

# Cheleby: An Internet Topology Mapping System



Hakan Kardeş

Talha Öz, David Shelly,  
and Mehmet H. Güneş

**ISMA 2011 AIMS-3**  
**Workshop on Active Internet Measurements**

La Jolla, CA


10 Feb 2011



# Outline

- Introduction
- Issues & Related Work
- Cheleby
- Experimental Results
- Conclusions and Future Work

- Web of interconnected networks
  - Grows with no central authority
  - Autonomous Systems optimize local communication efficiency
  - The building blocks are engineered and studied in depth
  - Global entity has not been well characterized



- Understand topological and functional characteristics of the Internet
  - Essential to design, implement, protect, and operate underlying network technologies, protocols, services, and applications
- Topology measurement studies require representative Internet maps
  - Such maps are not publicly available
- Direct and Indirect Probing



# Roadmap

- Introduction
- **Related Work**
- Cheleby
- Experimental Results
- Conclusions and Future Work



# Topology Collection Systems

- **Ark**
  - 53 Monitors around the world
  - Traces every /24 subnet
  - AS/Router Level
- **Dimes**
  - 20K Monitors (home users)
  - annotates the links with delay and loss statistics
  - PoP/AS Level
- **iPlane**
  - 200 PlanetLab nodes, 1000 destination
  - PoP/AS Level



# Roadmap

- Introduction
- Related Work
- **Cheleby**
- Experimental Results
- Conclusions and Future Work

- Ultimate goal of *Cheleby* is to generate Internet topology maps at varying levels
  - Topology information collection
  - Topology construction
  - Visualization



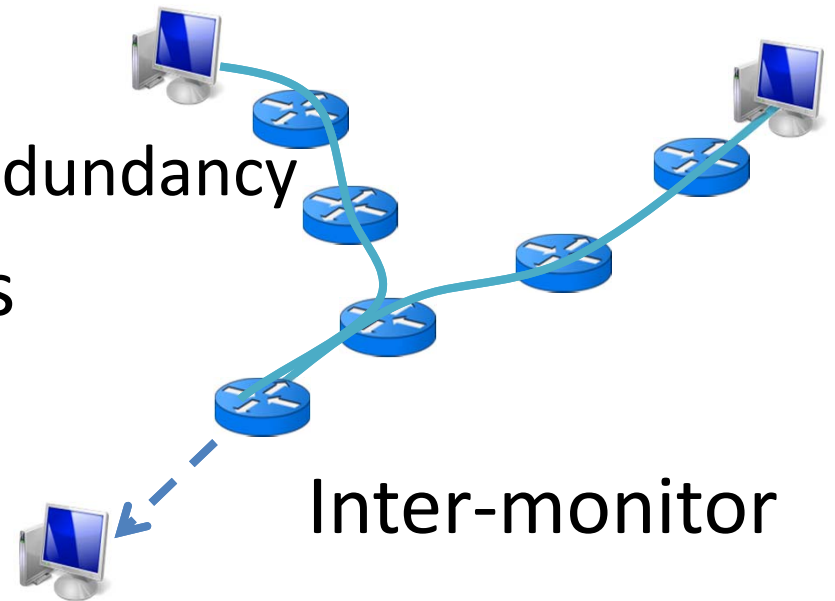
# N

# Challenges

- Infrastructural Issues
- Sampling
  - Vantage Points and Destination List
- Probing Overhead
  - Inter- and Intra-monitor Redundancy
- Responsiveness of Routers
  - ICMP, UDP, TCP
- Load Balancing Routers
  - Per destination, per flow, per packet



