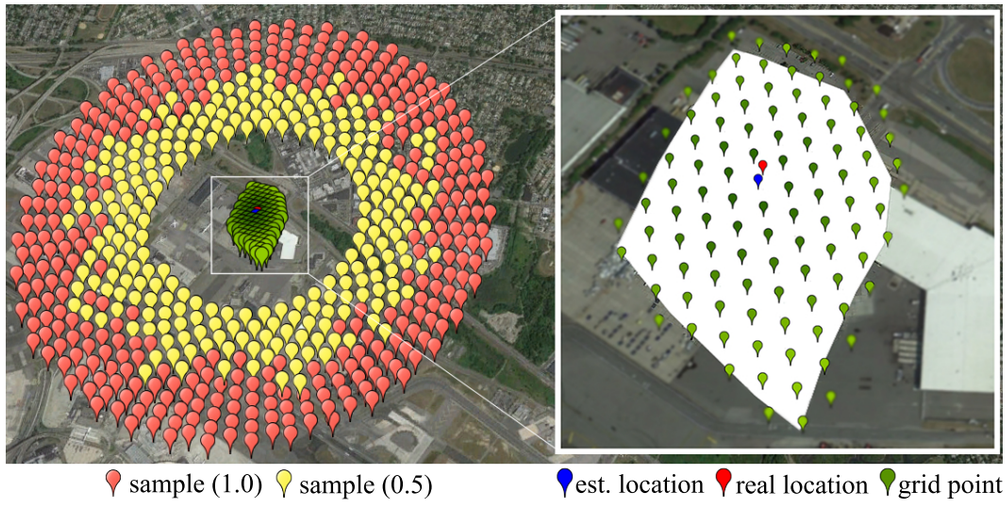

- 1400 queries per attack
- Multiple accounts for efficiency (no social connection required)

Skout: Locating Static Users

Wally is somewhere in NY

- 1400 queries per attack
- Multiple accounts for efficiency (no social connection required)

