

# Time's Forgotten: Using NTP to Understand Internet Latency

Ram Durairajan

University of Wisconsin-Madison

Sathiya K. Mani

University of Wisconsin-Madison

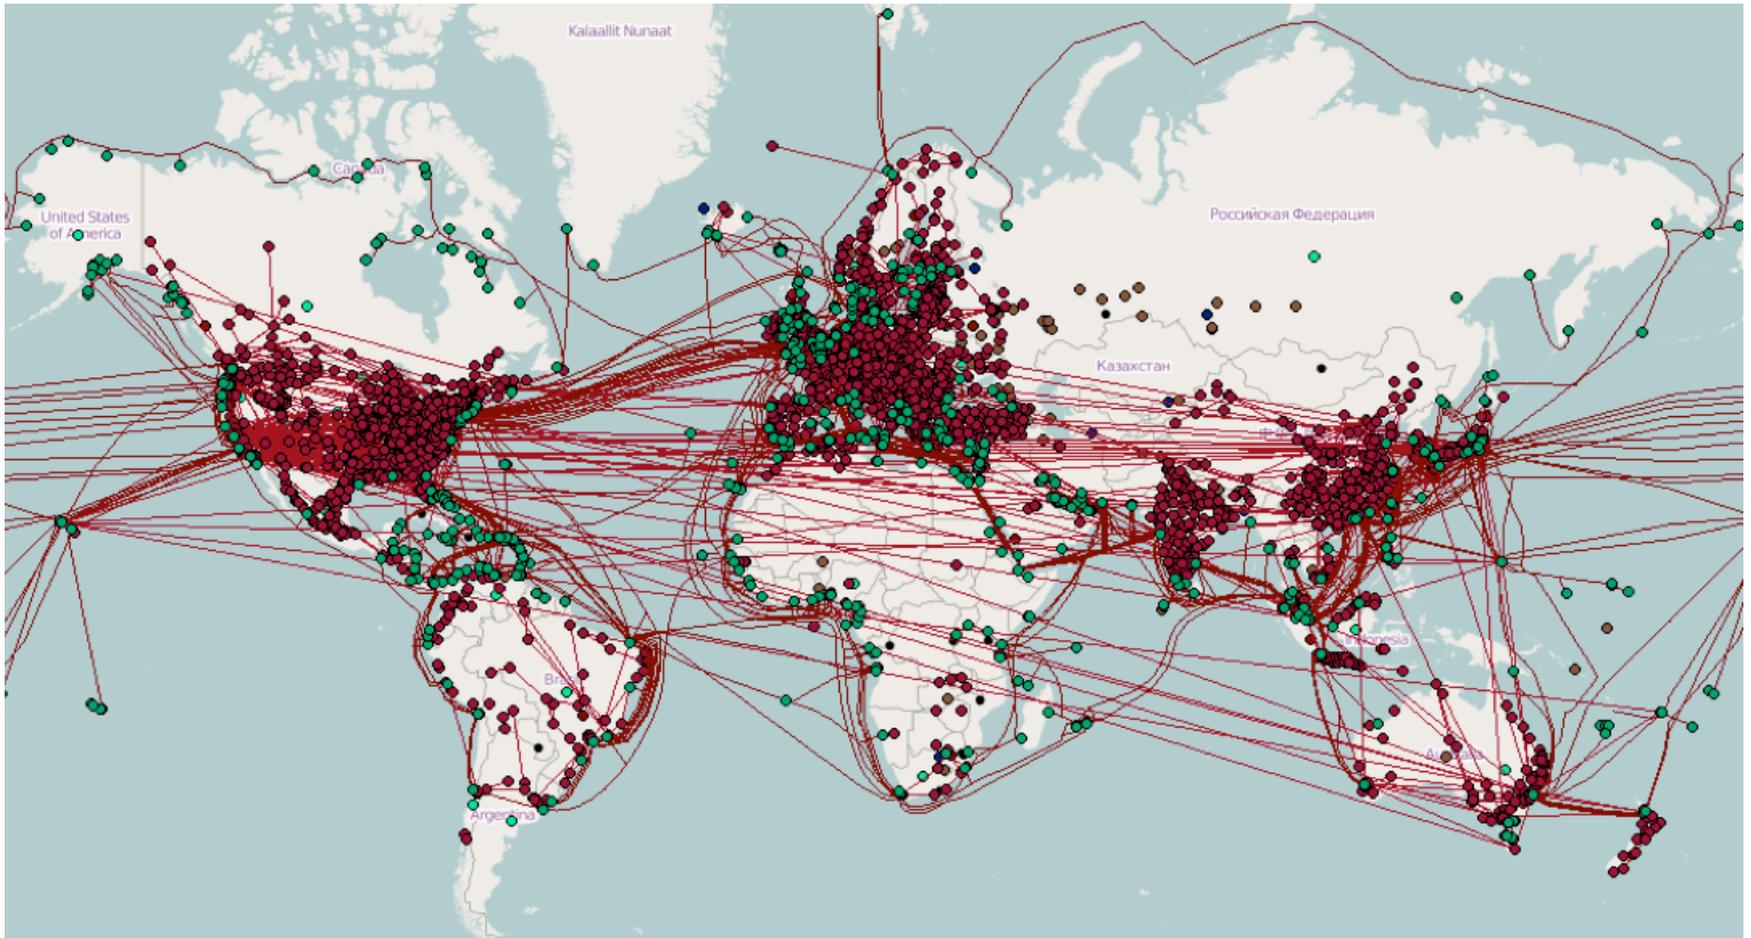
Joel Sommers

Colgate University

Paul Barford

University of Wisconsin-Madison &  
comScore, Inc.

# The Internet

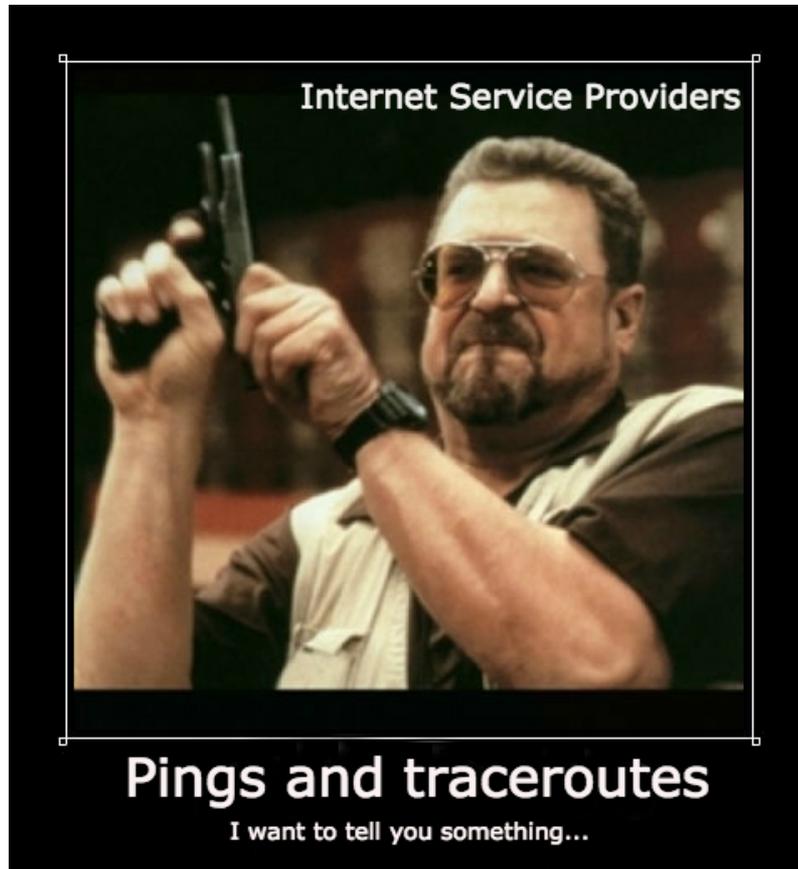


Broad understanding of latency is challenging due to Internet's scale and dynamics. [as  
.org](https://www.ietf.org)

# Active measurements to the rescue

- Understanding latency is (almost) always based on *ping* and/or *traceroute* measurements
- Other great problems
  - Outage quantification
  - SLA monitoring
  - Topology inference and modeling

# Problems with pings and traceroutes



Additional traffic

Coverage problem

Occasionally blocked

Management difficulties

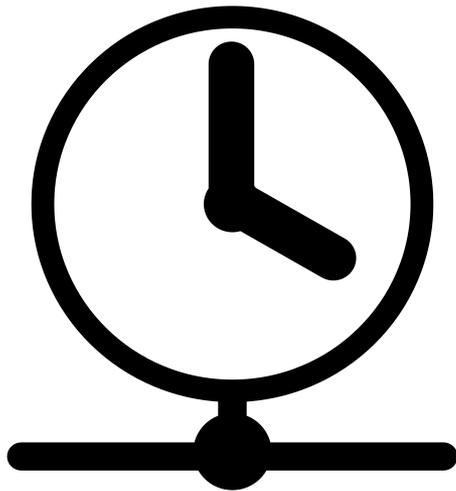
# Time's forgotten

- Why not use logs from NTP servers?