Intra-monitor



Inter-monitor



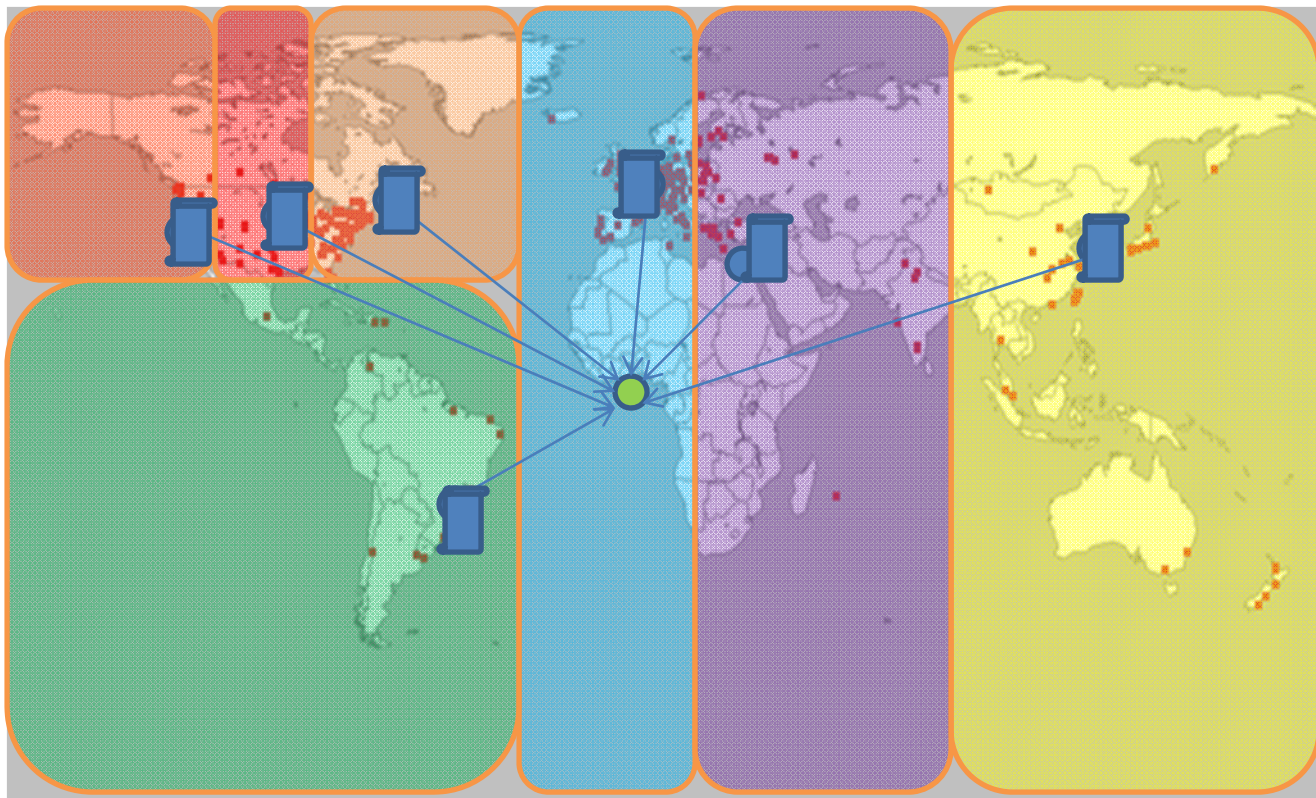
## Infrastructural Issues & VP sampling

- Current working (CoMon) + Paris Traceroutable
  - ~550 nodes can run Paris Traceroute
- 4 processes at a monitor
- Team Formation (inter-monitor redundancy)
- Task Assignment



# Inter-monitor Redundancy

- Team Formation (location info from PLCAPI)



- Generating initial set
  - A list of route advertisements (/24, A.B.C.1)
  - Each file has  $1200 \pm 200$  IP addresses
- Update after each run
  - Replace non-observed IPs with observed
  - Add /30 and /31 mates of observed IPs
  - To increase subnet completeness rDNS and Ping
    - Set increased from  $\sim 3.0$  M to  $\sim 3.8$  M IP addresses




# New Destination List

| 10% Subnet<br>(missing IPs) | rDNS    | Ping    | Total found | Not found |
|-----------------------------|---------|---------|-------------|-----------|
| 651,800                     | 273,244 | 229,497 | 368,935     | 282,865   |