RUDP&RANDUDP
avg distance 37.4m

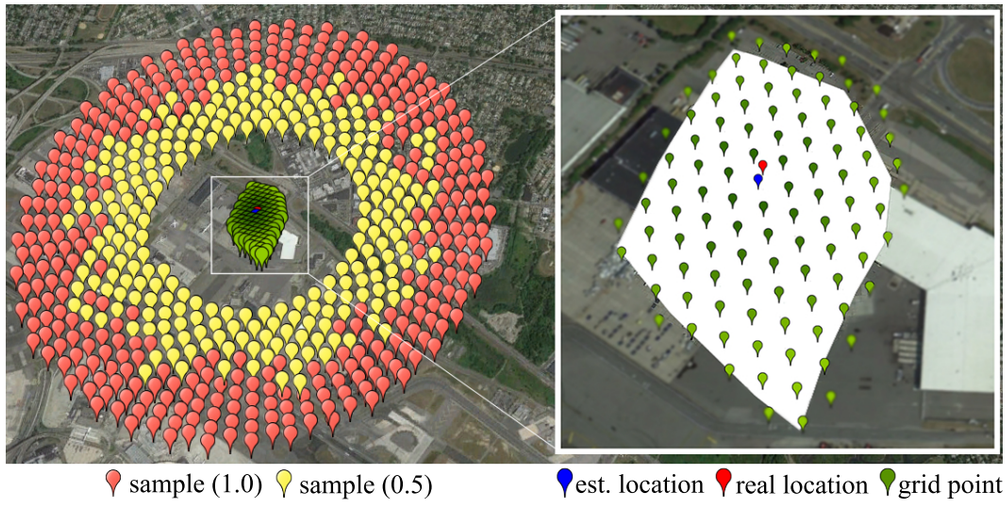

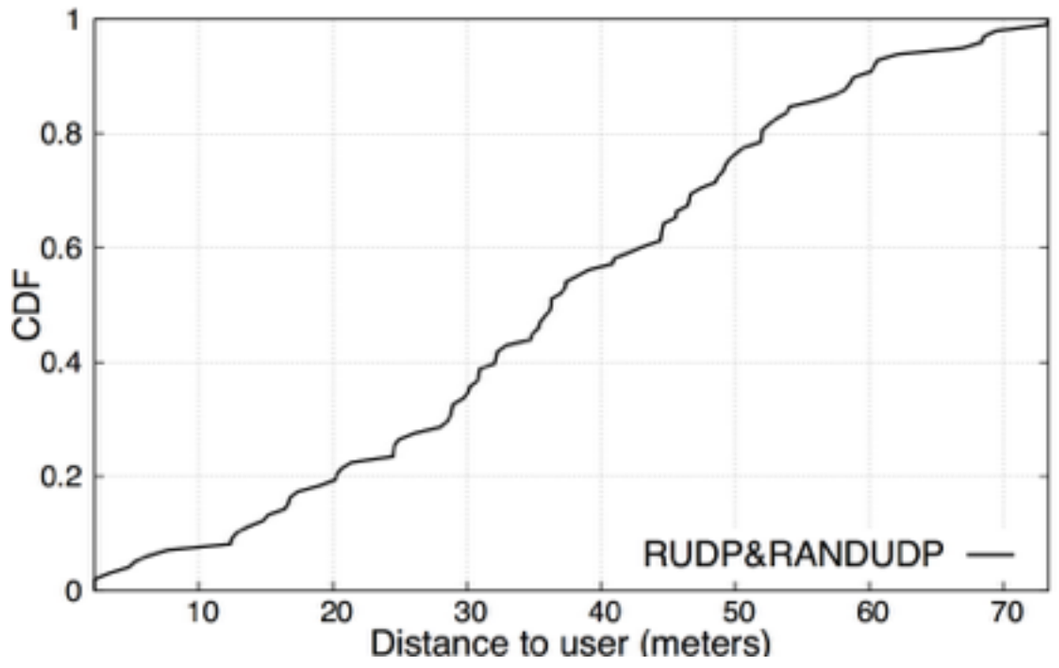

polakis@cs.columbia.edu

Preventing User Discovery

- Revisit s*patial cloaking*
  - Has been proposed for various applications
- Adapted to distance-based proximity services
- Offers nice trade-off between usability and privacy
  - Fits usability requirements of popular services

