

IP GEOLOCATION

Using Traceroute Location Propagation
and IP Range Location Interpolation

Ovidiu Dan, Vaibhav Parikh, Brian. D. Davison





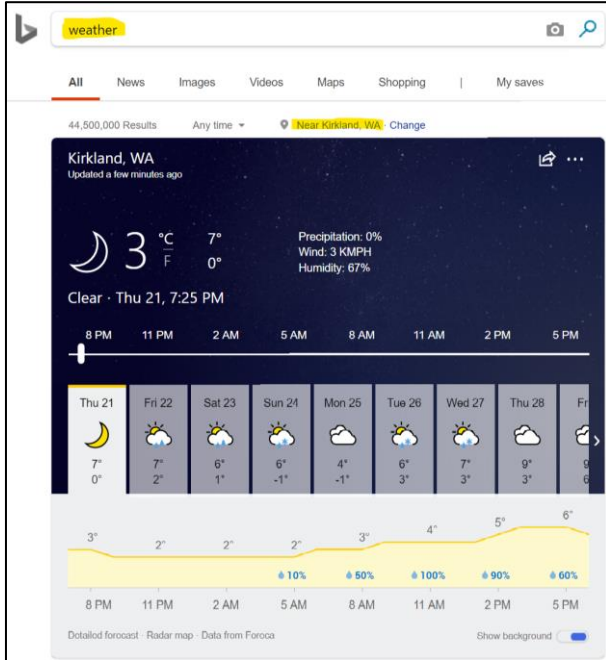
- » What is IP Geolocation
- » Commercial Geolocation Services
- » Motivation and Goal
- » Privacy and Datasets
- » IP Colocation
- » IP Interpolation
- » Traceroute Location Propagation
- » Related Work (Appendix)

1.

WHAT IS IP GEOLOCATION

Background, Motivation, Research Questions, Contributions





weather

All News Images Videos Maps Shopping | My saves

44,500,000 Results Any time Near Kirkland, WA Change

Kirkland, WA

Updated a few minutes ago

3°C 7°F
Precipitation: 0%
Wind: 3 KMPH
Humidity: 67%

Clear · Thu 21, 7:25 PM

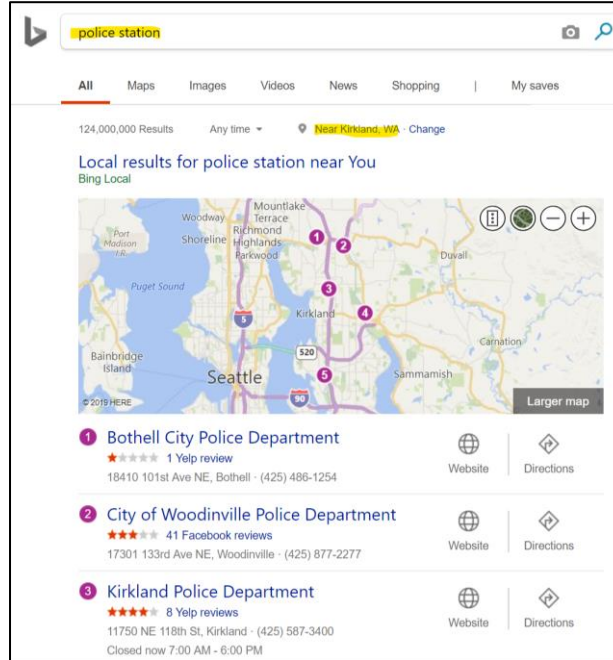
8 PM 11 PM 2 AM 5 AM 8 AM 11 AM 2 PM 5 PM

Thu 21	Fri 22	Sat 23	Sun 24	Mon 25	Tue 26	Wed 27	Thu 28	Fri 29
7° 0°	7° 2°	6° 1°	6° -1°	4° -1°	6° 3°	7° 3°	9° 3°	9° 6°

3° 2° 2° 2° 3° 4° 5° 6°

8 PM 11 PM 2 AM 5 AM 8 AM 11 AM 2 PM 5 PM

Detailed forecast · Radar map · Data from Forcasto Show background




police station

All Maps Images Videos News Shopping | My saves

124,000,000 Results Any time Near Kirkland, WA Change

Local results for police station near You

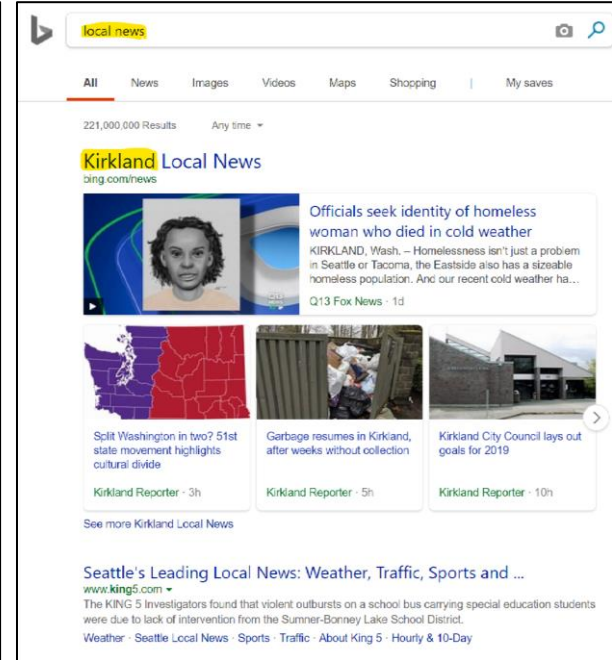
Bing Local



1 **Bothell City Police Department**
★★★★★ 1 Yelp review
18410 101st Ave NE, Bothell · (425) 486-1254 Website Directions

2 **City of Woodinville Police Department**
★★★★★ 41 Facebook reviews
17301 133rd Ave NE, Woodinville · (425) 877-2277 Website Directions

3 **Kirkland Police Department**
★★★★★ 8 Yelp reviews
11750 NE 118th St, Kirkland · (425) 587-3400
Closed now 7:00 AM - 6:00 PM Website Directions



local news

All News Images Videos Maps Shopping | My saves

221,000,000 Results Any time

Kirkland Local News

bing.com/news

Officials seek identity of homeless woman who died in cold weather
KIRKLAND, Wash. — Homelessness isn't just a problem in Seattle or Tacoma, the Eastside also has a sizeable homeless population. And our recent cold weather ha...
Q13 Fox News · 1d

Split Washington in two? 51st state movement highlights cultural divide
Kirkland Reporter · 3h

Garbage resumes in Kirkland, after weeks without collection
Kirkland Reporter · 5h

Kirkland City Council lays out goals for 2019
Kirkland Reporter · 10h

See more Kirkland Local News

Seattle's Leading Local News: Weather, Traffic, Sports and ...

www.king5.com

The KING 5 Investigators found that violent outbursts on a school bus carrying special education students were due to lack of intervention from the Sumner-Bonney Lake School District.