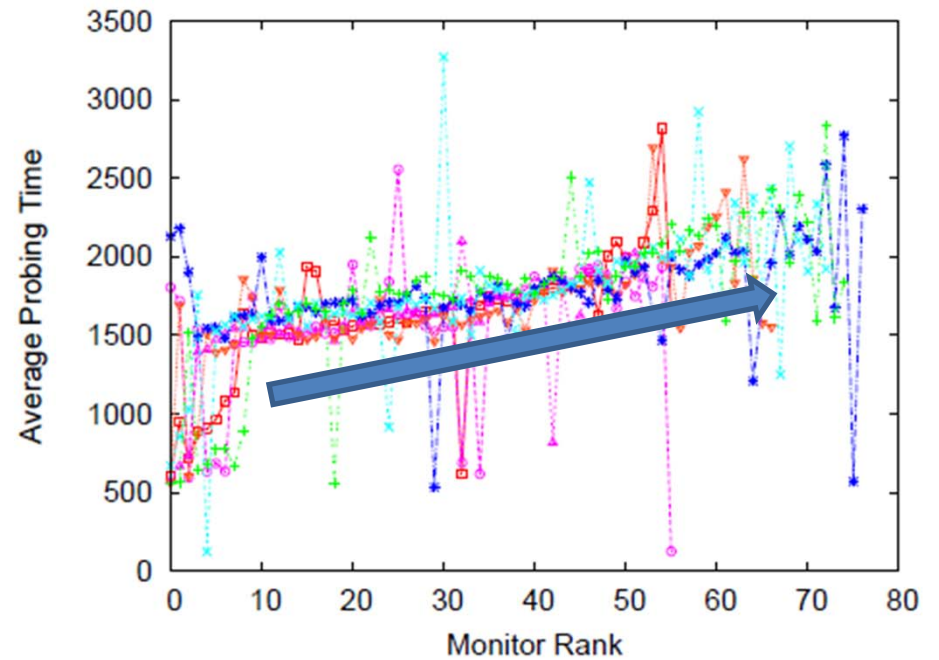
| Initial IPs | Total Observed<br>(with mates) | Newly<br>Observed | /30 mate | /31 mate | New Total<br>IPs |
|-------------|--------------------------------|-------------------|----------|----------|------------------|
| 3,09 M      | 1,259,298                      | 630,987           | 535,185  | 93,126   | 3,813,256        |



# Task Assignment



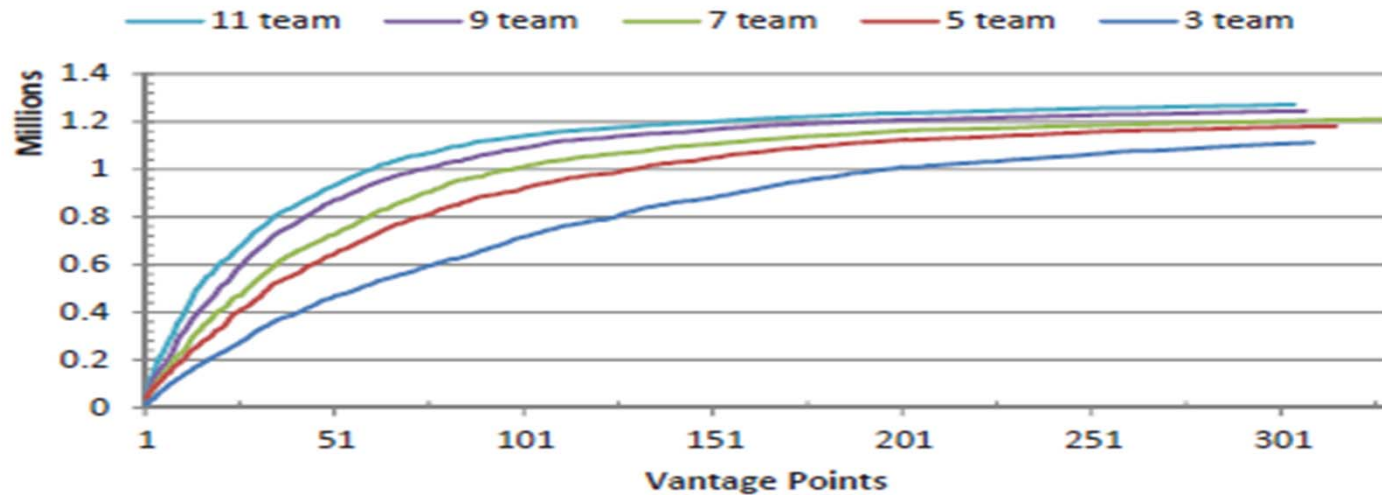
a) Average Number of Probes




b) Average Completion Time (sec)



