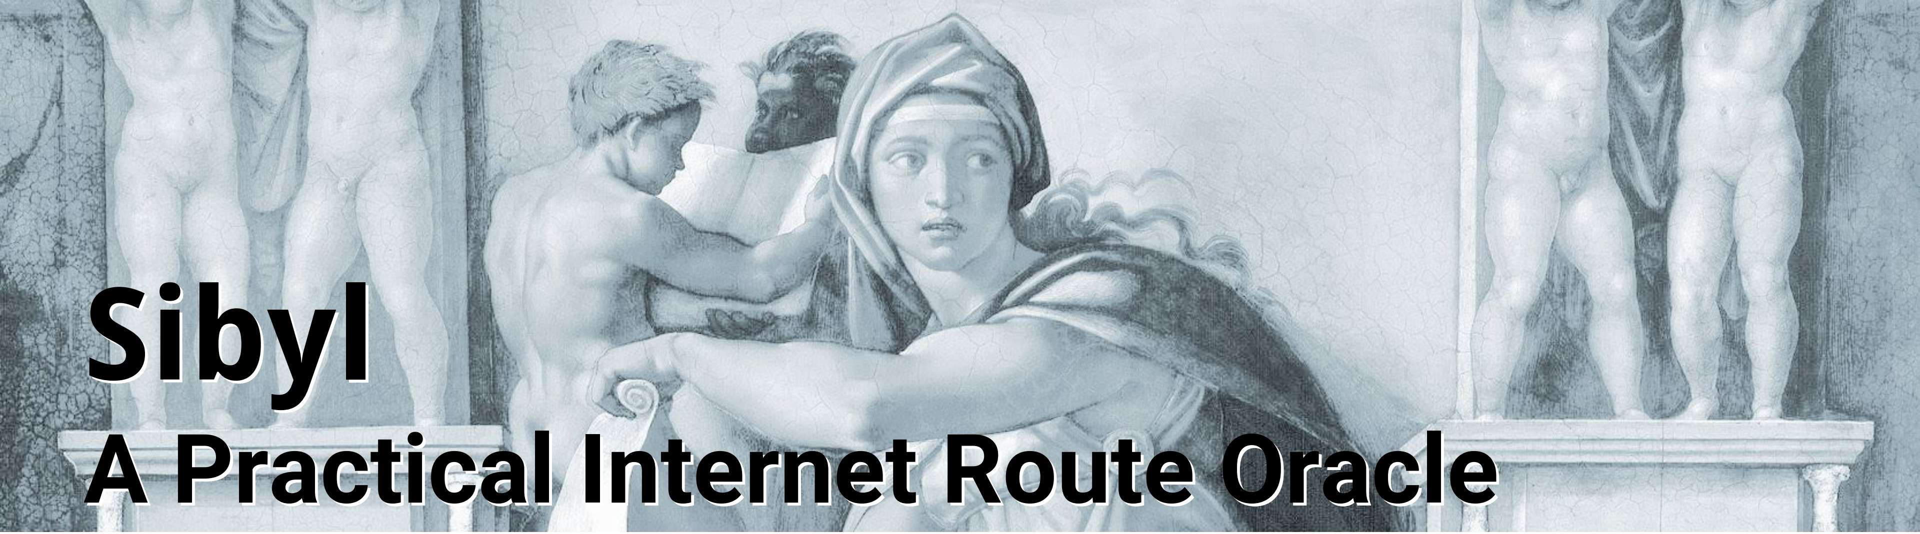

Sibyl
A Practical Internet Route Oracle

Ítalo Cunha[1], Pietro Marchetta[2], Matt Calder[3], Yi-Ching Chiu[3]
Brandon Schlinker[3], Bruno Machado[1], Antonio Pescapè[2]
Vasileios Giotsas[4], Harsha Madhyastha[5], Ethan Katz-Bassett[3]

UFMG
UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II
USC University of Southern California
caida
UNIVERSITY OF MICHIGAN

The Art of Network Troubleshooting

1. Operator has routing problem
2. Runs traceroute to the destination
3. Needs to identify where the problem is

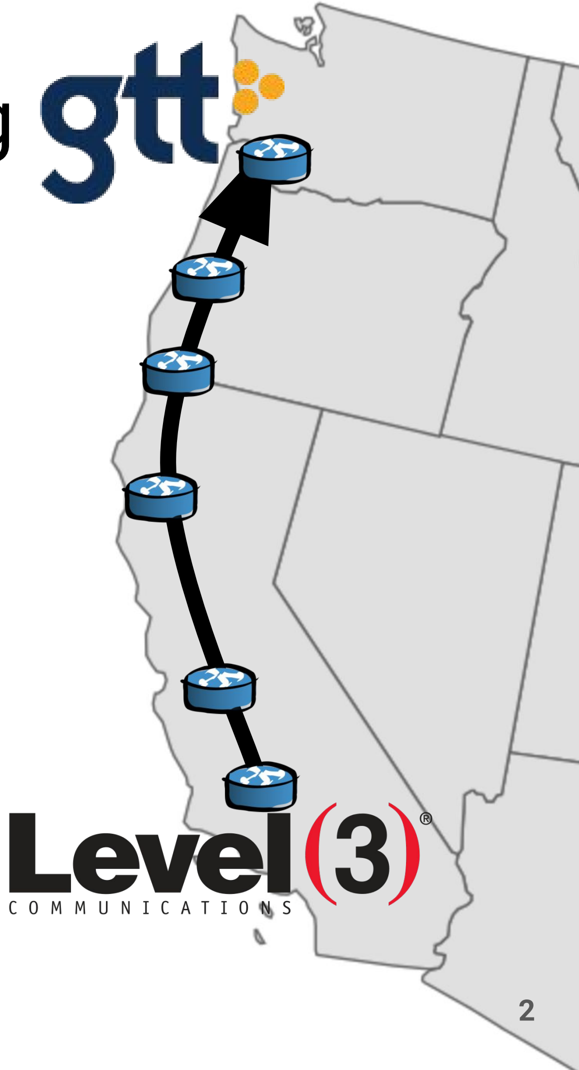

The Art of Network Troubleshooting

1. Operator has routing problem
2. Runs traceroute to the destination
3. Needs to identify where the problem is

"Are routes through GTT in Seattle experiencing problems?"

"Are routes through Level3 in LA experiencing problems?"

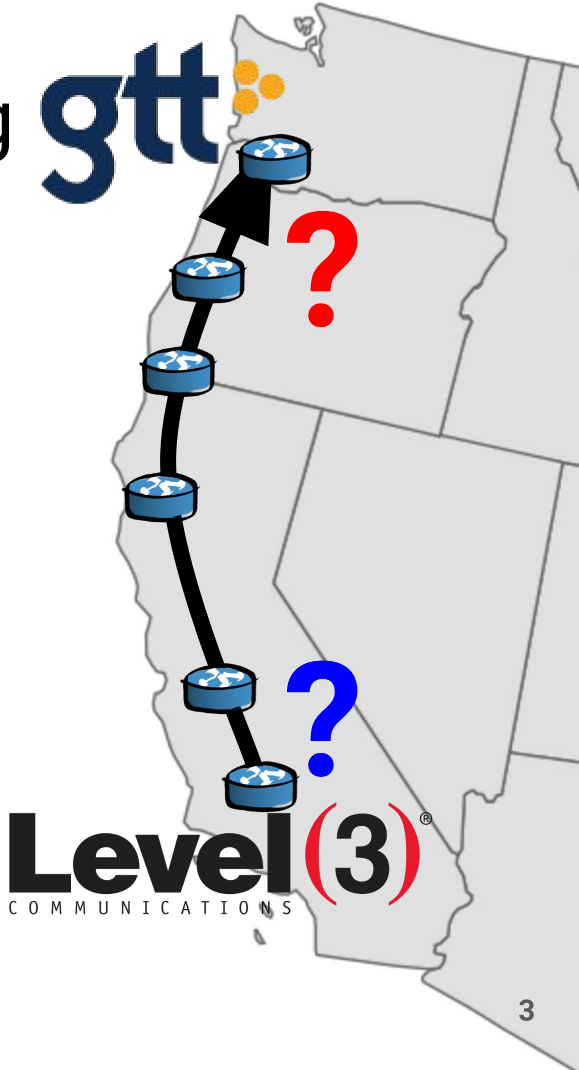

The Art of Network Troubleshooting

1. Operator has routing problem
2. Runs traceroute to the destination
3. Needs to identify where the problem is

“Are routes through GTT in Seattle experiencing problems?”

“Are routes through Level3 in LA experiencing problems?”

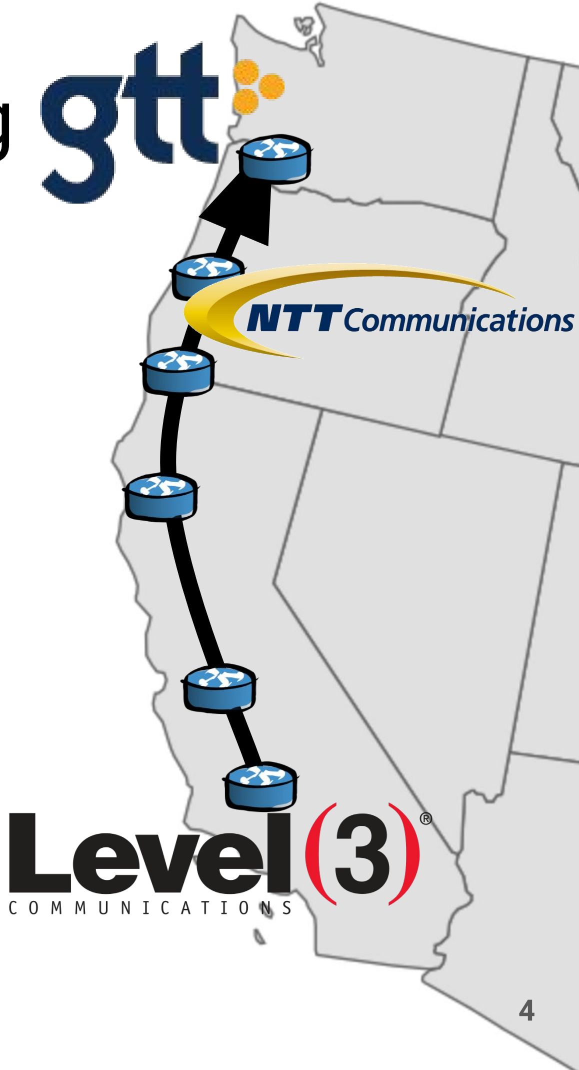

Operators Have Complex Questions

Someone suggests problem is on link between NTT and GTT in Seattle. Which routes use that link and could be impacted?

Operators Have Complex Questions

Someone suggests problem is on link between NTT and GTT in Seattle. Which routes use that link and could be impacted?

If route to destination D is impacted, which of my providers have a route that avoids that link?

Providers in my region with routes that avoid that link?

- [outages] GTT Peering issues *Hugh Smallwood*
    - [outages] GTT Peering issues *Andree Toonk*
        - [outages] GTT Peering issues *Adam Davenport*
    - [outages] GTT Peering issues *Adam Davenport*

2. Runs
3. Need

“Are r

“Are r

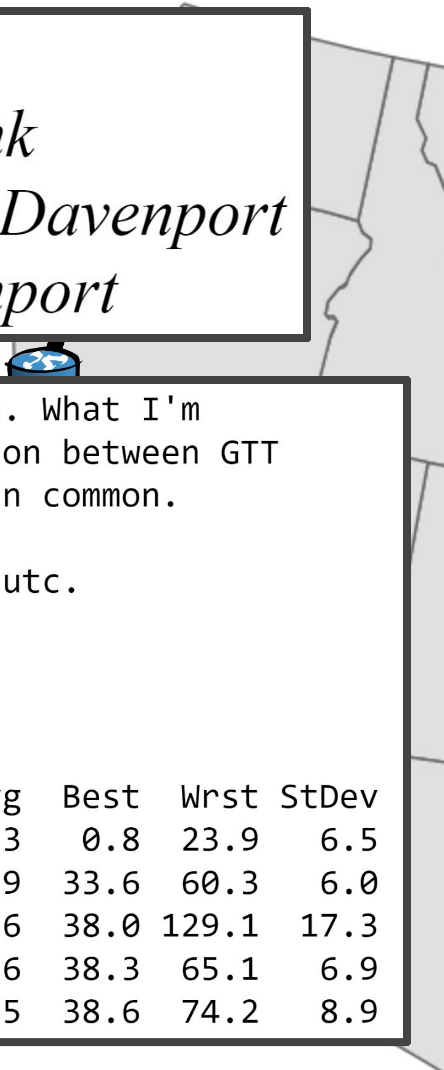

```
Seeing similar issues between a few US, Asia and EU sites. What I'm
seeing is what appears to be a congested peering connection between GTT
and NTT in Seattle as that is what all traceroutes have in common.

Our graphs indicate this started Jan 1st, around ~ 15:00 utc.

Example traceroute from Seattle (ntt) to Vancouver (gtt):


HOST: rtr1-re0.sea                 Loss%   Snt   Last   Avg  Best  Wrst StDev
  1. xe-0-0-0-34.r04.sttlwa01.us.  0.0%    60    4.1   8.3   0.8  23.9   6.5
  2. ae3.sea22.ip4.gtt.net        20.0%    60   44.6  41.9  33.6  60.3   6.0
  3. xe-1-2-0.van10.ip4.gtt.net   18.3%    60   48.9  51.6  38.0 129.1  17.3
  4. opendns-gw.ip4.gtt.net       20.0%    60   46.6  46.6  38.3  65.1   6.9
  5. rtr1.yvr.opendns.com         28.3%    60   44.8  47.5  38.6  74.2   8.9
```

Problems with the Current Approach

1. Takes time
2. Requires help from others
    a. Needs operator with right VP to respond
    b. Destination to measure may not be obvious
3. Limits analyses

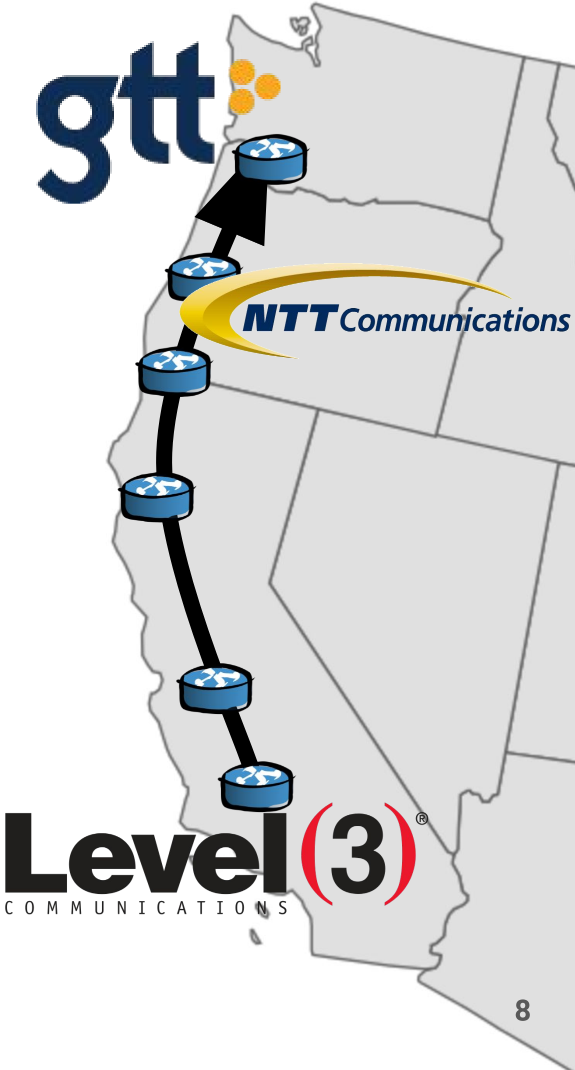

The Art of Internet Measurement

1. Researcher wants to study, for example, boomerang routes

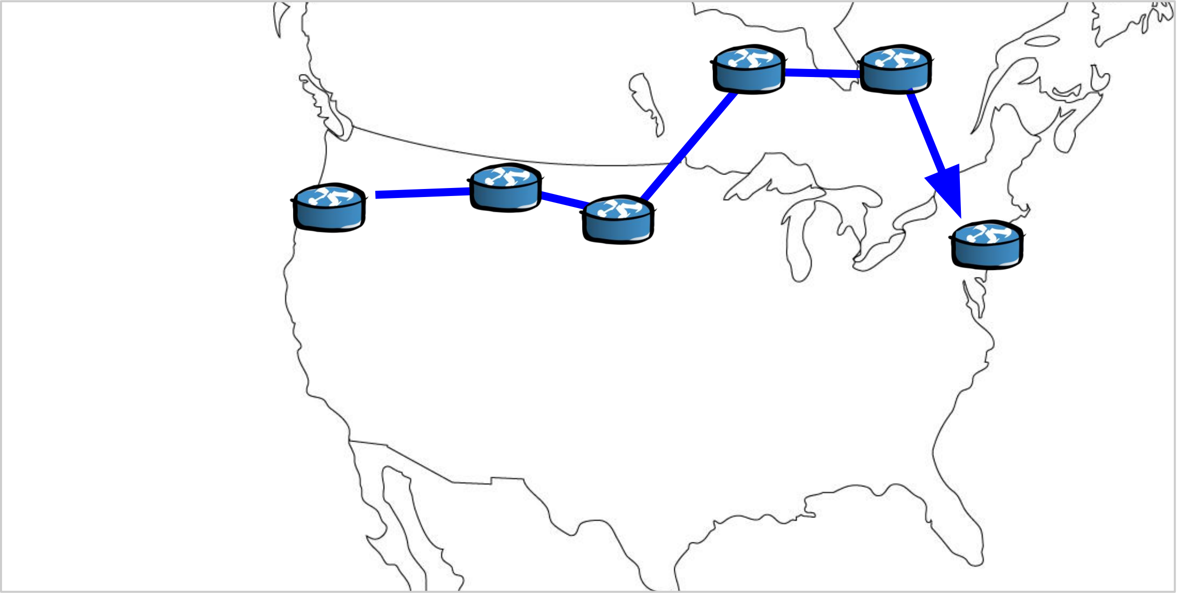

Researchers Have Complex Questions

What routes from the US to the US detour through Canada?

Researchers Have Complex Questions

What routes from the US to the US detour through Canada?

Eastward routes detouring through Canada?

Routes from Southern US through Canada? Mexico?

Routes from Africa to Africa through Europe?

The Art of Internet Measurement

1. Researcher wants to study, for example, boomerang routes
2. Collect (a lot of) traceroute measurements

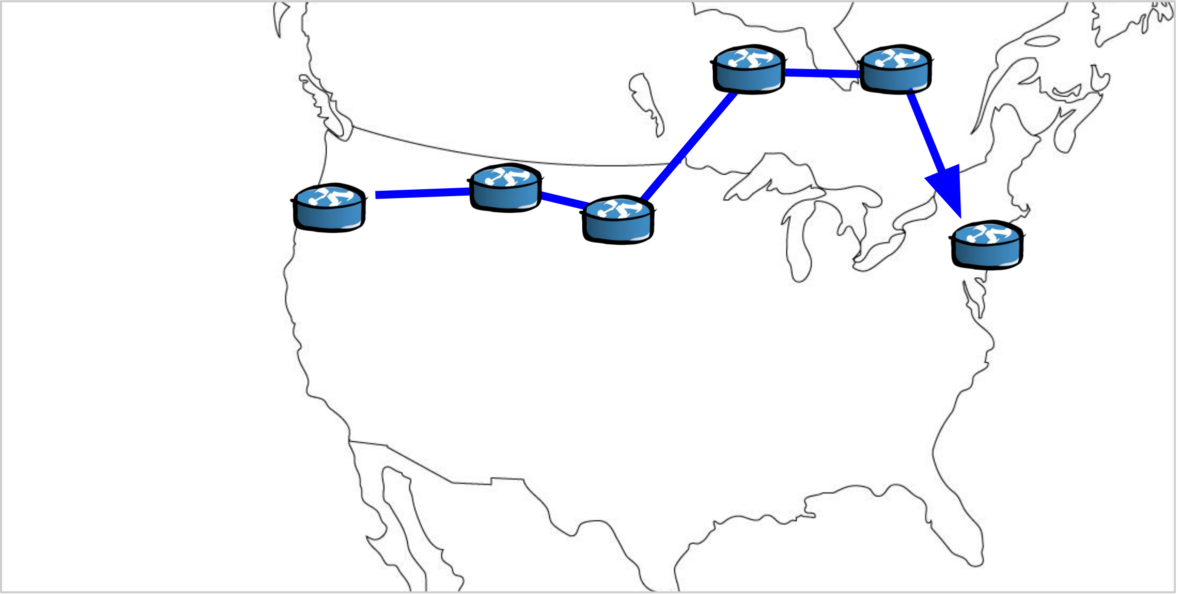

The Art of Internet Measurement

1. Resea
2. Collec

Index of /datasets/iplane-traceroutes/2016

Name	Last modified	Size	Description
Parent Directory		-	
traces_2016_01_01.tar.gz	02-Jan-2016 19:02	1.6G	
traces_2016_01_02.tar.gz	03-Jan-2016 19:03	1.6G	
traces_2016_01_03.tar.gz	04-Jan-2016 19:03	1.6G	
traces_2016_01_04.tar.gz	05-Jan-2016 19:03	1.6G	
traces_2016_01_05.tar.gz	06-Jan-2016 19:02	1.6G	
traces_2016_01_06.tar.gz	07-Jan-2016 19:03	1.6G	
traces_2016_01_07.tar.gz	08-Jan-2016 19:03	1.6G	

The Art of Internet Measurement

1. Resea
2. Collec

Index of /datasets/iplane-traceroutes/2016

Name	Last modified	Size	Description
traces_2016_01_02.tar.gz	03-Jan-2016 19:03	1.6G	
traces_2016_01_03.tar.gz	04-Jan-2016 19:03	1.6G	
traces_2016_01_04.tar.gz	05-Jan-2016 19:03	1.6G	
traces_2016_01_05.tar.gz	06-Jan-2016 19:02	1.6G	
traces_2016_01_06.tar.gz	07-Jan-2016 19:03	1.6G	
traces_2016_01_07.tar.gz	08-Jan-2016 19:03	1.6G	

PlanetLab lacks diversity. Other platforms cannot make exhaustive measurements.

“The number one go-to tool is traceroute.”

NANOG Troubleshooting Tutorial, 2009.
NANOG Traceroute Tutorial, 2014.

“The number one go-to tool is traceroute.”

NANOG Troubleshooting Tutorial, 2009.
NANOG Traceroute Tutorial, 2014.

But traceroute only answers one simple question:

“What is the path from vantage point s to destination d?”

“The number one go-to tool is traceroute.”

NANOG Troubleshooting Tutorial, 2009.
NANOG Traceroute Tutorial, 2014.

But traceroute only answers one simple question:

“What is the path from vantage point s to destination d?”