Weather · Seattle Local News · Sports · Traffic · About King 5 · Hourly & 10-Day

StartIP	EndIP	Country	State	City
1.0.16.0	1.0.16.255	JP	Tokyo	Tokyo
1.0.48.0	1.0.48.255	CN	Guangdong	Guangzhou
187.153.184.0	187.153.184.255	MX	Quintana roo	Cancun



- » Search Engine Personalization
- » Content Delivery Networks
- » Credit card fraud protection
- » Click fraud protection
- » Advertising network targeting
- » Cybercrime investigations
- » Location based content licensing
- » E-commerce
- » Organizations with regional offices
- » Automatic Language Selection
- » Network Security
- » Academic Research

- » Personalization
- » Cohort modeling

Add a rule for **blocking payments**
Payments matching this rule will be blocked unless they already matched an allow rule.

Block if

Example rules

```
Block if :charge_attempts_per_card_number_hourly: > 10 and  
:card_country: in @my_card_country_blocklist;
```

```
Block if :card_funding: = 'prepaid' and :amount_in_usd: > 1000.00
```

```
Block if :is_disposable_email: and :is_anonymous_ip:
```

```
Block if :card_country: in @my_card_country_blocklist and  
:ip_country: in @my_ip_country_blocklist;
```

Dashboard

Home Frontend **Backend** Pages Categories Tags Post types Search

iQ Block Country

Backend Options

Block visitors from visiting the backend (administrator) of your website:

Your IP address is [redacted]
Do NOT set the 'Block' You will NOT be able to [redacted]

Select the countries that should be blocked from visiting your backend:
Use the x behind the country to remove a country from this blacklist.

Inverse the selection above:
If you select this option only the countries that are selected are allowed.

Backend whitelist IP4 and/or IP6 addresses:
Use a semicolon (;) to separate IP addresses

Backend blacklist IP4 and/or IP6 addresses:
Use a semicolon (;) to separate IP addresses

Google **plumbers**

All Maps News Images Shopping More Settings Tools

About 697,000,000 results (0.53 seconds)

50+ plumbers nearby Sponsored

- Roto-Rooter Plumbing ...**
4.5 ★★★★★ · See reviews
GOOGLE GUARANTEED
Serves 98033
(206) 565-0312
Open 24/7
- Roto-Rooter Plumbing ...**
4.7 ★★★★★ · See reviews
GOOGLE GUARANTEED
Serves 98033
(425) 606-2932
Open 24/7
- Mr Rooter Plumbing of ...**
4.7 ★★★★★ · See reviews
GOOGLE GUARANTEED
Serves 98033
(206) 889-6788
Open 24/7


→ More plumbers in Kirkland

AAV MEMBERSHIP TRAVEL INSURANCE DISCOUNTS AUTOMOTIVE MAPS STORES

Locations Events

WASHINGTON

- Battle Ground (Insurance Only)
- Bellevue
- Bellingham
- Bothell (Insurance Only)
- Bremerton
- Issaquah
- Longview (Insurance Only)
- Lynnwood
- Seattle
- Seattle - Westlake (Insurance Only)
- Spokane
- Tacoma
- Tri-Cities
- Tukwila
- Vancouver
- Wenatchee



NETFLIX

Vezi ce urmează.

VIZIONEAZĂ ORIUNDE. ANULEAZĂ ORICÂND.

INCEARCĂ O LUNĂ GRATUIT!



2.

COMMERCIAL GEOLOCATION

Current commercial location services and their accuracy





- » **Akamai EdgeScape** (since 2000)
- » **Neustar IP Intelligence** (since 2000)
- » **IP2Location** (since 2001)
- » **IPelligence** (since 2006)
- » **MaxMind** (since 2002)
- » **Digital Element's NetAcuity** (since 1999)





» **Data and techniques are proprietary**

- » Direct partnerships with ISPs
- » Mobile phone traffic
- » Data from weather, news apps
- » Crowdsourcing
- » Network delay (ping)
- » Network topology (traceroute)
- » WHOIS databases
- » Reverse DNS hostnames



- » MaxMind and IP2Location reveal city-level accuracy

- » Numbers self-reported, not audited

- » **Maxmind**

- » Mean unweighted accuracy for 100 countries is 76.5% at 50 km

- » USA: accuracy of 86% at 50 km

- » **IP2Location**

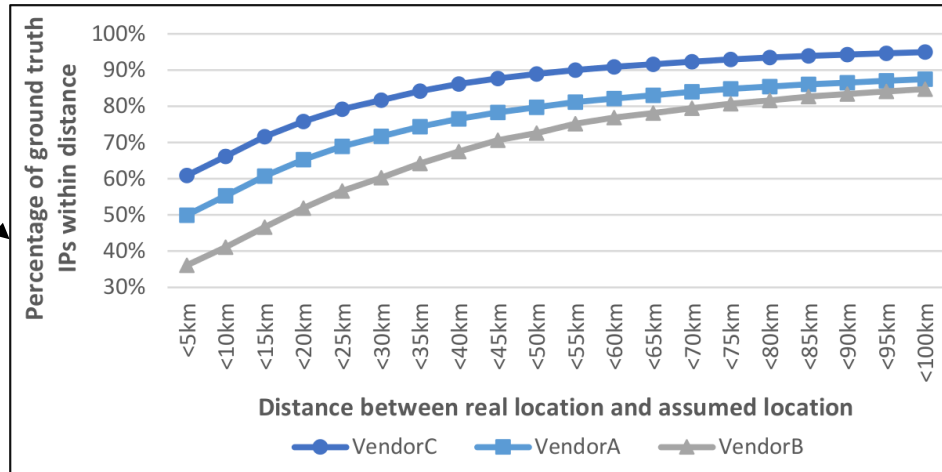

- » Mean unweighted accuracy 77% for error smaller than 50 kilometers

- » USA: accuracy of 99.52% at 50 km





- » Ground truth: 8.4 million IPs
- » Three commercial services
- » Exact city accuracy
- » Error distance in USA



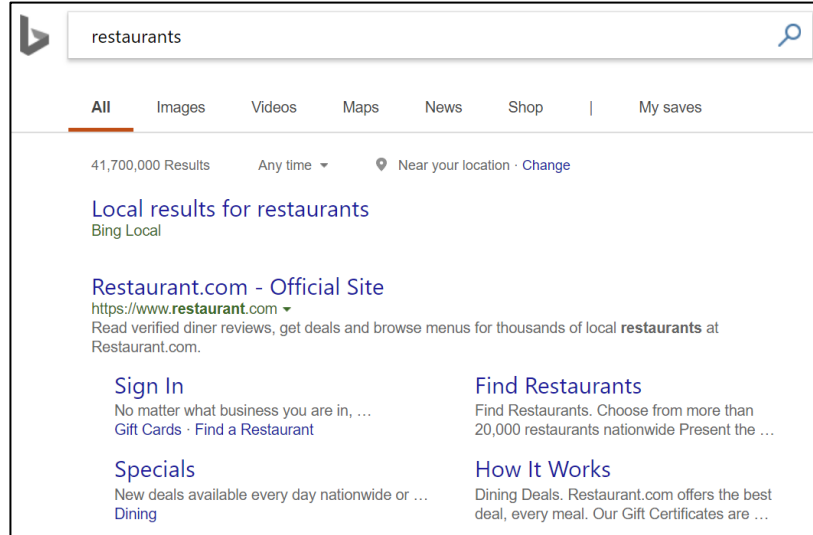
3.