	Active Measurements	NTP	
Additional traffic	YES	NO	} We are relying on existing time sync. procedure.
Management difficulties	YES	NO	
Occasionally blocked	YES	NO	} No coordination & I/O blocking issues.
Coverage problems	YES	NO	} Not blocked.
			} Widely used in routers, DCs, desktops, etc.

# NTP Background

- Hierarchical organization of time sources
  - Stratum-1, stratum-2, etc.
- Clock discipline algorithm

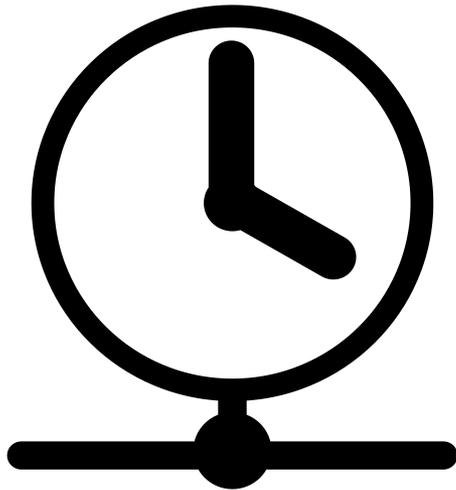


Rapid polling initially



# NTP Background

- Hierarchical organization of time sources
  - Stratum-1, stratum-2, etc.
- Clock discipline algorithm

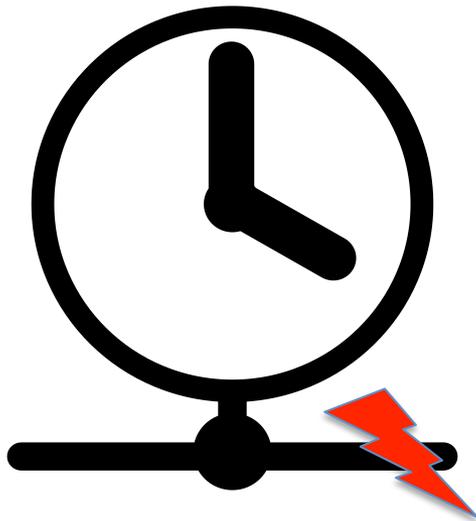


Decreased polling after  
synchronization



# NTP Background

- Hierarchical organization of time sources
  - Stratum-1, stratum-2, etc.
- Clock discipline algorithm



Rapid polling again



# NTP Background

- Hierarchical organization of time sources
  - Stratum-1, stratum-2, etc.
- Clock discipline algorithm
- Four timestamps are generated due to polling
  - Time when request is sent by the client
  - Time when request is received by the server
  - Time when response is sent by the server
  - Time when response is received by the receiver

# NTP Background

- Hierarchical organization of time sources
  - Stratum-1, stratum-2, etc.
- Clock discipline algorithm
- Four timestamps are generated due to polling
  - Time when request is sent by the client
  - Time when request is received by the server
  - Time when response is sent by the server
  - Time when response is received by the receiver

# NTP to the rescue

- Goal: Understand basic characteristics of Internet latency
- Analysis of logs from 10 NTP servers for a day

Server Location
Wisconsin
Utah
California

# Challenges in using latencies from NTP

- Invalid measurements
  - Malformed headers
  - Packet errors
  - Missing timestamps
  - Negative latency
- Client's synchronization stage with server?
  - Starting up? Fully synchronized?
  - No explicit information in the logs

# Filtering invalid measurements

- Two step filtering to remove bad latencies
  - Simple filtering to remove invalid packets
    - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!

