DisCarte: A Disjunctive Internet Cartographer

Rob Sherwood, Adam Bender,
Neil Spring

University of Maryland

Introduction

- Network maps are useful
- Existing maps are small, stale, or inaccurate
- Mapping well requires:
  - lots of measurements (lots of error)
  - many techniques (lots of disagreement)

Contributions

- Show how to use Record Route
- DisCarte
  - Cross-validates measurements to reduce error
- Introduce disjunctive logic programming (DLP)
  - Develop divide and conquer scheme to scale DLP to Internet-sized topologies
- Validate DisCarte against public topologies

Contributions

- Show how to use Record Route
- DisCarte
  - Cross-validates measurements to reduce error
- Introduce disjunctive logic programming (DLP)
  - Develop divide and conquer scheme to scale DLP to Internet-sized topologies
- Validate DisCarte against public topologies

Topology Inference

Topology Inference

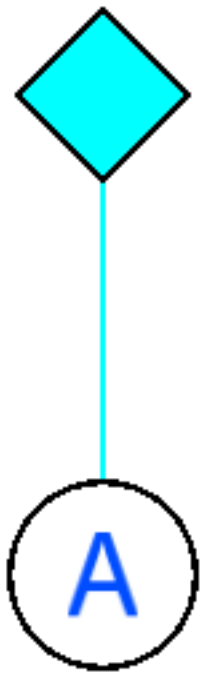

Topology Inference

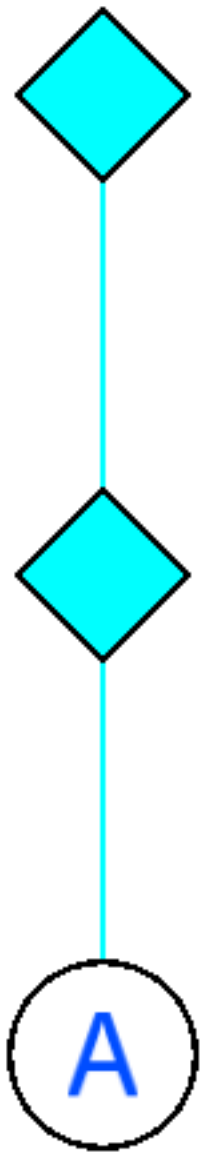

Topology Inference

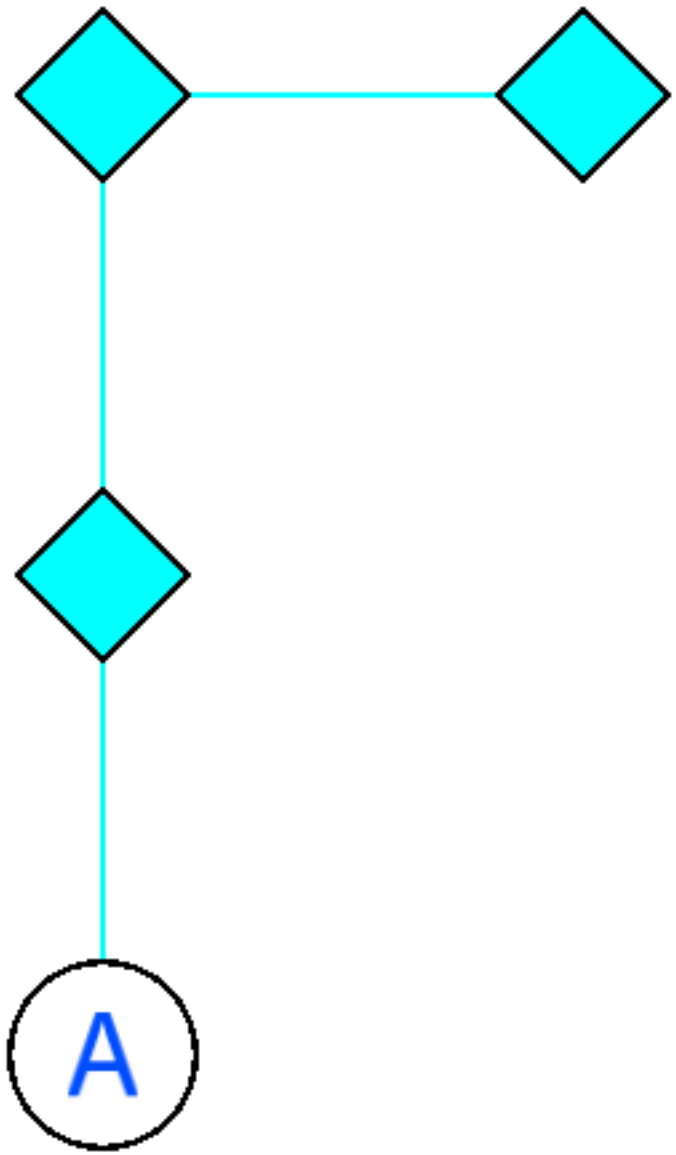