MOTIVATION

Impact of inaccurate IP geolocation in the context of search engines



User Location: **Unknown**



restaurants

All Images Videos Maps News Shop | My saves

41,700,000 Results Any time Near your location · Change

Local results for restaurants
Bing Local

Restaurant.com - Official Site
<https://www.restaurant.com>
Read verified diner reviews, get deals and browse menus for thousands of local **restaurants** at Restaurant.com.

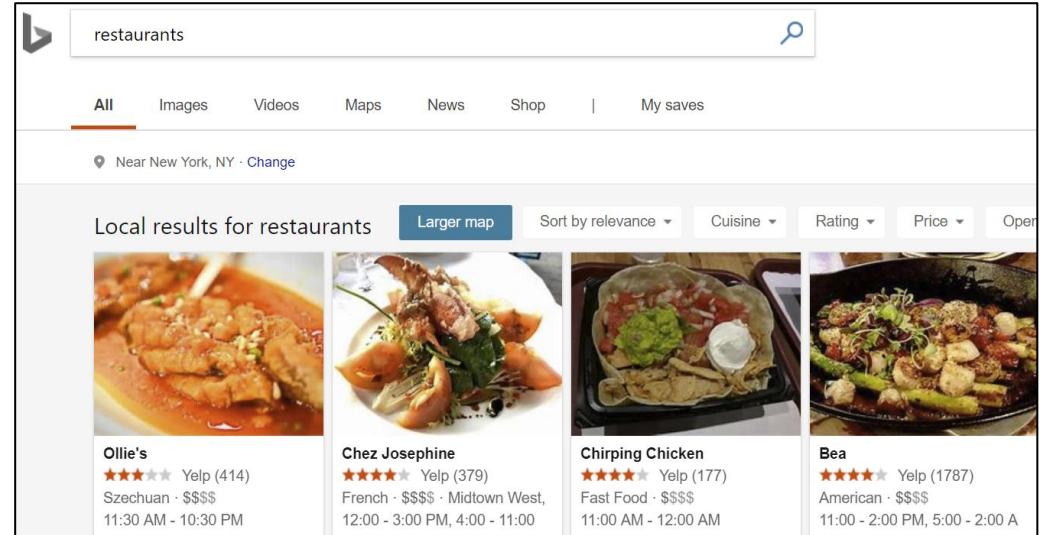
Sign In
No matter what business you are in, ...
[Gift Cards](#) · [Find a Restaurant](#)

Specials
New deals available every day nationwide or ...
[Dining](#)

Find Restaurants
Find Restaurants. Choose from more than 20,000 restaurants nationwide Present the ...

How It Works
Dining Deals. Restaurant.com offers the best deal, every meal. Our Gift Certificates are ...

User Location: **New York City**







restaurants

All Images Videos Maps News Shop | My saves

Near New York, NY · Change

Local results for restaurants [Larger map](#) Sort by relevance Cuisine Rating Price Open

			
Ollie's ★★★★★ Yelp (414) Szechuan · \$\$\$\$ 11:30 AM - 10:30 PM	Chez Josephine ★★★★★ Yelp (379) French · \$\$\$\$ · Midtown West, 12:00 - 3:00 PM, 4:00 - 11:00	Chirping Chicken ★★★★★ Yelp (177) Fast Food · \$\$\$\$ 11:00 AM - 12:00 AM	Bea ★★★★★ Yelp (1787) American · \$\$\$\$ 11:00 - 2:00 PM, 5:00 - 2:00 A

- » IP Geolocation is not a solved problem
- » Commercial geolocation services are proprietary black boxes
- » Previous research lacks accuracy and scale

Getting wrong GEO Location - How can I fix it?

- ▲ So I did multiple ip searches on multiple websites and the Geographical location is Sri Lanka. It shows IPS is DataShack, LC
- 1 The real location should be Kansas City, MO. USA My server's ip is 199.168.103.234 or play.fluixstudios.com
- ▼ We are renting the server from www.wholesaleinternet.net
- ★

Is it possible to correct my GeoIP location

- ▲ Numerous services and [advertising](#) do GeoIP lookups. They get the country correct, but often the city is wrong. Are there places where I can go and correct this, letting people know the actual location of my IP address. I have a static IP, so it's not going to go to another user in a different location.
- 4

services advertisements geolocation geoip

Why does my IP location put me in the wrong major city.

What Is My IP? » Questions » Why does my IP location put me in the wrong major city.



Daniel Barch asked 2 years ago

Netflix thinks I'm in a different country

If TV shows and movies are not available in your language, or if you are missing content that your local friends see, Netflix may think you are in a different country than where you are currently located.

📌 If you signed up for Netflix in a European Union member country but are currently watching Netflix in a different European Union member country, please see [Traveling or Relocating with Netflix](#) for more

When Google Gets Your Location Wrong!

Lauren
26 Sep 2017

Views: 18537

Recently, Google's desktop news began showing me the weather and local news for Detroit in the state of Michigan, rather than for my current location.

I switched my location but YouTube is still limited due to geo-restrictions. What can I do?



Melissa
1 year ago · Updated

Geo-restricted sites can use various data to recognize your real location such as IP tracking and if you've visited the site before, using cookies on your device. ZenMate allows you to bypass some of these methods by providing a different IP address in the corresponding country.

» Correct location vs incorrect location

- » Mined 7 days of search engine logs in US market from Oct 2014
- » Intersect with separate unused ground truth
- » Compare ground truth location with Bing IP geolocation
 - » Impressions with error $> 15\text{km}$ vs error $\leq 15\text{km}$

Metrics	Local Intent
Overall Click-Through rate (any link)	-4.3%
Algorithmic Click-Through rate (on search results)	-1.1%
Algorithmic Click Success	-6.1%
Ads Click-Through Rate	-17.9%
Ads Click Success	-15.2%
Ad Revenue	-40.3%



- » **Overall goal:** Create an IP geolocation database from scratch
- » Part of a larger research project
- » Approaches
 - » Reverse DNS
 - » Geographic Clicks
 - » GPS
 - » Web Index
 - » Query Logs
 - » Interpolation
 - » Traceroutes
 - » WHOIS
- » Some approaches used in production today in Bing

4.