# Filtering invalid measurements

- Two step filtering to remove bad latencies
  - Simple filtering to remove invalid packets
    - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!
  - Filtering by leveraging polling behaviors
    - Exhibited by the clock discipline algorithm

# Filtering invalid measurements

- Two step filtering to remove bad latencies
  - Simple filtering to remove invalid packets
    - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!
  - Filtering by leveraging polling behaviors
    - Exhibited by the clock discipline algorithm
      - Monotonically increasing polling values



# Filtering invalid measurements

- Two step filtering to remove bad latencies
  - Simple filtering to remove invalid packets
    - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!
  - Filtering by leveraging polling behaviors
    - Exhibited by the clock discipline algorithm
      - Monotonically increasing polling values
      - Monotonically decreasing polling values



# Filtering invalid measurements

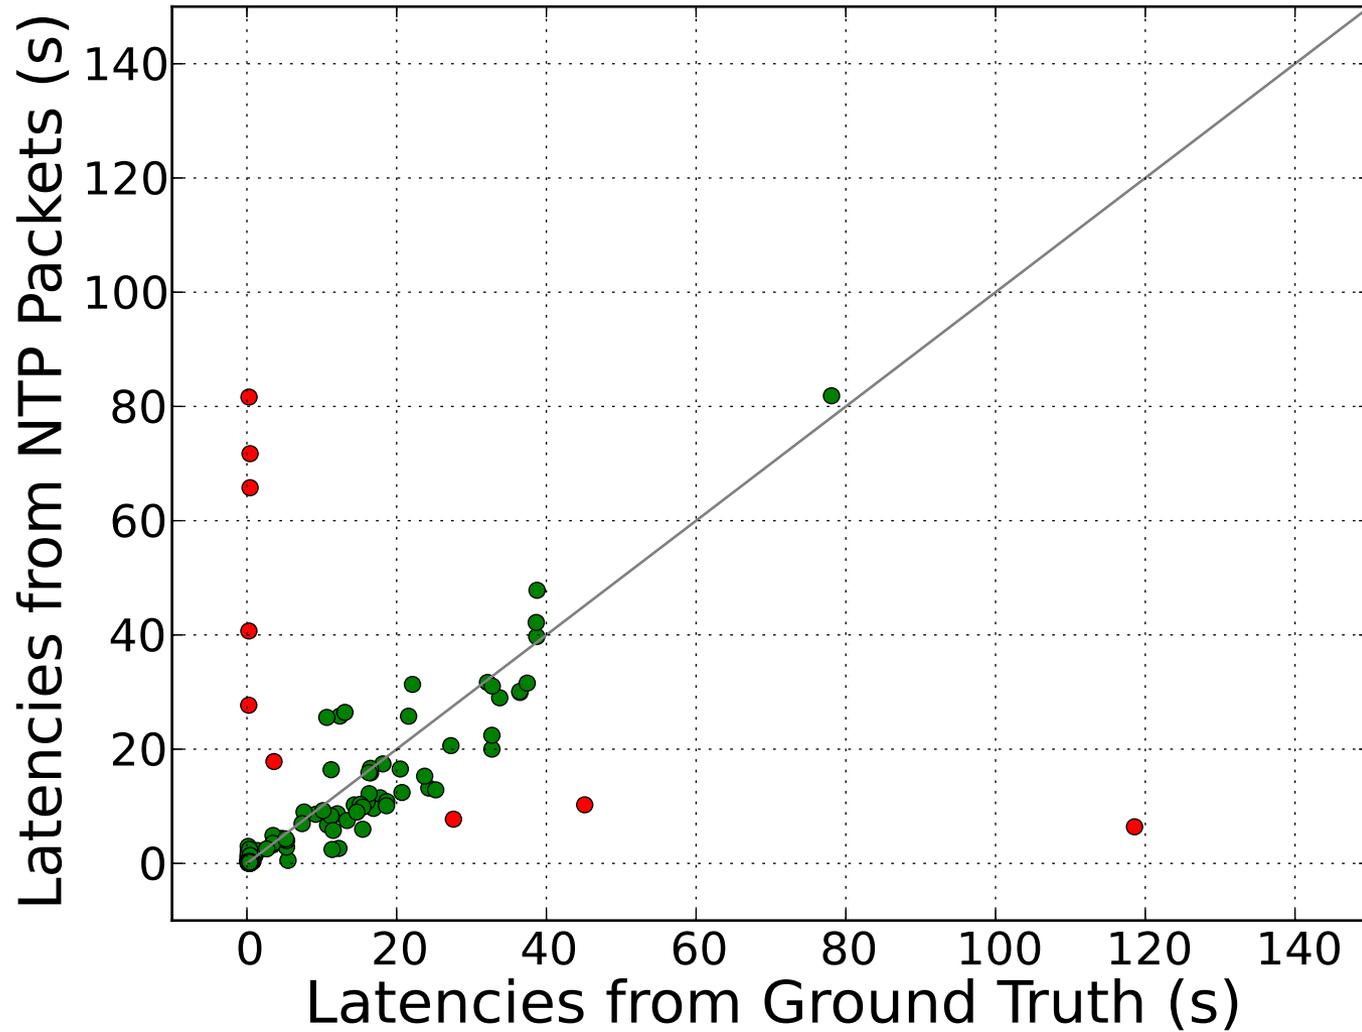
- Two step filtering to remove bad latencies
    - Simple filtering to remove invalid packets
      - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!
    - Filtering by leveraging polling behaviors
      - Exhibited by the clock discipline algorithm
        - Monotonically increasing polling values
        - Monotonically decreasing polling values
        - Constant polling values
-