Topology Inference

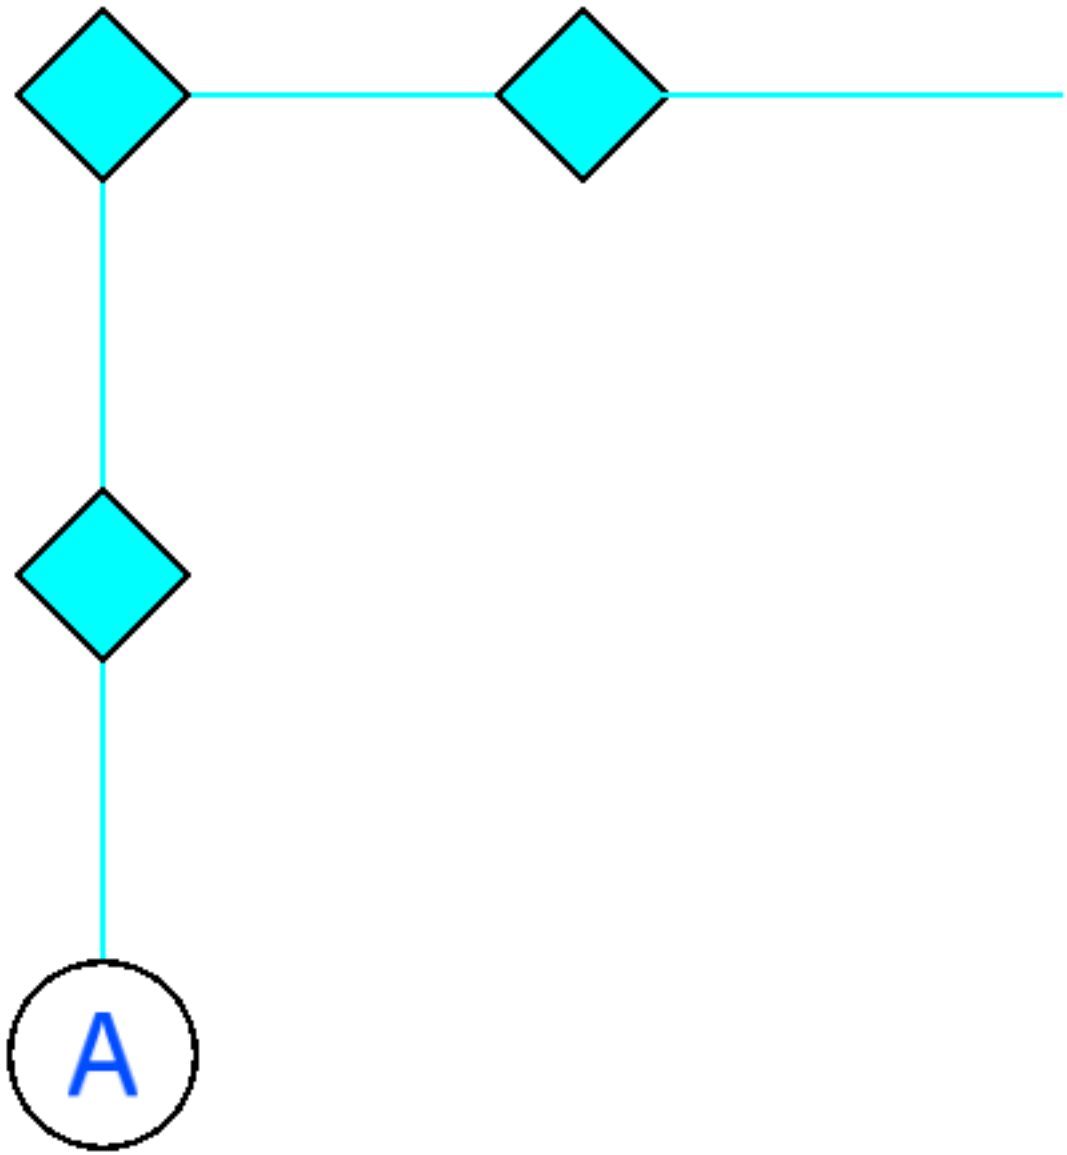

Topology Inference

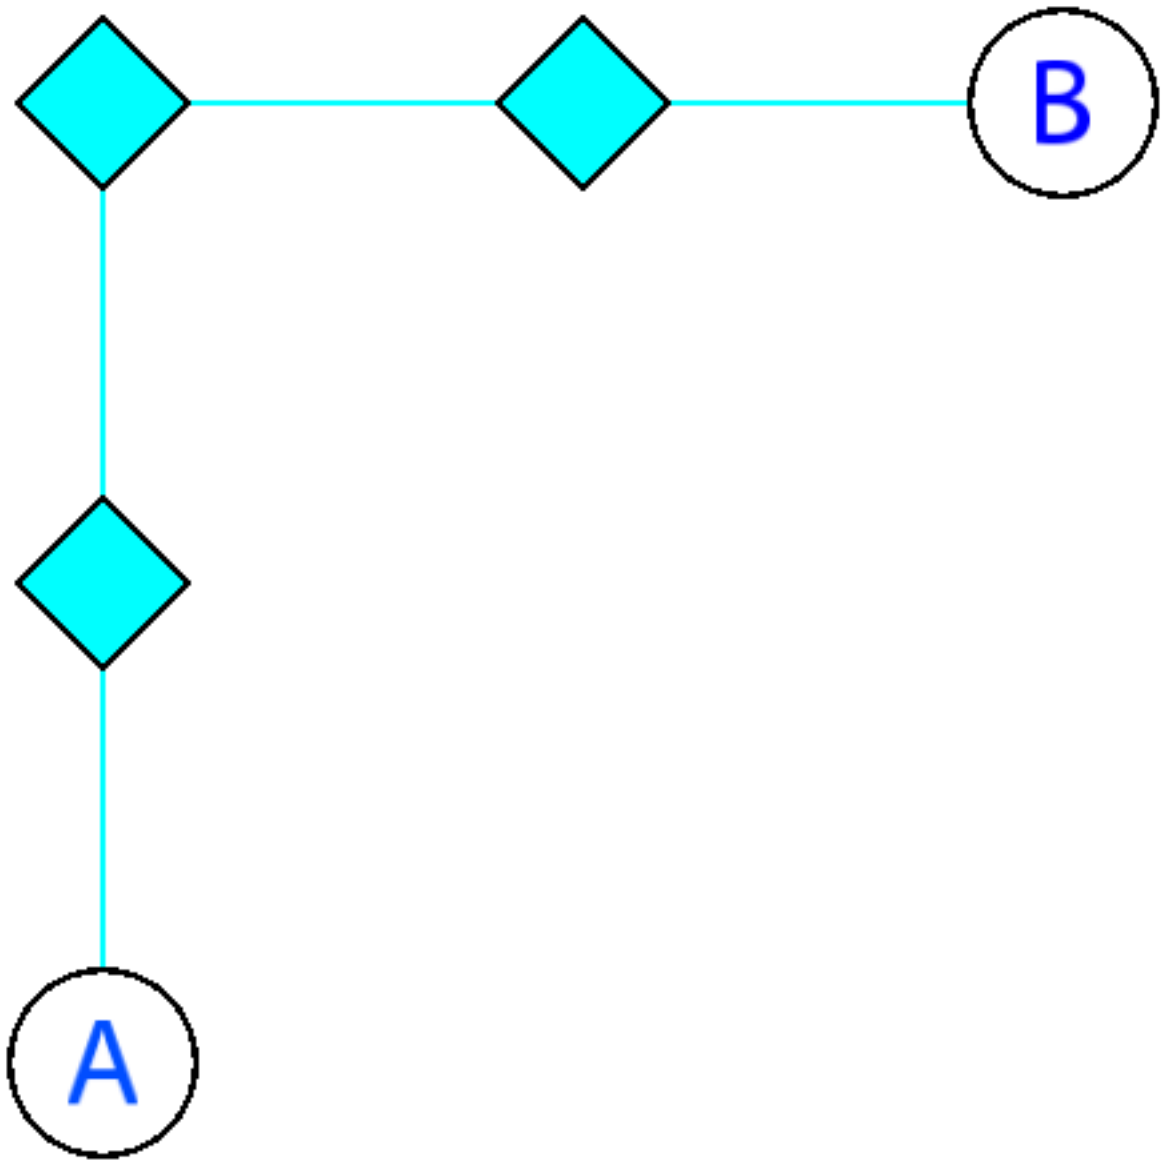

Topology Inference

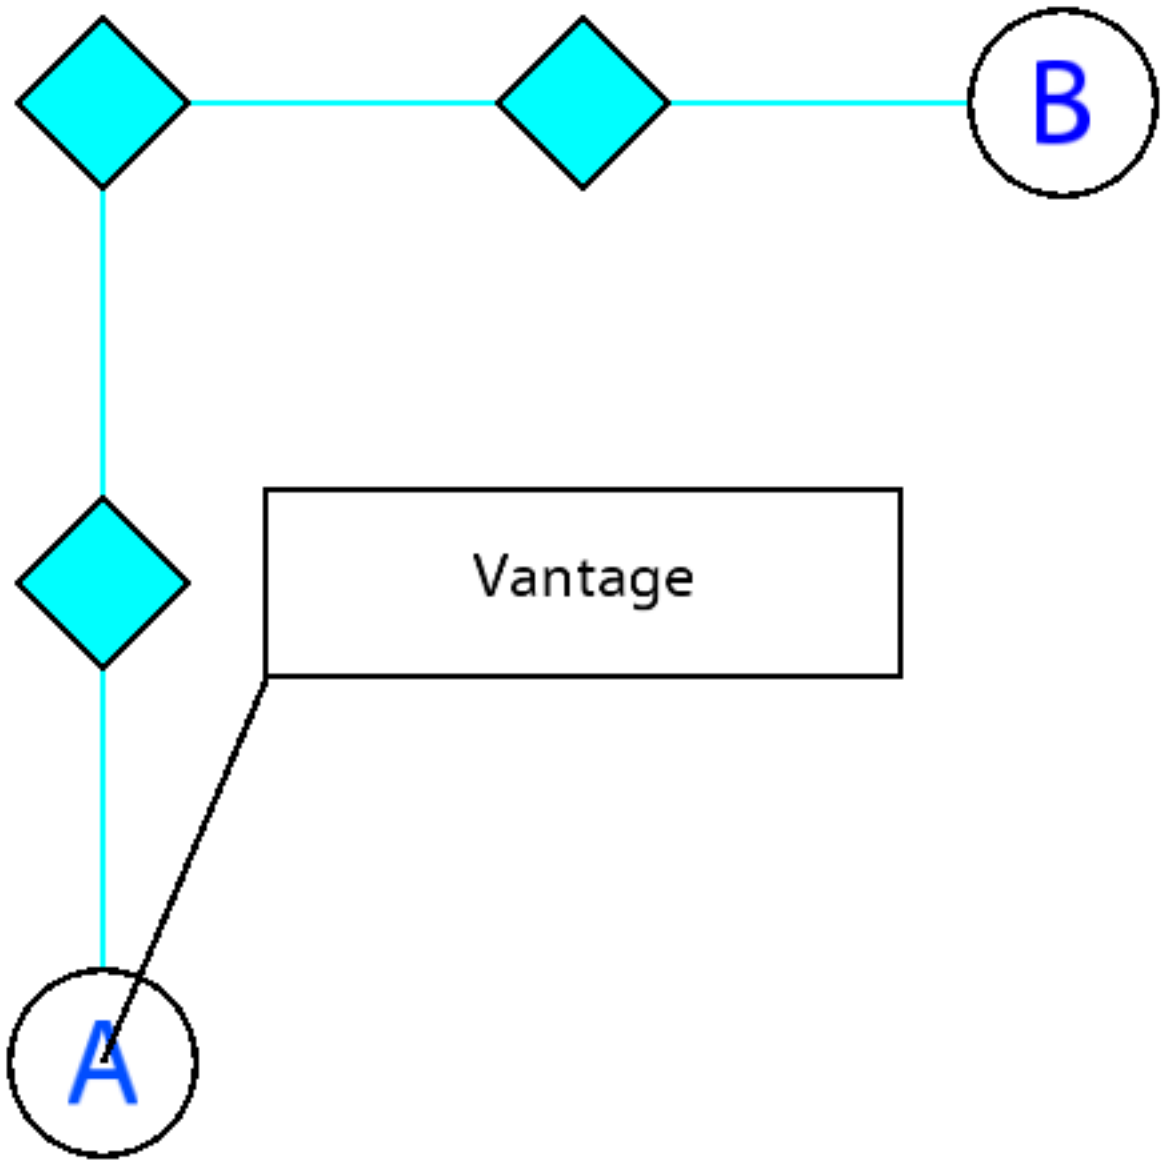

Topology Inference

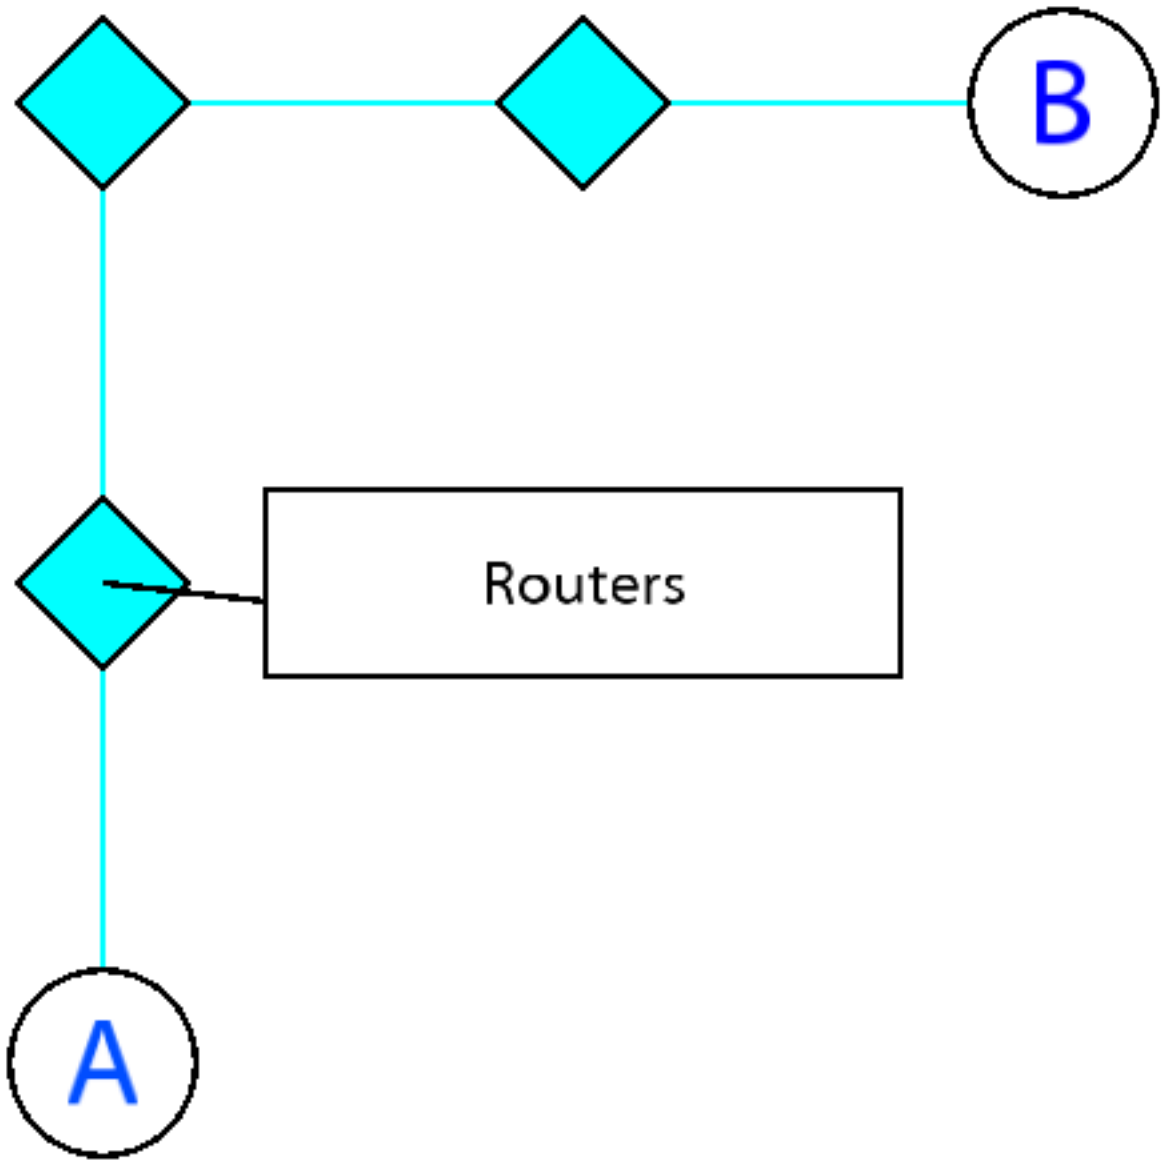

Topology Inference

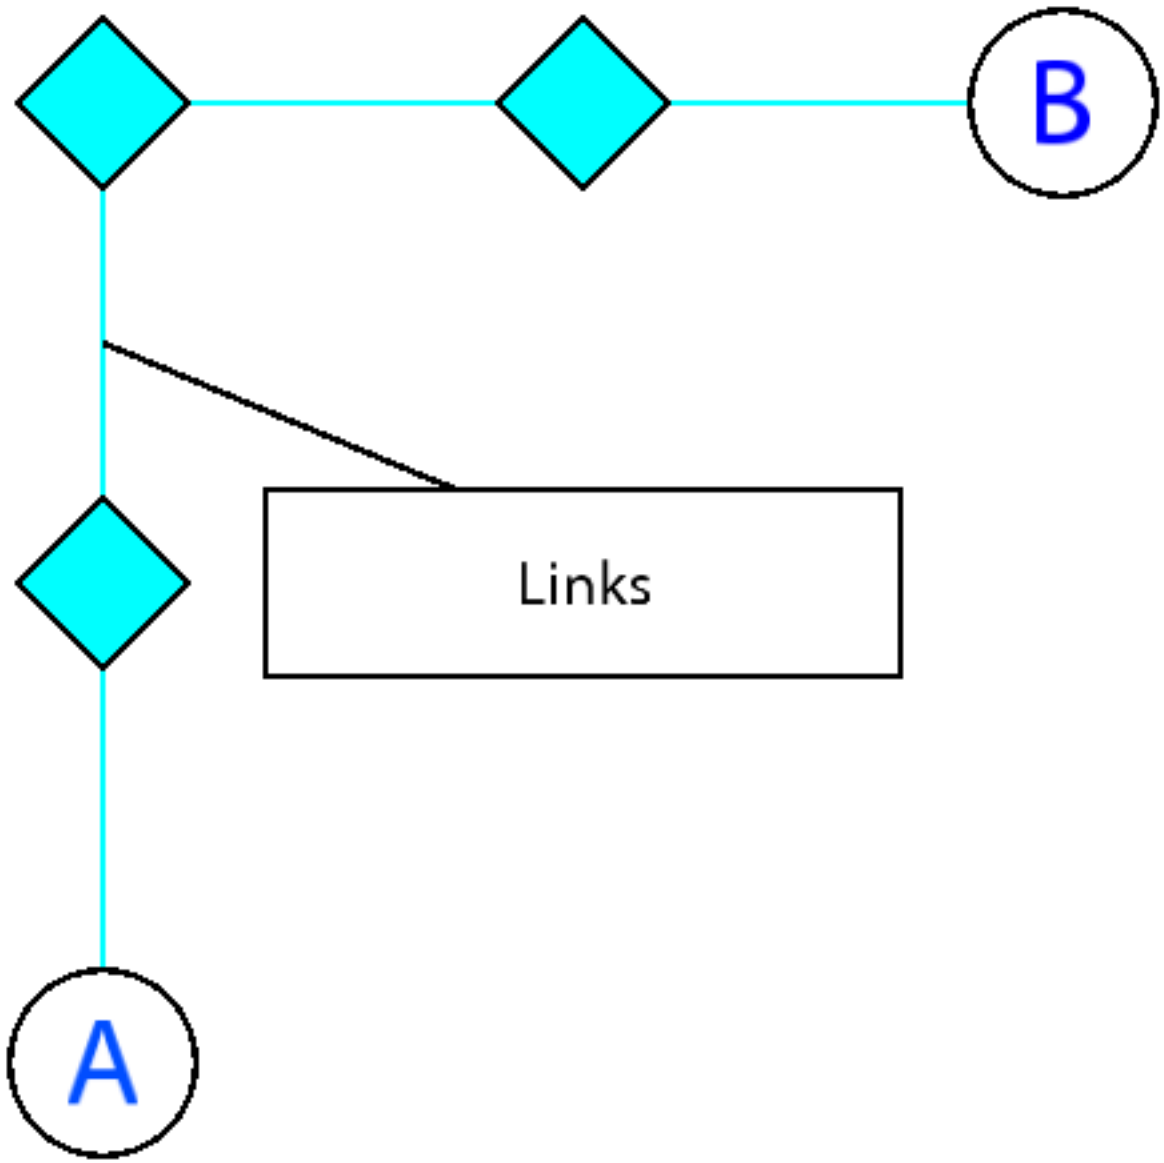

Topology Inference

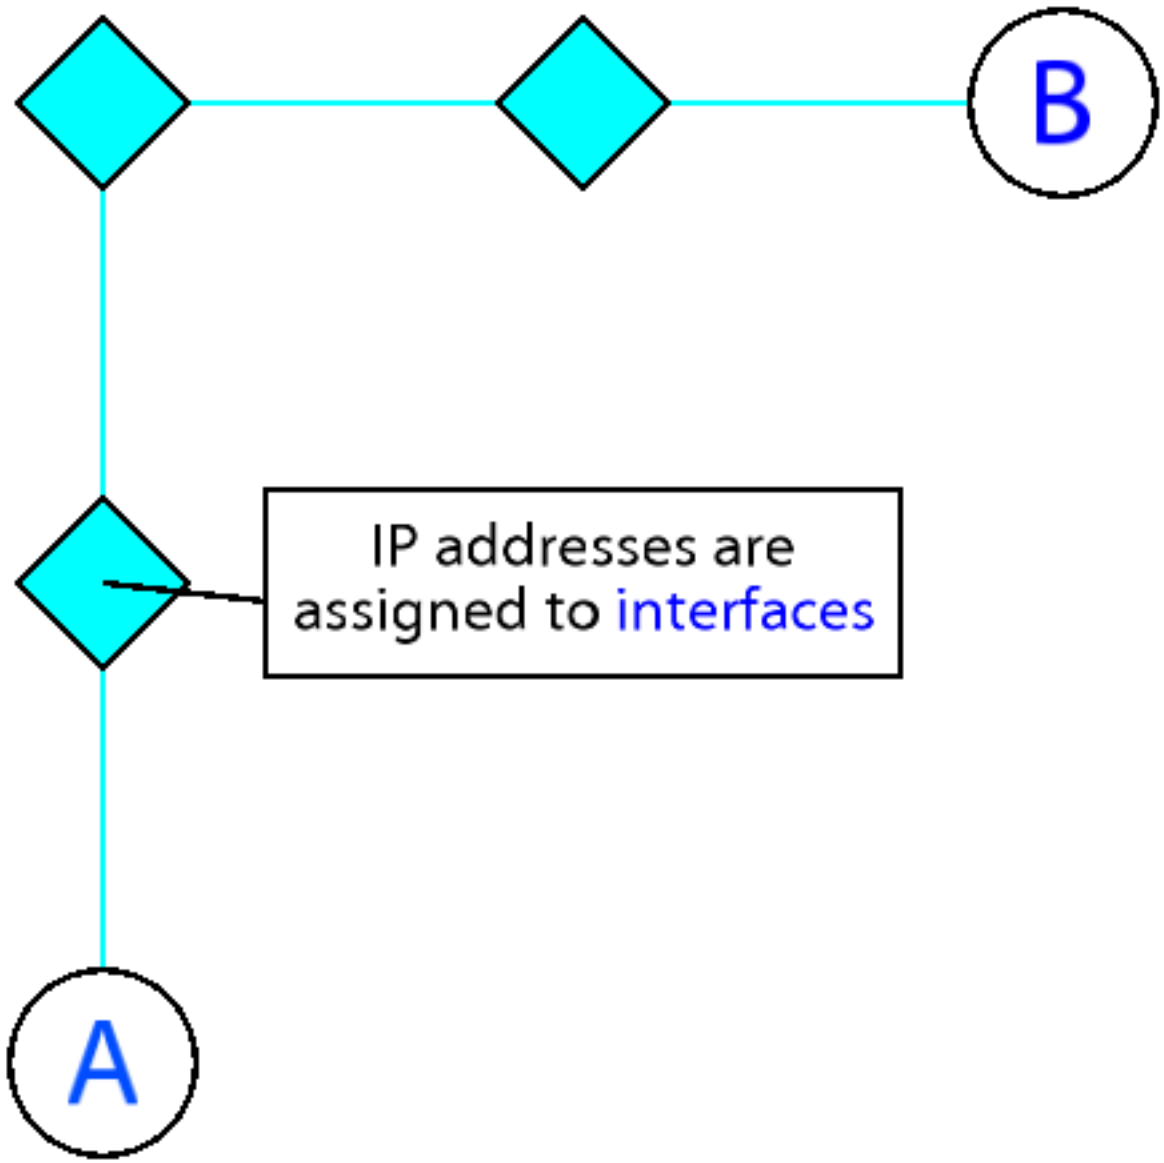

IP Addresses Are Interfaces

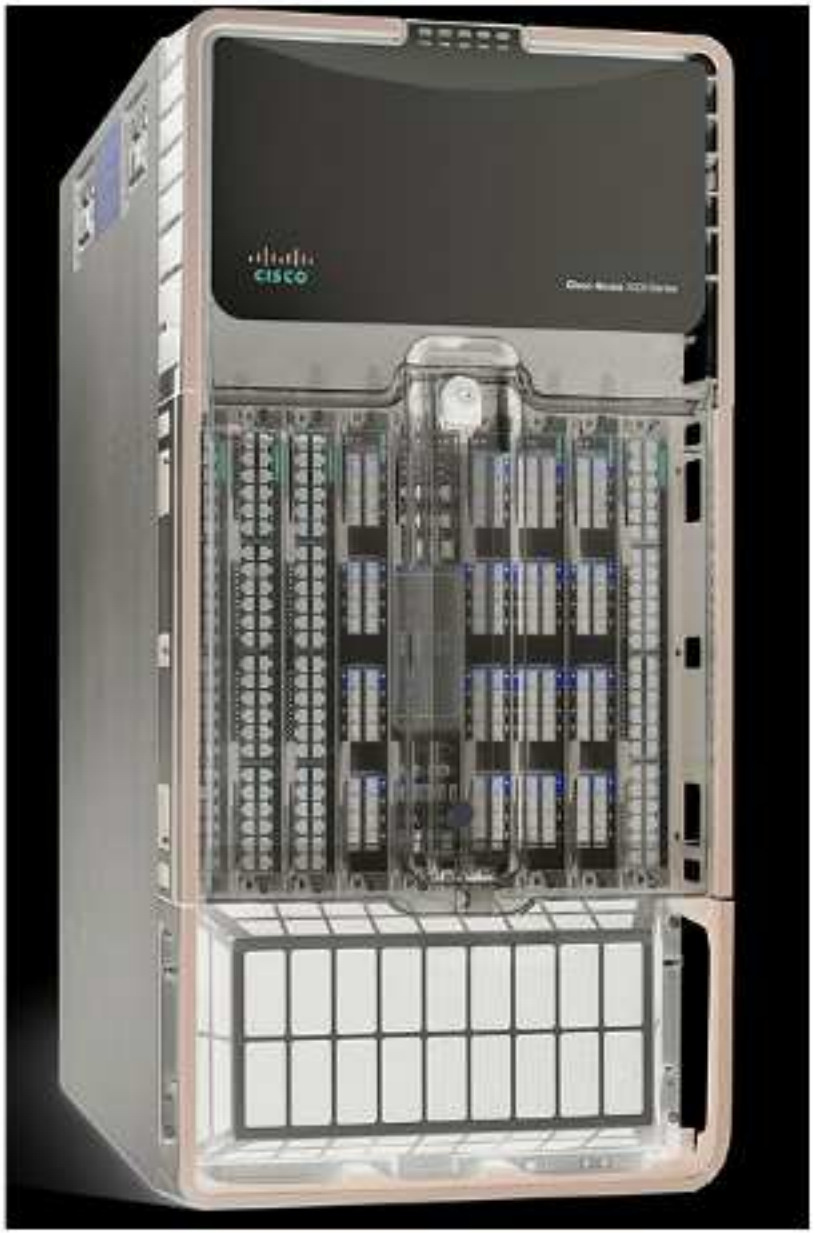

- All routers have multiple interfaces
- Many-to-one mapping from interfaces to routers