PRIVACY AND DATASETS

Privacy-Conscious IP Geolocation





Your Apps Know Where You Were Last Night, and They're Not Keeping It Secret

Dozens of companies use smartphone locations to help advertisers and even hedge funds. They say it's anonymous, but the data shows how personal it is.

Apps sending location, secretly.

Posted by [Joel Reardon](#) on [May 14, 2018](#)



SAFE GRAPH



» Protecting User Privacy

- » Ground Truth Sets
- » Training data





- » IP Geolocation
 - » Does not track individual's user location
 - » Uses IP addresses and not accounts
 - » Is coarse, it maps IP ranges to cities
 - » Arguably, more privacy conscious than user-level GPS tracking



» 9 billion traceroutes

» Collected between January and November 2017

» From Archipelago (Ark) Measurement Infrastructure of CAIDA

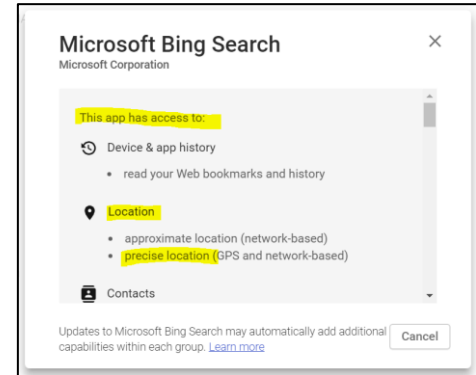
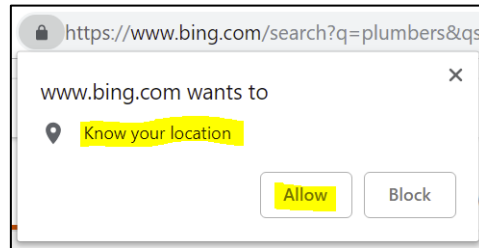
» 208 servers located in 63 countries

» Every 48 hours a random IP address is chosen in each /24 prefix, then the chosen IP addresses are individually probed by random Ark servers.



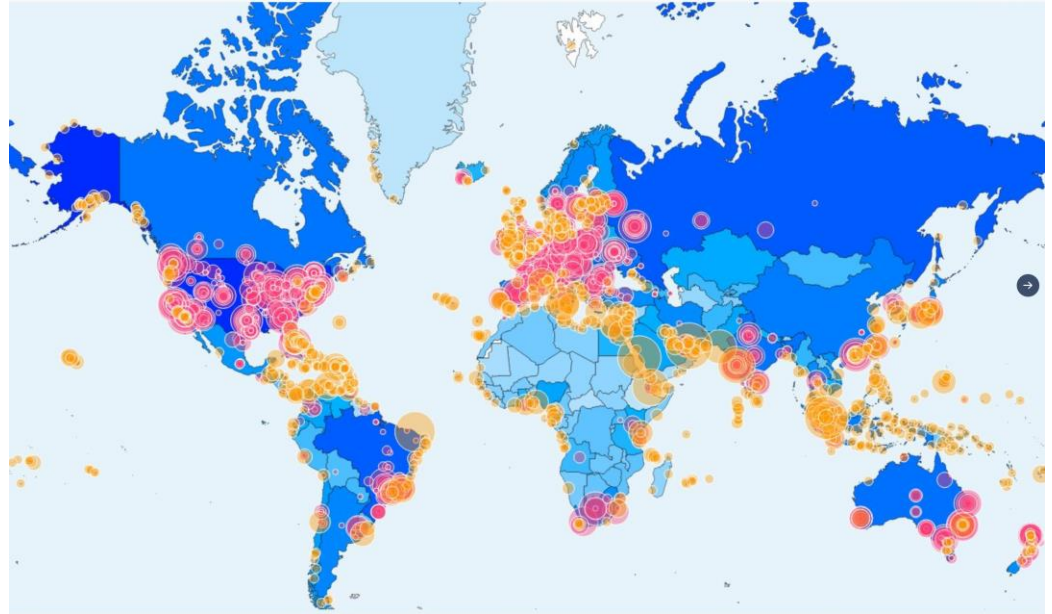
» 8.9 million IP addresses with known location

- » From Bing query logs: 28-day period ending on December 1st, 2017
- » Opt-in and anonymized data from browsers and apps
- » IPs from wired, wireless, and cell networks
- » From all countries in the world
- » Only used in training and testing, individual data points are not part of the final geolocation database





- » Smaller alternative to our proprietary ground truth set
- » 400 IP ranges spanning 128,000 IP addresses, along with geographic coordinates
- » Disadvantage: heavily infrastructure focused vs end users




5.