We need a new tool to answer our questions

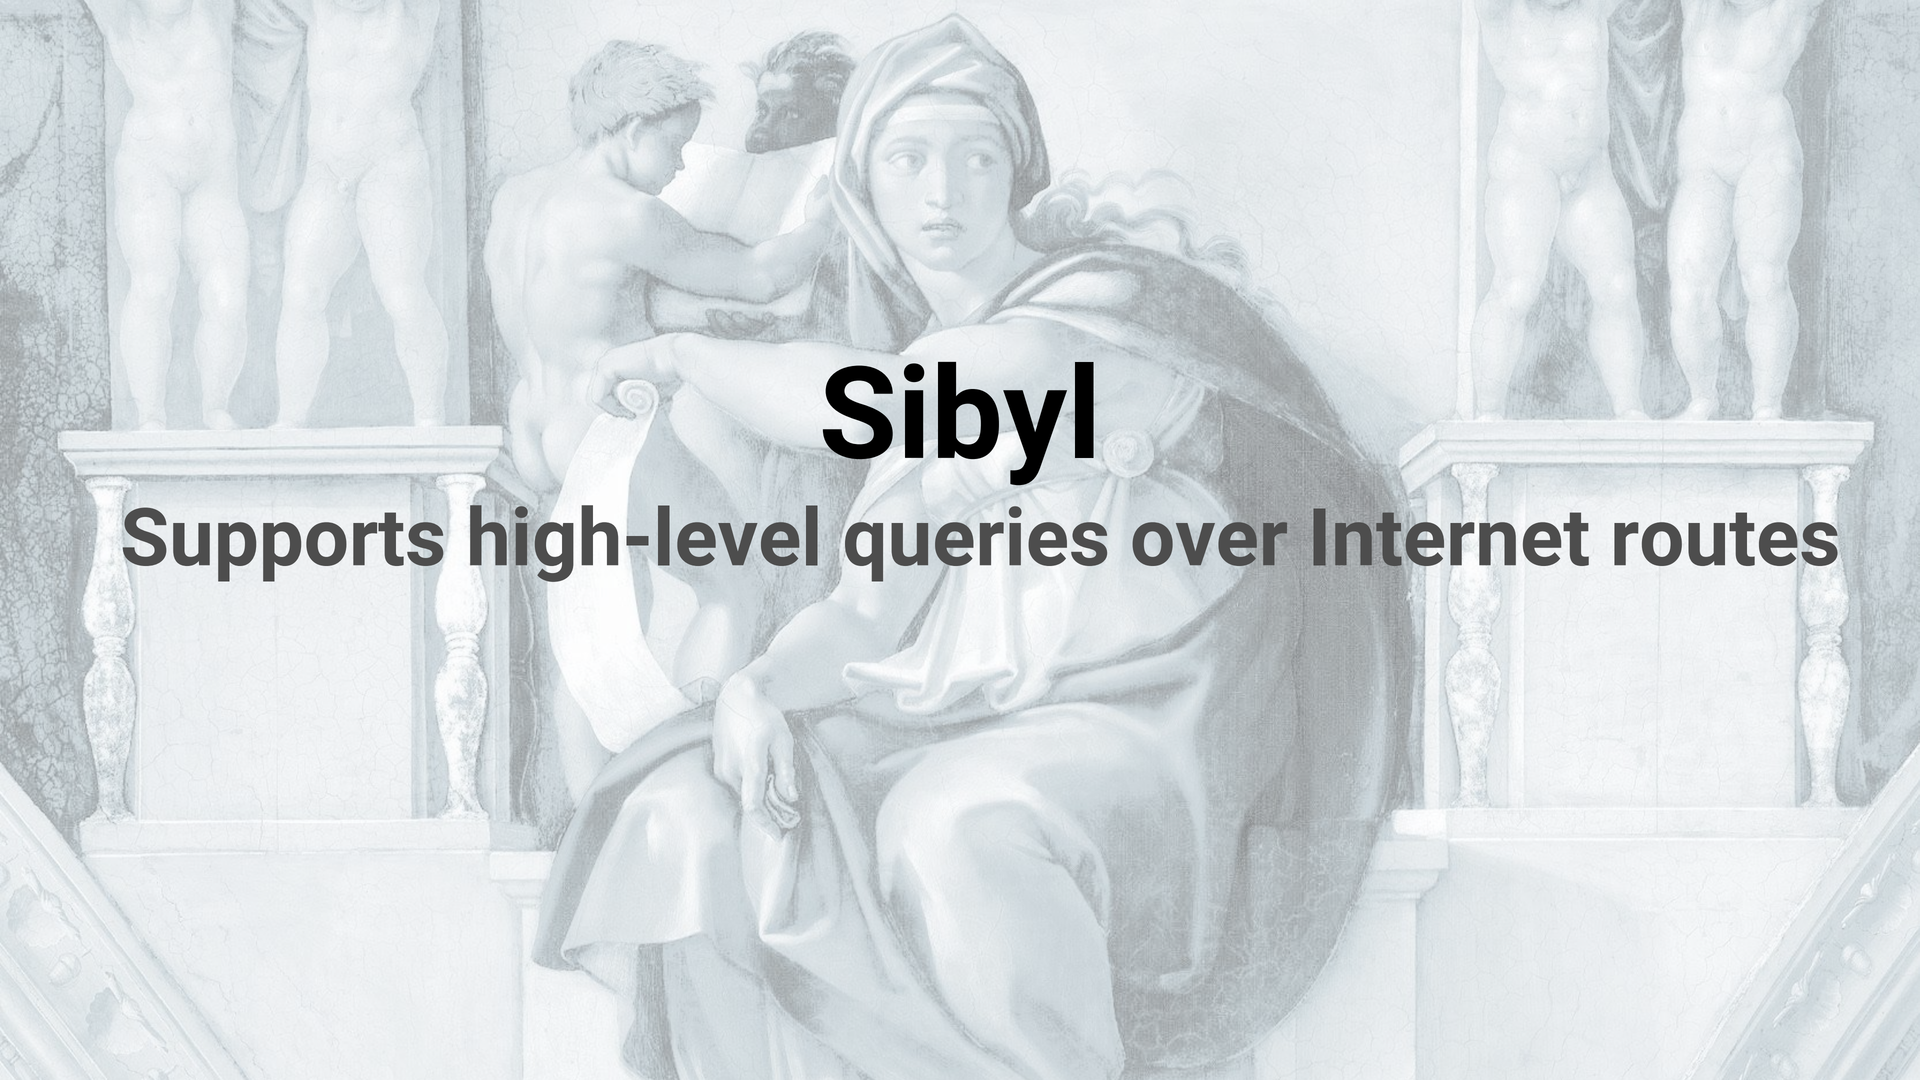

Sibyl

Supports high-level queries over Internet routes

Goal: Serve High-Level Routing Queries

Users submit
high-level queries

Queries

Sibyl

Traceroutes
that match queries

Query Internet Routes Using Regular Expressions

Symbols representing boolean expressions that IP addresses must satisfy

`Level3&LA` `GTT&Seattle`

Query Internet Routes Using Regular Expressions

Symbols representing boolean expressions
that IP addresses must satisfy

`Level3&LA`   `GTT&Seattle`

Routes traversing Level3 in LA and GTT in Seattle?

```
^ .* Level3&LA .* GTT&Seattle .* $
```

Regular Expressions Capture Operational Questions

Someone suggests problem is on link between NTT and GTT in Seattle. Which routes use that link and could be impacted?

```
^ .* NTT&Seattle GTT&Seattle .* $
```

Regular Expressions Capture Operational Questions

Someone suggests problem is on link between NTT and GTT in Seattle. Which routes use that link and could be impacted?

```
^ .* NTT&Seattle GTT&Seattle .* $
```

If route to destination D is impacted, which of my providers have a route that avoids that link?

```
^ USC (?!NTT GTT&Seattle).* D $
```

Providers in my region with routes that avoid that link?

```
^ .* LA (?!NTT GTT&Seattle).* D $
```

Regular Expressions Capture Research Questions

What routes from the US to the US detour through Canada?

```
^ US* Canada* US* $
```

Eastward routes detouring through Canada?

```
^ EastCoast .* Canada .* WestCoast $
```

Routes from Southern US through Canada? Mexico?

```
^ SouthernUS .* (Canada|Mexico) .* US $
```

Routes from Africa to Africa through Europe?

```
^ Africa .* Europe .* Africa $
```

How do we achieve route coverage to satisfy complex and diverse queries?

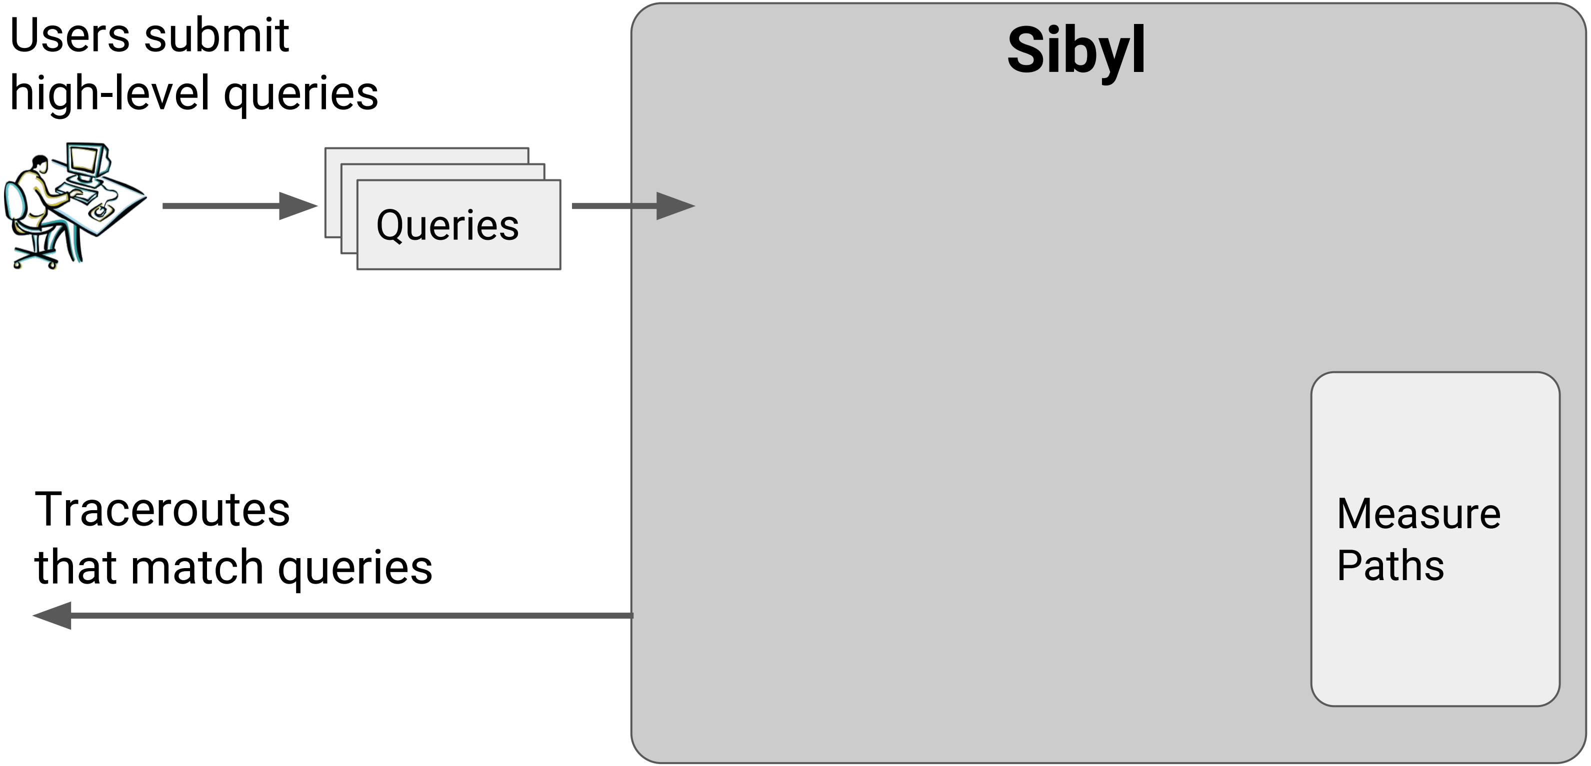

Combining Platforms Provides Great Coverage

Vantage points in 100% of very large networks
and 50% of small and regional networks

Combining Platforms Provides Great Coverage

Vantage points in 100% of very large networks and 50% of small and regional networks

Measurement Platform	Vantage Points	Vantage Point ASes
PlanetLab	~400	~250
Traceroute servers	~500	~500
RIPE Atlas	~9000	~2700

Vantage Points Are Resource Constrained

Want to use network coverage to answer queries,
but cannot issue all measurements

Measurement Platform	Vantage Points	Vantage Point ASes	Traces per Day
PlanetLab	~400	~250	**16,000K**
Traceroute servers	~500	~500	**145K**
RIPE Atlas	~9000	~2700	**35K**

Vantage Points Are Resource Constrained

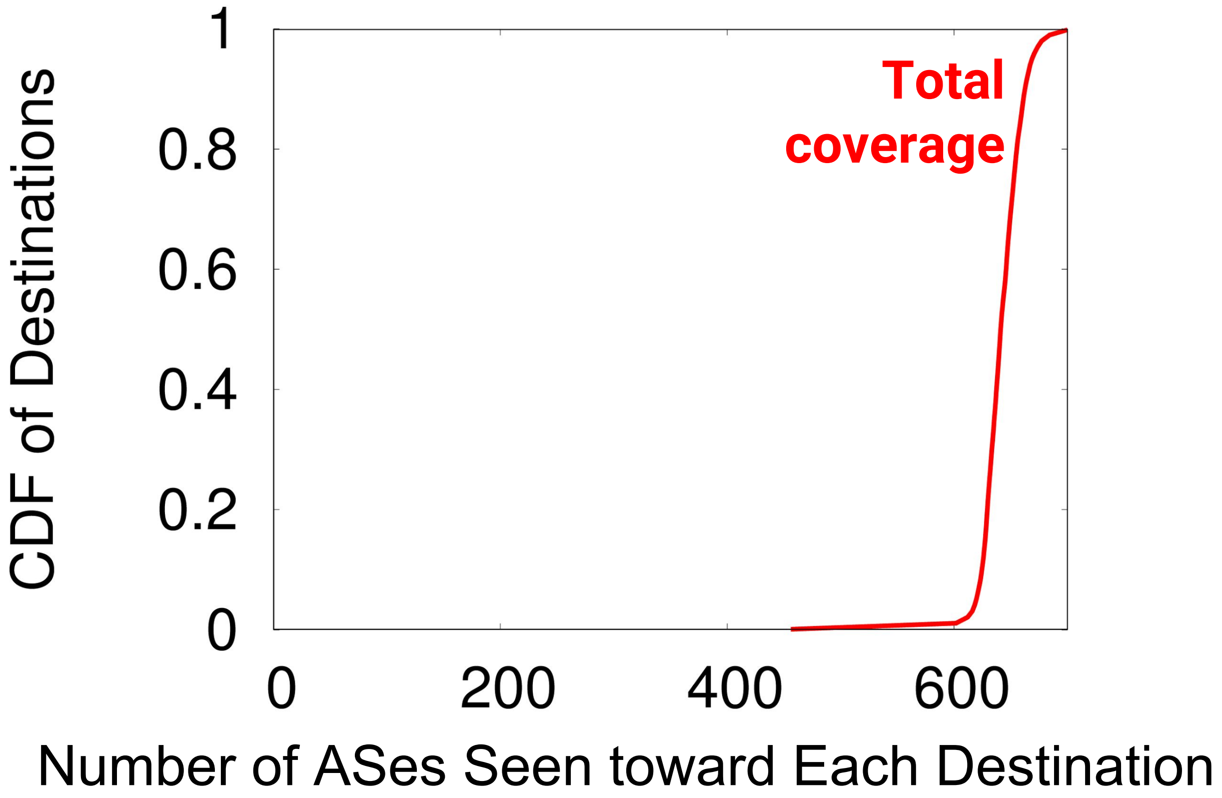

Resource Constraints Limit Visibility

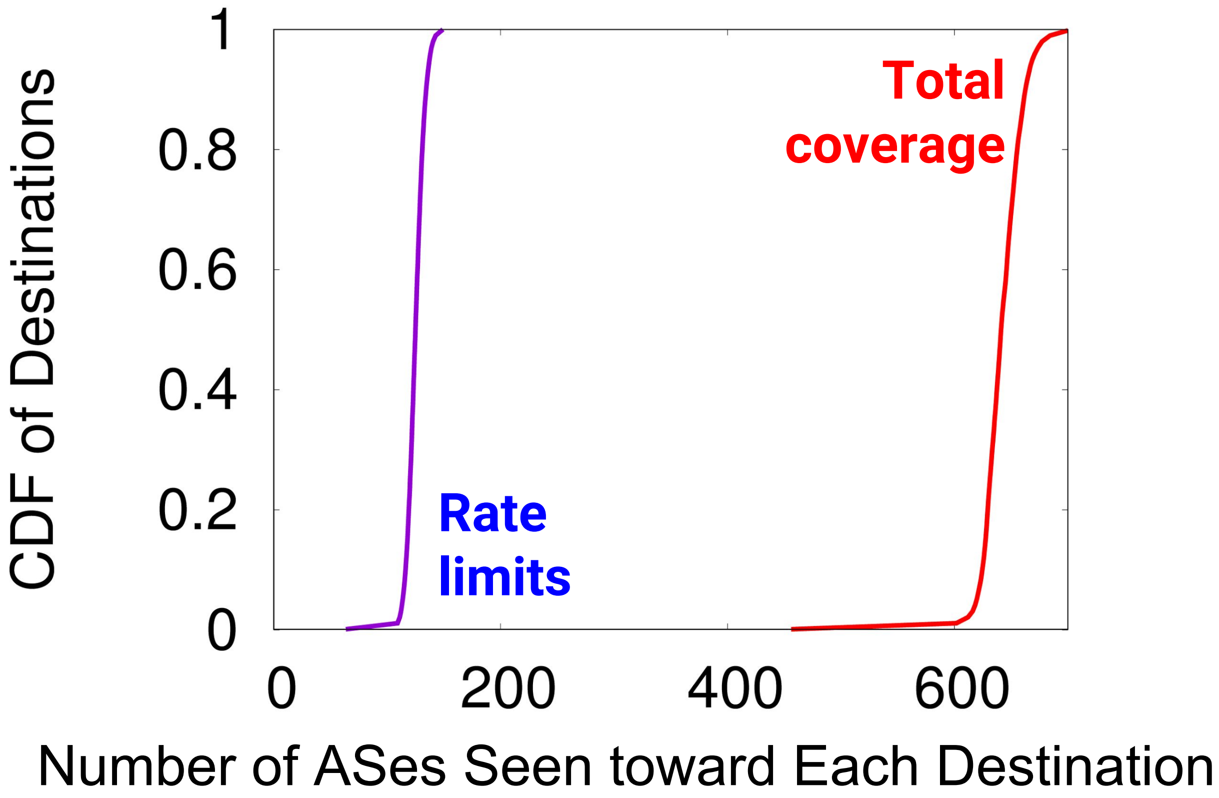

Resource Constraints Limit Visibility

Goal: satisfy queries as if we have total coverage, despite rate limits

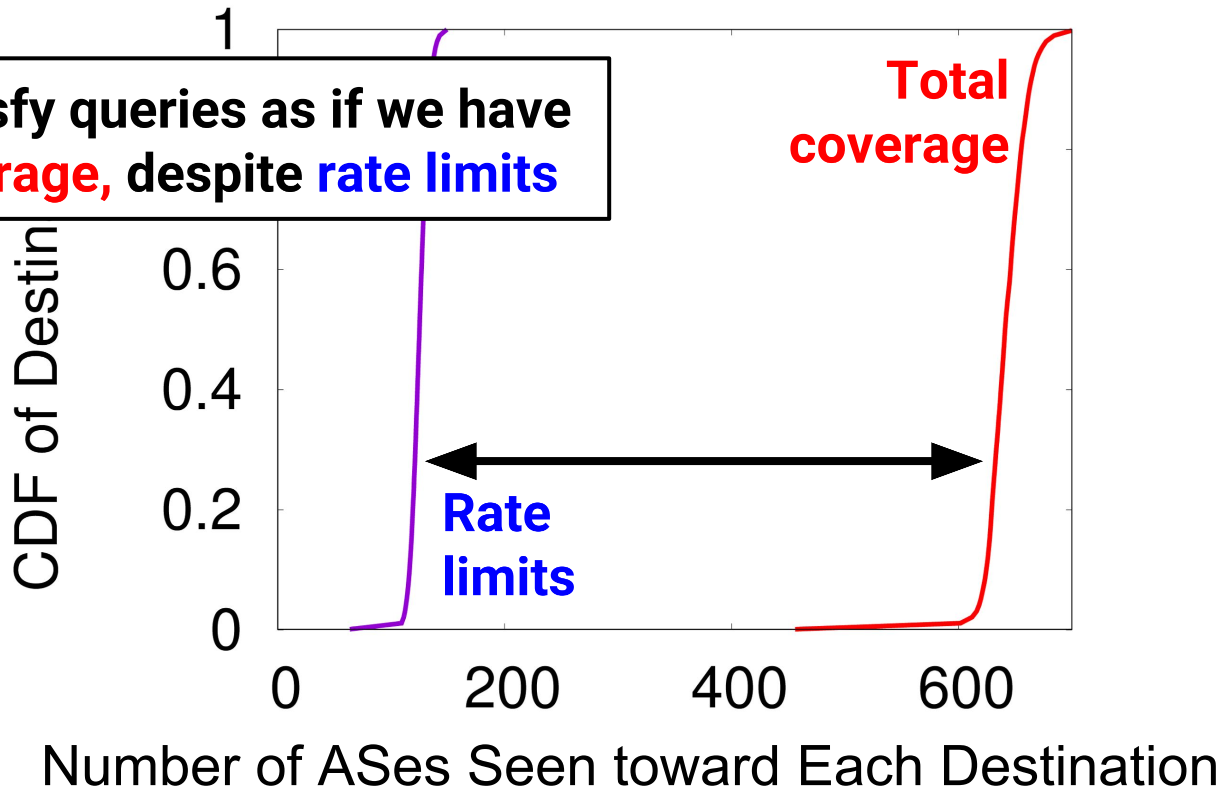

Resource Constraints Limit Visibility

Goal: satisfy queries as if we have total coverage, despite rate limits

Approach: carefully choose measurements

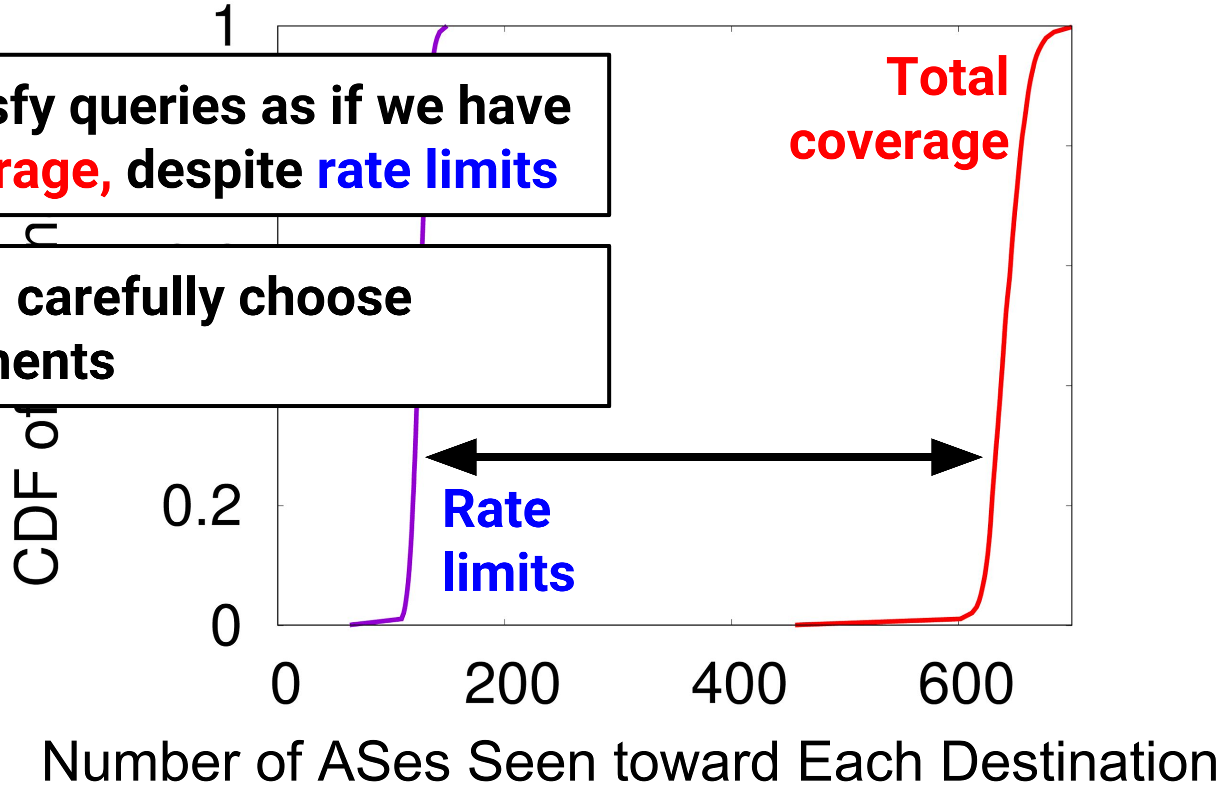

Too many possible measurements. How do we prune the set of measurements to consider?