IP Addresses Are Interfaces

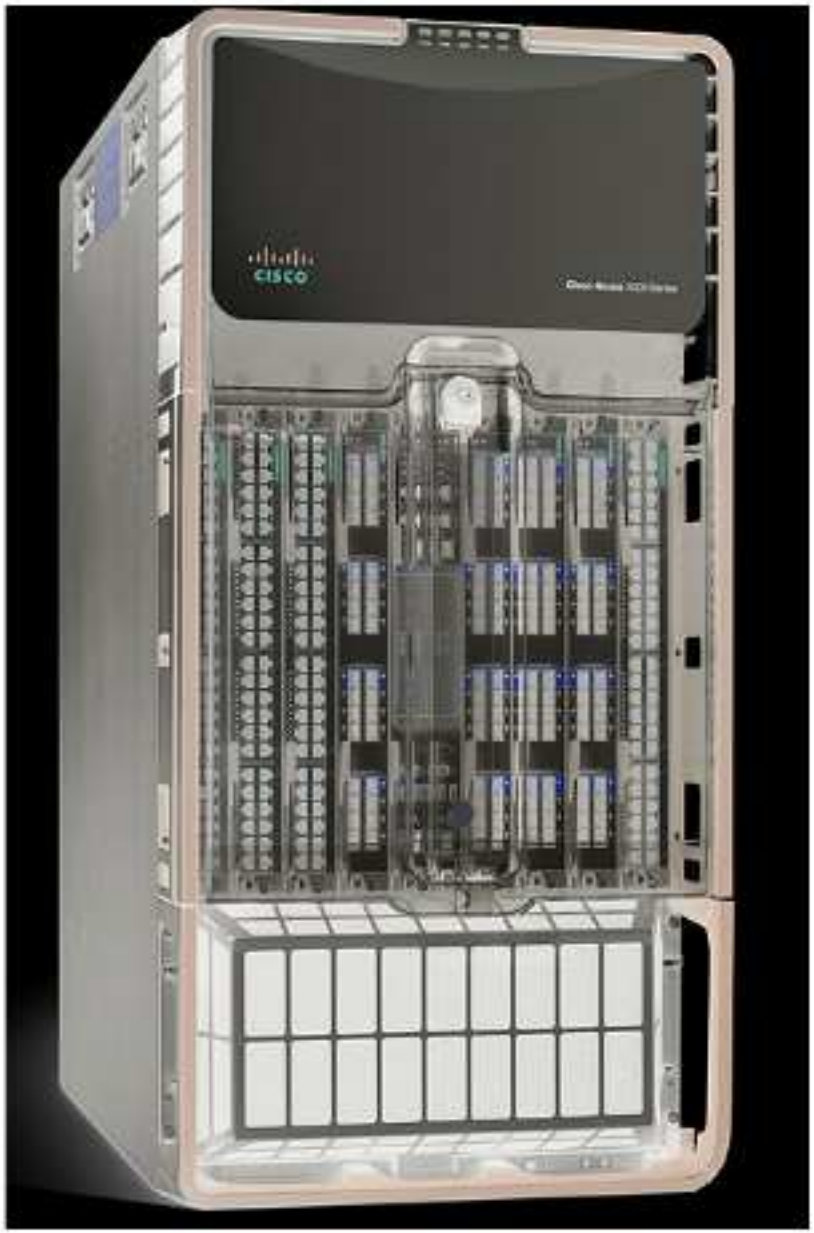

- All routers have multiple interfaces
- Many-to-one mapping from interfaces to routers