IP LOCATION INTERPOLATION

Extrapolating location of entire IP range from a few individual IPs

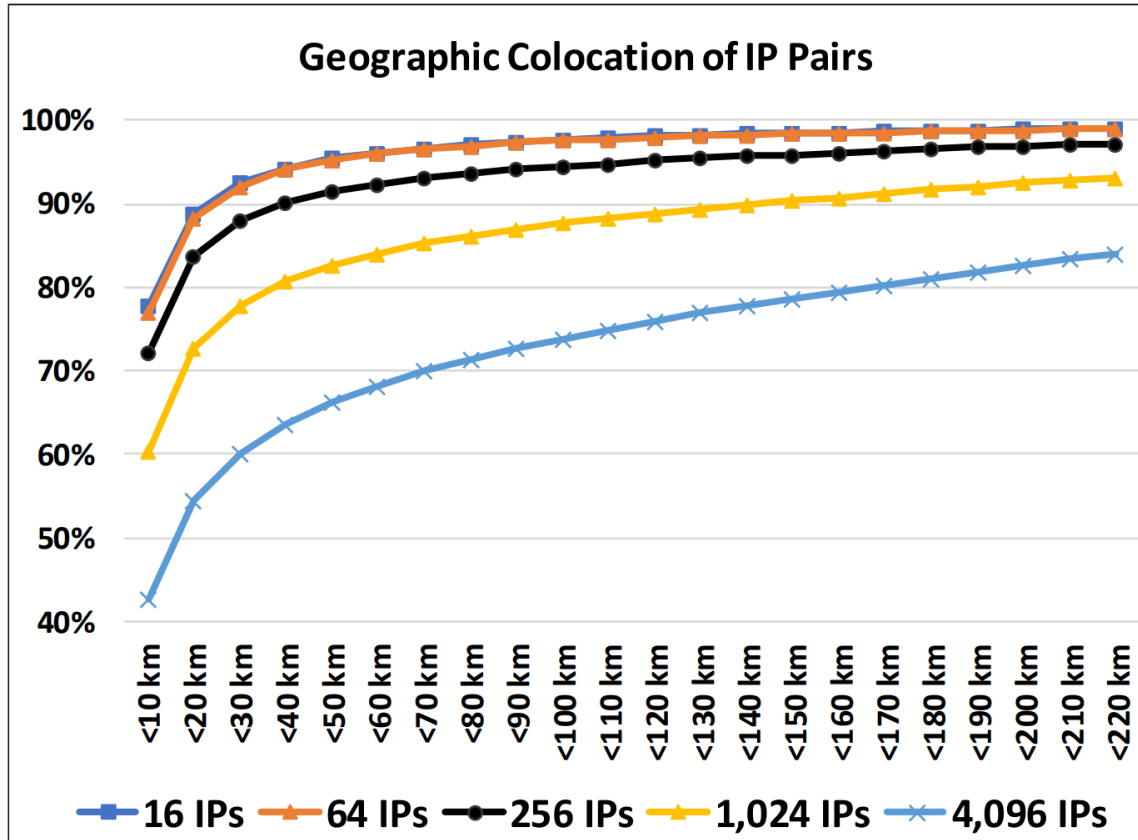


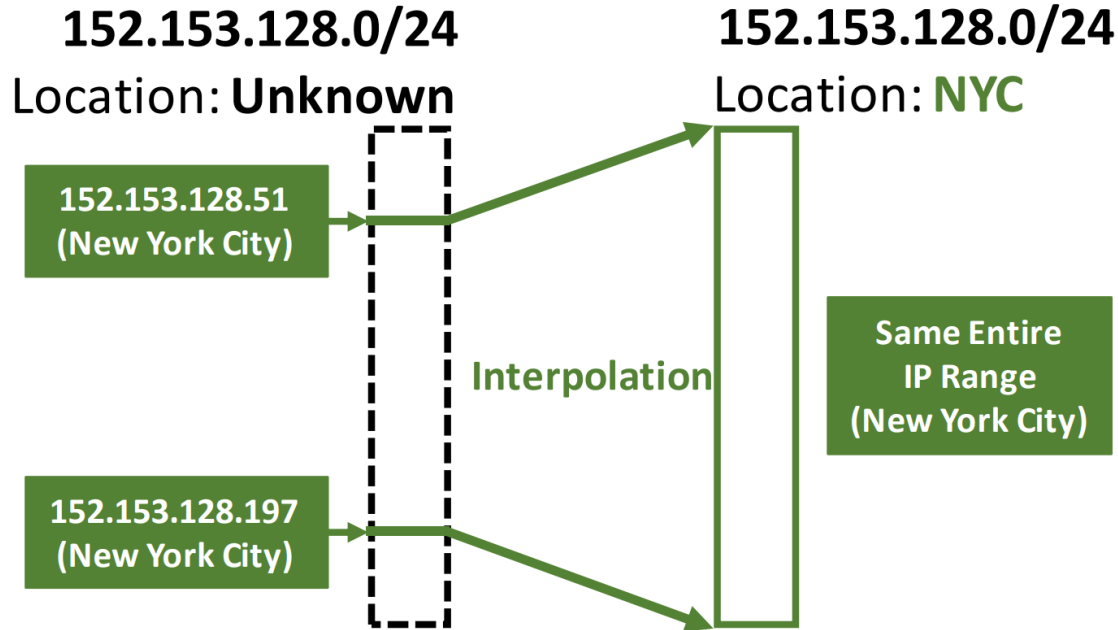
152.153.128.0/24



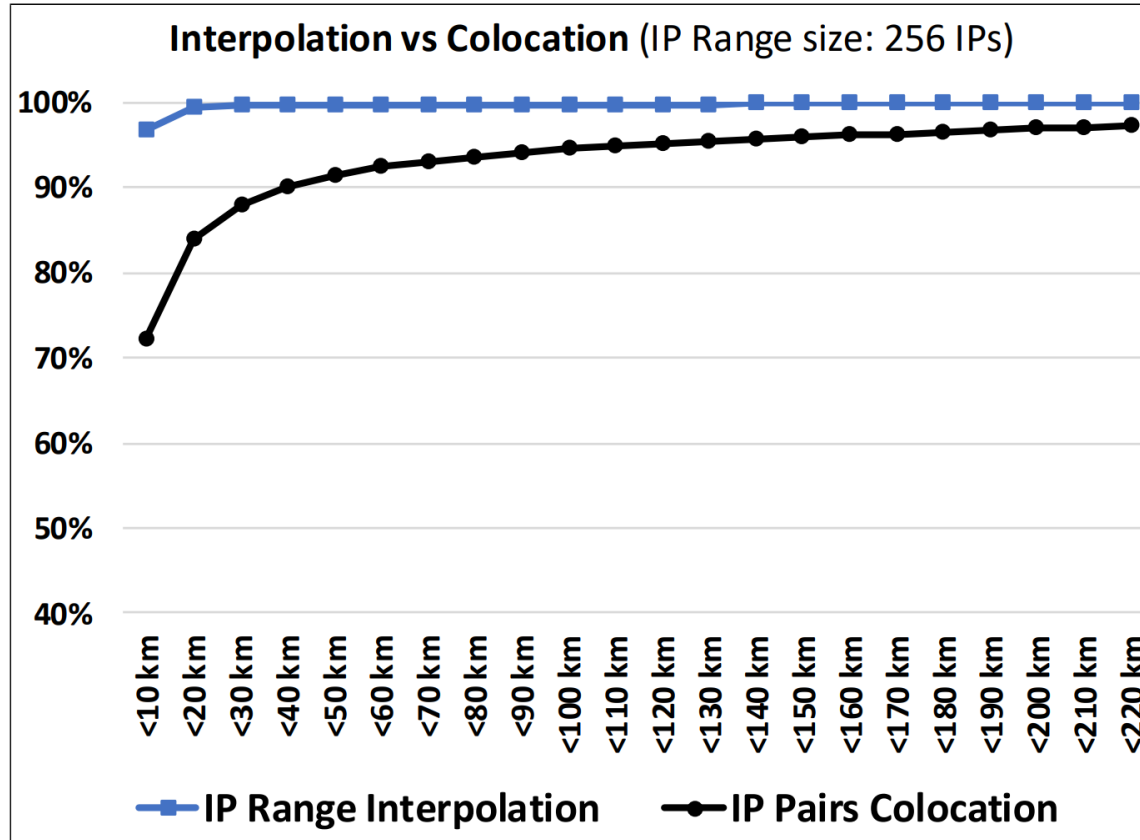
	Hosts	Netmask
/30	4	255.255.255.252
/29	8	255.255.255.248
/28	16	255.255.255.240
/27	32	255.255.255.224
/26	64	255.255.255.192
/25	128	255.255.255.128
/24	256	255.255.255.0
/23	512	255.255.254.0
/22	1024	255.255.252.0
/21	2048	255.255.248.0
/20	4096	255.255.240.0
/19	8192	255.255.224.0
/18	16384	255.255.192.0
/17	32768	255.255.128.0
/16	65536	255.255.0.0

- » Two IP addresses are collocated if they are in the same IP range.
- » IP ranges can have different sizes
- » In example above the 2 IPs are also geographically collocated





Example of IP range interpolation. Since IP range 152.153.128.0/24 contains two IP addresses with the same known location (coordinates in New York City), we propagate that location to the entire IP range.