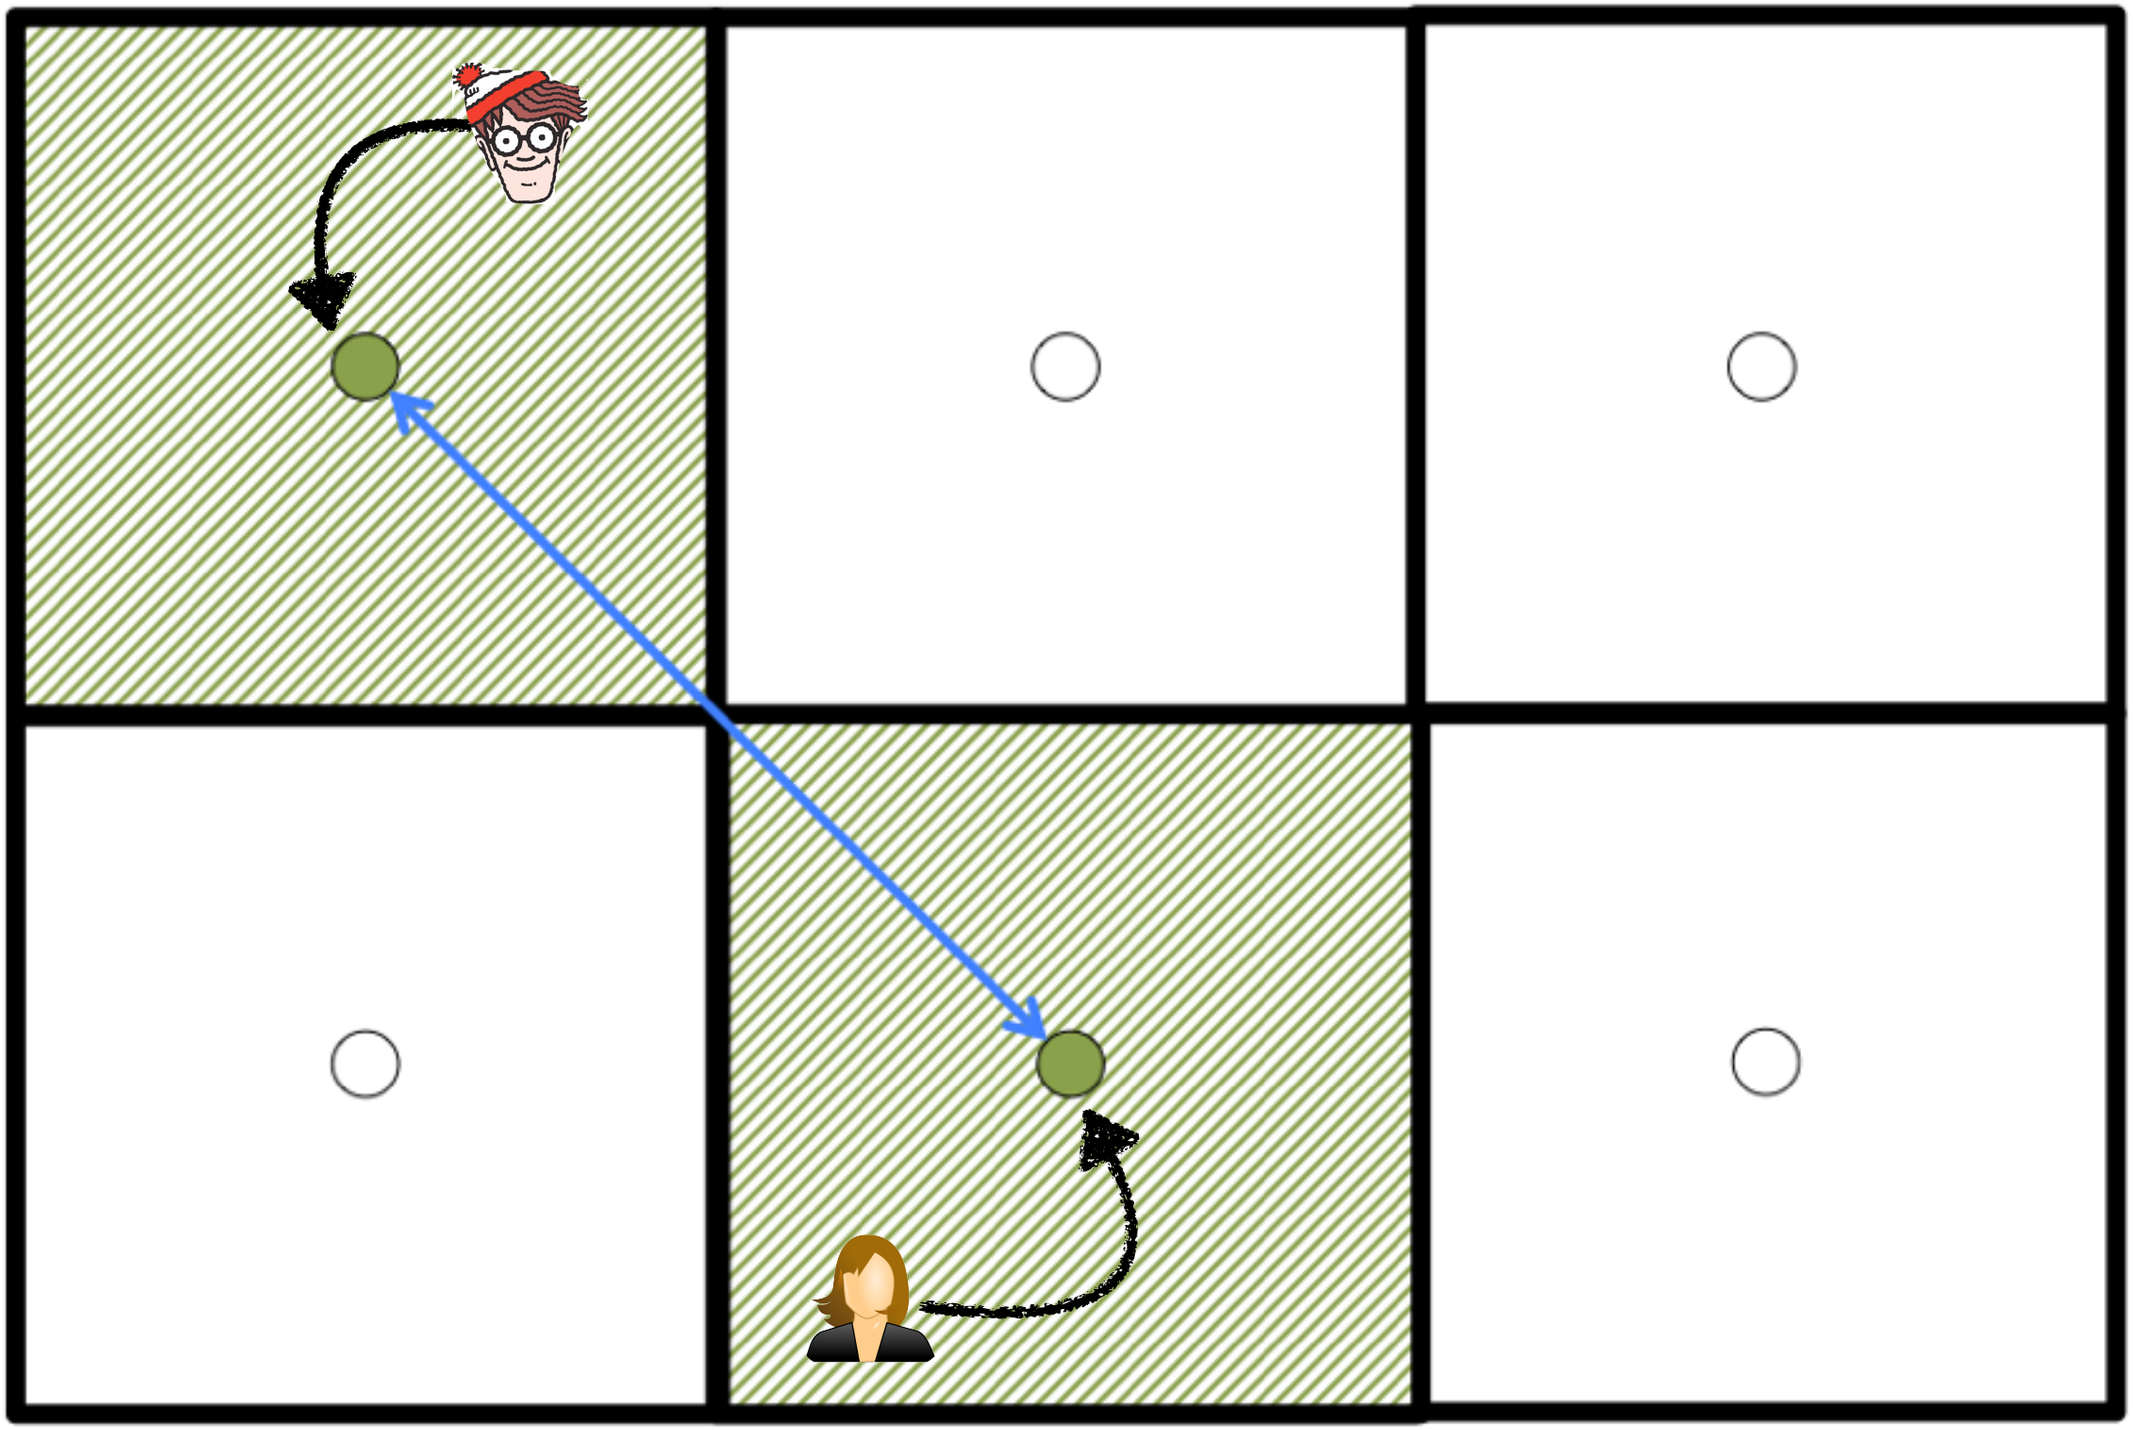

Wally's Location
Alice's Location
Distance
Grid Point
Cloaked Location
Cloaking Region

Auditing Framework

- Many apps and services remain un-tested (vulnerable?)
  - Too many for us to tackle!!
- We have released the **LBSProximityAuditor** framework
  - Implements **DUDP** and **RUDP**
  - Supports custom proximity oracles (e.g., Grindr)

More info:
www.cs.columbia.edu/~polakis/

Responsible Disclosure

Responsible Disclosure

- Contacted the 4 services we attacked
  - Details of our attacks
  - Recommendations for spatial cloaking construction

Responsible Disclosure

- Contacted the 4 services we attacked
  - Details of our attacks
  - Recommendations for spatial cloaking construction

☑ Facebook and Foursquare verified our findings and adopted spatial cloaking

Responsible Disclosure

- Contacted the 4 services we attacked
  - Details of our attacks
  - Recommendations for spatial cloaking construction

☑ Facebook and Foursquare verified our findings and adopted spatial cloaking

☑ Facebook used our framework to test their defense

Future Work & Conclusions

Future Work & Conclusions

- Demonstrated practical and effective attacks against popular services with various defenses
  - No industry standard for ensuring privacy (ad-hoc)
  - **Many** apps still vulnerable

Future Work & Conclusions

- Demonstrated practical and effective attacks against popular services with various defenses
  - No industry standard for ensuring privacy (ad-hoc)
  - **Many** apps still vulnerable
- Disclosure resulted in Facebook and Foursquare adopting spatial cloaking

Future Work & Conclusions

- Demonstrated practical and effective attacks against popular services with various defenses
  - No industry standard for ensuring privacy (ad-hoc)
  - **Many** apps still vulnerable
- Disclosure resulted in Facebook and Foursquare adopting spatial cloaking
- Spatial cloaking not a panacea!
  - Moving victims [Ghinita et al., GIS ’09]
  - Attacks leveraging prior locations of user [Theodorakopoulos et al., WPES ’14]

Future Work & Conclusions

source: Walker Book Ltd

Future Work & Conclusions

source: Walker Book Ltd

Questions?

more info:

www.cs.columbia.edu/~polakis/

polakis@cs.columbia.edu