# Filtering invalid measurements

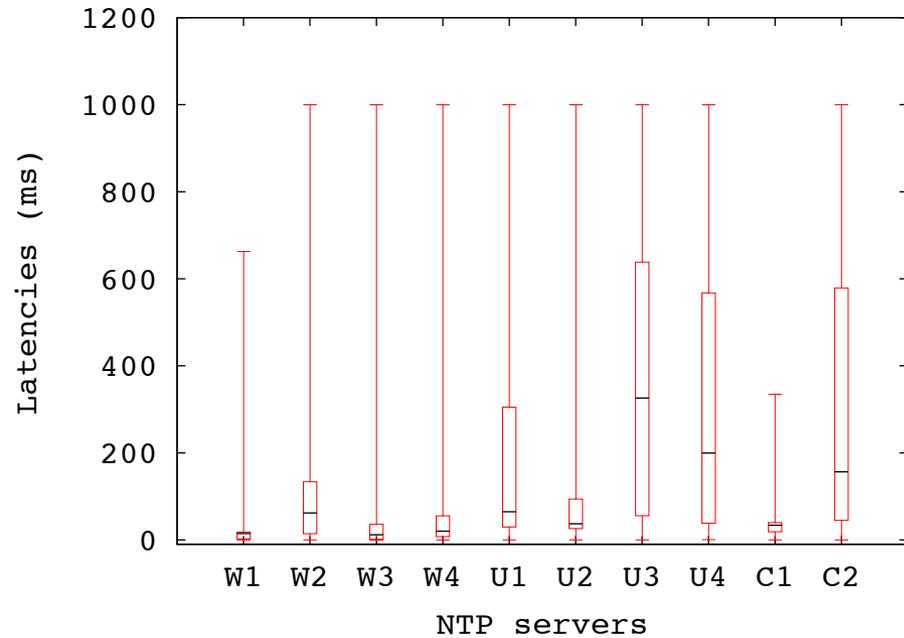
- Two step filtering to remove bad latencies
  - Simple filtering to remove invalid packets
    - 48.86 M (out of 73.83 M) packets filtered leaving us with about 25 M packets!
  - Filtering by leveraging polling behaviors
    - Exhibited by the clock discipline algorithm
      - Monotonically increasing polling values
      - Monotonically decreasing polling values
      - Constant polling values
      - Varying (non-monotonic) polling values