6.

TRACEROUTE LOCATION PROPAGATION

Background, Motivation, Research Questions, Contributions





```
Command Prompt

C:\Users\zmart>ping www.lehigh.edu

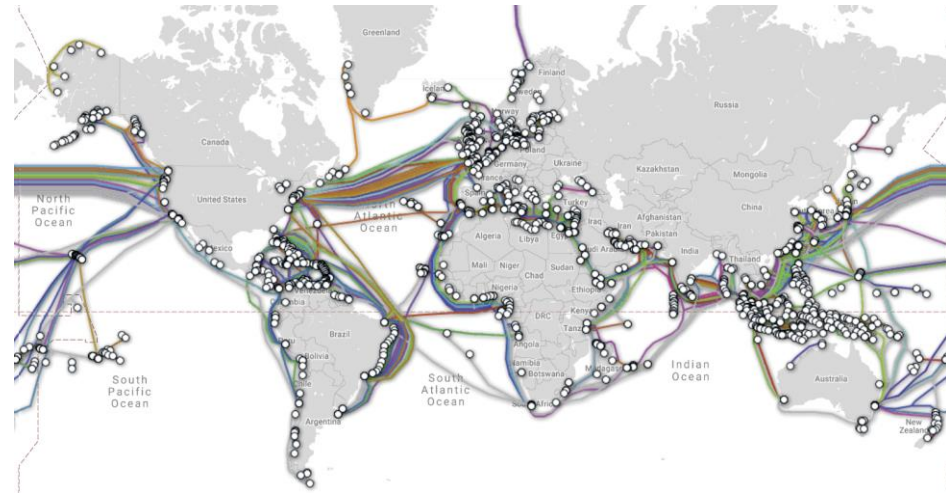
Pinging www.lehigh.edu [128.180.1.16] with 32 bytes of data:
Reply from 128.180.1.16: bytes=32 time=79ms TTL=50
Reply from 128.180.1.16: bytes=32 time=79ms TTL=50
Reply from 128.180.1.16: bytes=32 time=79ms TTL=50
Reply from 128.180.1.16: bytes=32 time=79ms TTL=50
```

```
Command Prompt

C:\Users\zmart>tracert www.lehigh.edu

Tracing route to www.lehigh.edu [128.180.1.16]
over a maximum of 30 hops:


  0  1 ms    1 ms    1 ms    Soul.ziggystardust.selfhost.corp.microsoft.com [10.0.10.2]
  1  <1 ms   <1 ms   <1 ms   10.0.10.1
  2  2 ms    1 ms    1 ms    50-46-216-1.evrt.wa.frontiernet.net [50.46.216.1]
  3  2 ms    2 ms    2 ms    172.76.21.177
  4  3 ms    3 ms    3 ms    ae3---0.cor02.sttl.wa.frontiernet.net [74.40.1.101]
  5  3 ms    3 ms    3 ms    ae1---0.cbr01.sttl.wa.frontiernet.net [74.40.5.126]
  6  3 ms    3 ms    3 ms    74.43.94.7
  7  77 ms   77 ms   78 ms   216.156.16.80.ptr.us.xo.net [216.156.16.80]
  8  75 ms   76 ms   82 ms   207.88.12.228.ptr.us.xo.net [207.88.12.228]
  9  76 ms   76 ms   75 ms   207.88.12.144.ptr.us.xo.net [207.88.12.144]
 10  78 ms   76 ms   77 ms   207.88.12.190.ptr.us.xo.net [207.88.12.190]
 11  77 ms   77 ms   78 ms   207.88.12.189.ptr.us.xo.net [207.88.12.189]
 12  78 ms   78 ms   77 ms   207.88.12.164.ptr.us.xo.net [207.88.12.164]
 13  77 ms   76 ms   76 ms   207.88.12.213.ptr.us.xo.net [207.88.12.213]
 14  80 ms   78 ms   78 ms   207.88.12.214.ptr.us.xo.net [207.88.12.214]
 15  78 ms   78 ms   77 ms   207.88.12.207.ptr.us.xo.net [207.88.12.207]
 16  79 ms   77 ms   78 ms   216.156.16.45.ptr.us.xo.net [216.156.16.45]
 17  79 ms   79 ms   79 ms   mail.tfains.net [209.118.23.114]
 18  80 ms   80 ms   79 ms   128.180.18.202
 19  79 ms   79 ms   79 ms   ws1iv.cc.lehigh.edu [128.180.1.16]
```



Speed of light in fiber

» 2/3 speed of light: 199 km/ms (124 miles/ms)

» 4/9 speed of light: 133 km/ms (83 miles/ms)






Combine traceroute with IP interpolation

- » Extract latency neighbors
- » Propagate locations through neighbors
- » Apply interpolation to increase coverage
- » Evaluate against ground truth

	Median error	% Err <10km	RMSE in km
Traceroute-GPS-HighAcc	<u>4.3 km</u>	<u>67.7%</u>	<u>329.3</u>
Traceroute-GPS-HighCov	10.1 km	50.5%	423.6
Traceroute-PeeringDB	8.4 km	61.1%	2124.9
Commercial Provider A	11.1 km	47.2%	545.9
Commercial Provider B	16.7 km	36.7%	545.3





- » Traceroute Dataset Parsing Library
 - » <https://github.com/zmarty/ScamperTracerouteParser>
 - » C# library that can parse the text traceroute dumps extracted from binary warts files using sc_analysis_dump (part of the CAIDA Scamper project).
- » PeeringDB parsing and generation library
 - » <https://github.com/zmarty/PeeringDBToTSV>
 - » Tiny C# library which reads the PeeringDB SQLite databases and can output the data in TSV format more suitable for IP geolocation research

The screenshot shows a GitHub commit page for the repository 'ScamperTracerouteParser / ScamperTracerouteParser'. The commit is by user 'zmarty' and is titled 'Add maxRTT'. The commit message is 'Add maxRTT'. The commit includes four files: 'ScamperTracerouteParser.Console' (Added support to detect latency neighbors), 'ScamperTracerouteParser.UnitTests' (Add support for maxHopDifference), 'ScamperTracerouteParser' (Add maxRTT), and 'ScamperTracerouteParser.sln' (Initial checkin).