# Number of Teams



a) Number of Known Nodes



b) Number of Edges




# Number of Teams

| Num of Teams | Time (min) | Observed IPs | Percentage | Difference | Per min IPs | Per min Time diff | Num Traces |
|--------------|------------|--------------|------------|------------|-------------|-------------------|------------|
| 3            | 540        | 1,110,647    | 79.26%     |            | 2,056.75    |                   | 9,000,000  |
| 5            | 630        | 1,180,886    | 84.28%     | 70,239     | 1,874.42    | 182.33            | 15,000,000 |
| 7            | 770        | 1,209,289    | 86.30%     | 28,403     | 1,570.50    | 303.91            | 21,000,000 |
| 9            | 1220       | 1,244,742    | 88.83%     | 35,453     | 1,020.28    | 550.22            | 27,000,000 |
| 11           | 1540       | 1,271,041    | 90.71%     | 26,299     | 825.35      | 194.92            | 33,000,000 |
| All          |            | 1,401,197    |            |            |             |                   |            |





# Round Trip Time Experiment





# Team Statistics

|  | Team 1   | Team 2   | Team 3   | Team 4   | Team 5   | Team 6   | Team 7   |
|--|----------|----------|----------|----------|----------|----------|----------|
| <b>Monitors</b>                                  | 56.63    | 53.88    | 55.5     | 56.75    | 77.25    | 73.63    | 76.75    |
| <b>Incomplete</b>                                | 7.43     | 30.28    | 24.03    | 35.72    | 12.85    | 12.35    | 12.15    |
| <b>Completed total</b>                           | 3,453    | 3,430    | 3,436    | 3,424    | 3,447    | 3,448    | 3,448    |
| <b>Completed total %</b>                         | 99.80%   | 99.13%   | 99.31%   | 98.96%   | 99.62%   | 99.65%   | 99.65%   |
| <b>1<sup>st</sup> trial completed</b>            | 3,436.38 | 3,366.63 | 3,395.38 | 3,363.88 | 3,420.25 | 3,424.25 | 3,421.75 |
| <b>1<sup>st</sup> trial completed %</b>          | 99.32%   | 97.30%   | 98.13%   | 97.22%   | 98.85%   | 98.97%   | 98.89%   |
| <b>2<sup>nd</sup> trial completed</b>            | 16.2     | 63.1     | 40.6     | 60.4     | 26.9     | 23.4     | 26.1     |
| <b>2<sup>nd</sup> trial completed %</b>          | 68.59%   | 67.58%   | 62.83%   | 62.84%   | 67.67%   | 65.45%   | 68.24%   |
| <b>Destination probing time (average in sec)</b> | 1,463.46 | 1,324.97 | 1,546.25 | 1,592.54 | 1,744.10 | 1,744.67 | 1,546.60 |

(Average of 8 rounds)



# Intra-monitor Redundancy

- Partial-traceroute
  - Keep pairs
    - IP observed right before the destination AS
    - Its hop distance
  - Start probing with the minimum observed hop distance
  - Check if the new pair was observed already