Topology Inference

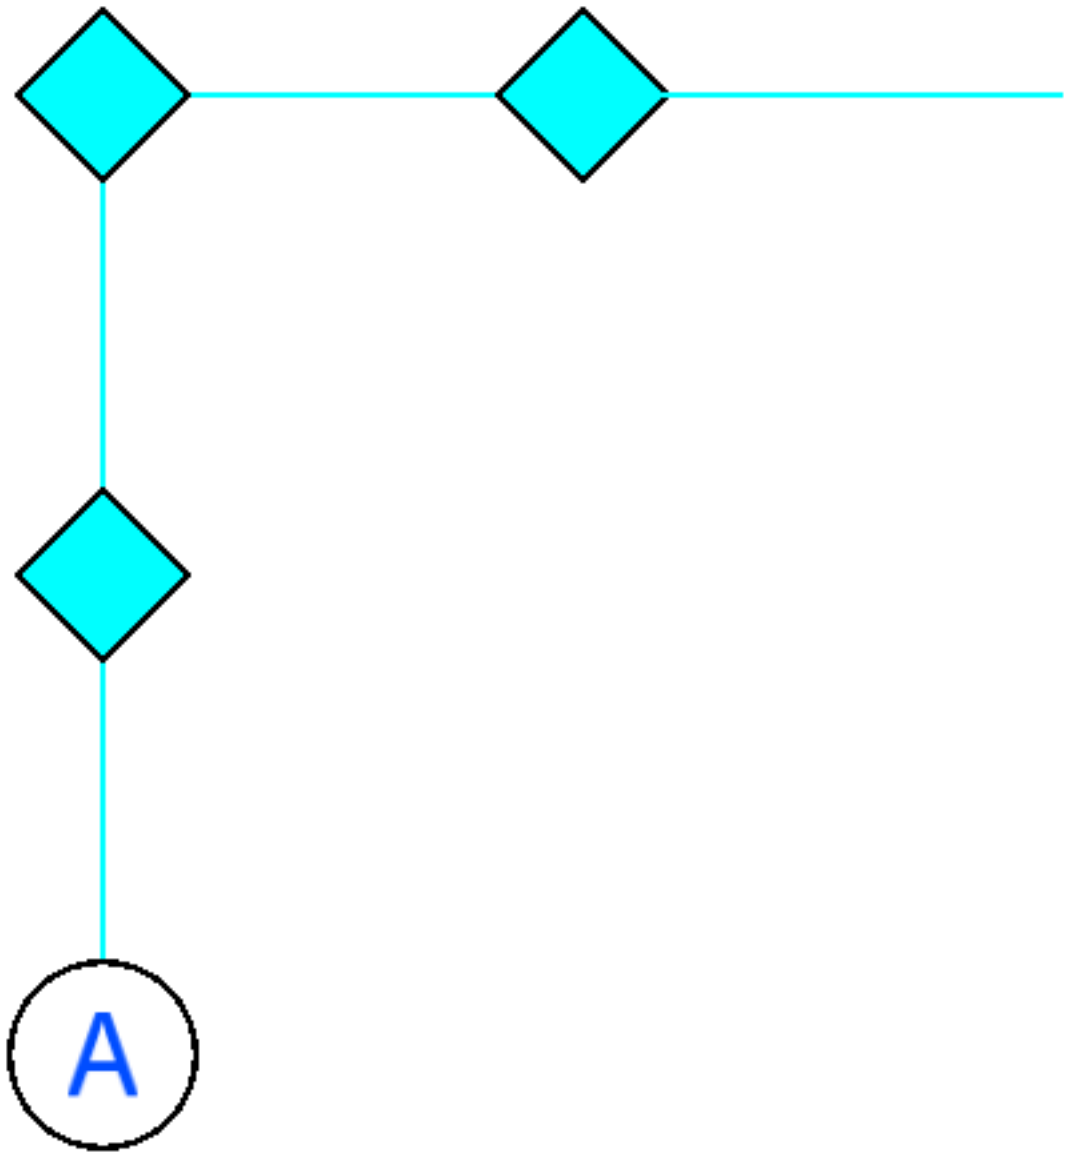

Topology Inference

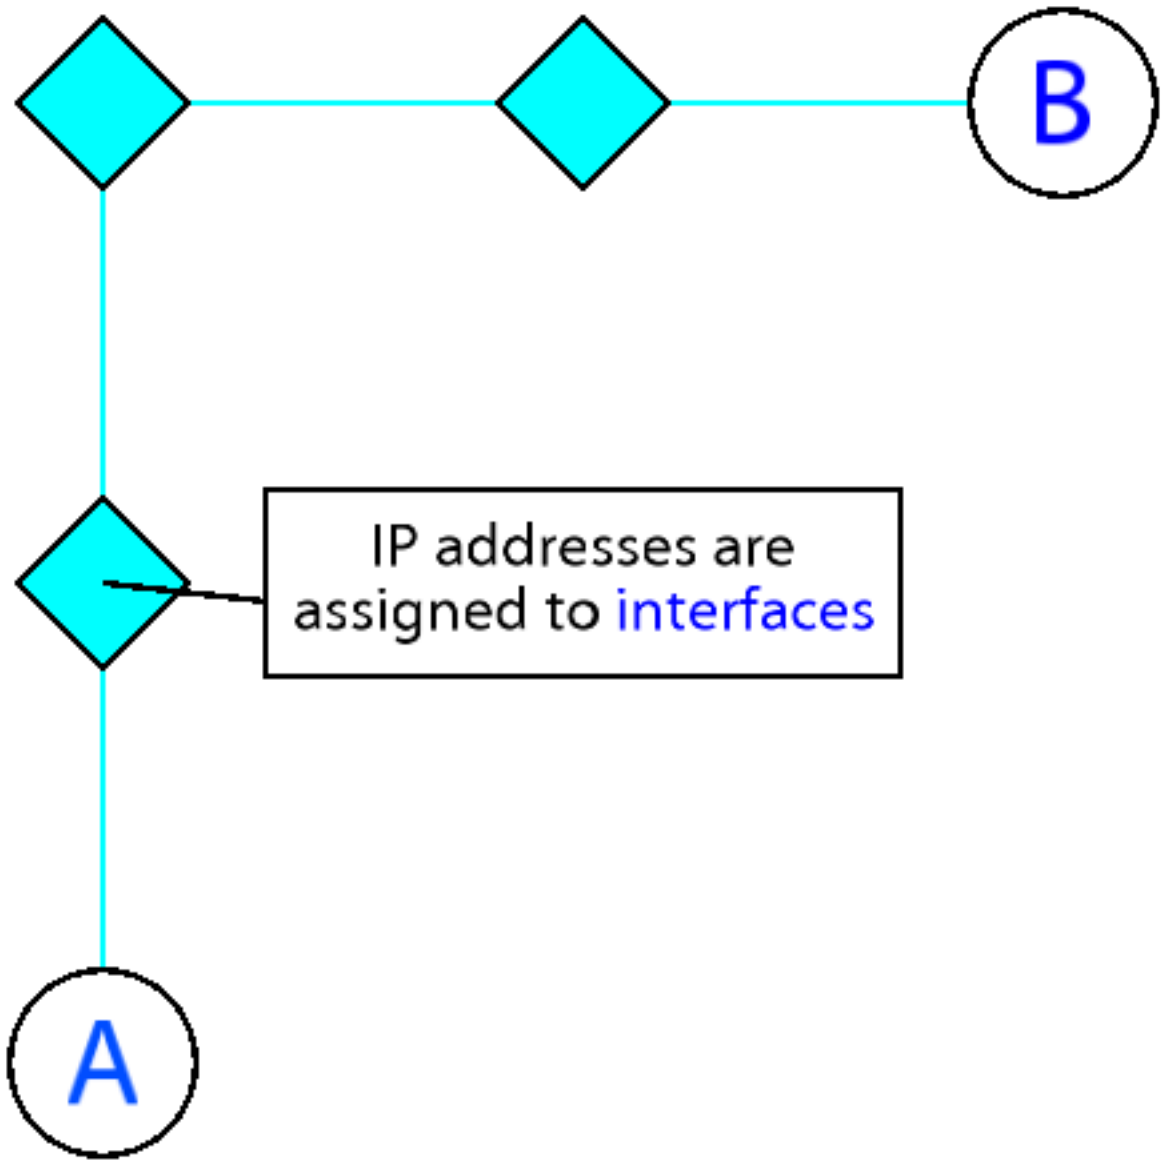

Topology Inference

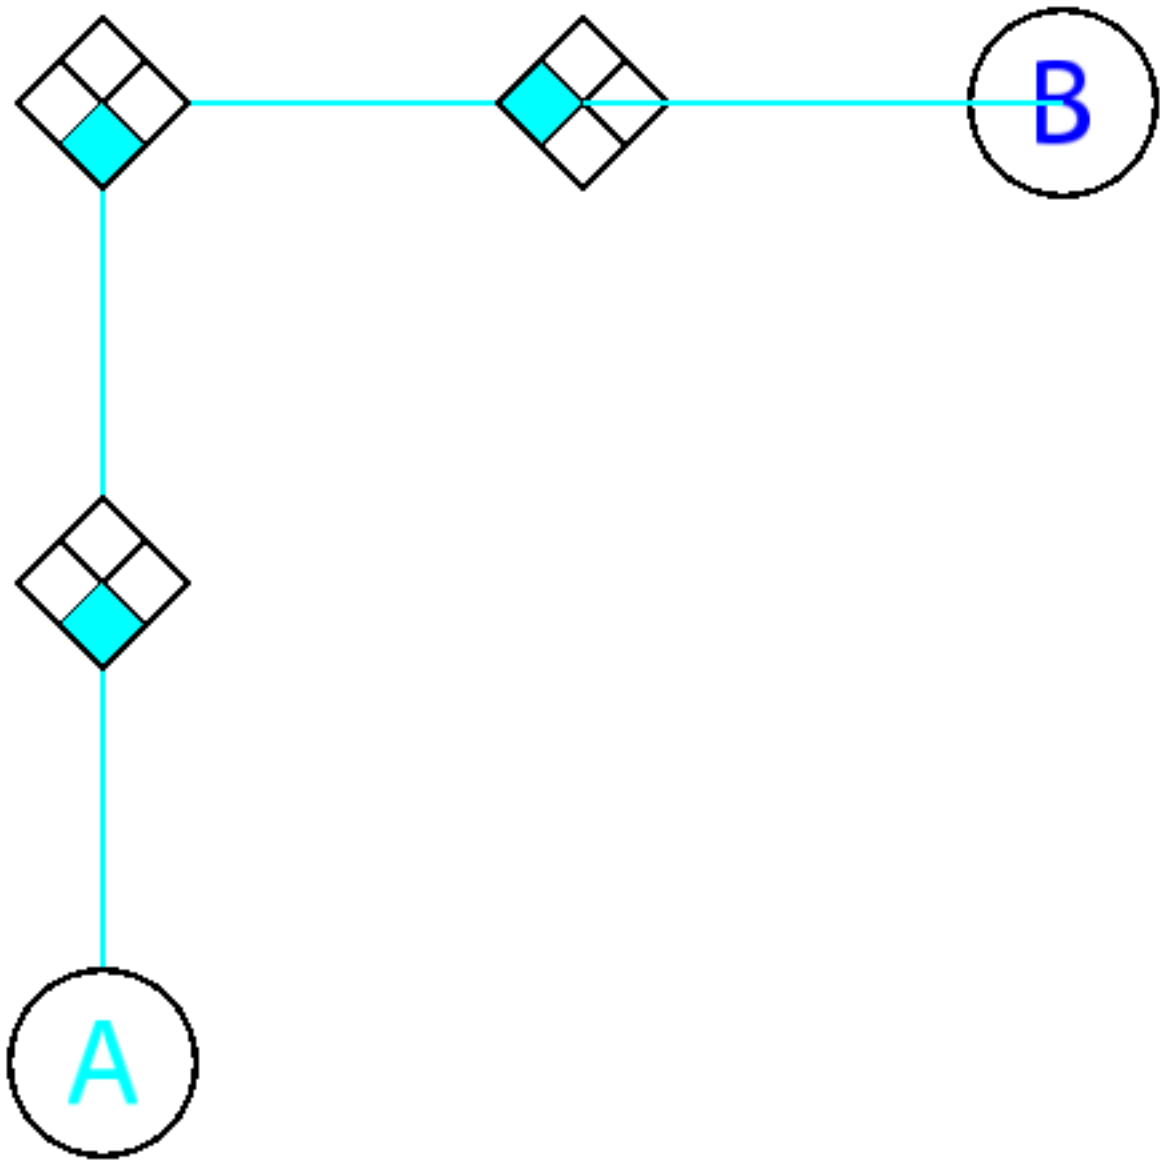