File	Description
ScamperTracerouteParser.Console	Added support to detect latency neighbors
ScamperTracerouteParser.UnitTests	Add support for maxHopDifference
ScamperTracerouteParser	Add maxRTT
ScamperTracerouteParser.sln	Initial checkin

THANK YOU

Ovidiu Dan, Vaibhav Parikh, Brian. D. Davison



7.

RELATED WORK

Brief summary of two decades of related work in IP geolocation






- » **Network delay**
- » **Network Topology**
 - » Traceroute
 - » BGP / ASN
- » **Reverse DNS**
- » **IP Interpolation**
- » **WHOIS databases**
- » **Social networks**
- » **Web Mining**
- » **Crowdsourcing**

		Ping	Trace- route	Reverse DNS	Inter- polation	WHOIS	BGP AS	Crowd Social	Other
2000	NetGeo [112]			✓		✓			✓
2001	GeoTrack [69, 70]		✓	✓					
2001	GeoPing [69, 70]	✓							
2001	GeoCluster [69, 70]				✓		✓		
2002	undns [84]			✓					
2004	Ziviani [71]	✓							
2004	CBG [2, 40]	✓							
2004	PBE [82]	✓	✓						
2004	RBE [82]	✓	✓						
2006	TBG [83]		✓	✓					
2006	GeoBuD [89]	✓	✓						
2007	Octant [31]	✓	✓	✓		✓			
2009	SG [72]	✓							
2009	Structon [32]		✓		✓		✓		✓
2010	SocialGraph [101]							✓	
2010	WBG [113]		✓			✓			
2010	Geo-RhOL [88]	✓	✓						
2010	NB-LHP [90]	✓	✓						✓
2011	Spotter [30]	✓							
2011	WebCBG [100]	✓	✓						✓
2011	HawkEyes [118]	✓	✓						
2012	SDP [75]	✓							
2012	Posit [91, 92]	✓	✓						
2013	GeoGet [114]								✓
2013	Checkin-Geo [63]				✓			✓	✓
2013	PCFL [95]	✓	✓						✓
2013	RUEL [95]	✓	✓	✓					✓
2013	PathAudit [93]			✓					
2013	AdvancedGeo [122]		✓	✓		✓			✓
2014	DRoP [97]		✓	✓					
2015	Dragoon [119-121]	✓	✓						
2015	Alidade [33]	✓	✓	✓	✓	✓	✓		✓
2016	GeoSpeed [60]				✓			✓	
2016	AIG [116]	✓							
2016	Neural-RBF [76]	✓							
2017	HLOC [99]	✓		✓					
2018	CRLB [77]	✓							

		Ping	Trace-route	Reverse DNS	Interpolation	WHOIS	BGP AS	Crowd Social	Other
2000	NetGeo [112]			✓		✓			✓
2001	GeoTrack [69, 70]		✓	✓					
2001	GeoPing [69, 70]	✓							
2001	GeoCluster [69, 70]				✓		✓		
2002	undns [84]			✓					
2004	Ziviani [71]	✓							
2004	CBG [2, 40]	✓							
2004	PBE [82]	✓	✓						
2004	RBE [82]	✓	✓						
2006	TBG [83]		✓	✓					
2006	GeoBuD [89]	✓	✓						
2007	Octant [31]	✓	✓	✓		✓			
2009	SG [72]	✓							
2009	Structon [32]		✓		✓		✓		✓
2010	SocialGraph [101]							✓	
2010	WBG [113]		✓			✓			
2010	Geo-RhOL [88]	✓	✓						
2010	NB-LHP [90]	✓	✓						✓
2011	Spotter [30]	✓	✓						
2011	WebCBG [100]	✓	✓						✓
2011	HawkEyes [118]	✓	✓						
2012	SDP [75]	✓							
2012	Posit [91, 92]	✓	✓						
2013	GeoGet [114]								✓
2013	Checkin-Geo [63]				✓			✓	✓
2013	PCFL [95]	✓	✓						✓
2013	RUEL [95]	✓	✓						✓
2013	PathAudit [93]			✓					
2013	AdvancedGeo [122]		✓	✓		✓			✓
2014	DRoP [97]		✓	✓					
2015	Dragoon [119-121]	✓	✓						
2015	Alidade [33]	✓	✓	✓	✓	✓			✓
2016	GeoSpeed [60]				✓			✓	
2016	AIG [116]	✓							
2016	Neural-RBF [76]	✓							
2017	HLOC [99]	✓		✓					
2018	CRLB [77]	✓							



Network Delay and Network Topology

- » Ping, Traceroute, BGP, ASN
- » Active probes
- » Router table dumps
- » Most studied



GeoPing (2001)

- » Measure ping from each probe servers to all the others
- » For a new target IP, measure ping from all servers to it
- » Pick probe server vector with smallest latency
- » Assign probe server location to target IP

Active Probe Servers

Location: **Known**

	Known Location IP 1	Known Location IP 2	Known Location IP 3	Known Location IP N
Known Location IP 1	0ms	13ms	43ms	12ms
Known Location IP 2	11ms	0ms	9ms	207ms
Known Location IP 3	42ms	17ms	0ms	4ms
Known Location IP N	132ms	43ms	0.5ms	0ms

Pick vector w/
smallest distance



New target IP
Location: **Unknown**


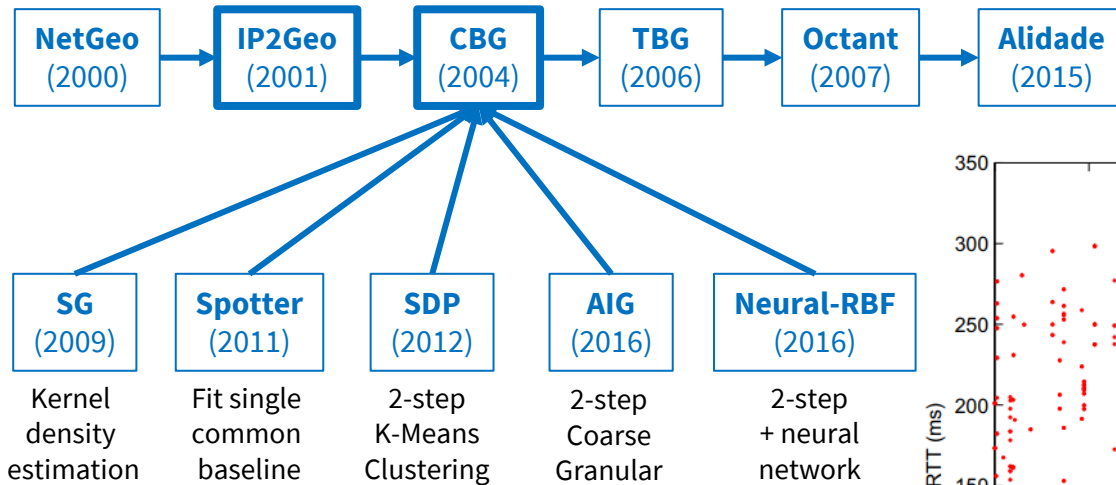
	IP 1	IP 2	IP 3	IP N
Unknown Location	10ms	0.1ms	9ms	204ms