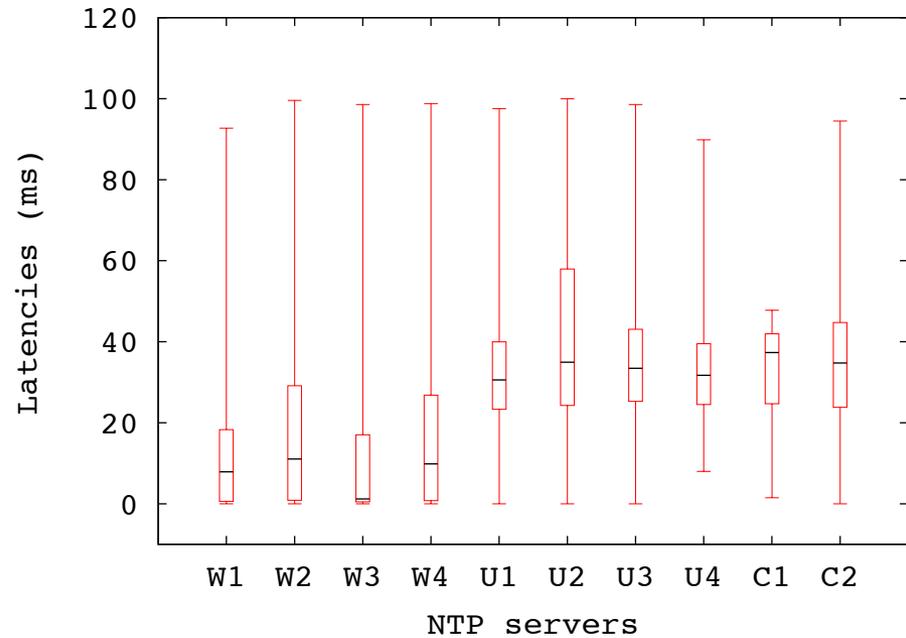
# Filtering results



# Latency characteristics



Clients distributed worldwide



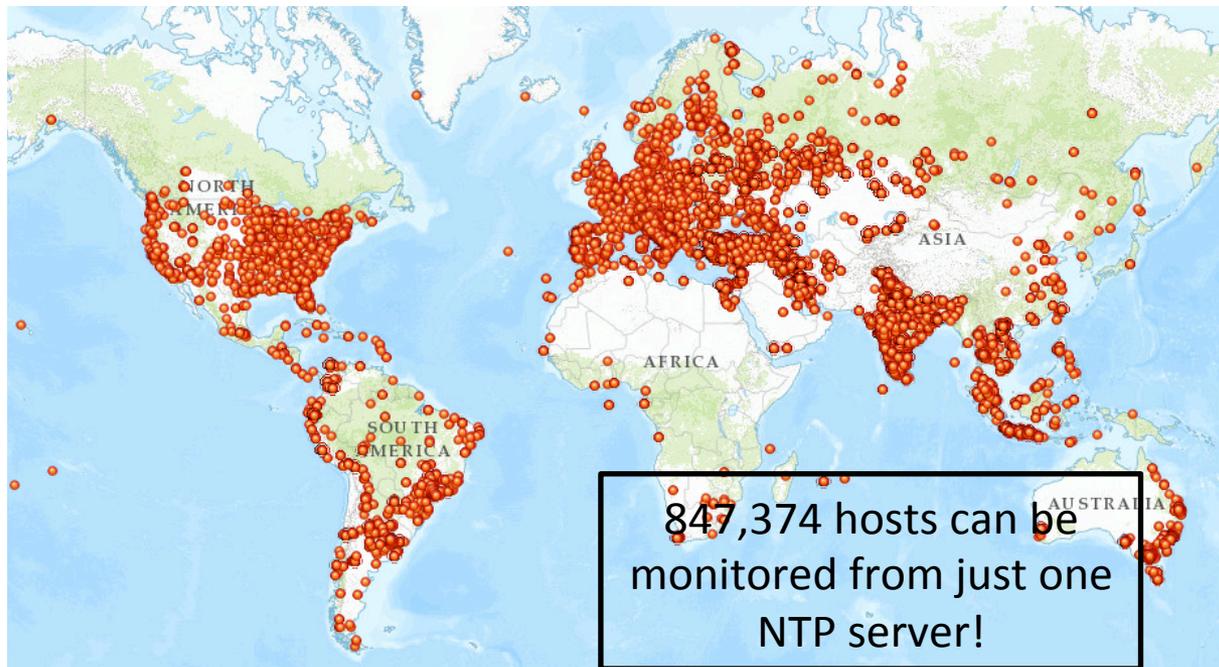
Clients distributed in the US

# Latency characteristics

	1999	2015
Client latencies	90% of clients had latencies < 100ms	99% of clients had latencies < 100ms
Bottlenecks	Stratum-1 servers were bottlenecked	Stratum-1 servers are not bottlenecked anymore!

# Future work

- Opens up many new opportunities
  - Internet monitoring without traceroutes/pings



Distribution of clients talking to *only one* stratum-2 NTP server at UW-Madison

# Future work

- Opens up many new opportunities
  - Internet monitoring without traceroutes/pings
  - Replicate previous efforts by leveraging NTP logs
    - E.g., can we find outage characteristics without Thunderpings?

# Thank you!

## Acknowledgements

- Murray Anderegg, James Babb, Tim Czerwonka, John Ricketts, and Aaron Topence for providing NTP server logs.
- Judah Levine and Dave Plonka for all the helpful discussions.



# Backup

# Research Question

- Can we understand latency without pings and traceroutes?
  - Can we also extend coverage?
  - Can we also improve accuracy?