Topology Inference

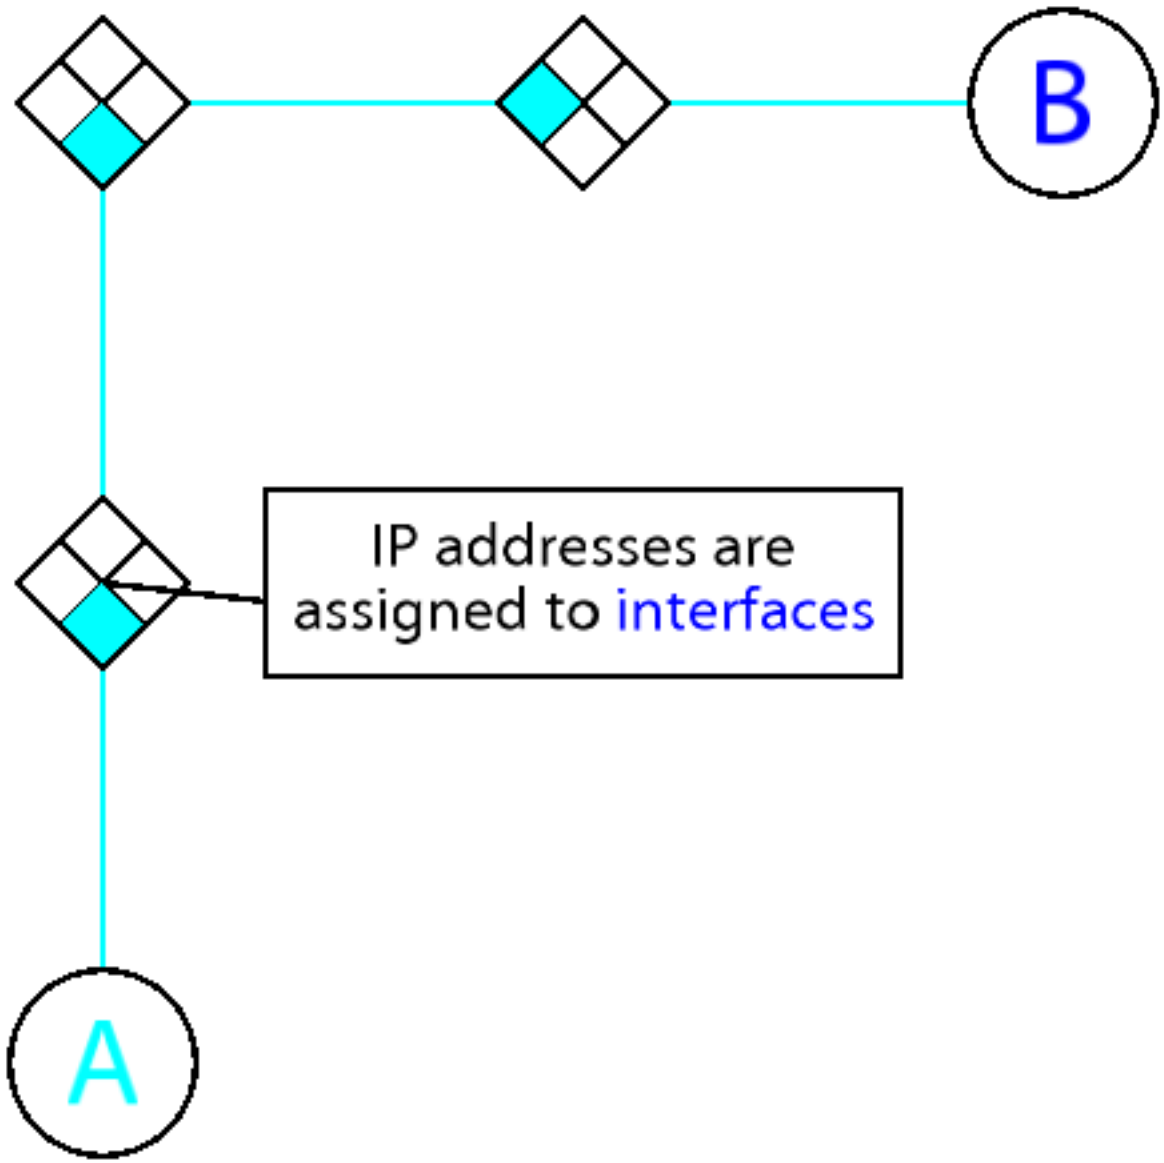

Topology Inference

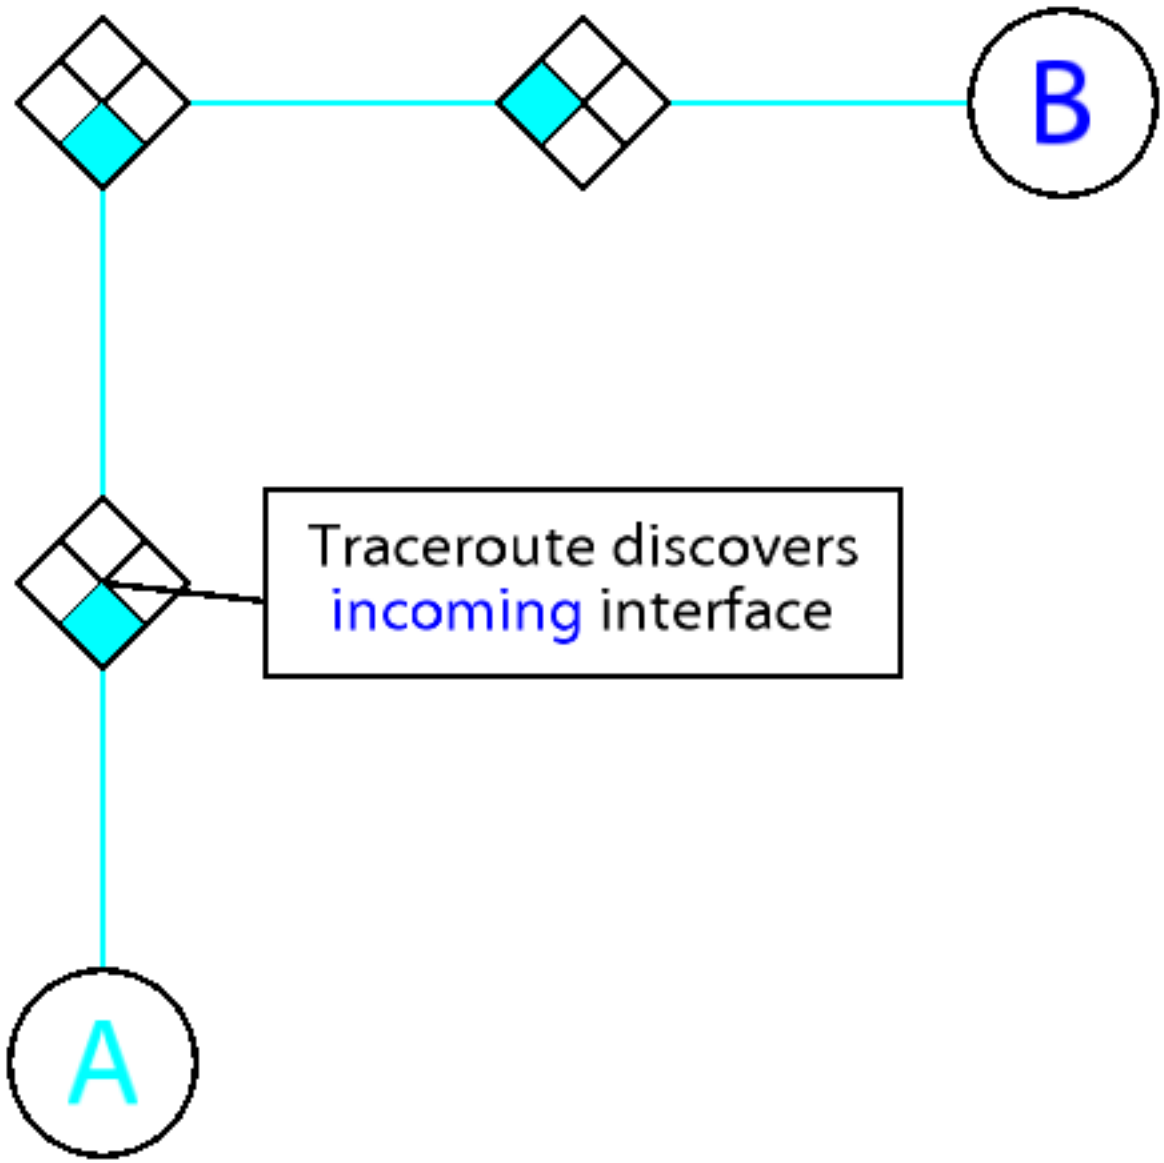

Topology Inference

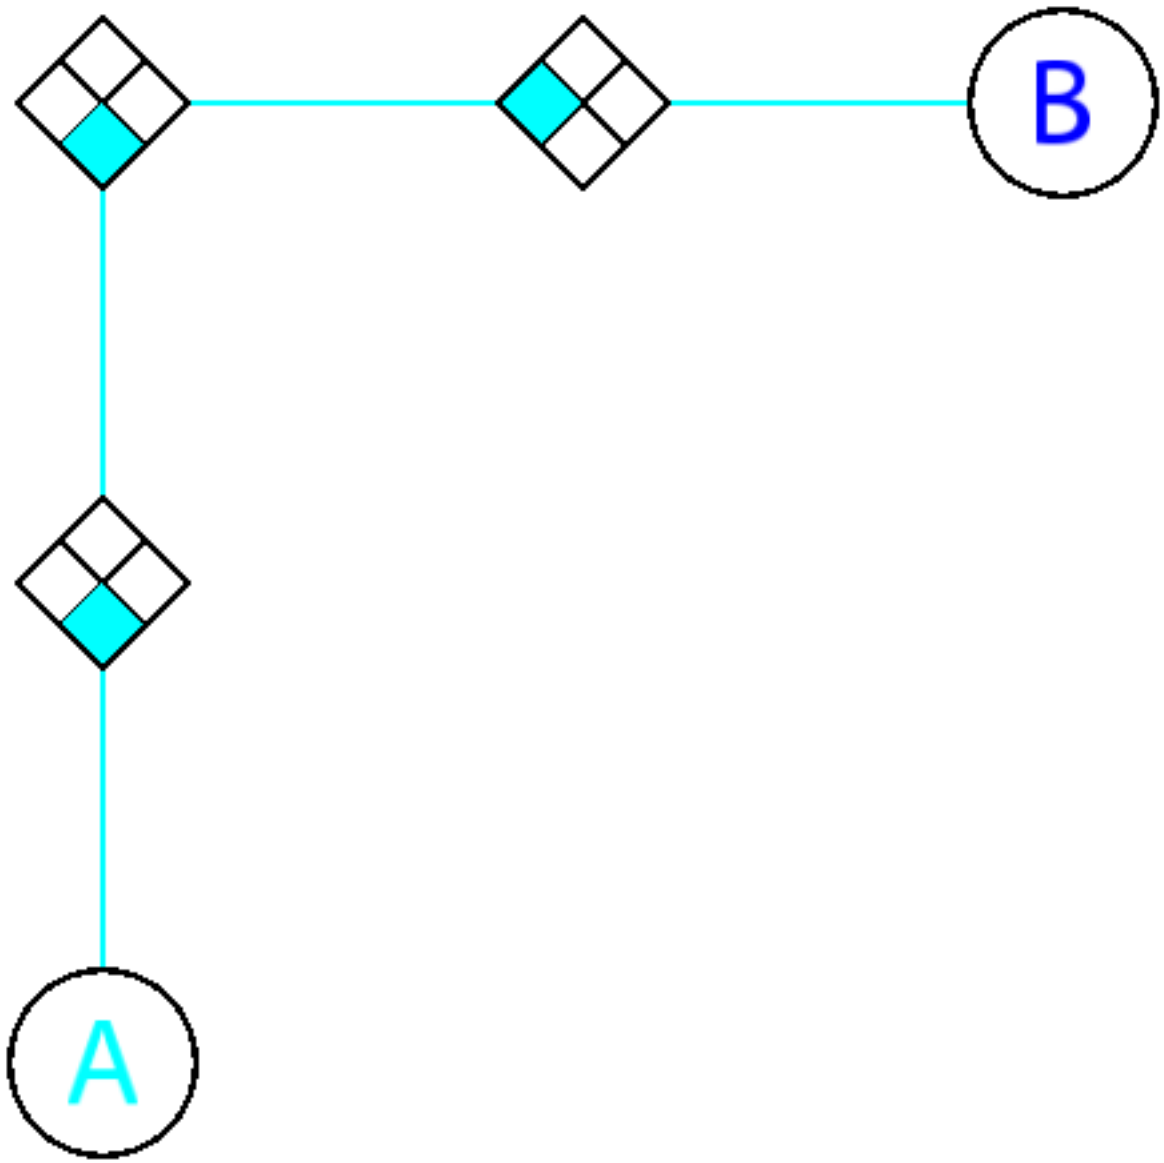

Topology Inference