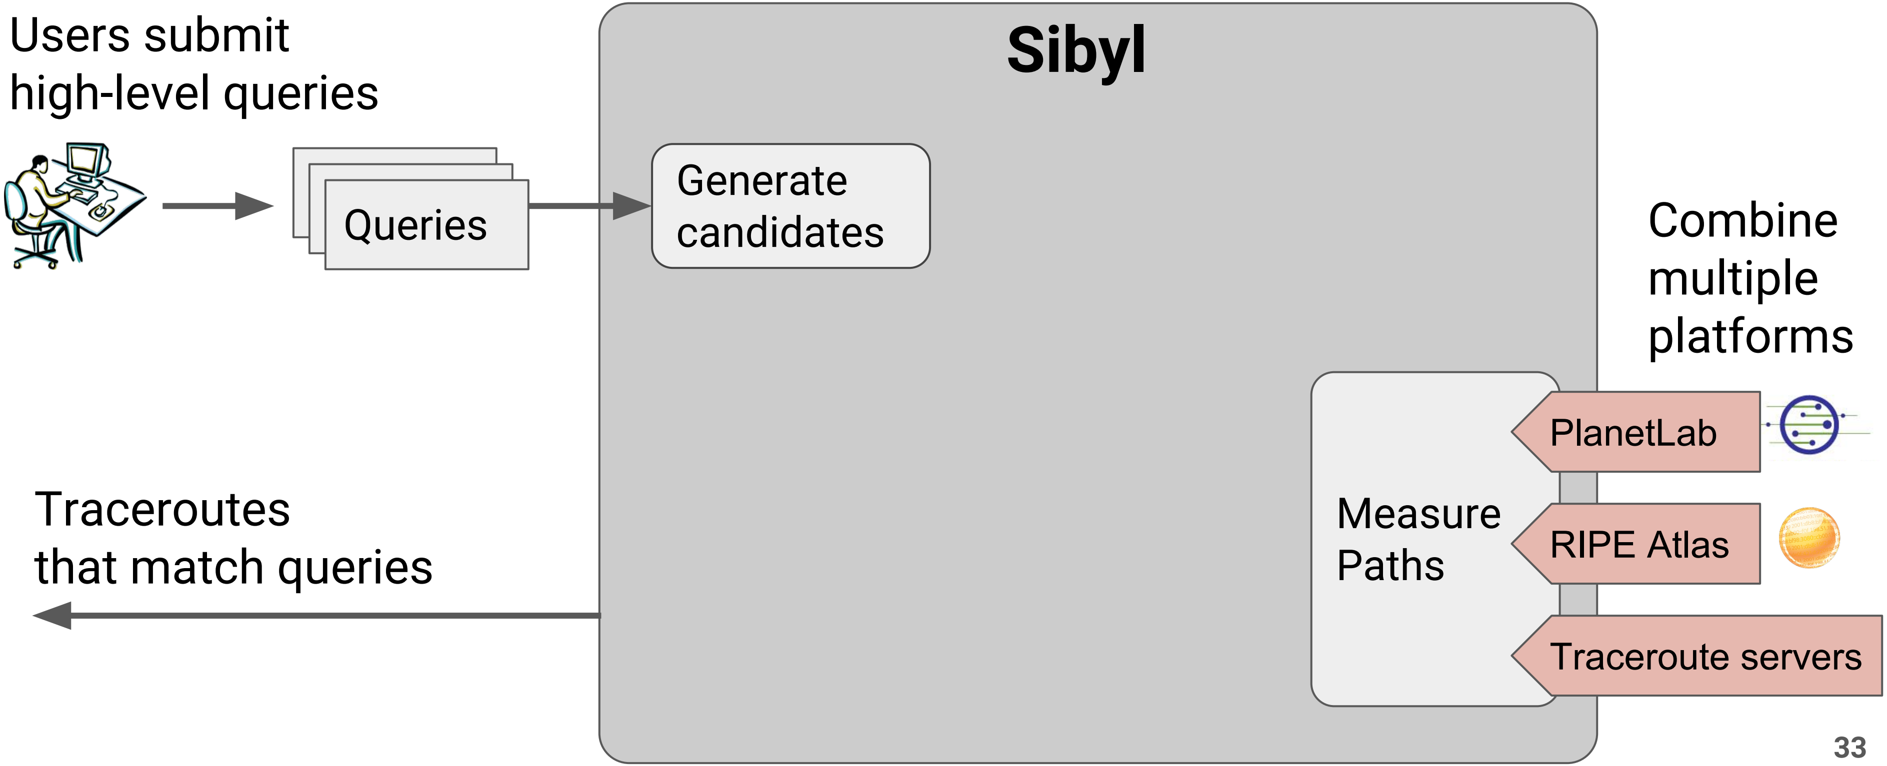

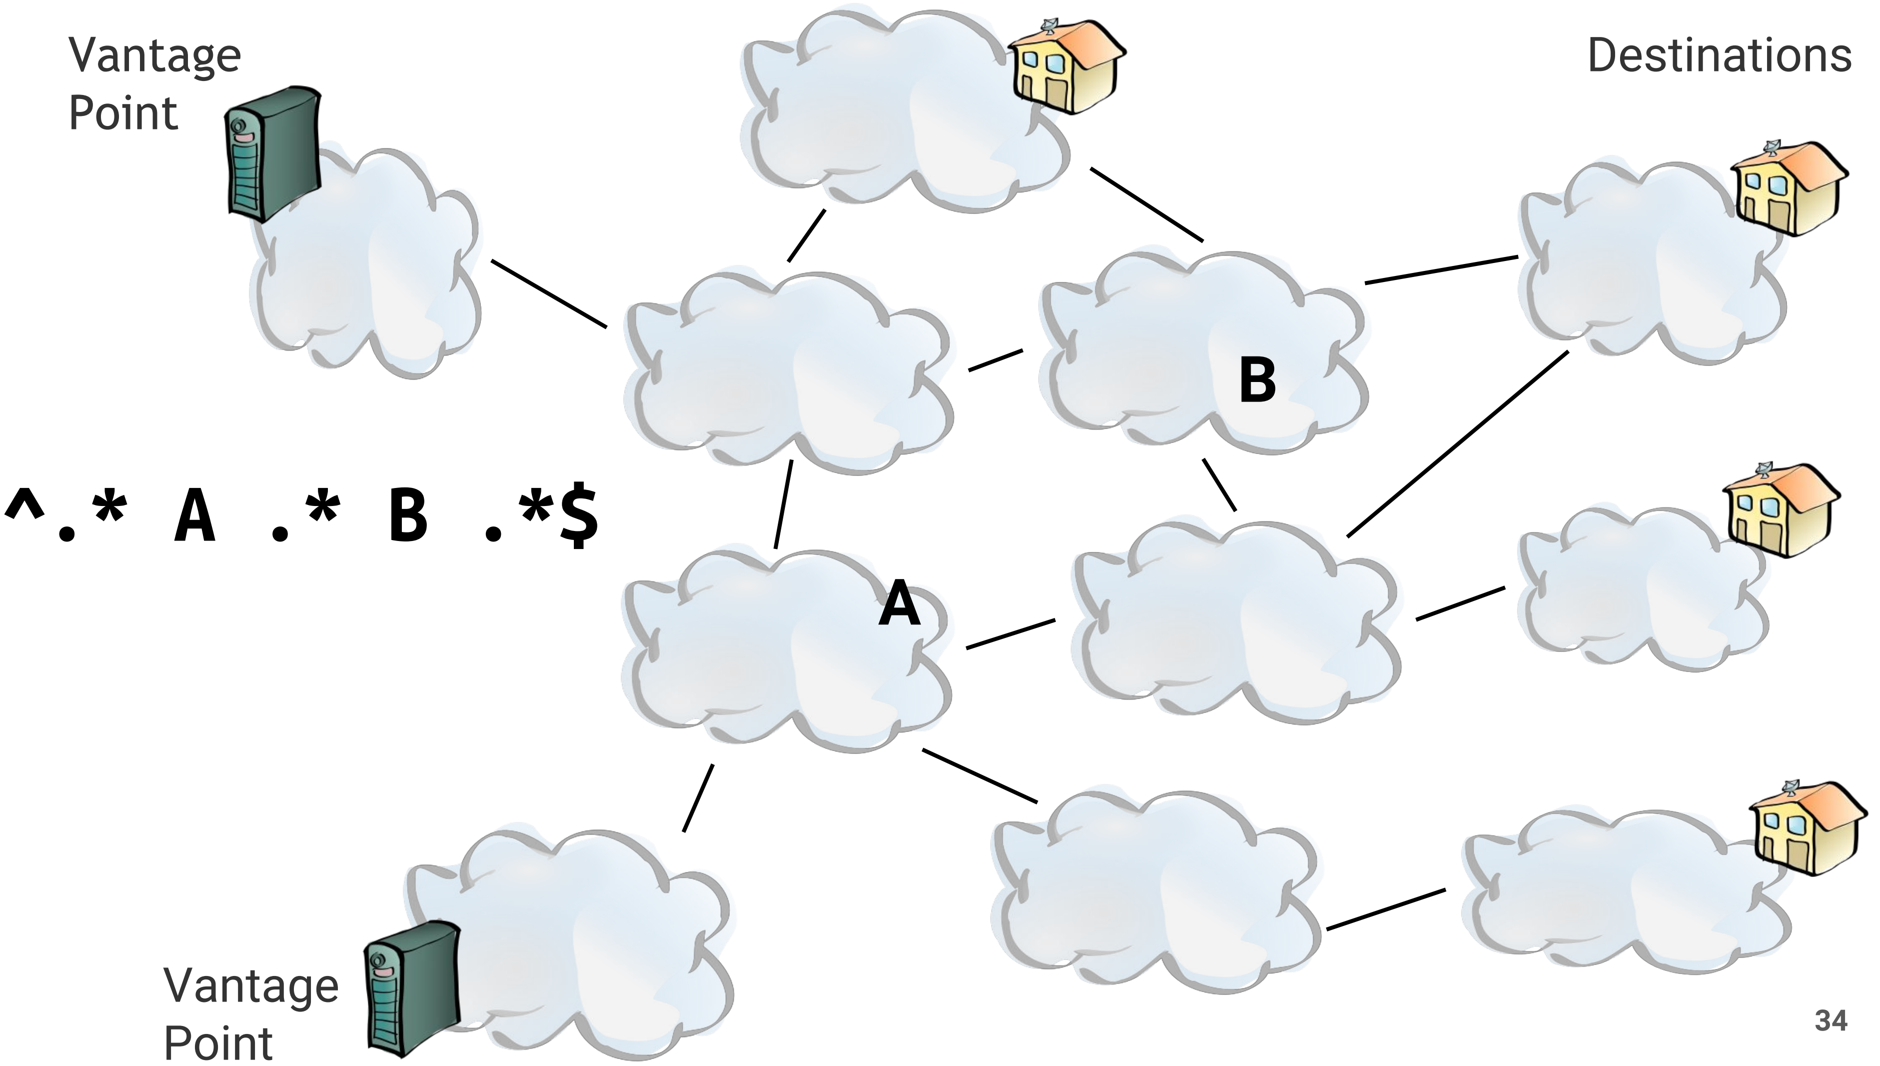
Vantage Point
Destinations
B
^.* A .* B .*$
A
Vantage Point