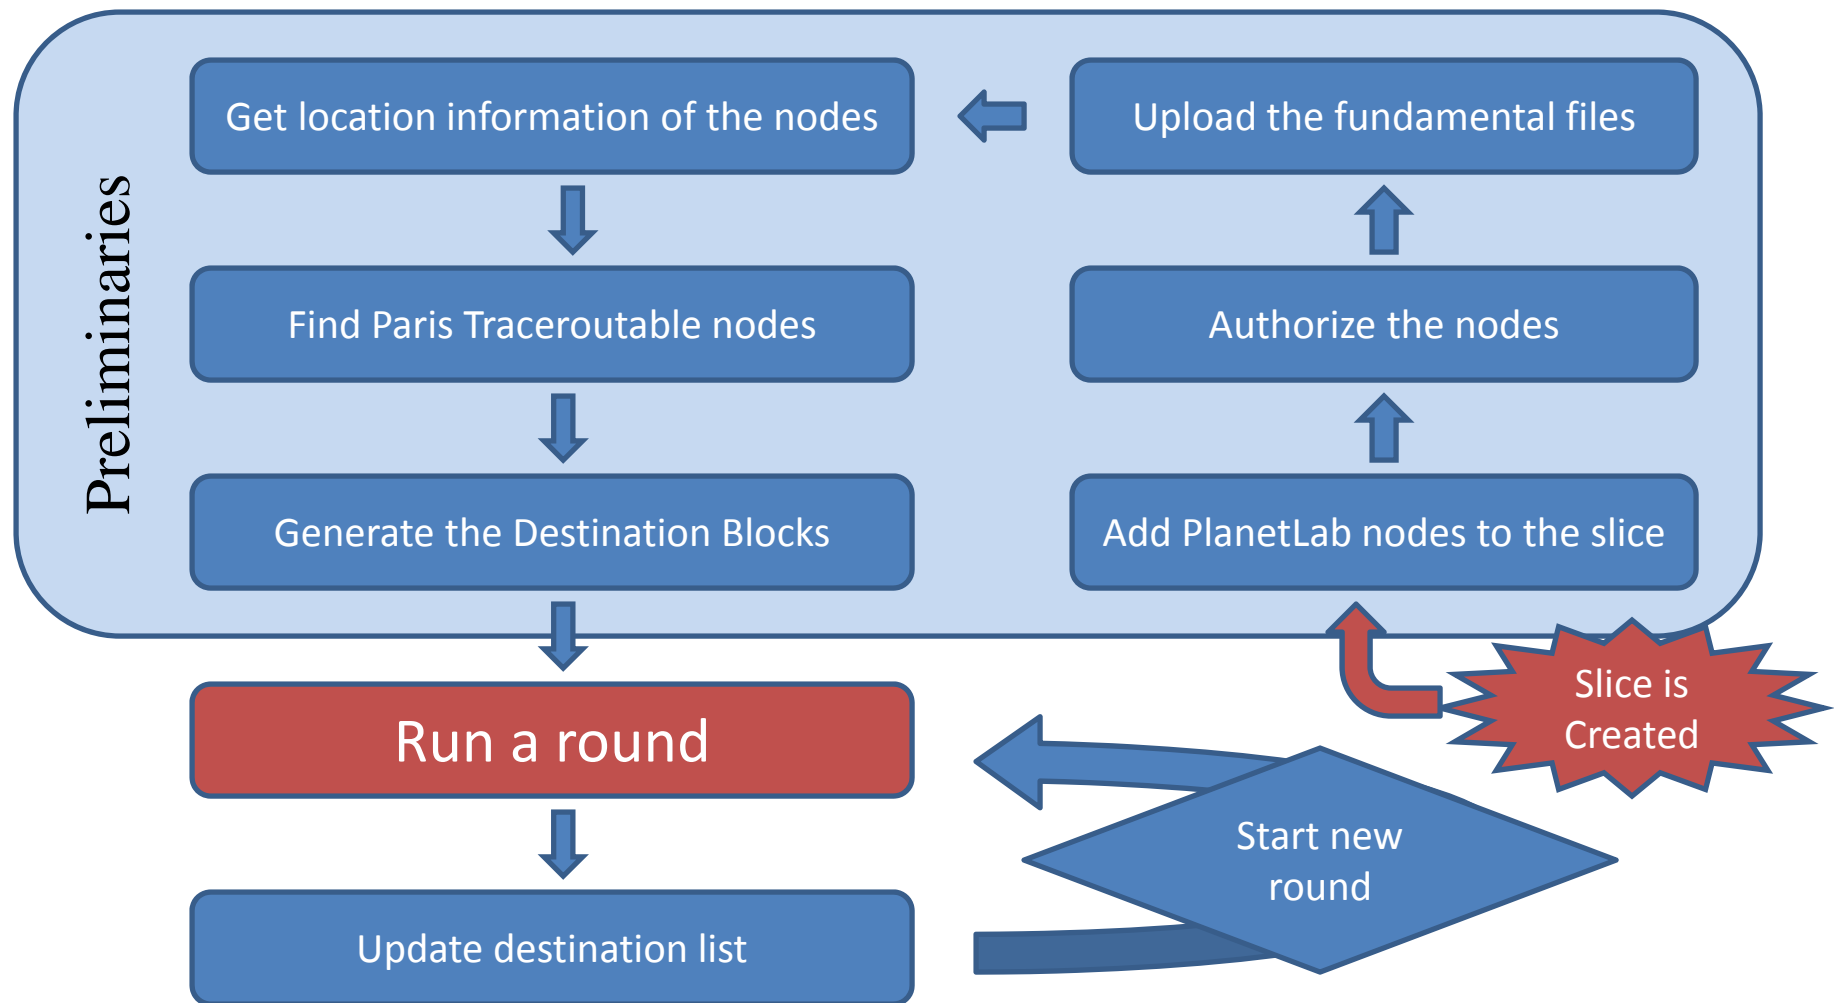


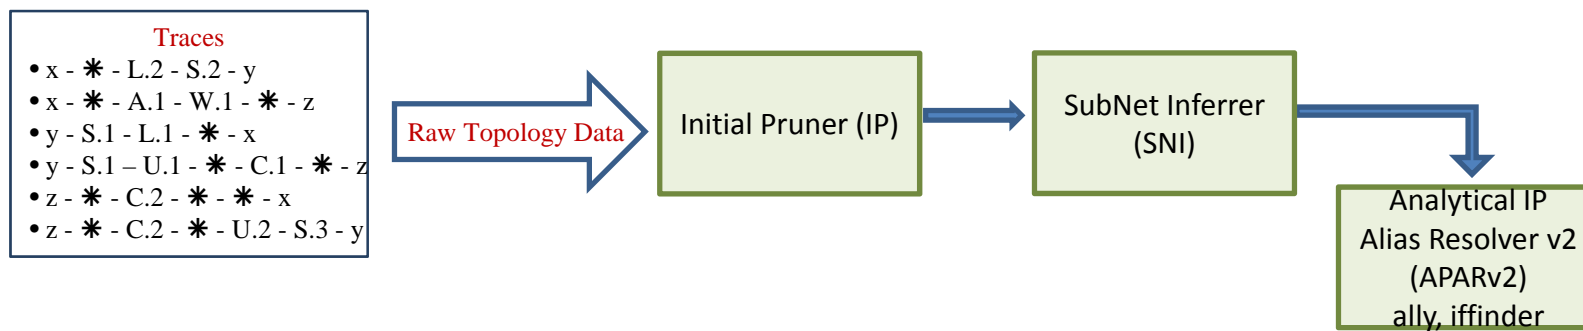
# Partial Traceroute Gain

| Experiment with 7 teams | Traces (M) | Partial (%) | Saved probes (M) |
|-------------------------|------------|-------------|------------------|
| Round 1                 | 22.39      | 35.8        | 65.23            |
| Round 2                 | 22.42      | 36.2        | 66.14            |
| Round 3                 | 22.42      | 35.9        | 67.68            |
| Round 4                 | 22.40      | 35.1        | 66.32            |
| Round 5                 | 22.42      | 34.2        | 63.98            |
| Round 6                 | 22.41      | 35.6        | 67.90            |
| Round 7                 | 22.42      | 35.3        | 66.19            |
| Round 8                 | 22.03      | 35.4        | 65.98            |