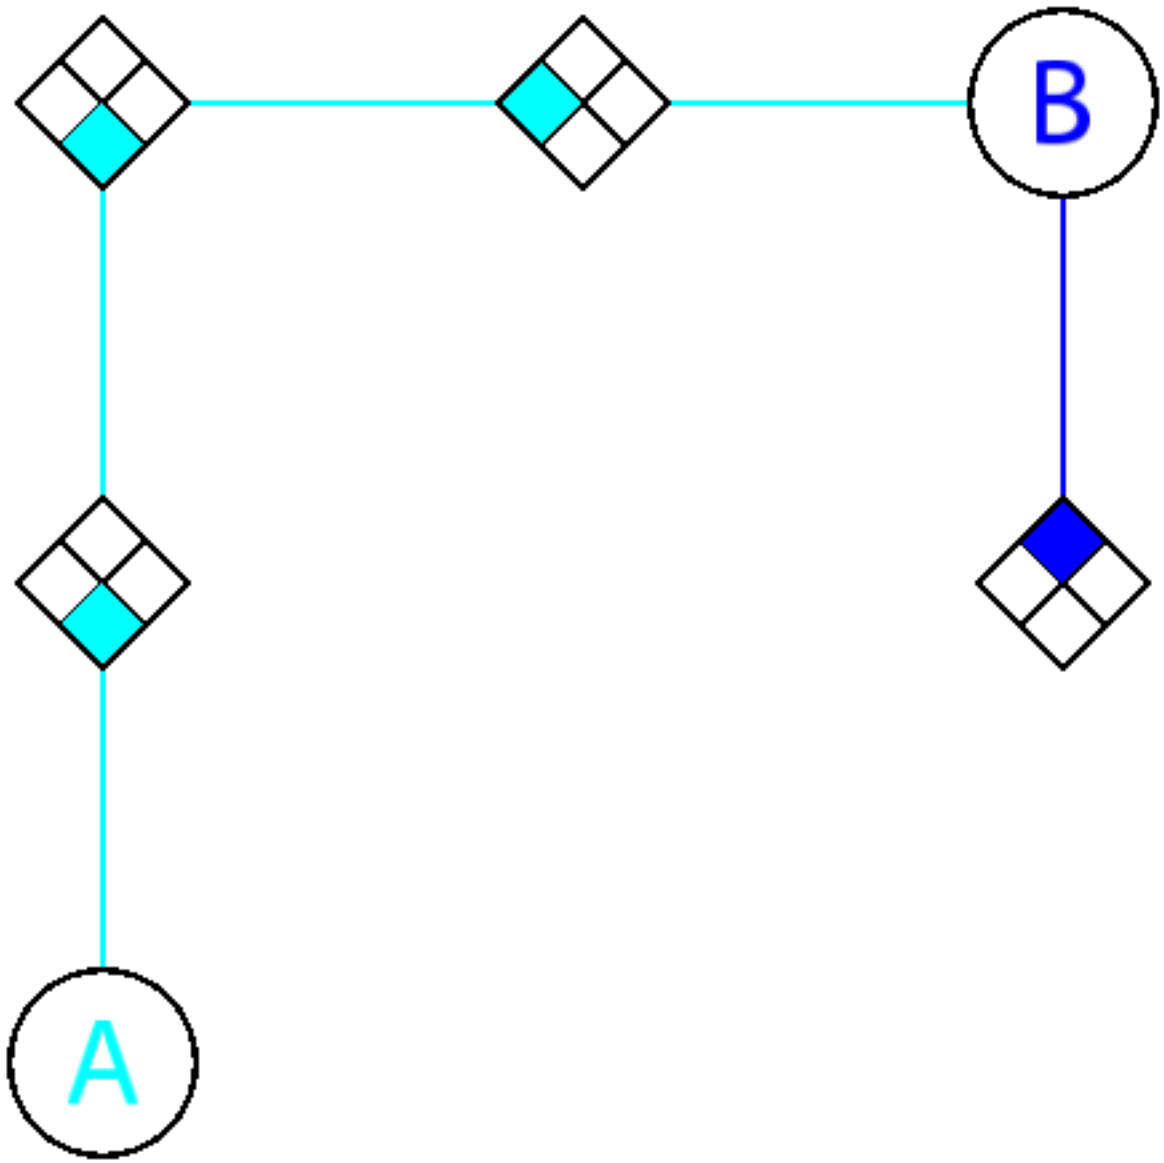

Topology Inference

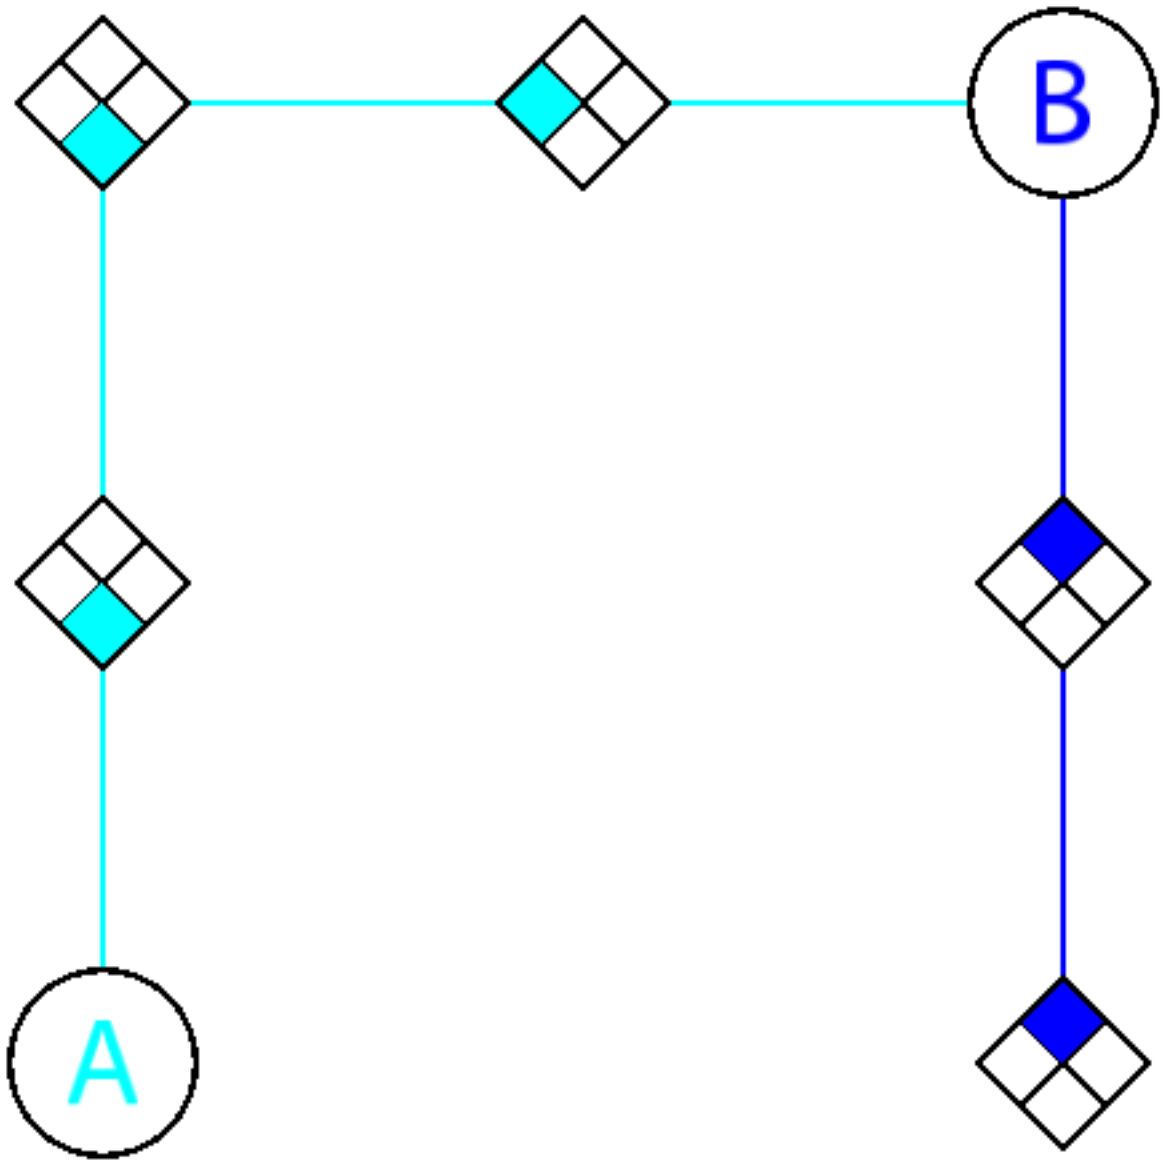

Topology Inference

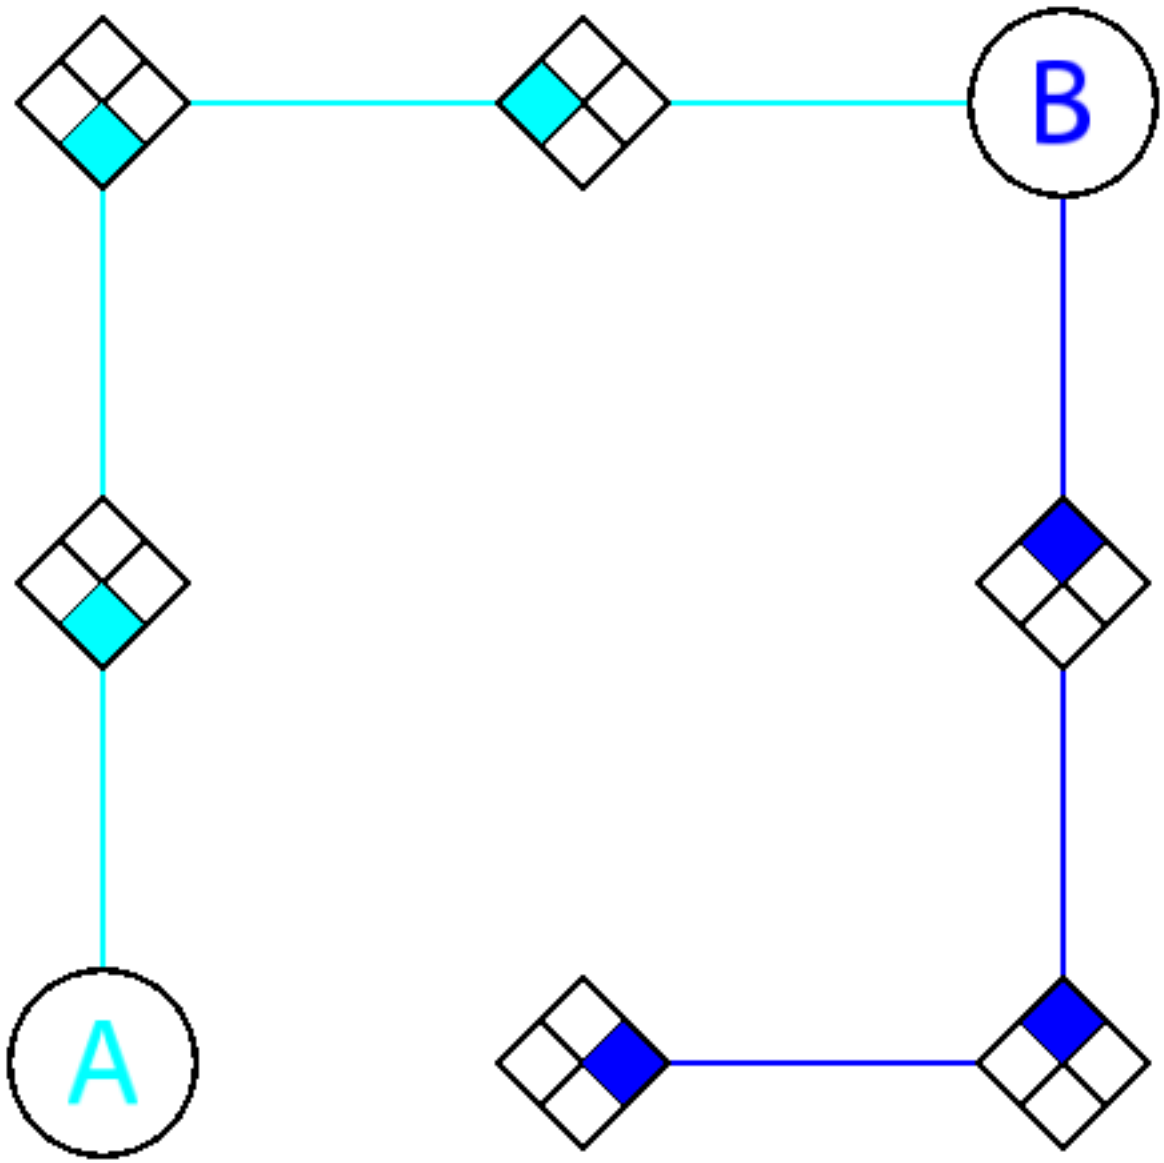

Topology Inference