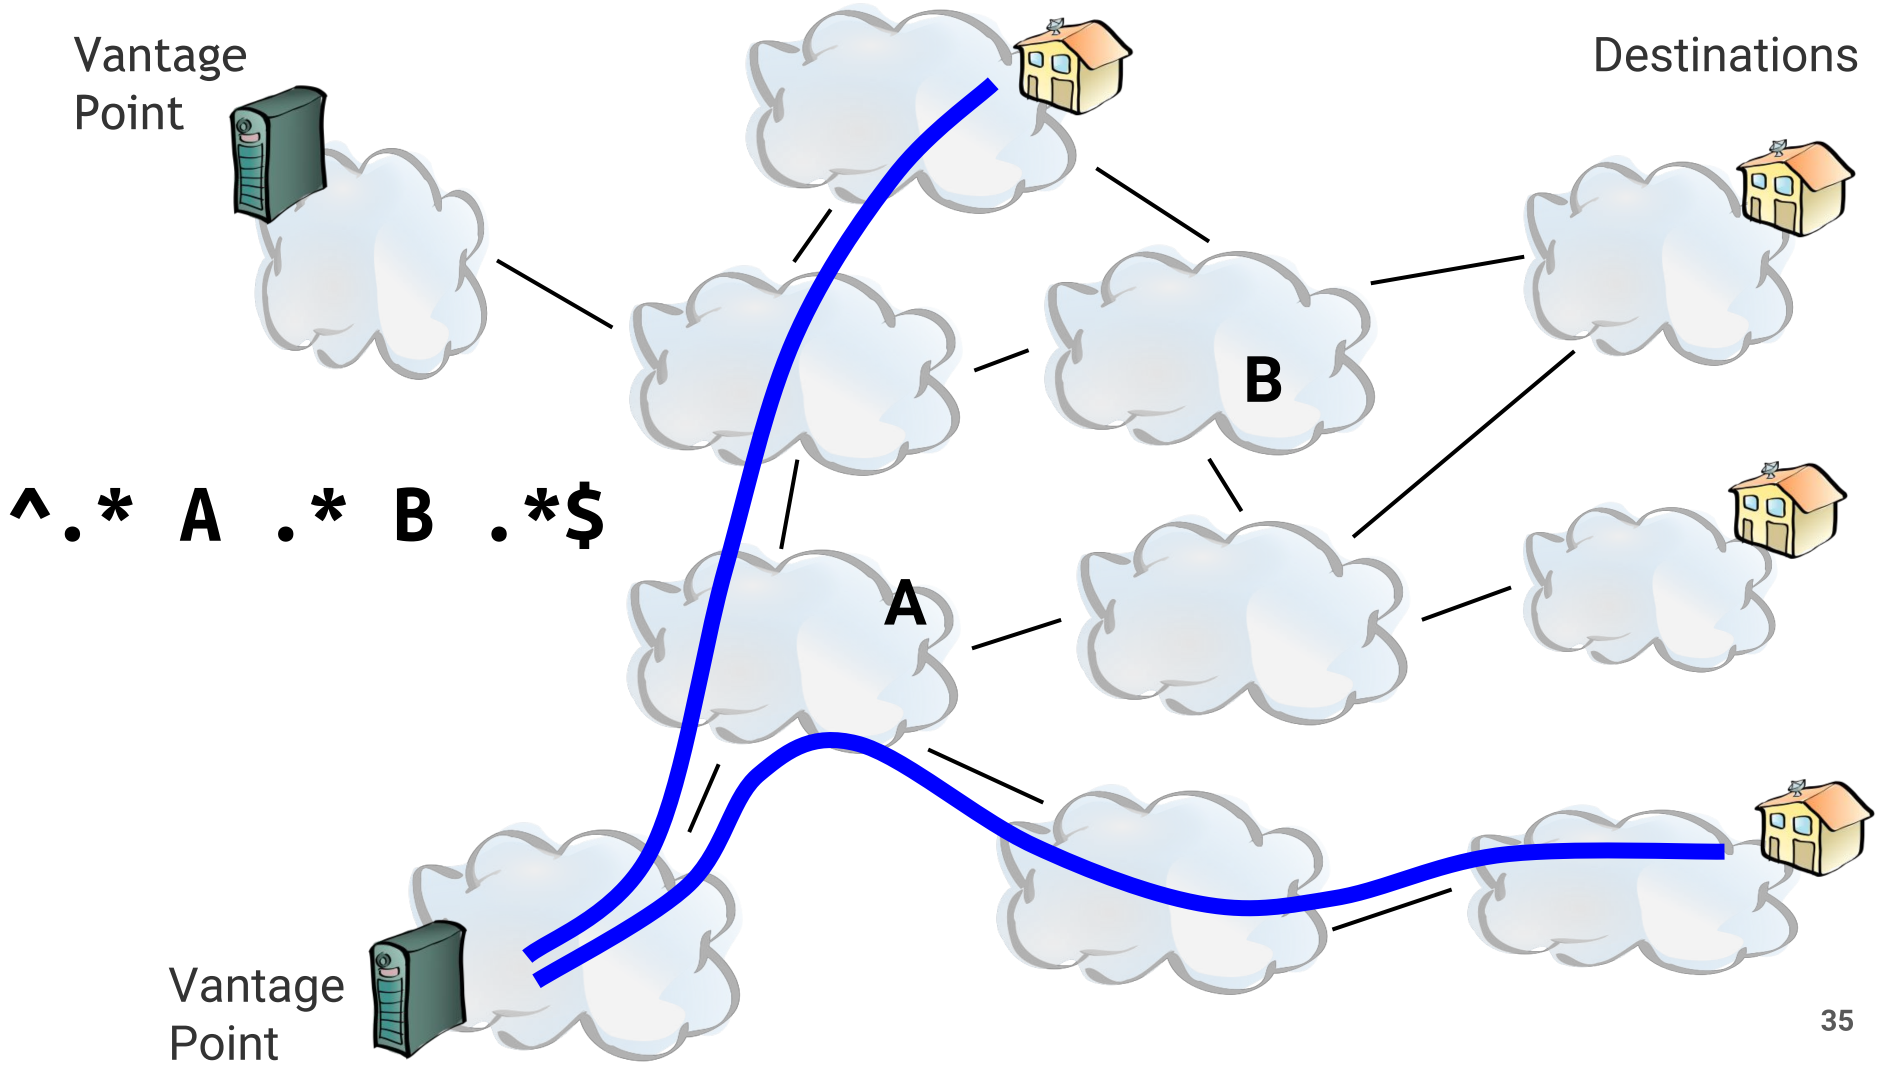

Vantage Point
Destinations
B
^.* A .* B .*$
A
Vantage Point

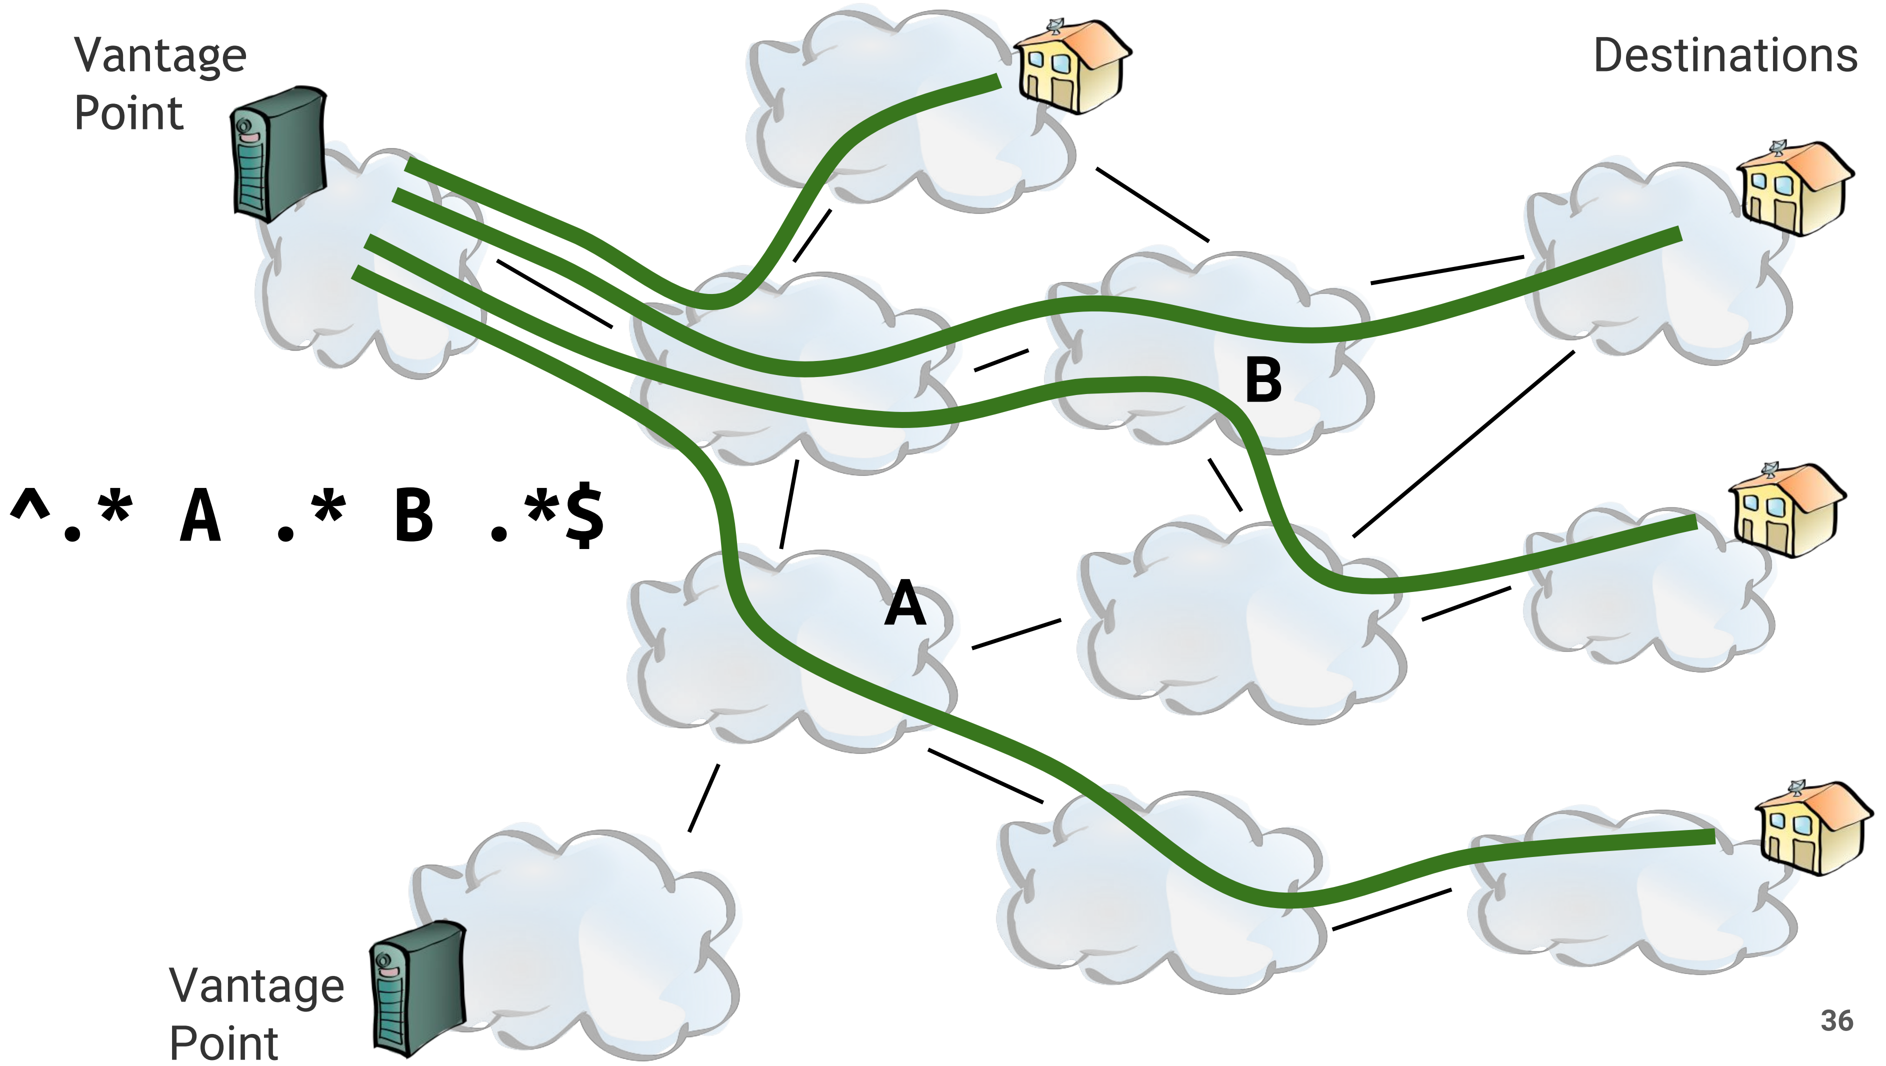

Vantage Point
Destinations
B
A
^.* A .* B .*$
Vantage Point

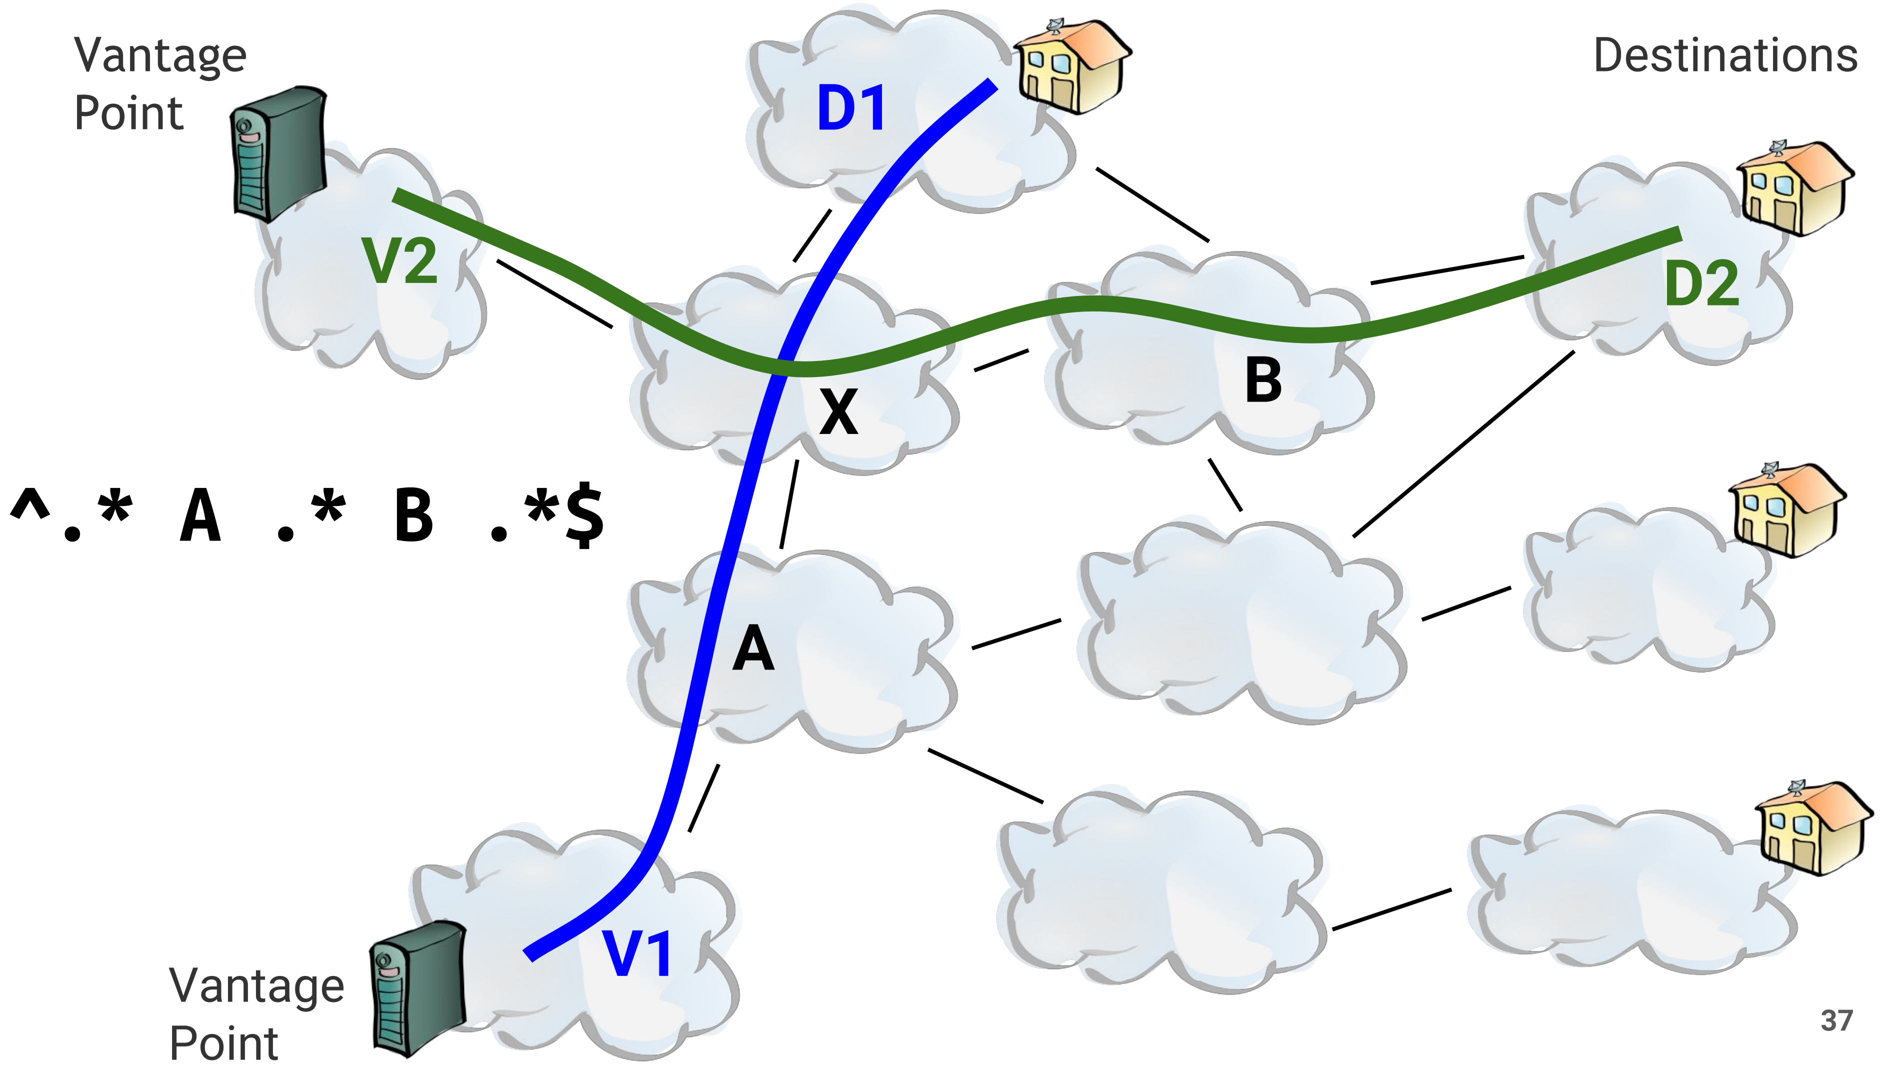
Vantage Point
Destinations
D1
V2
D2
X
B
^.* A .* B .*$
A
V1
Vantage Point

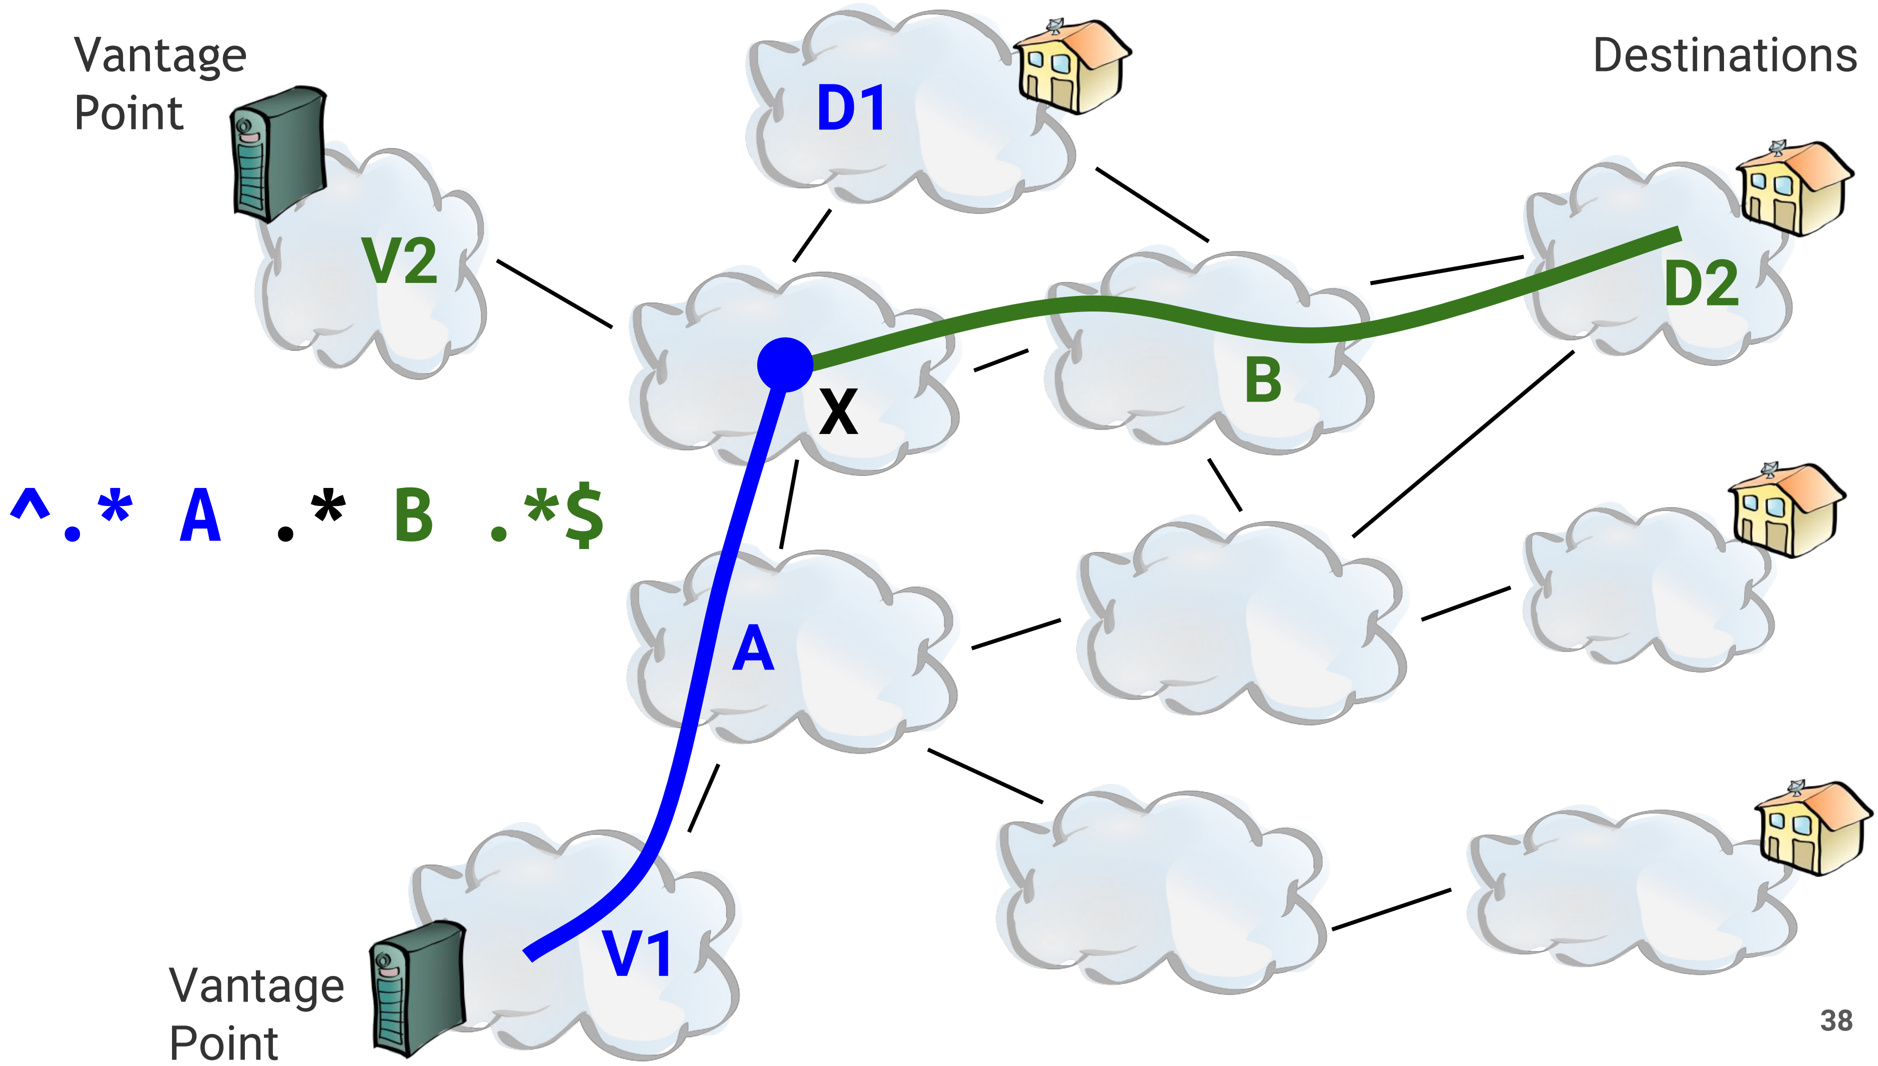

Vantage Point
D1
Destinations
V2
D2
X
B
^.* A .* B .*$
A
V1
Vantage Point

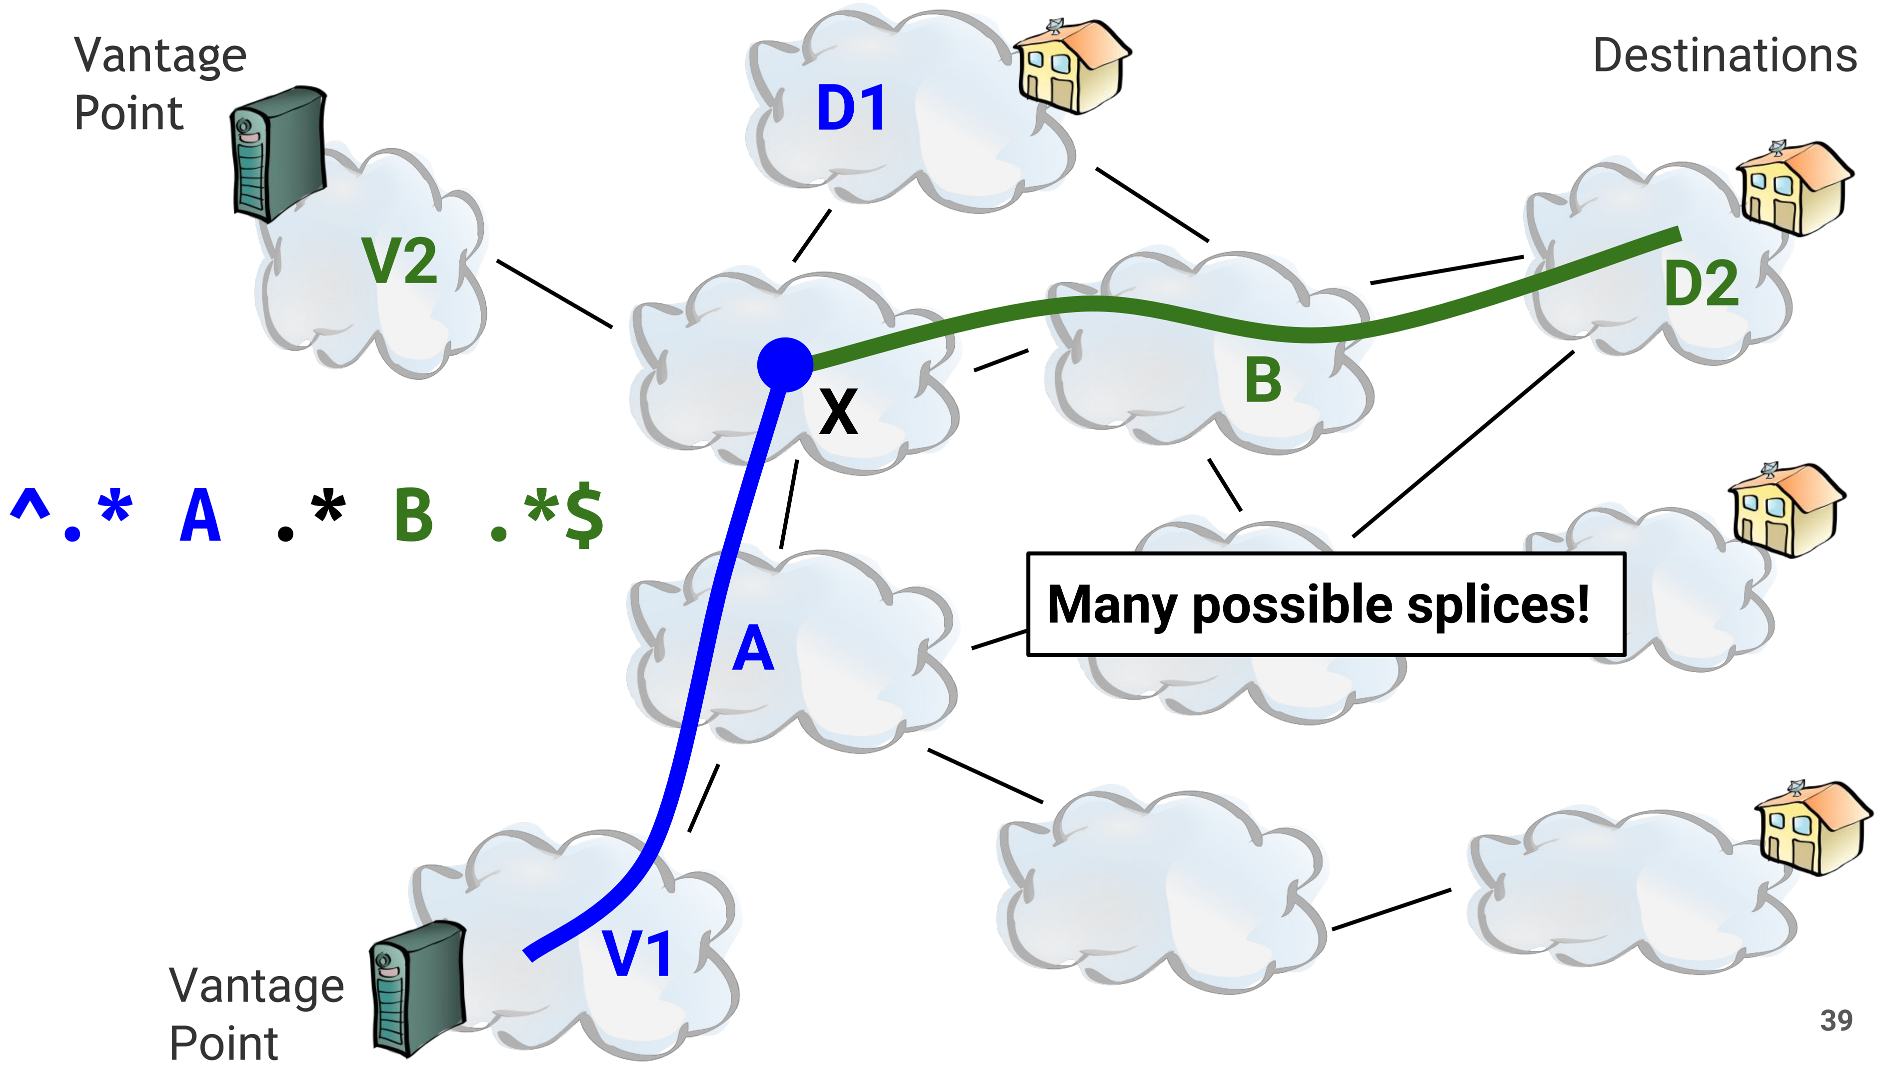

Vantage Point
D1
Destinations
V2
D2
X
B
^.* A .* B .*$
A
Many possible splices!
V1
Vantage Point

How to estimate which candidate measurements will match queries?

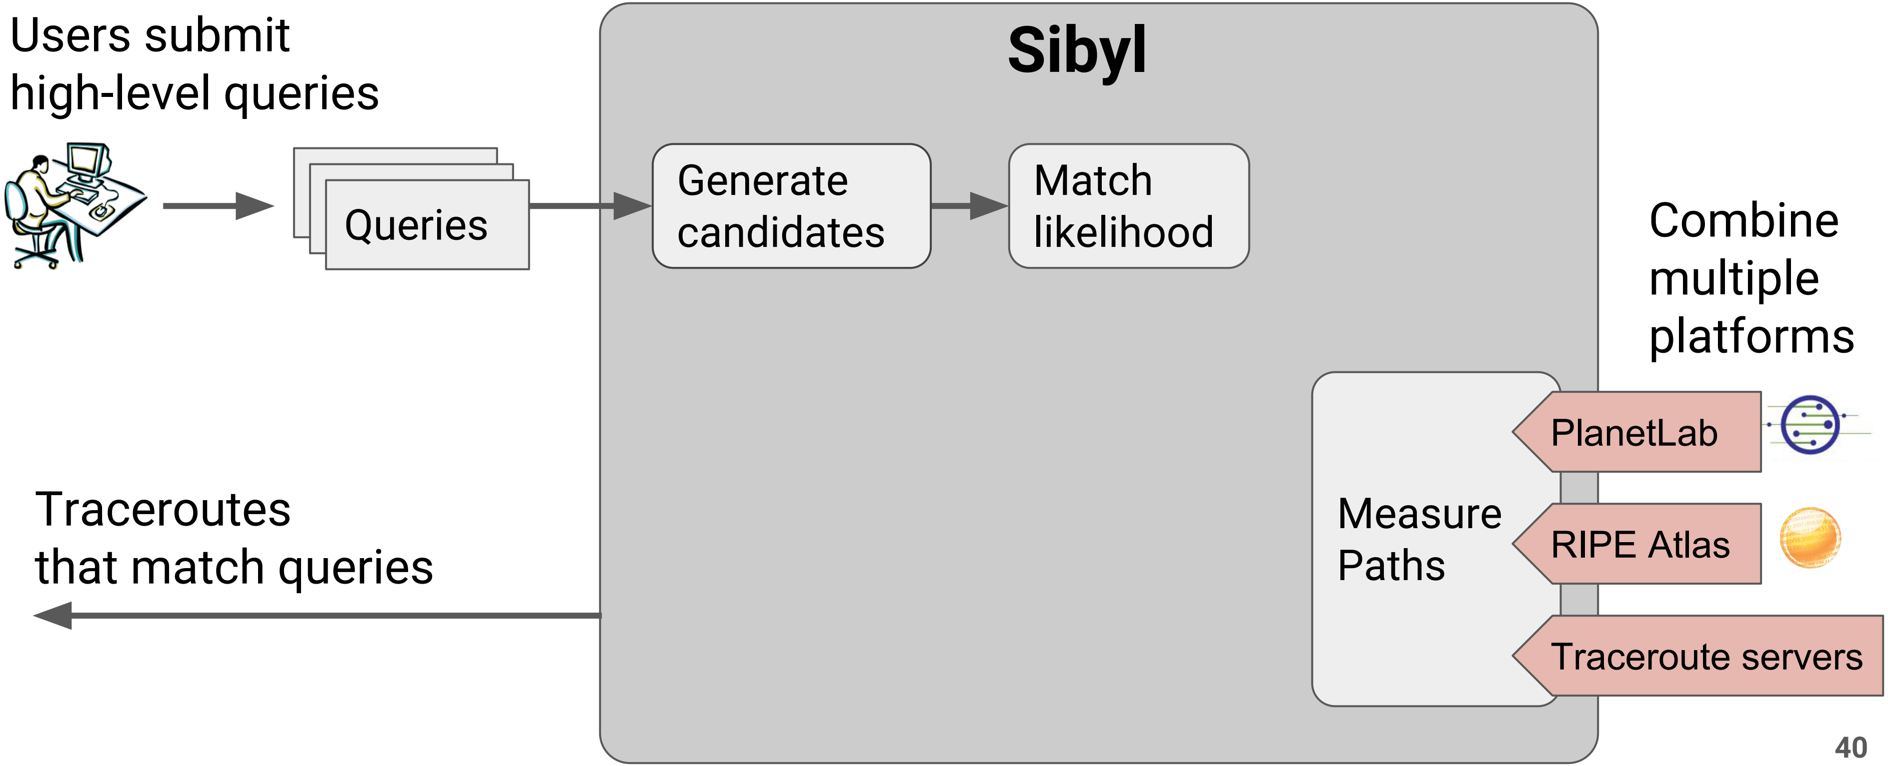

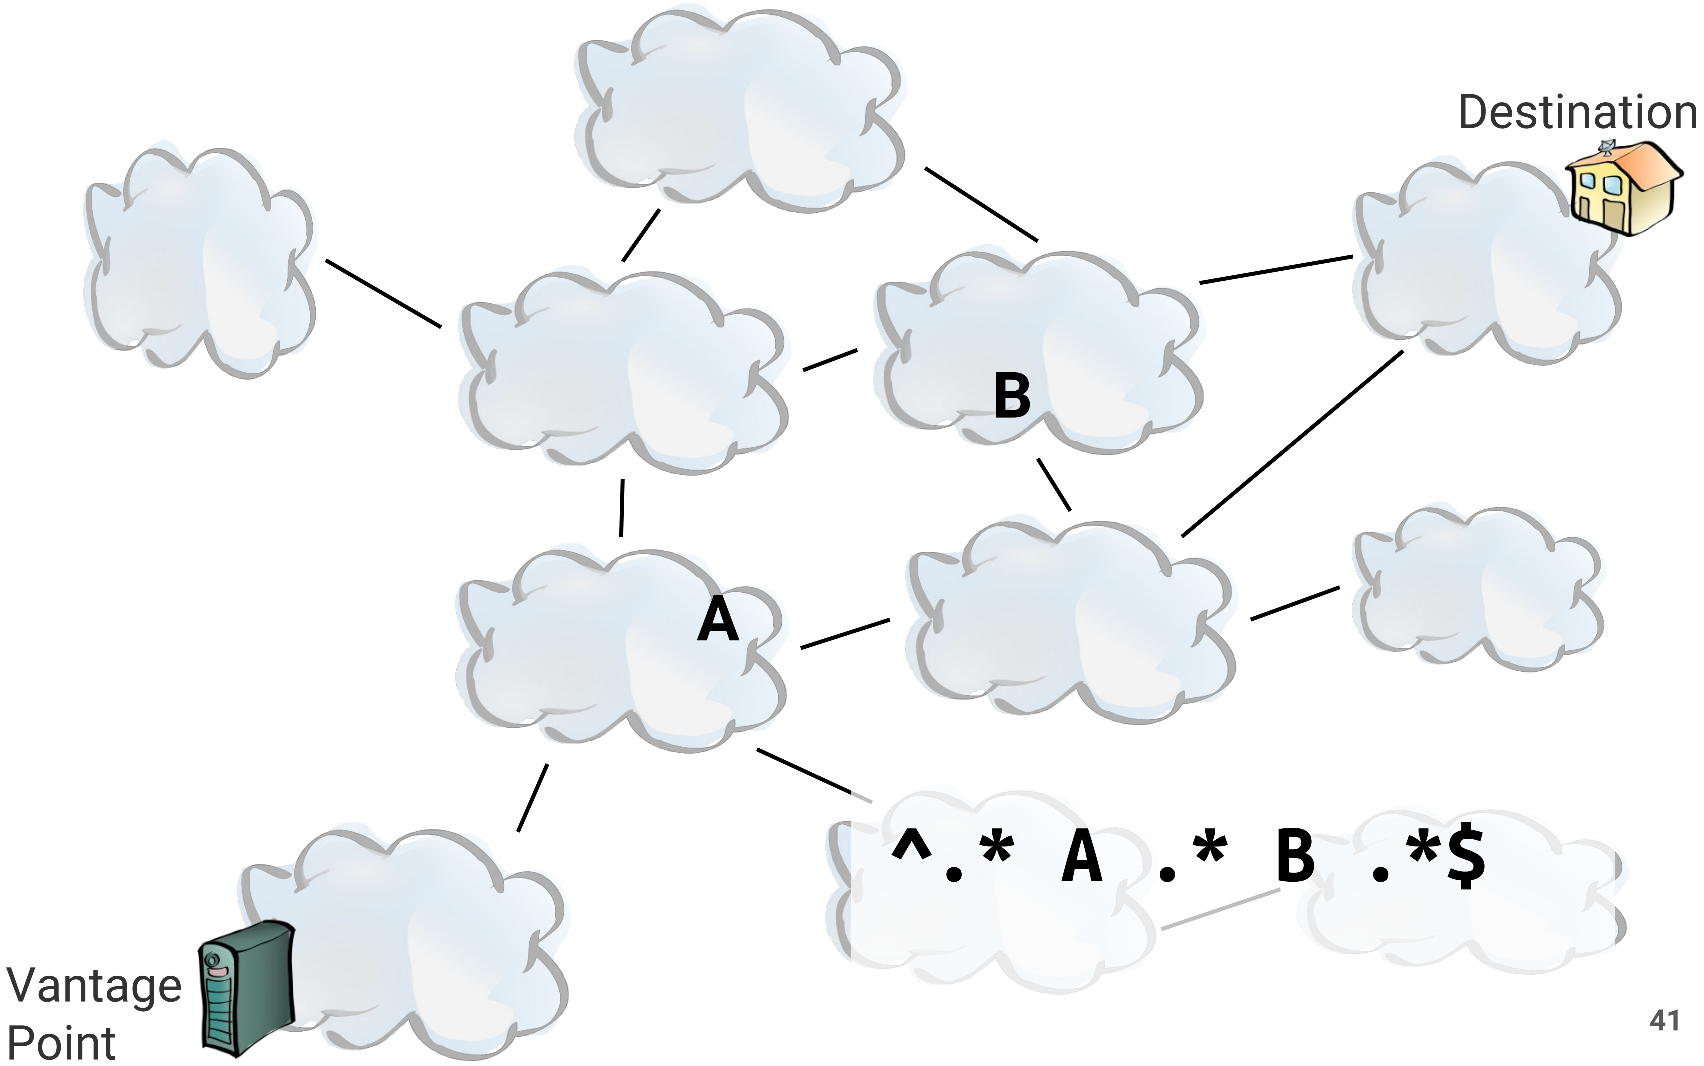
Destination
B
A
^.* A .* B .*$
Vantage
Point