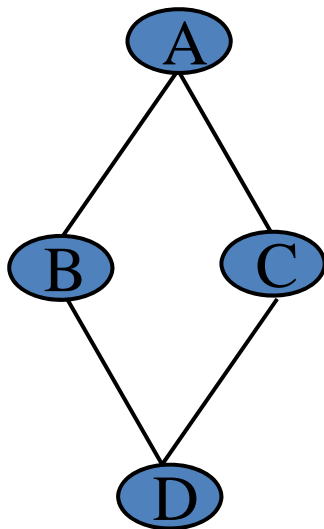


# Data Collection System

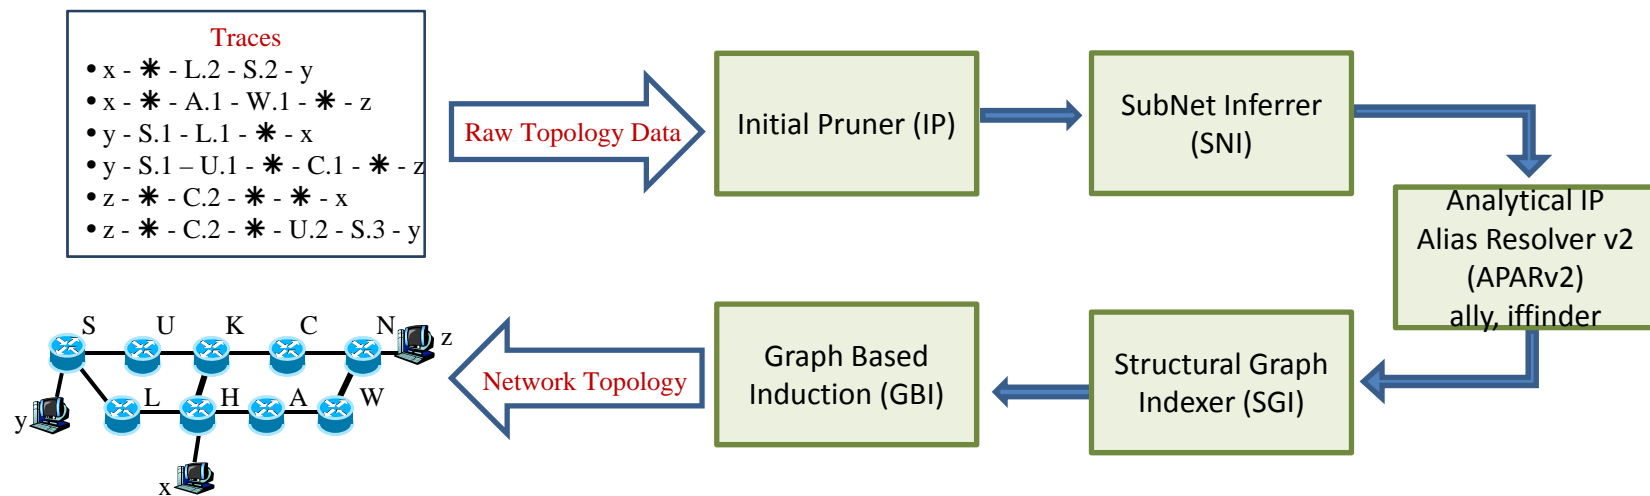




•  
•  
•



•  
•  
•







# Roadmap

- Introduction
- Related Work
- Cheleby
- **Experimental Results**
- Conclusions and Future Work



# Subnet Resolution

| Subnet Size  | \24   | \25   | \26   | \27    | \28    | \29     | \30    | \31    |
|--------------|-------|-------|-------|--------|--------|---------|--------|--------|
| Count        | 4     | 36    | 184   | 1,294  | 8,836  | 93,110  | 20,543 | 37,468 |
| Completeness | 26.3% | 30.0% | 28.3% | 27.7%  | 28.0%  | 39.9%   | 100%   | 100%   |
| All IPs      | 268   | 1,359 | 3,228 | 10,767 | 34,587 | 219,745 | 41,086 | 74,936 |



# IP Alias Resolution

| Resolver               | Alias Sets | Aliased IPs |
|------------------------|------------|-------------|
| APARv2                 | 38,012     | 128,495     |
| Ally (path traces)     | 32,860     | 65,720      |
| Ally (common neighbor) | 32,595     | 65,190      |
| Ally (Subnet)          | 25,436     | 50,872      |
| Ally (Combined)        | 55,027     | 110,054     |
| iffinder               | 305        | 610         |
| Combined               | 82,962     | 216,628     |



# Unresponsive Router Resolution

| Initial   | I. Pruner | Rate Lim. | Triangle | Bipartite | Star    | Final *s |
|-----------|-----------|-----------|----------|-----------|---------|----------|
| 7,207,885 | 6,137,750 | 51,279    | 2,858    | 143,880   | 619,204 | 252,915  |



# Roadmap

- Introduction
- Related Work
- Cheleby
- Experimental Results
- Conclusions and Future Work

- We presented an Internet topology constructor system
  - takes raw Internet topology data and by using efficient algorithms
  - produces router level and link-level (i.e. subnet) maps which are ready to visualize
  - Cheleby: An Internet Topology Mapping System  
<http://cheleby.cse.unr.edu>



# Visualization

- Demo

# Questions & Comments





**Thank you**

A dark blue square containing a white, stylized, blocky letter 'N' with a double outline.

**N**