CBG (2004)

- » Determine *bestline* for each probe server → delay-to-distance graph
- » Draw circles around each active probe server
- » Intersect circles and find center of intersection → *multilateration*
- » Assign center location to target





SDP (2012)

- » Cluster bestline data points
- » Only use probes in closest cluster

TBG (2006)

- » Use intermediate routers on the traceroute path to constrain distances from latency measurements
- » Treat geolocation as an optimization problem on a graph
 - » Vertices – active probe servers and targets
 - » Edges – traceroute paths with network delay





		Ping	Trace-route	Reverse DNS	Inter-polation	WHOIS	BGP AS	Crowd Social	Other
2000	NetGeo [112]			✓		✓			✓
2001	GeoTrack [69, 70]		✓	✓					
2001	GeoPing [69, 70]	✓							
2001	GeoCluster [69, 70]				✓		✓		
2002	undns [84]			✓					
2004	Ziviani [71]	✓							
2004	CBG [2, 40]	✓							
2004	PBE [82]	✓	✓						
2004	RBE [82]	✓	✓						
2006	TBG [83]		✓	✓					
2006	GeoBuD [89]	✓	✓						
2007	Octant [31]	✓	✓	✓		✓			
2009	SG [72]	✓							
2009	Structon [32]		✓		✓		✓		✓
2010	SocialGraph [101]							✓	
2010	WBG [113]		✓			✓			
2010	Geo-RhOL [88]	✓	✓						
2010	NB-LHP [90]	✓	✓						✓
2011	Spotter [30]	✓	✓						✓
2011	WebCBG [100]	✓	✓						✓
2011	HawkEyes [118]	✓	✓						✓
2012	SDP [75]	✓							✓
2012	Posit [91, 92]	✓	✓						✓
2013	GeoGet [114]								✓
2013	Checkin-Geo [63]				✓			✓	✓
2013	PCFL [95]	✓	✓						✓
2013	RUEL [95]	✓	✓						✓
2013	PathAudit [93]			✓					✓
2013	AdvancedGeo [122]		✓	✓		✓			✓
2014	DRoP [97]		✓	✓					✓
2015	Dragoon [119-121]	✓	✓						✓
2015	Alidade [33]	✓	✓	✓	✓	✓	✓		✓
2016	GeoSpeed [60]				✓			✓	
2016	AIG [116]	✓							
2016	Neural-RBF [76]	✓							
2017	HLOC [99]	✓		✓					
2018	CRLB [77]	✓							

hostname	ccr21.	par01.	atlas.	cogentco.com
position	2	1	0	
hint	ccr	par	atlas	
location			Salas Atlas, ES	
		Paris, FR		
	Concord, CA			

» undns (2002)

- » Manual rules
- » US centric

» DRoP (2014)

- » Automatic parsing rules
- » Rules validated with network delay

» DDec (2015)

- » Combines undns with DRoP
- » Gives precedence to undns

» HLOC (2017)

- » Use location hints directly

		Ping	Trace-route	Reverse DNS	Interpolation	WHOIS	BGP AS	Crowd Social	Other
2000	NetGeo [112]			✓		✓			✓
2001	GeoTrack [69, 70]		✓	✓					
2001	GeoPing [69, 70]	✓							
2001	GeoCluster [69, 70]				✓		✓		
2002	undns [84]			✓					
2004	Ziviani [71]	✓							
2004	CBG [2, 40]	✓							
2004	PBE [82]	✓	✓						
2004	RBE [82]	✓	✓						
2006	TBG [83]		✓	✓					
2006	GeoBuD [89]	✓	✓						
2007	Octant [31]	✓	✓	✓		✓			
2009	SG [72]	✓							
2009	Structon [32]		✓		✓		✓		✓
2010	SocialGraph [101]							✓	
2010	WBG [113]		✓			✓			
2010	Geo-RhOL [88]	✓	✓						
2010	NB-LHP [90]	✓	✓						✓
2011	Spotter [30]	✓							
2011	WebCBG [100]	✓	✓						✓
2011	HawkEyes [118]	✓	✓						
2012	SDP [75]	✓							
2012	Posit [91, 92]	✓	✓						
2013	GeoGet [114]								✓
2013	Checkin-Geo [63]				✓			✓	✓
2013	PCFL [95]	✓	✓						✓
2013	RUEL [95]	✓	✓						✓
2013	PathAudit [93]			✓					
2013	AdvancedGeo [122]		✓	✓		✓			✓
2014	DRoP [97]		✓	✓					
2015	Dragoon [119-121]	✓	✓						
2015	Alidade [33]	✓	✓	✓		✓			✓
2016	GeoSpeed [60]				✓			✓	
2016	AIG [116]	✓							
2016	Neural-RBF [76]	✓							
2017	HLOC [99]	✓		✓					
2018	CRLB [77]	✓							

» Mine websites

- » Chinese websites - Structon (2009)
- » Yellow Pages + CBG - WebCBG (2011)

» Mine social graphs

- » Extrapolate location from friends - SocialGraph (2010)
- » Locations from social checkins - Checkin-Geo (2014)

» WHOIS databases


- » Combine WHOIS with traceroutes - WBG (2010)

» HTTP Network Delay

- » Use GET as "ping" for multilateration - GeoGet (2013)

» Crowdsourcing

- » Crowdsourced Internet speed test - Speed (2016)



» General

- » Small ground truth set
- » Insufficient geographic diversity
- » Lack of networking environment diversity
- » Poor city accuracy
- » Non-existent or vague evaluation
- » Usage of commercial IP geolocation databases for training or testing

» Network Delay and Network Topology

- » Require access to geographically diverse infrastructure
- » Active probes, ICMP is often filtered out
- » Not scalable, as locating even a single IP takes a long time (billions of IPs)
- » Cramér–Rao lower bound theoretical limit – best-case error of 20 kilometers