1. Predict all paths for candidate measurements

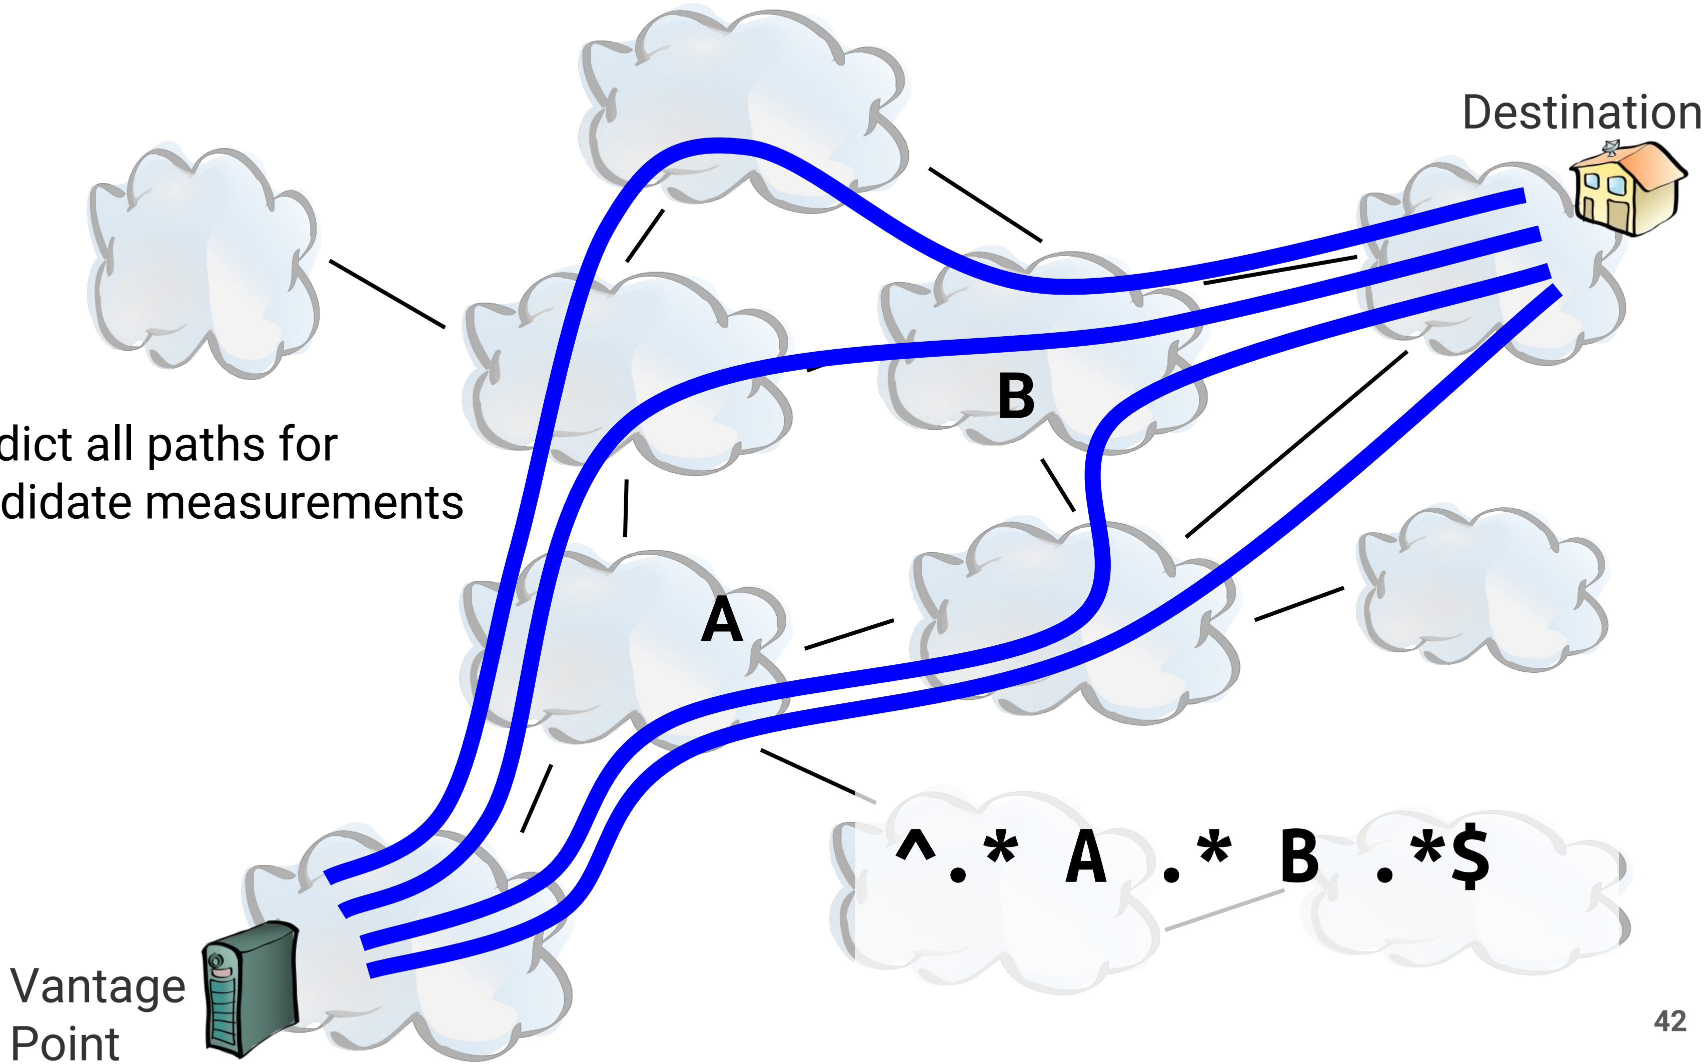

1. Predict all paths for candidate measurements
2. Compute likelihood of prediction being correct

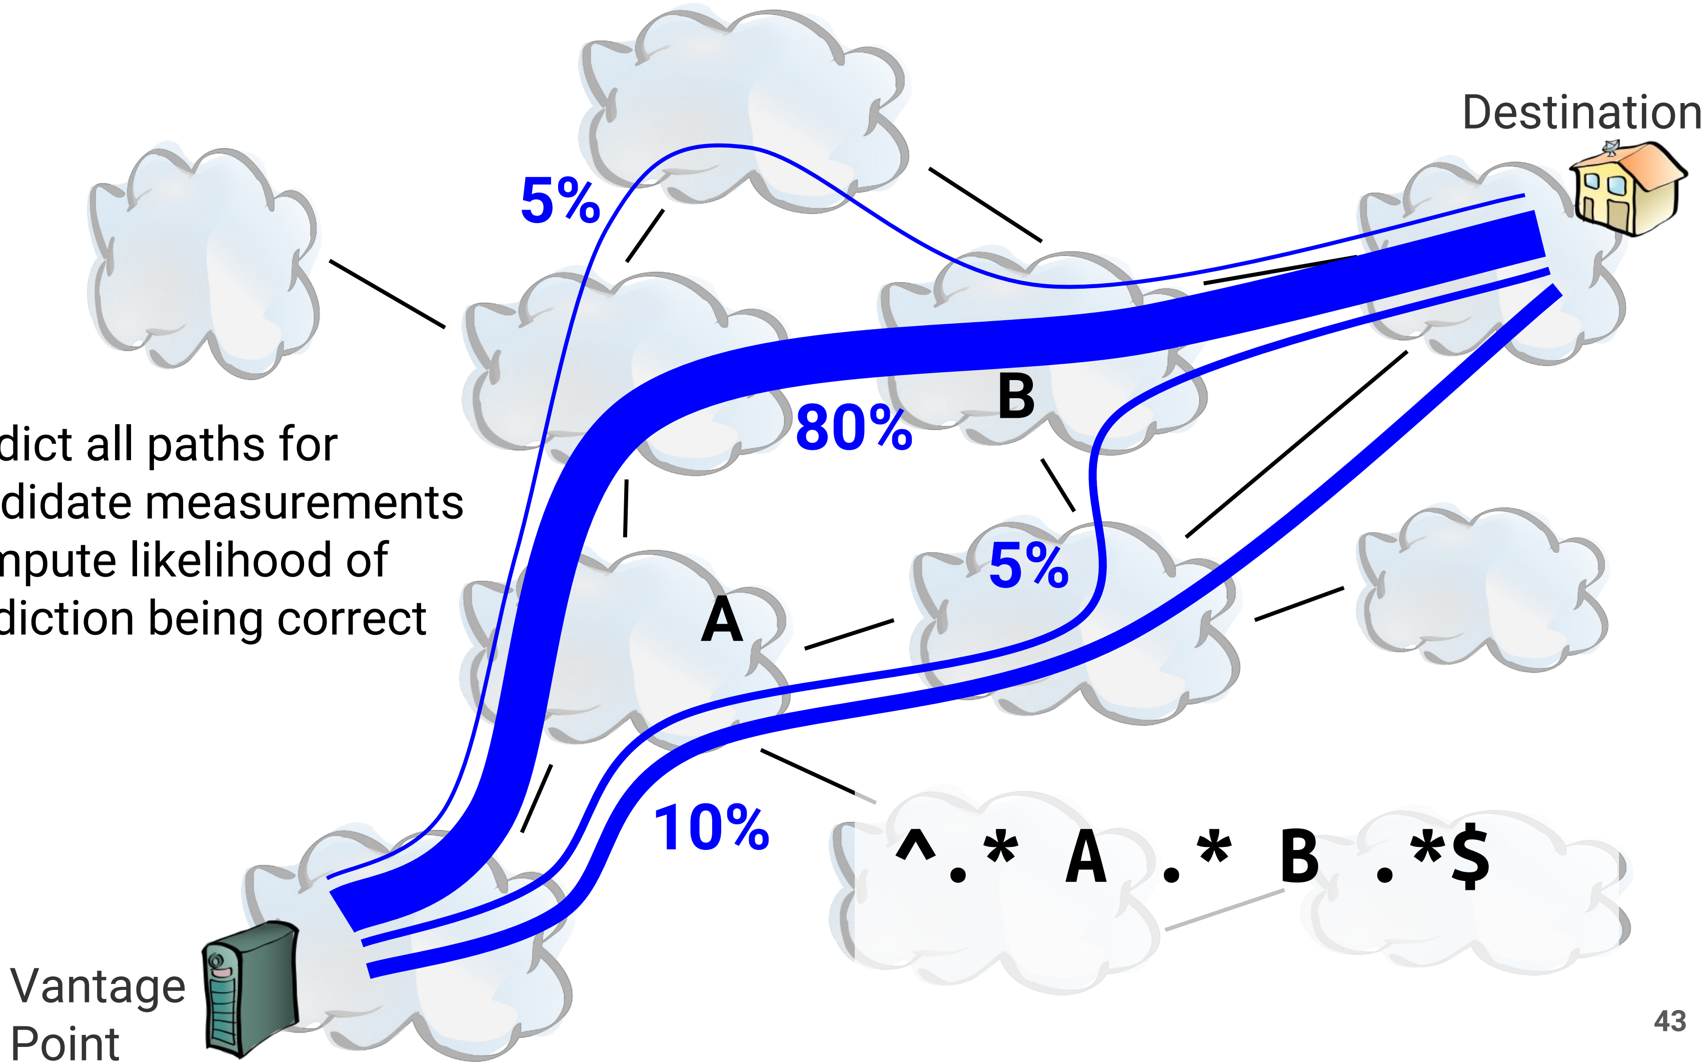

Computing Likelihood that Prediction is Correct

Input → set of predicted paths between source and destination

Output → likelihood of each prediction being correct

Approach → Compute route features and use machine learning

Computing Likelihood that Prediction is Correct

Input → set of predicted paths between source and destination

Output → likelihood of each prediction being correct

Approach → Compute route features and use machine learning

- Most important features include
  - AS peering relationships
  - Route length

1. Predict all paths for candidate measurements
2. Compute likelihood of prediction being correct
3. Compute likelihood actual path matches

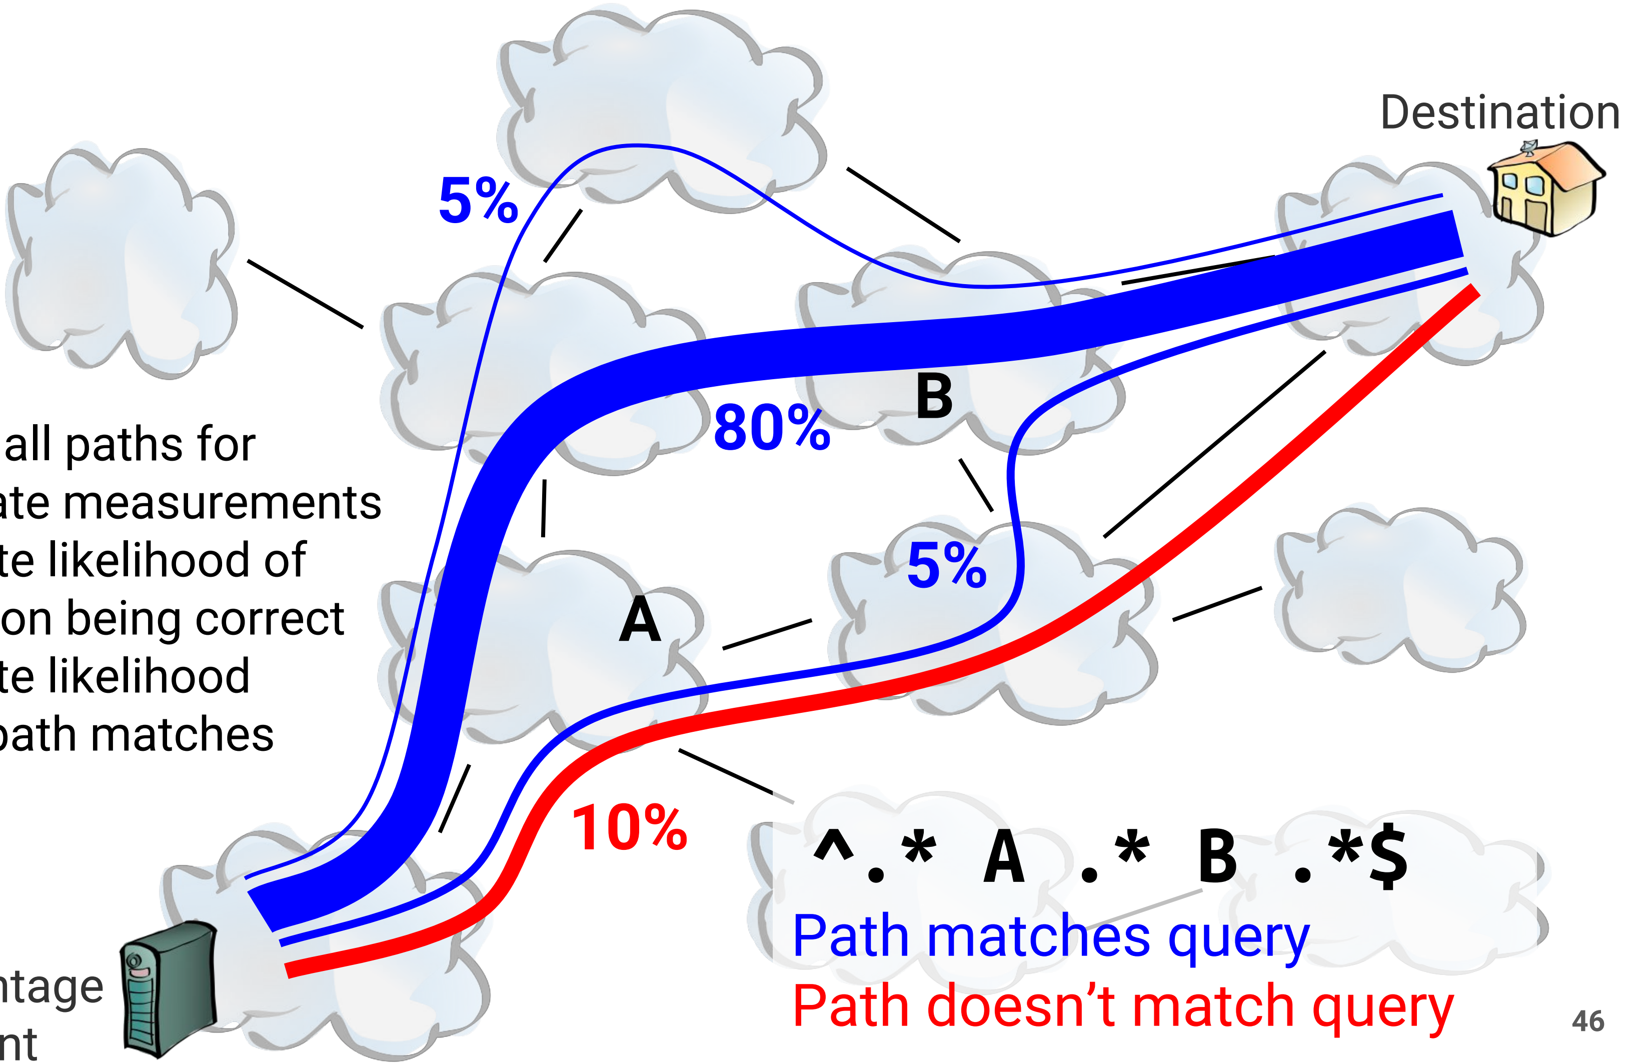

1. Predict all paths for candidate measurements
2. Compute likelihood of prediction being correct
3. Compute likelihood actual path matches

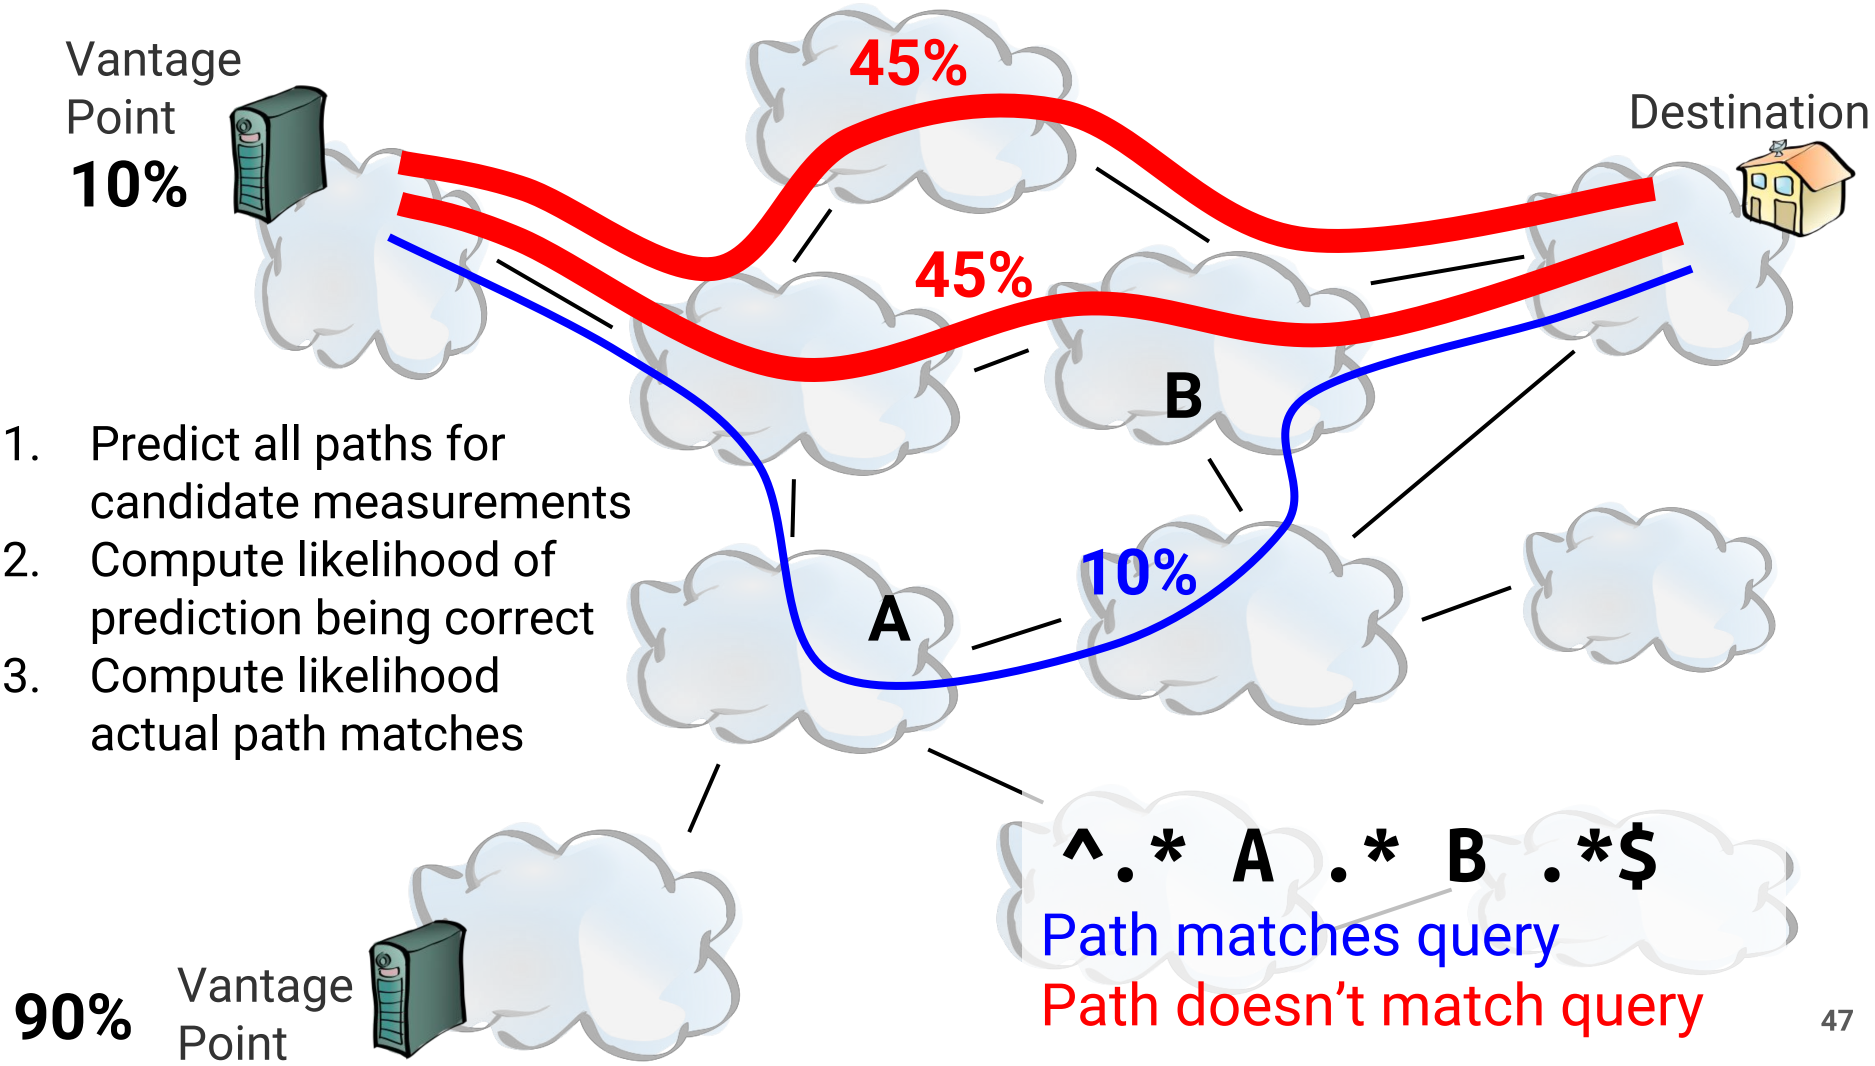

Match Likelihood Estimation Is Accurate

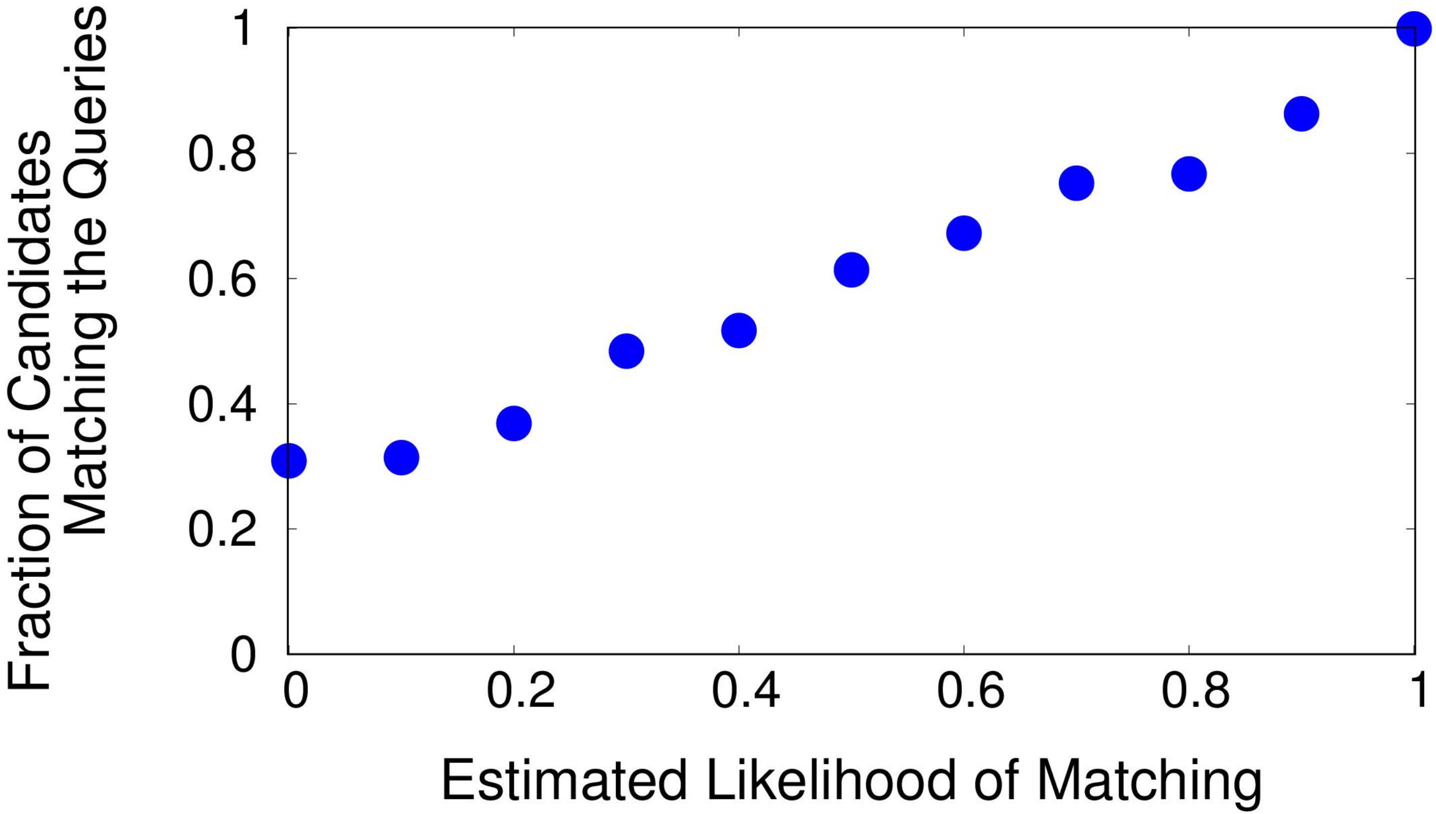

Match Likelihood Estimation Is Accurate

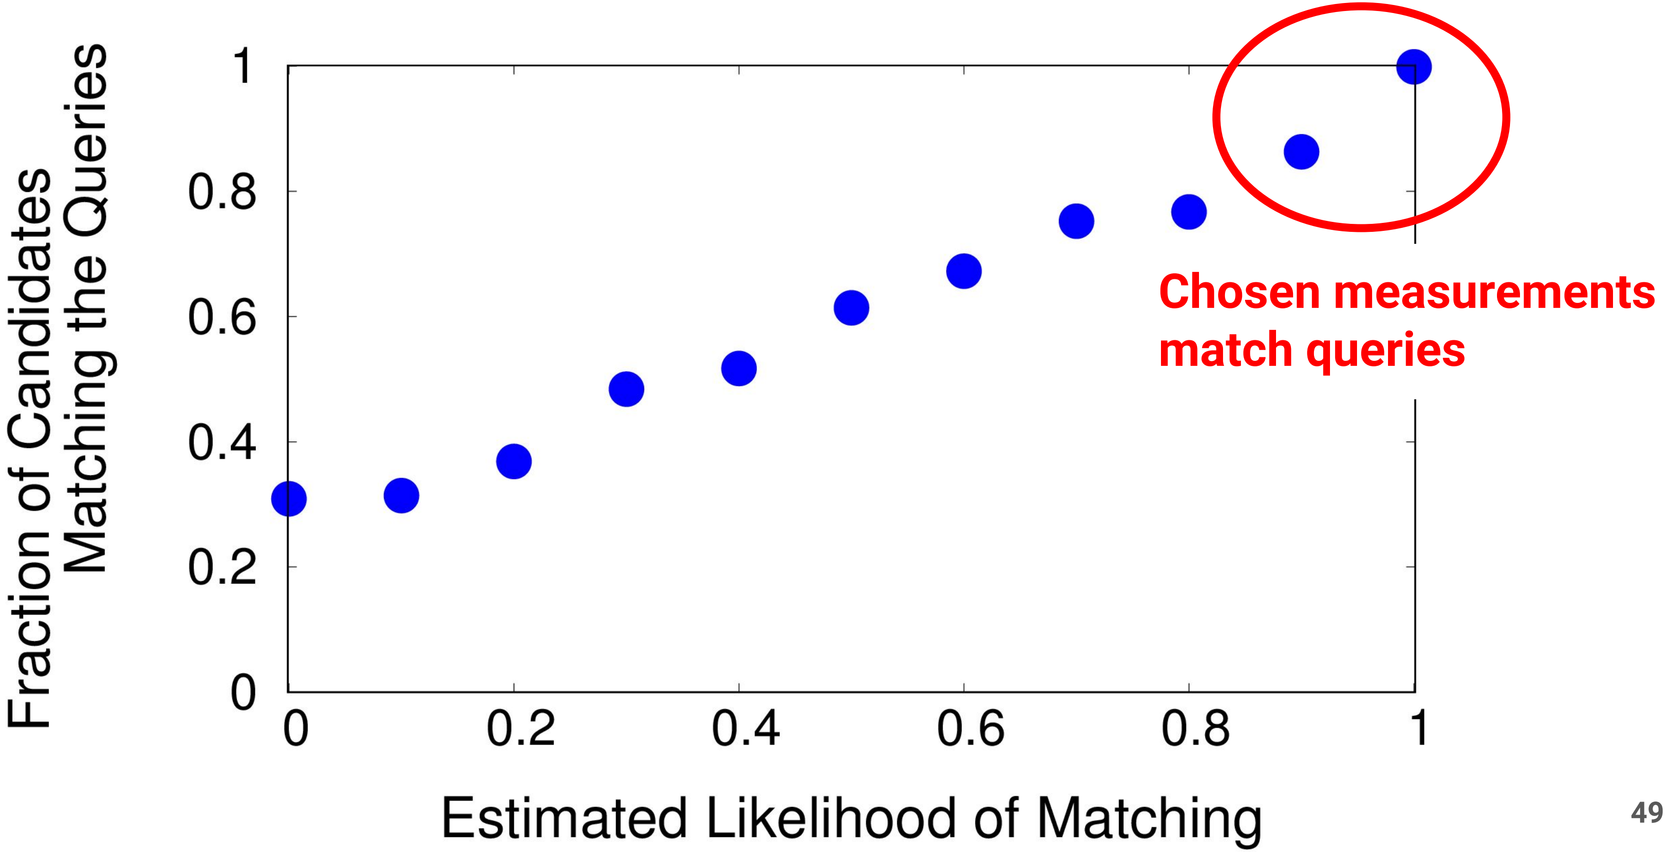

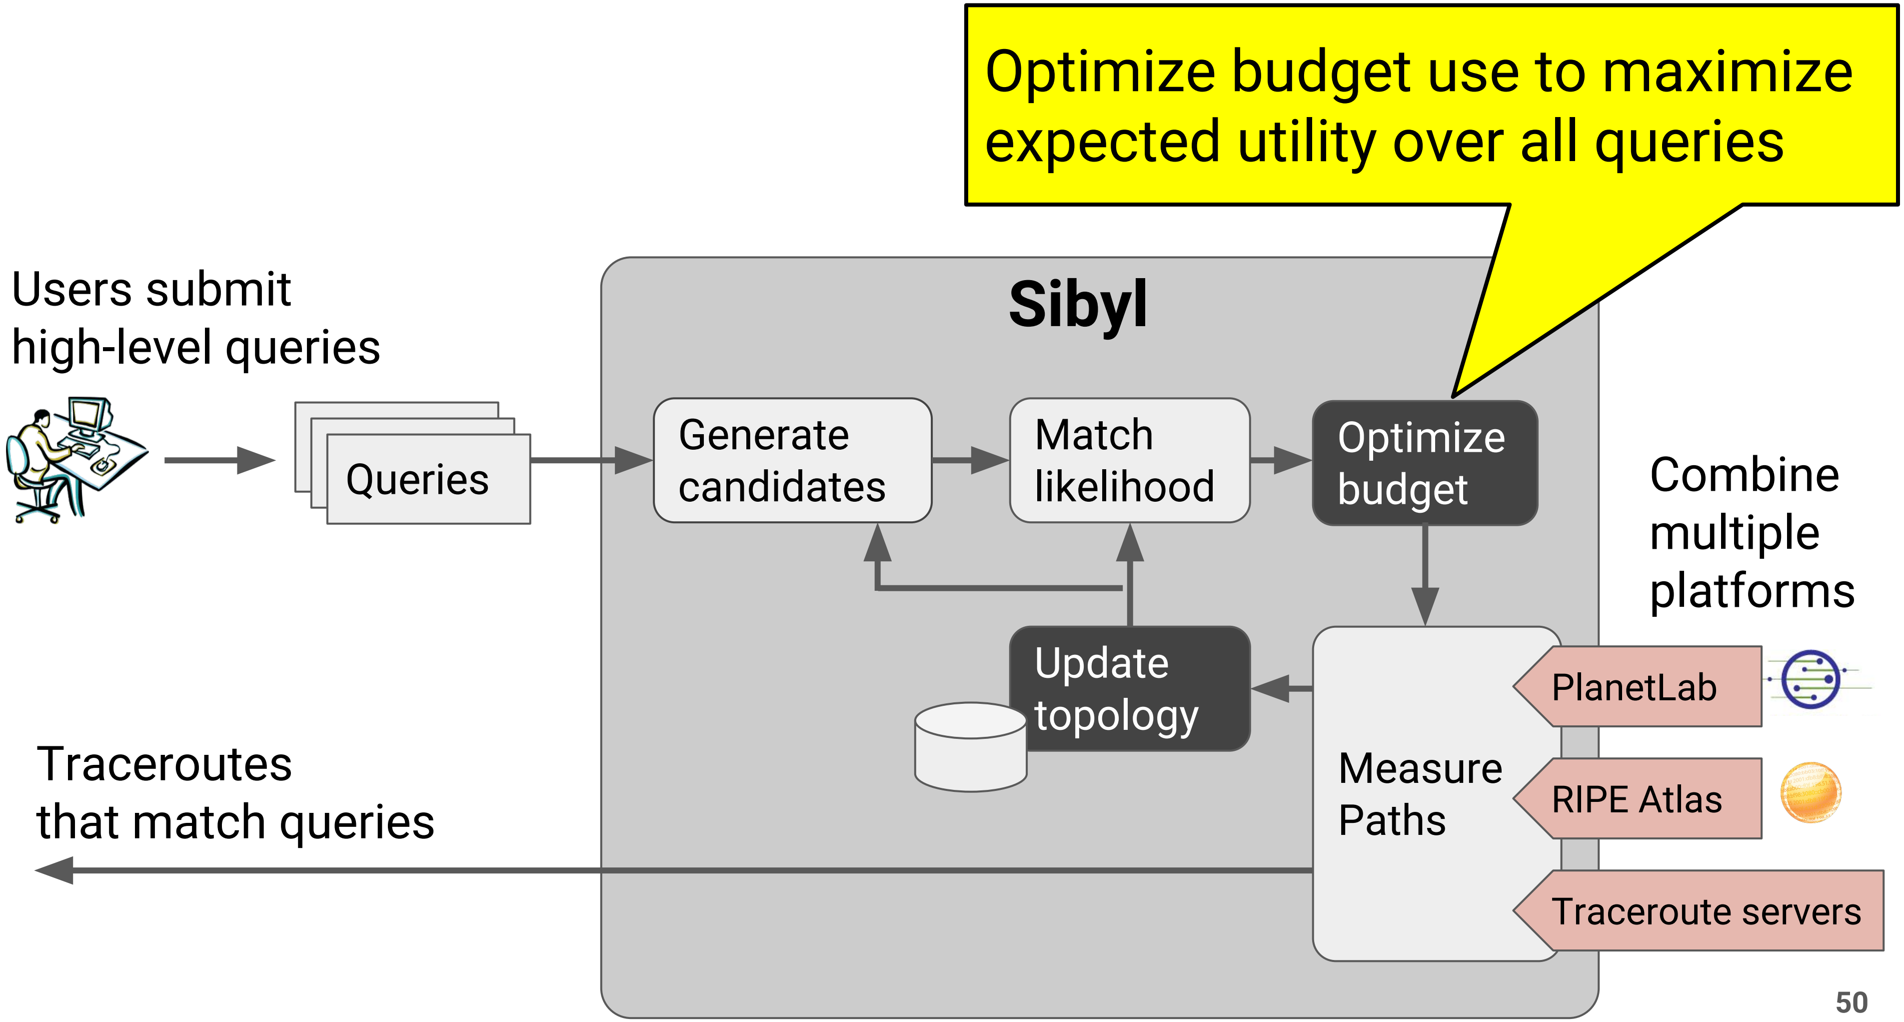

Optimize budget use to maximize expected utility over all queries
Users submit high-level queries
Queries
Sibyl
Generate candidates
Match likelihood
Optimize budget
Update topology
Measure Paths
PlanetLab
RIPE Atlas
Traceroute servers
Combine multiple platforms
Traceroutes that match queries

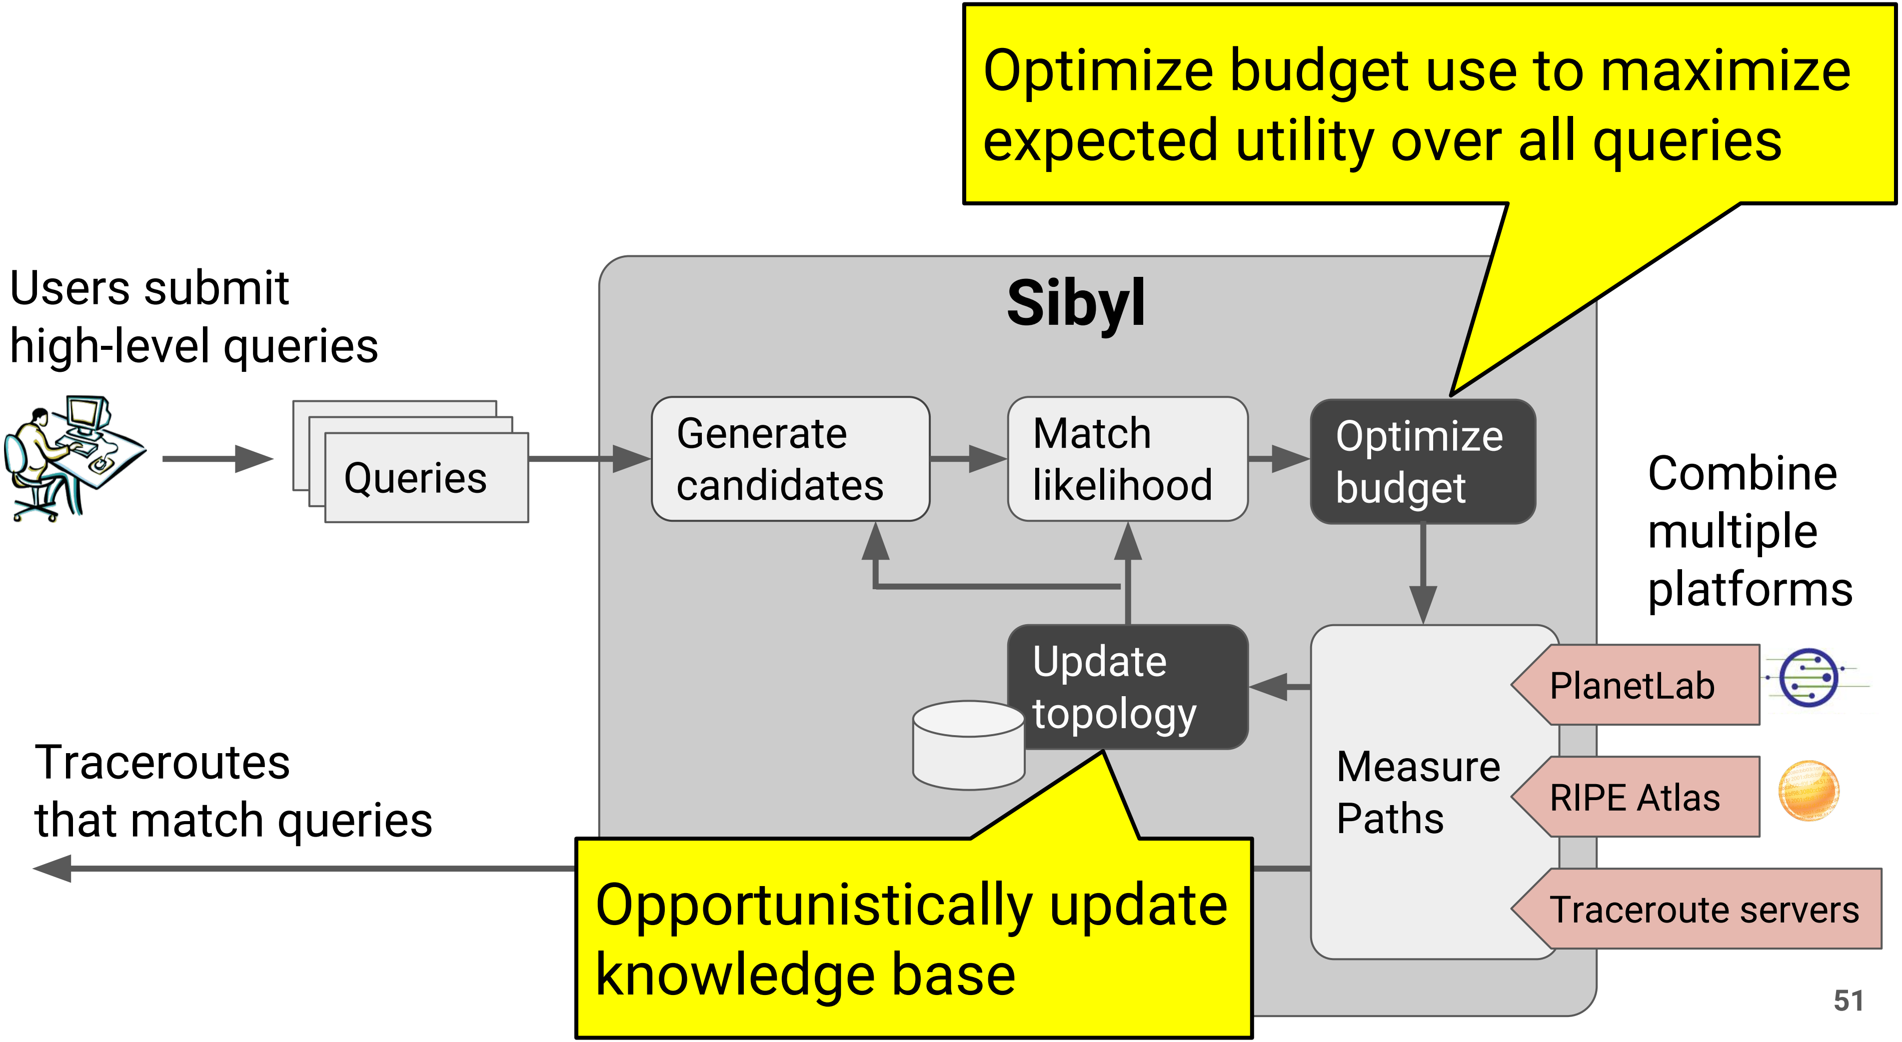
Optimize budget use to maximize expected utility over all queries
Users submit high-level queries
Sibyl
Queries
Generate candidates
Match likelihood
Optimize budget
Combine multiple platforms
Update topology
PlanetLab
Measure Paths
RIPE Atlas
Traceroutes that match queries
Traceroute servers
Opportunistically update knowledge base

Evaluation

Evaluated Sibyl over multiple rounds. Each round:

- System has fixed, limited probing budget to allocate
- Generate random (but satisfiable) queries
- Evaluate the fraction of queries Sibyl can satisfy

How Accurate Is Sibyl?

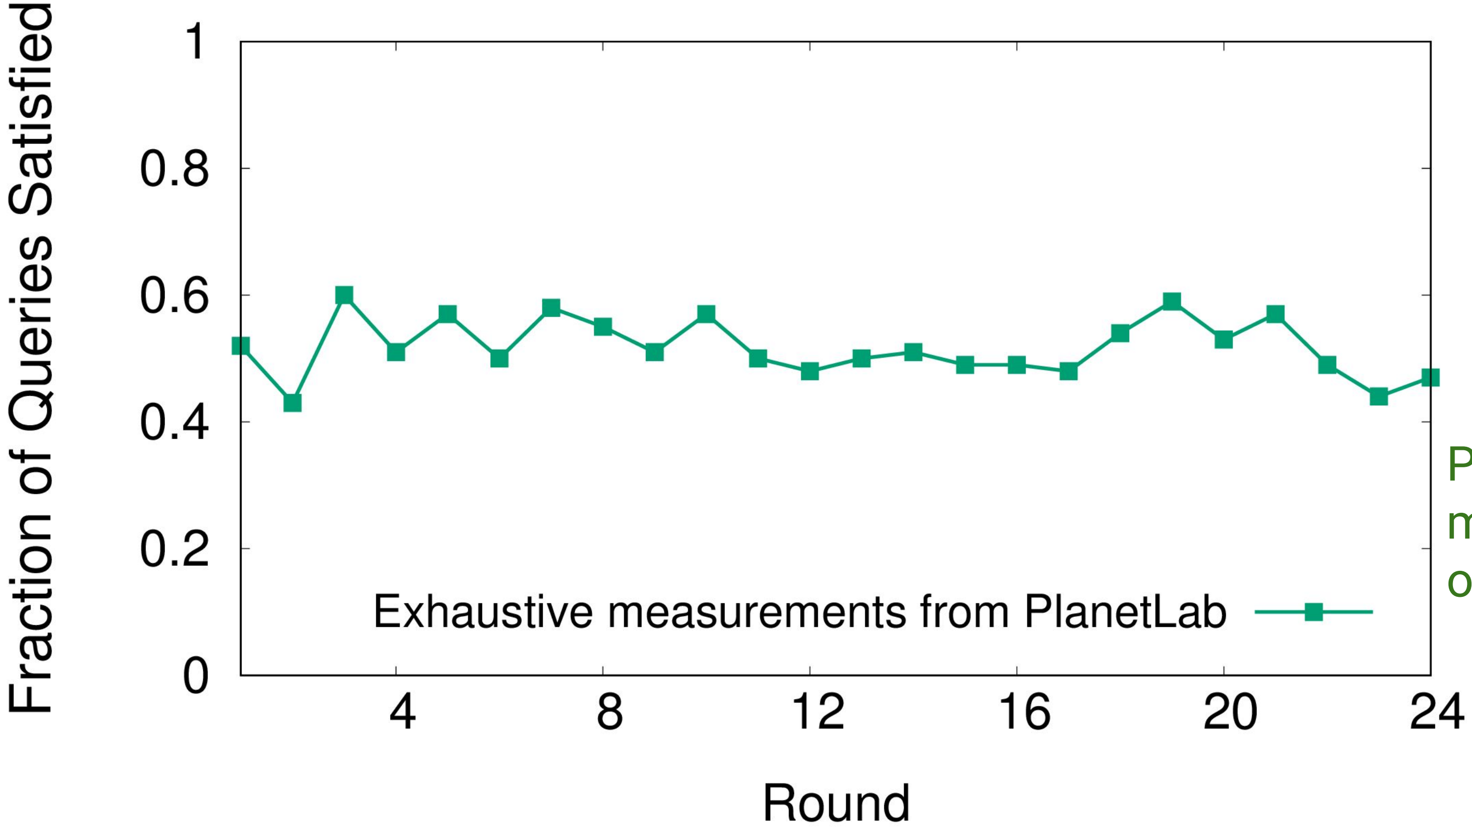

PlanetLab only matches 52% of queries

How Accurate Is Sibyl?

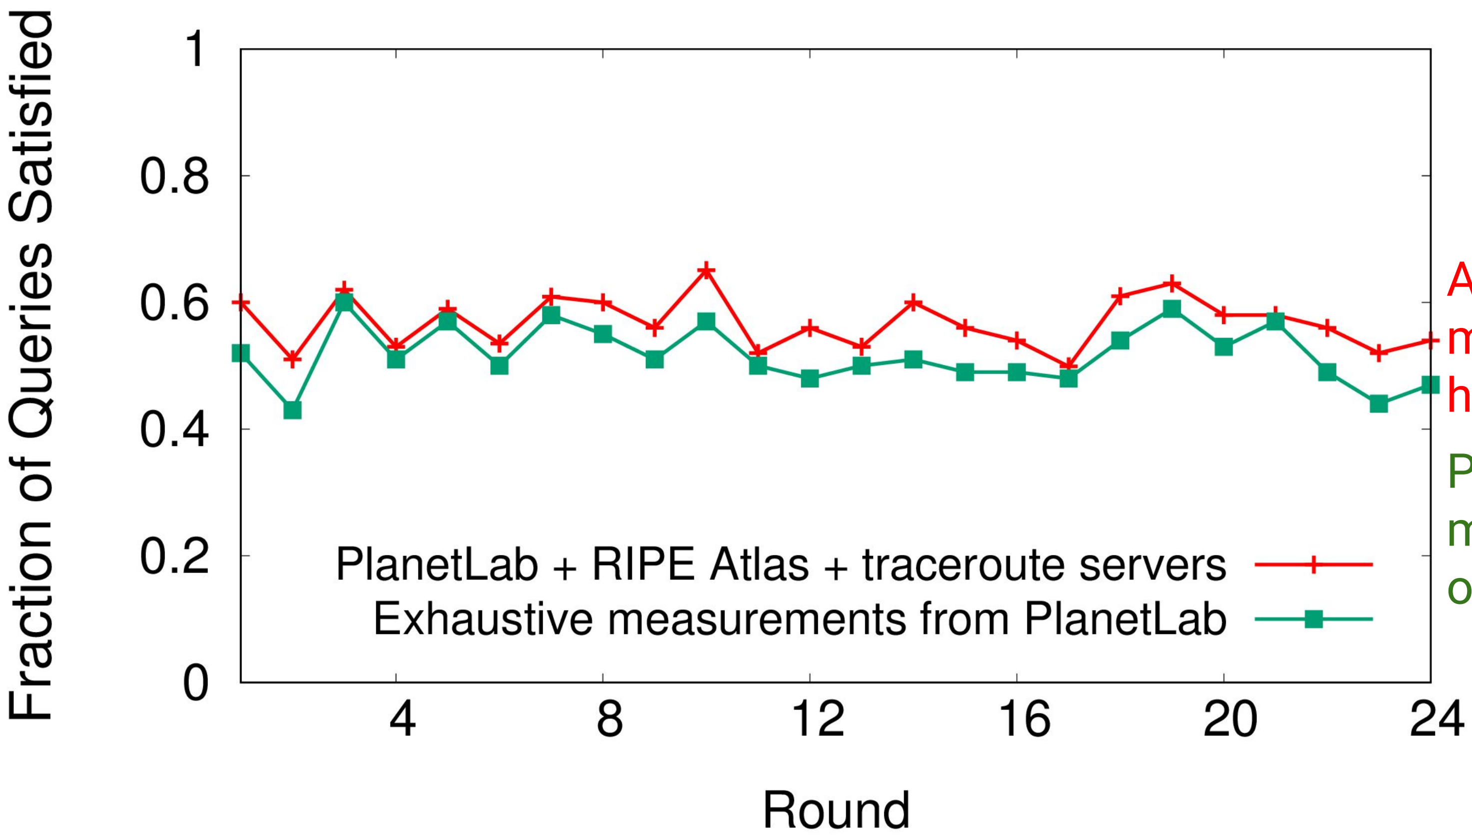

Adding random measurements helps little (56%)

PlanetLab only matches 52% of queries

How Accurate Is Sibyl?

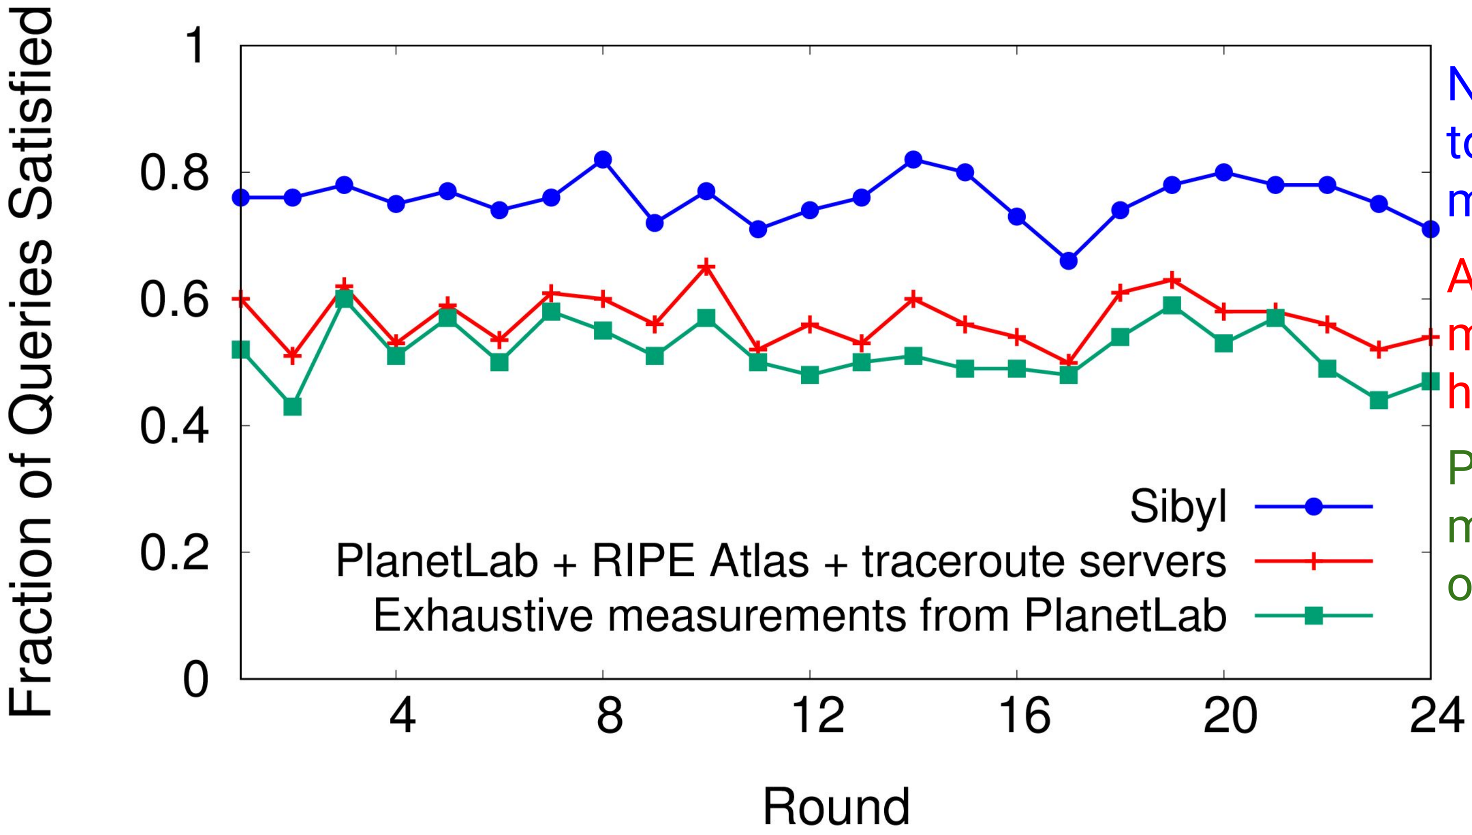

Need intelligence to effectively match (75%)

Adding random measurements helps little (56%)

PlanetLab only matches 52% of queries

Conclusions

- Supports high-level queries over Internet routes
- Combines platforms to improve coverage and budget
- Smartly chooses which measurements to issue
  - Overcome probing budget constraints

Future Work

- Improve path prediction and likelihood estimation
- Long-term budget allocation
  - Balance satisfying current queries vs benefit to serve future queries
  - Adapt probing rate as a function of query load
- Unify queries over historical and live data
  - “Give me paths that looked like X but now look like Y”

Query Language

We Have Complex Monitoring Systems

iSpy: diverse routes to a given prefix?

Reverse traceroute: routes through Level3 to monitoring node?

RocketFuel: routes through Verizon between Chicago and LA?

Regular Expressions Support Existing Systems

iSpy: diverse routes to a given prefix?

```
^ {•*} 184.164.224.0/19 $ by AS
```

Reverse traceroute: routes through Level3 to monitoring node?

```
^ •* Level3 .* USC $
```

RocketFuel: route through Verizon between Chicago and LA?

```
^ •* Verizon&Chicago •* Verizon&LA •* $
```

Internal Representation of Regular Expressions

Paths that go through Sprint's Chicago PoP on the way to Brazil:

```
^.*S.*U$ where S := Sprint&Chicago
               U := USC
```

From NANOG: Problem between Level3 in LA and GTT in Seattle:

```
^.*L.*G.*$ where L := Level3&LA
                 G := GTT&Seattle
```

Utility Optimization

Maximize Global Utility

$$\text{let } T = \bigcup_{p \in P} T_p$$

The union of the set of traceroutes chosen from each platform

Maximize Global Utility

$$\text{let } T = \bigcup_{p \in P} T_p$$

The union of the set of traceroutes chosen from each platform

$$\max_{T} \quad U\ (T)$$

Choose the set of traceroutes that maximizes utility over all queries

Maximize Global Utility

$$\text{let } T = \bigcup_{p \in P} T_p$$

The union of the set of traceroutes chosen from each platform

$$\max_T \sum_{q \in Q} U_q(T)$$

Choose the set of traceroutes that maximizes utility over all queries

Maximize Global Utility

$$\text{let } T = \bigcup_{p \in P} T_p$$

The union of the set of traceroutes chosen from each platform

$$\max_{T} \sum_{q \in Q} U_q(T)$$

Choose the set of traceroutes that maximizes utility over all queries

$$\text{subject to } |T_p| \leq B_p \ \forall p \in P$$

Subject to budget constraints of each platform

Utility functions

Existence queries → find one path that matches a query

$$\mathbb{E}\left[f_q(T)\right] = 1 - \prod_{t \in T} [1 - \mathbb{P}(t \in q)]$$

Utility functions

Existence queries → find one path that matches a query

$$\mathbb{E}\left[f_q(T)\right] = 1 - \prod_{t \in T}\left[1 - \mathbb{P}(t \in q)\right]$$

Diversity queries → find all different paths that match a part of the query

$$\mathbb{E}\left[f_q(T)\right] = \quad 1 - \prod_{t \in T}\left[1 - \mathbb{P}(t \in q \quad)\right]$$

Utility functions

Existence queries $\rightarrow$ find one path that matches a query

$$\mathbb{E}\left[f_q(T)\right] = 1 - \prod_{t \in T} [1 - \mathbb{P}(t \in q)]$$

Diversity queries $\rightarrow$ find all different paths that match a part of the query

$$\mathbb{E}\left[f_q(T)\right] = \sum_{h} \left\{ 1 - \prod_{t \in T} [1 - \mathbb{P}(t \in q \wedge h \in t)] \right\}$$

RuleFit and Likelihood Estimation

Most Important Features

Spliced Path Feature	Importance
1. PoP-level similarity with the other paths	1
2. PoP-level path length inflation vs iPlane's top-ranked path	.90
3. Total number of PoP splice points	.60
4. Total number of AS splice points	.55
5. AS splice point type	.52
6. AS splice point relationship with neighbors	.49
7. Number of PoPs in iPlane's top-ranked path	.44
Other features	$\leq .34$

RuleFit Jaccard Index Predictions Correlate

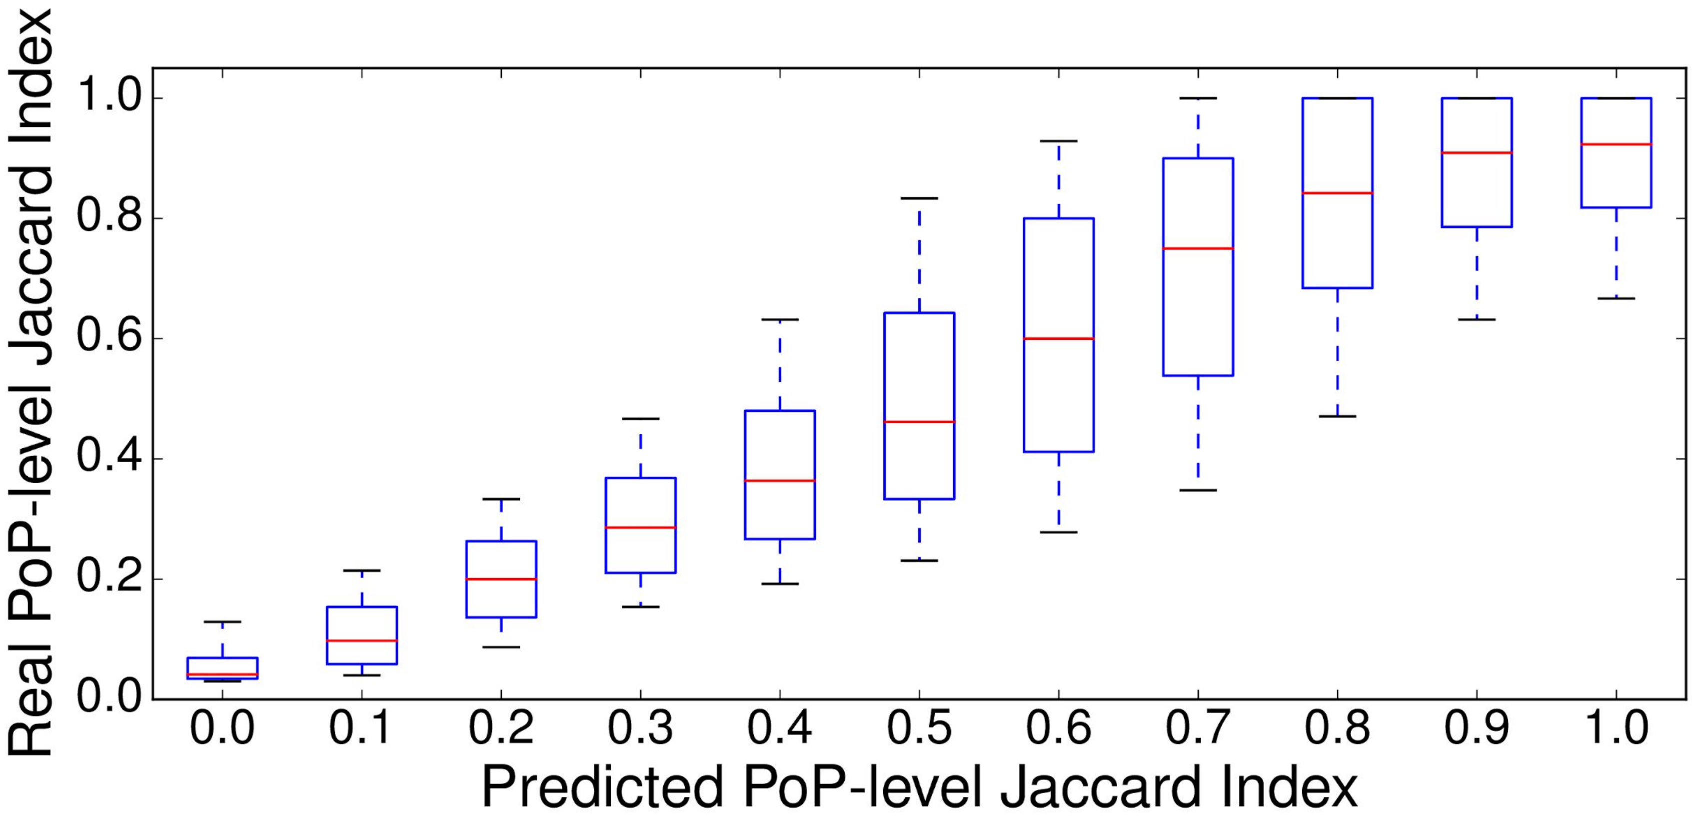

SFSA Splicing

Query FSA Splicing to Generate Candidates

^.* A .* B .*$

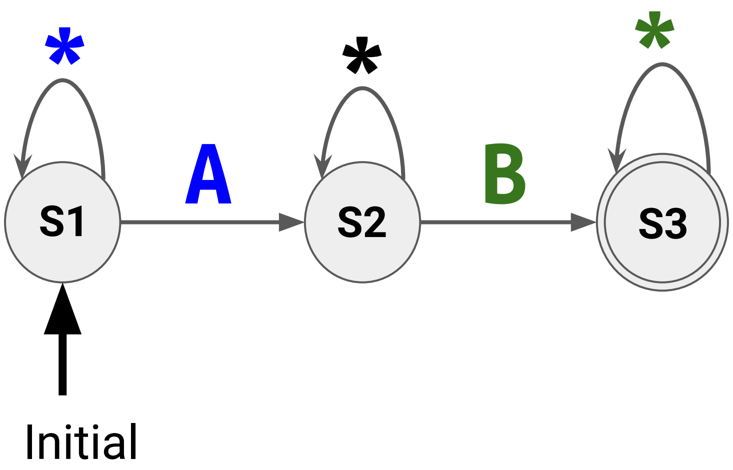

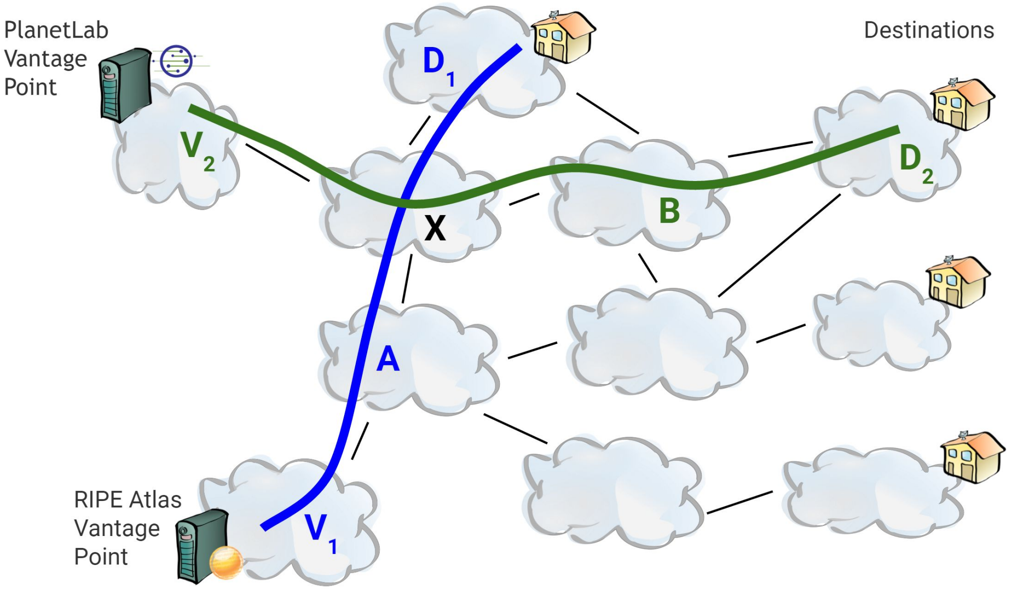

Query FSA Splicing to Generate Candidates

`^.*` `A` `.*` `B` `.*$`

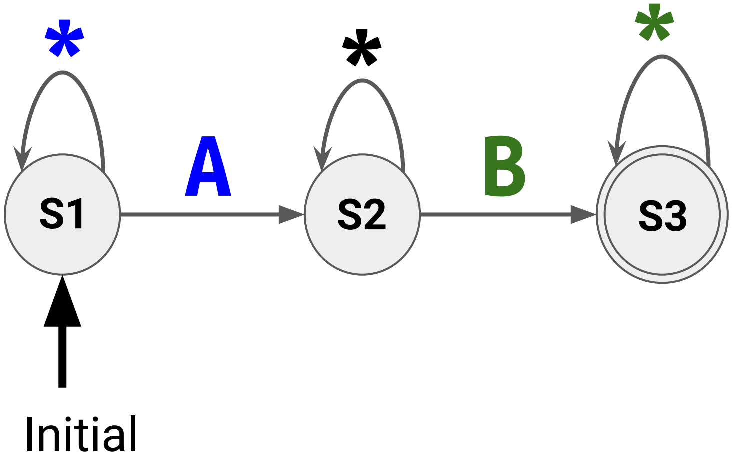

Hop	Transition	Hop	Transition
V1	$S1 \xrightarrow{*} S1$	V2	
A	$S1 \xrightarrow{A} S2$	X	
X	$S2 \xrightarrow{*} S2$	B	
D1	$S2 \xrightarrow{*} S2$	D2	

Query FSA Splicing to Generate Candidates

^.* A .* B .*$

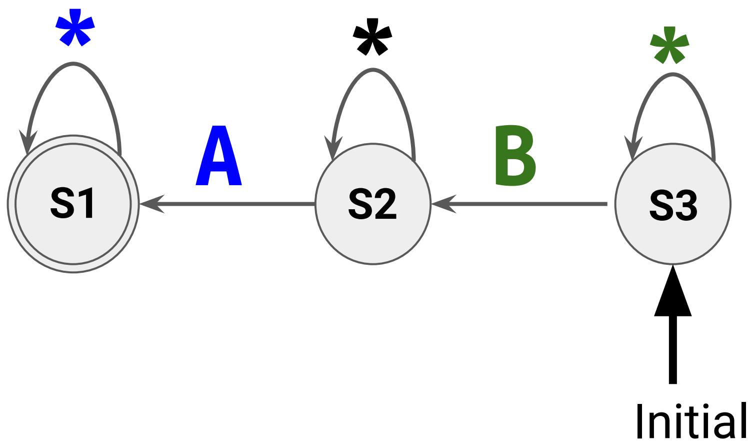

Hop	Transition	Hop	Transition
V1	$S1 \xrightarrow{*} S1$	V2	$S2 \xrightarrow{*} S2$
A	$S1 \xrightarrow{A} S2$	X	$S2 \xrightarrow{*} S2$
X	$S2 \xrightarrow{*} S2$	B	$S3 \xrightarrow{B} S2$
D1	$S2 \xrightarrow{*} S2$	D2	$S3 \xrightarrow{*} S3$

Query FSA Splicing to Generate Candidates

`^.* A .* B .*$`

`V1 A X B D2`

Hop	Transition	Hop	Transition
V1	$S1 \xrightarrow{*} S1$	V2	$S2 \xrightarrow{*} S2$
A	$S1 \xrightarrow{A} S2$	X	$S2 \xrightarrow{*} S2$
X	$S2 \xrightarrow{*} S2$	B	$S3 \xrightarrow{B} S2$
D1	$S2 \xrightarrow{*} S2$	D2	$S3 \xrightarrow{*} S3$

Evaluation

Queries Types Used in Evaluation

- Traverse AS toward destination
  - `^ .* A .* B $`
  - `^ [^A]+ A+ .* B $`
- Traverse link toward destination
  - `^ .* A B .* D $`
- Traverse three different locations in sequence
  - `^ .* A .* B .* C .* $`

How Much Better Can We Do?

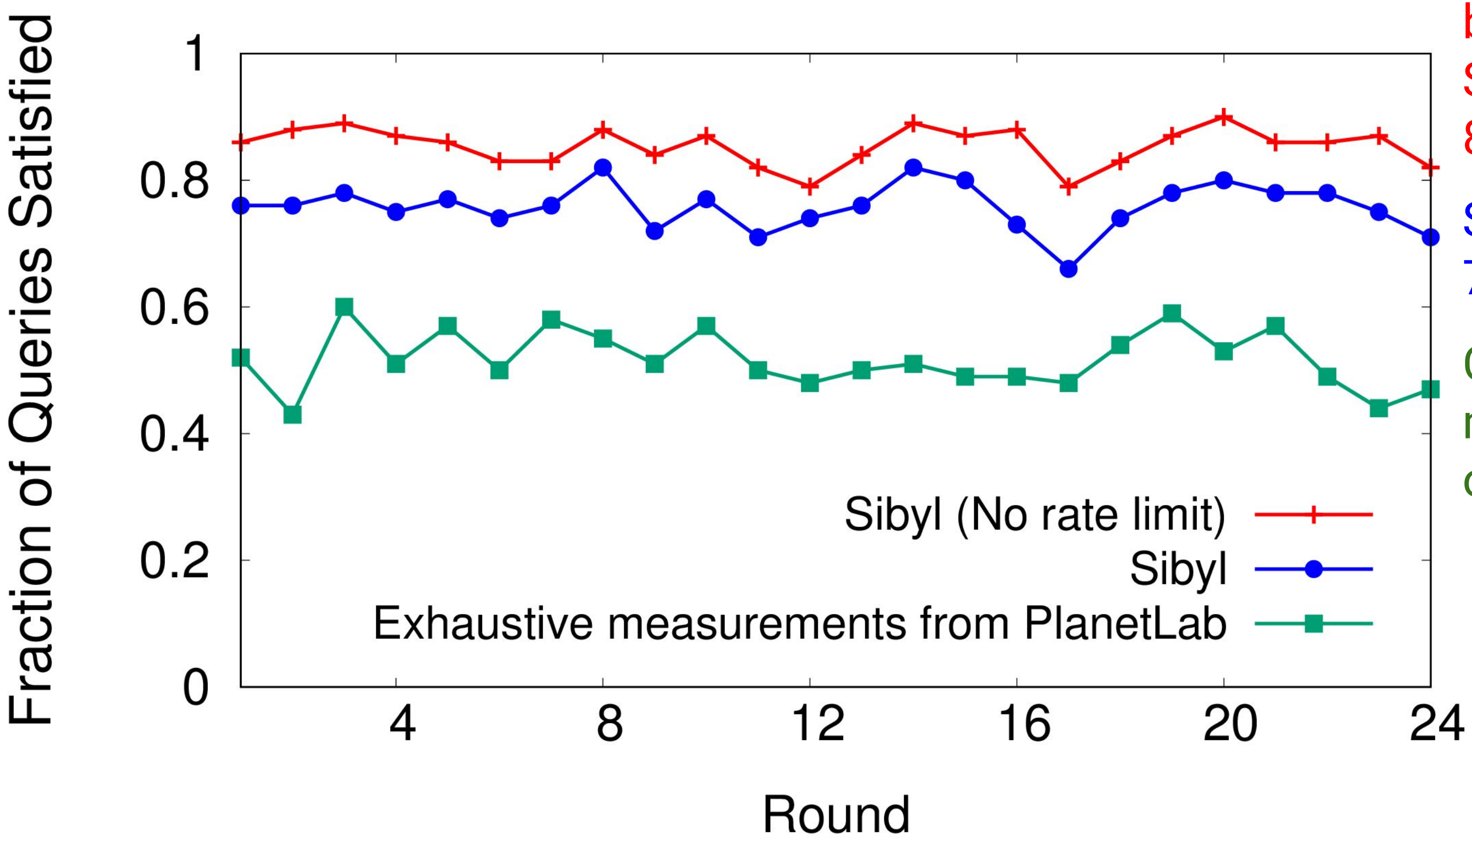

Unlimited budget would Sibyl to match 88% of queries

Sibyl matches 75% of queries

Current approach matches 52% of queries

How Much Does Each Module Contribute

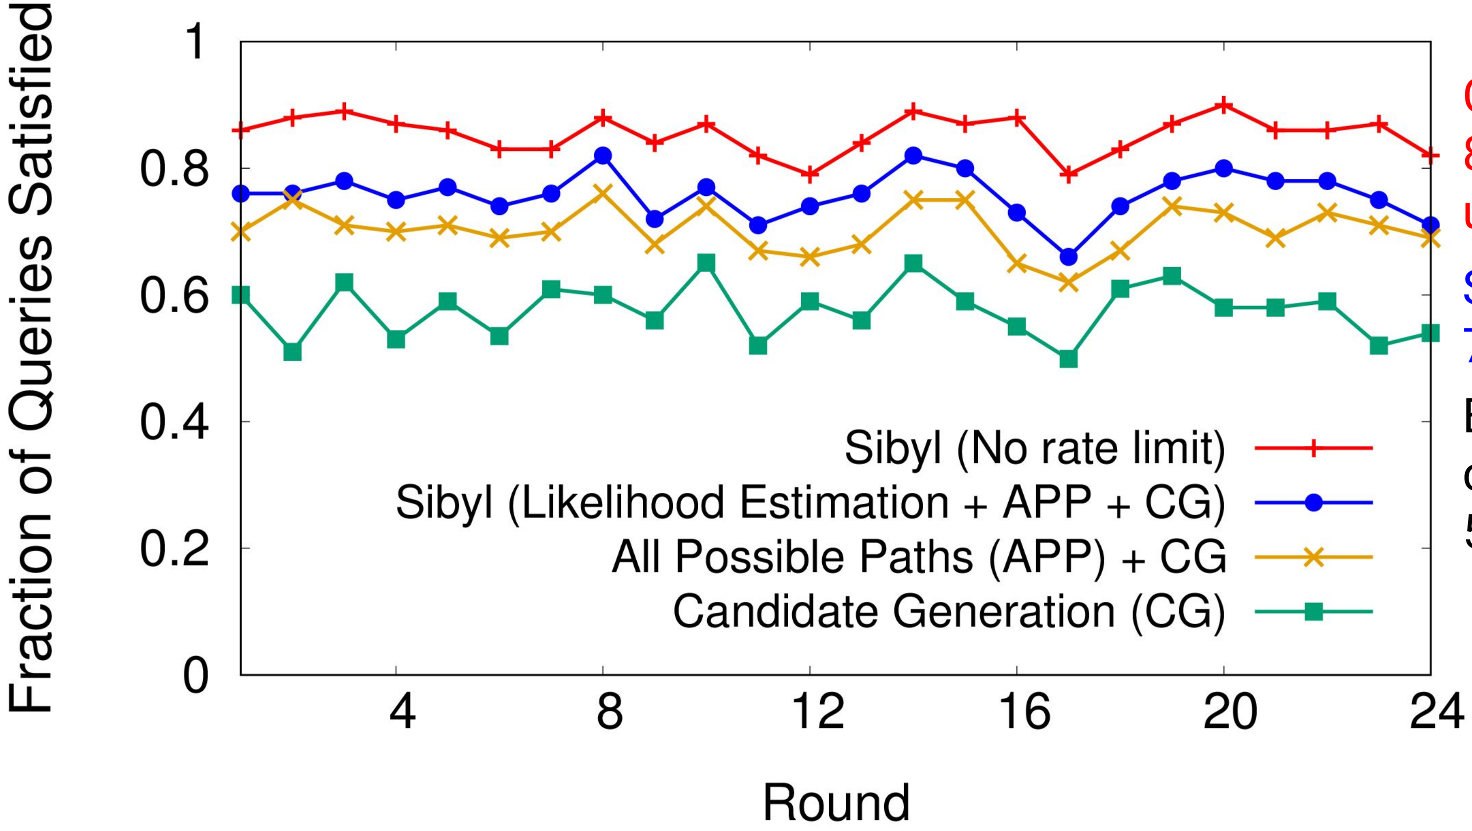

Could satisfy 84% with unlimited budget

Sibyl satisfies 75% of queries

Each module contributes 5-8% accuracy

Sibyl Is Effective at Different Granularities

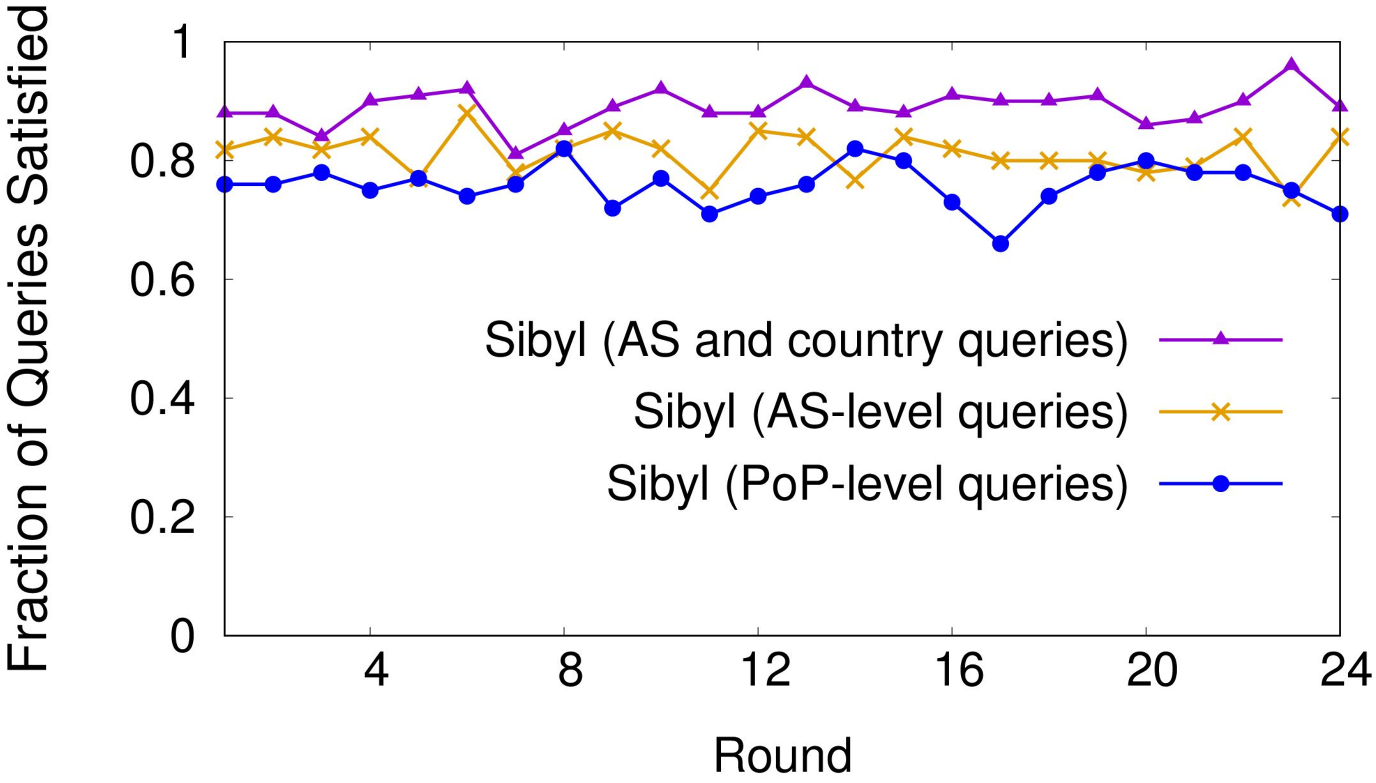

Combining Platforms Helps Satisfy More Queries

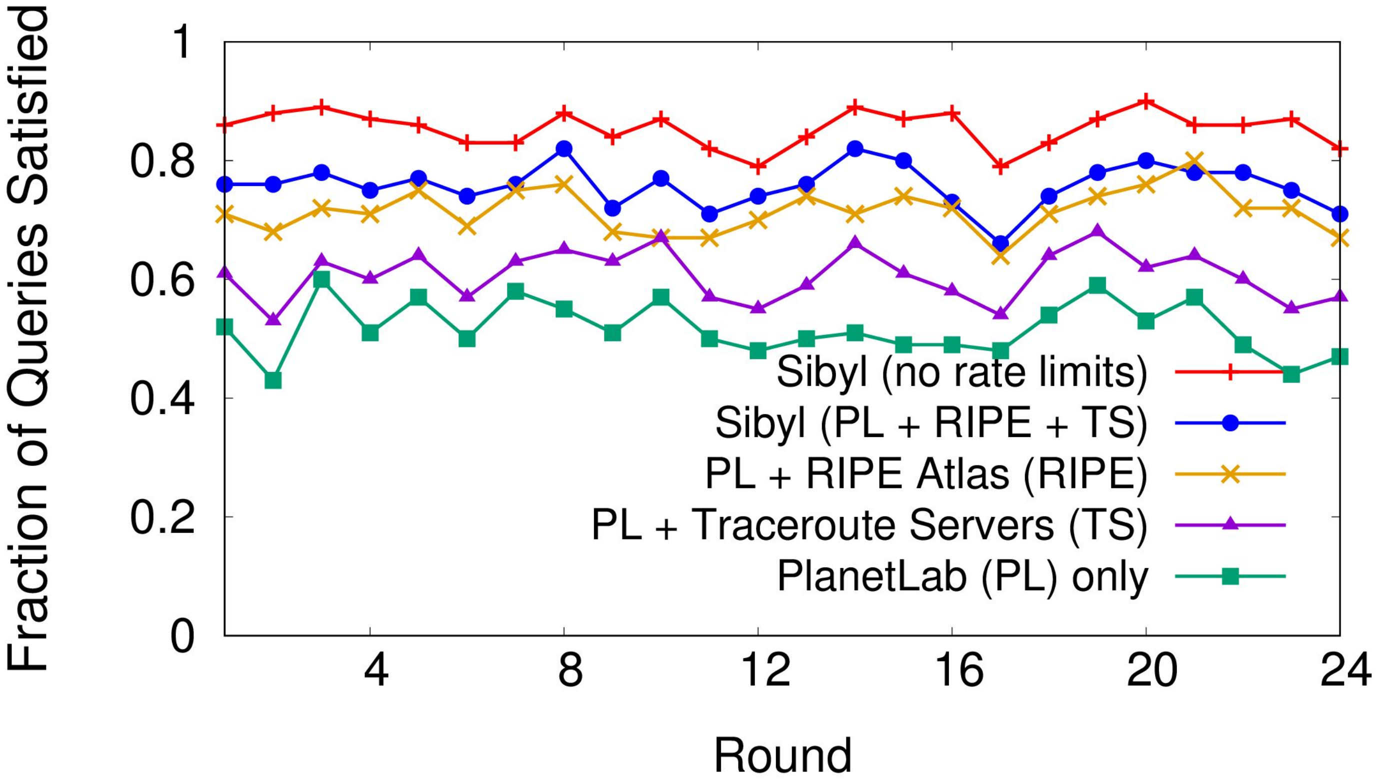

Sibyl Uncovers More Path Diversity

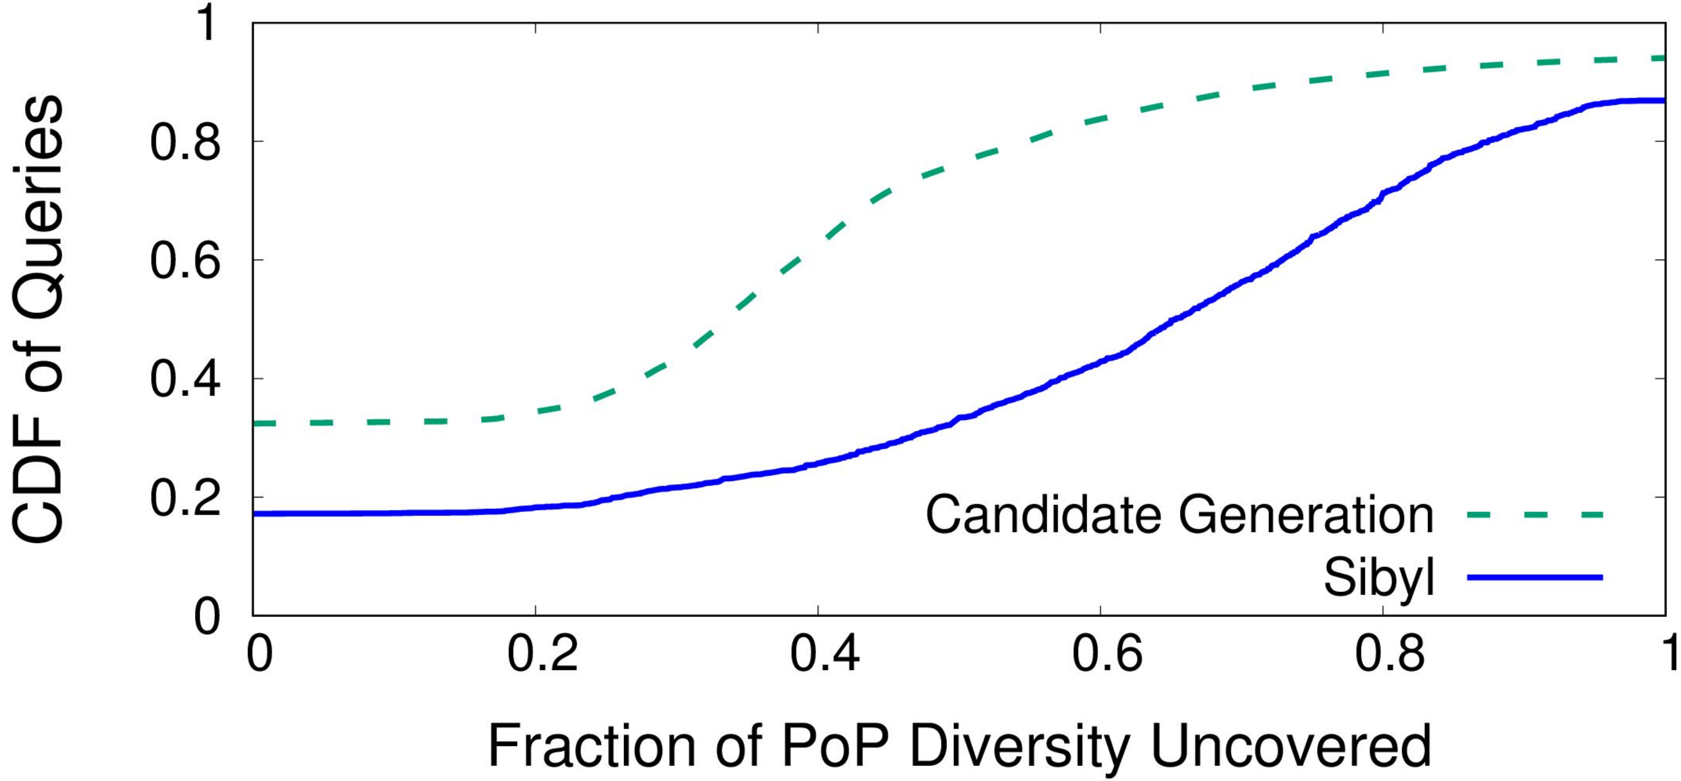

Staleness

Patching and Pruning Stale Measurements

New measurements may detect path changes

- Whenever we detect a change toward a destination, update all paths to that destination that overlap
- Whenever we detect a path change from a source, update all paths from that source that overlap

Staleness Has Small Impact

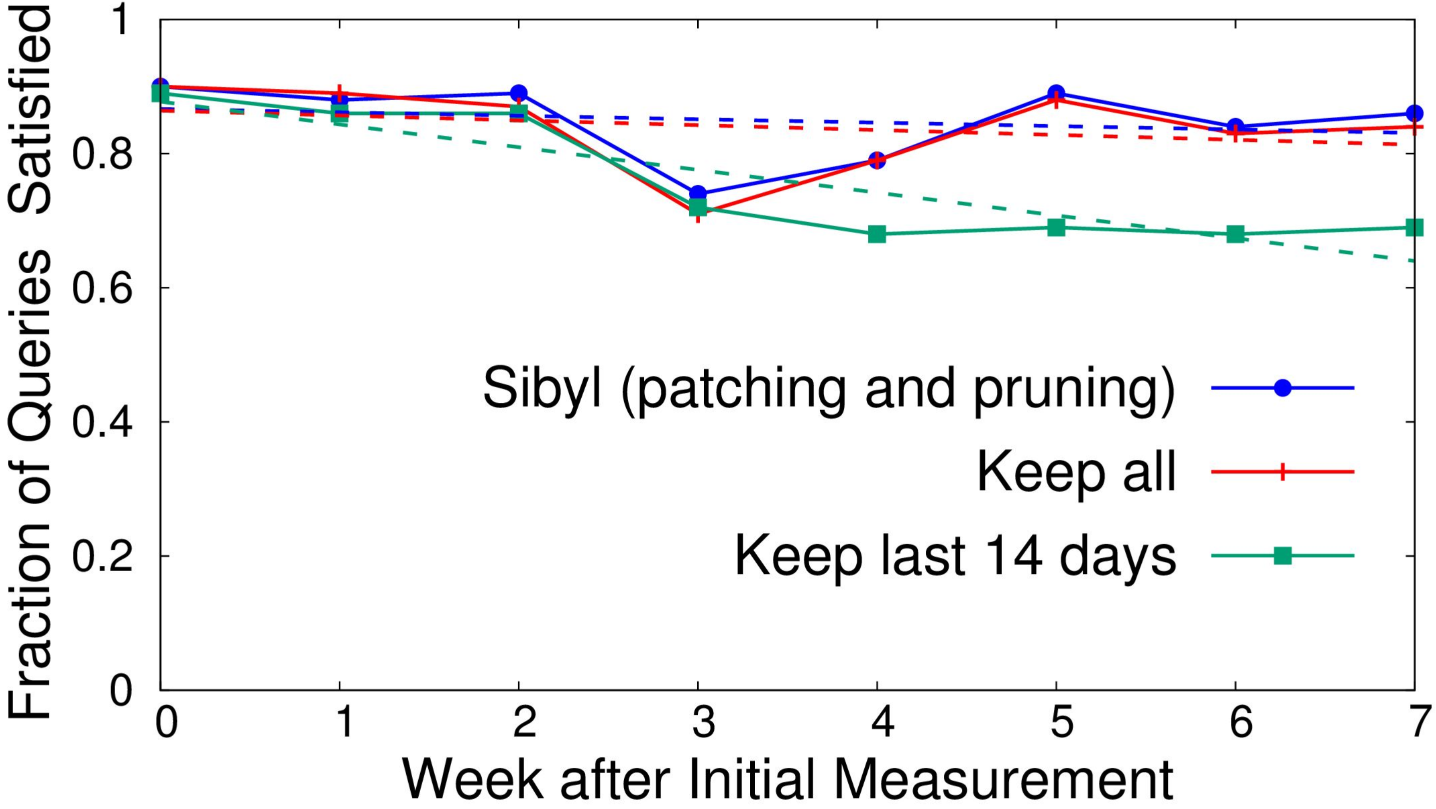

Path Change Properties

Uphill Path Changes Are Less Likely