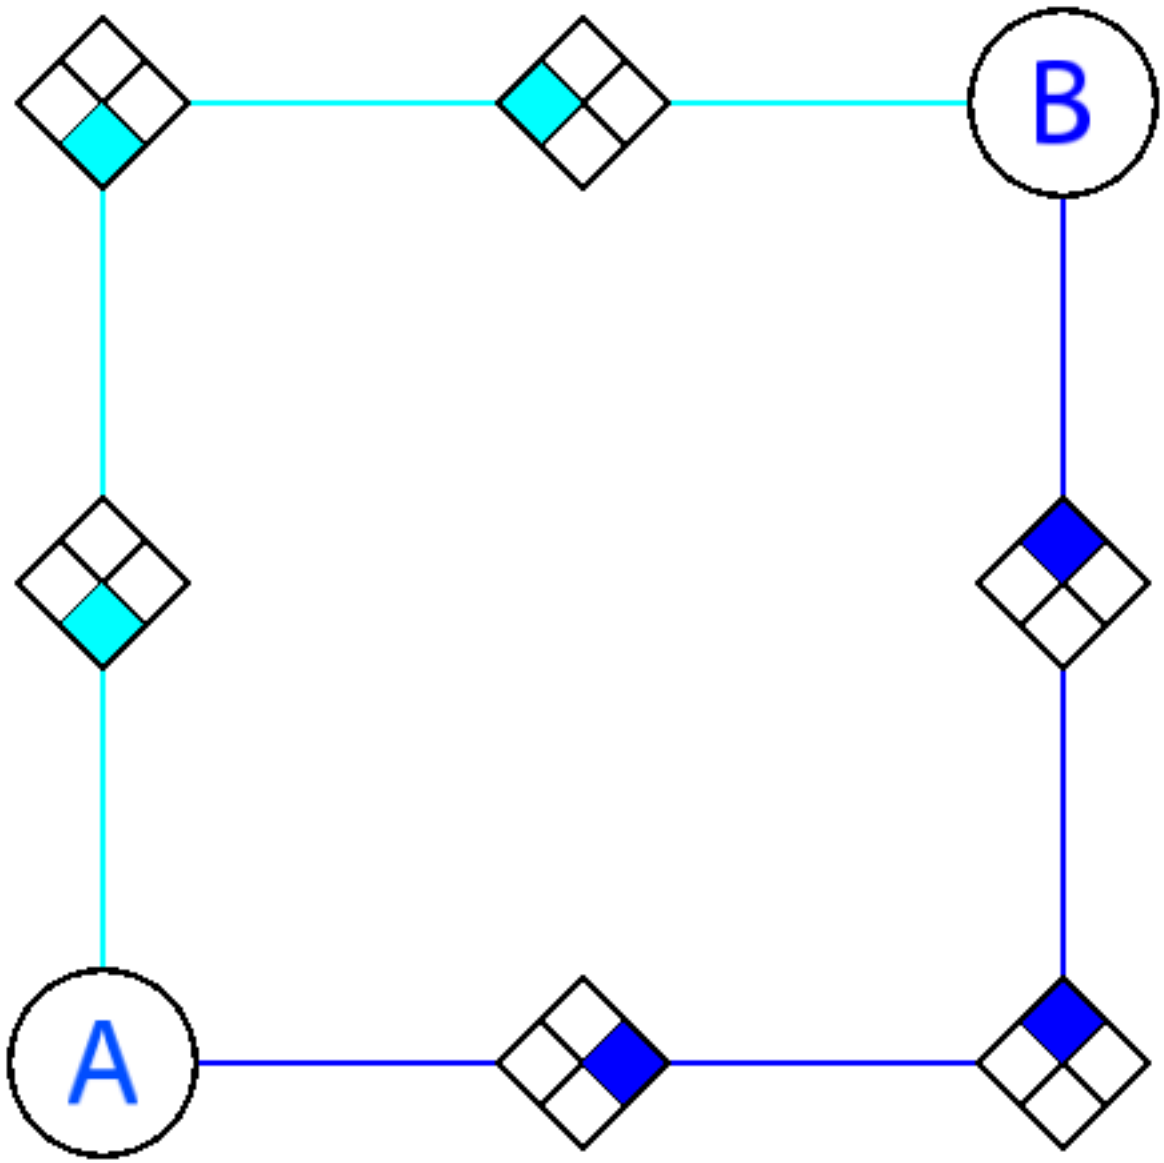

Topology Inference

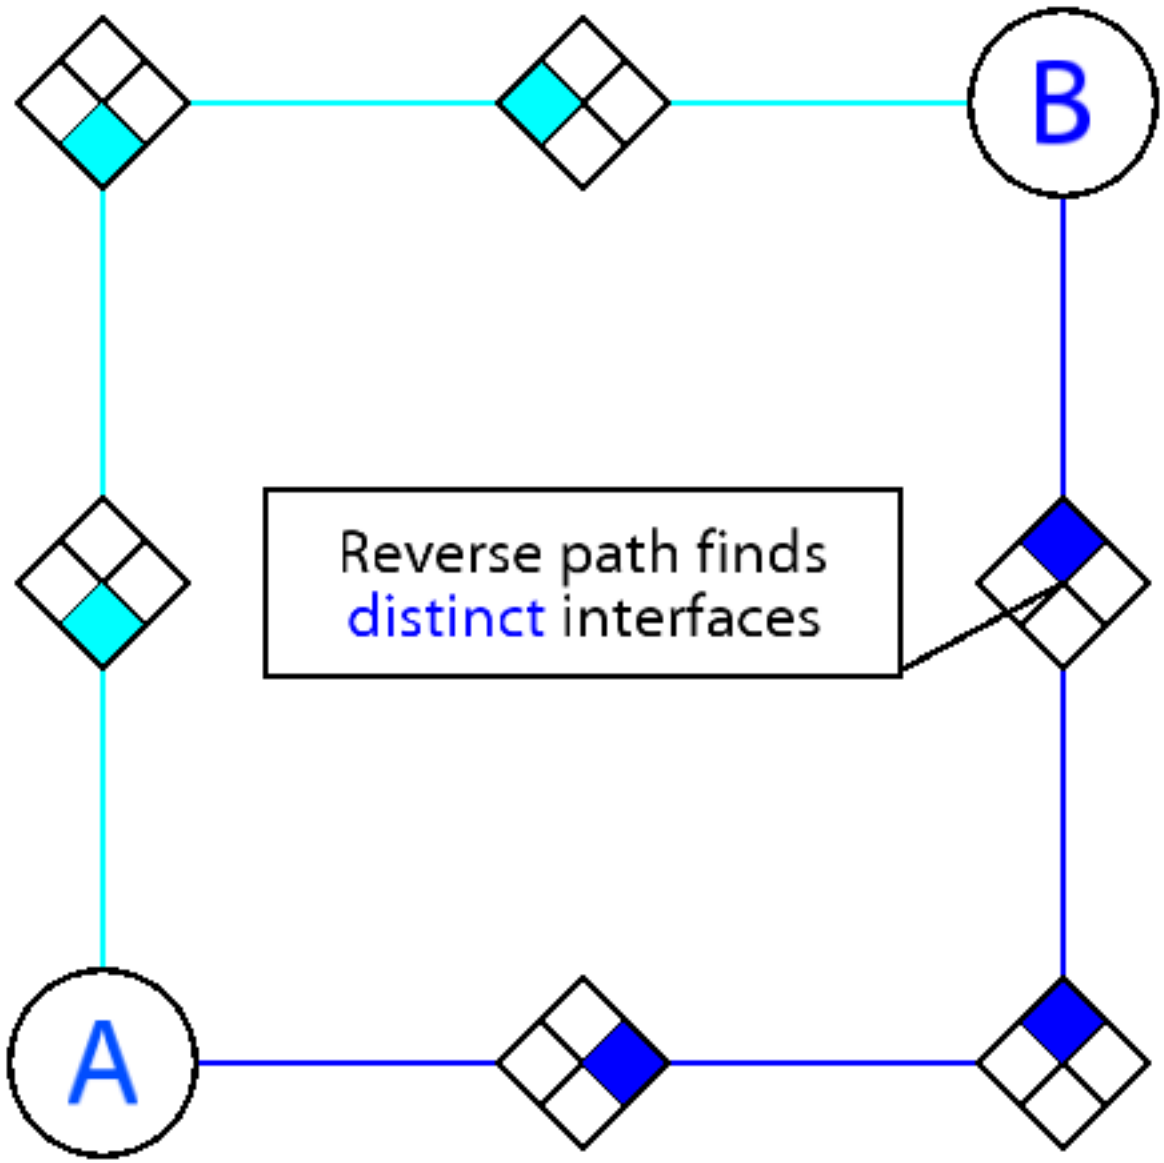

Ambiguous Topologies

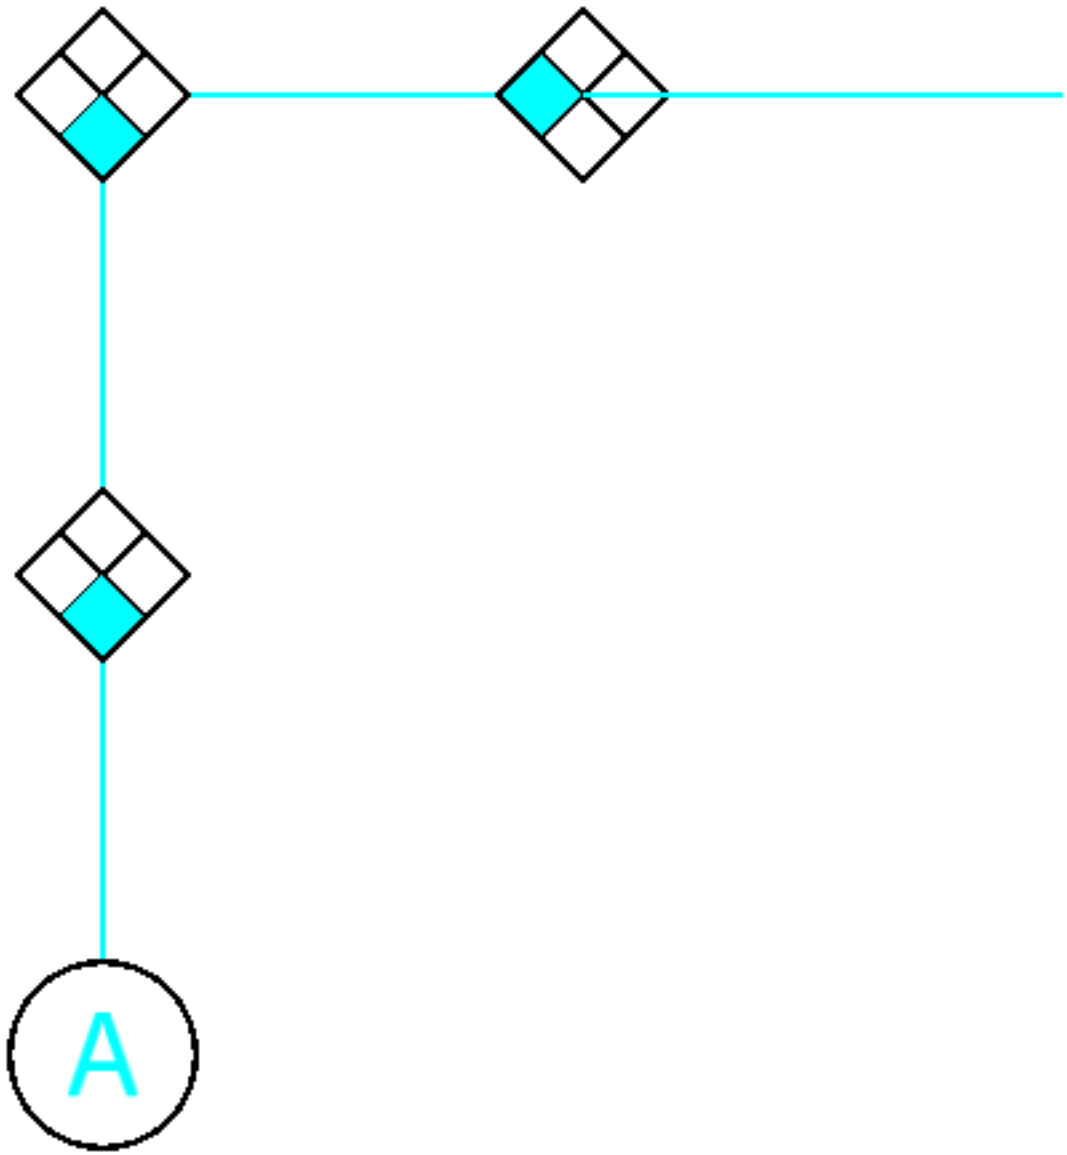

Ambiguous Topologies

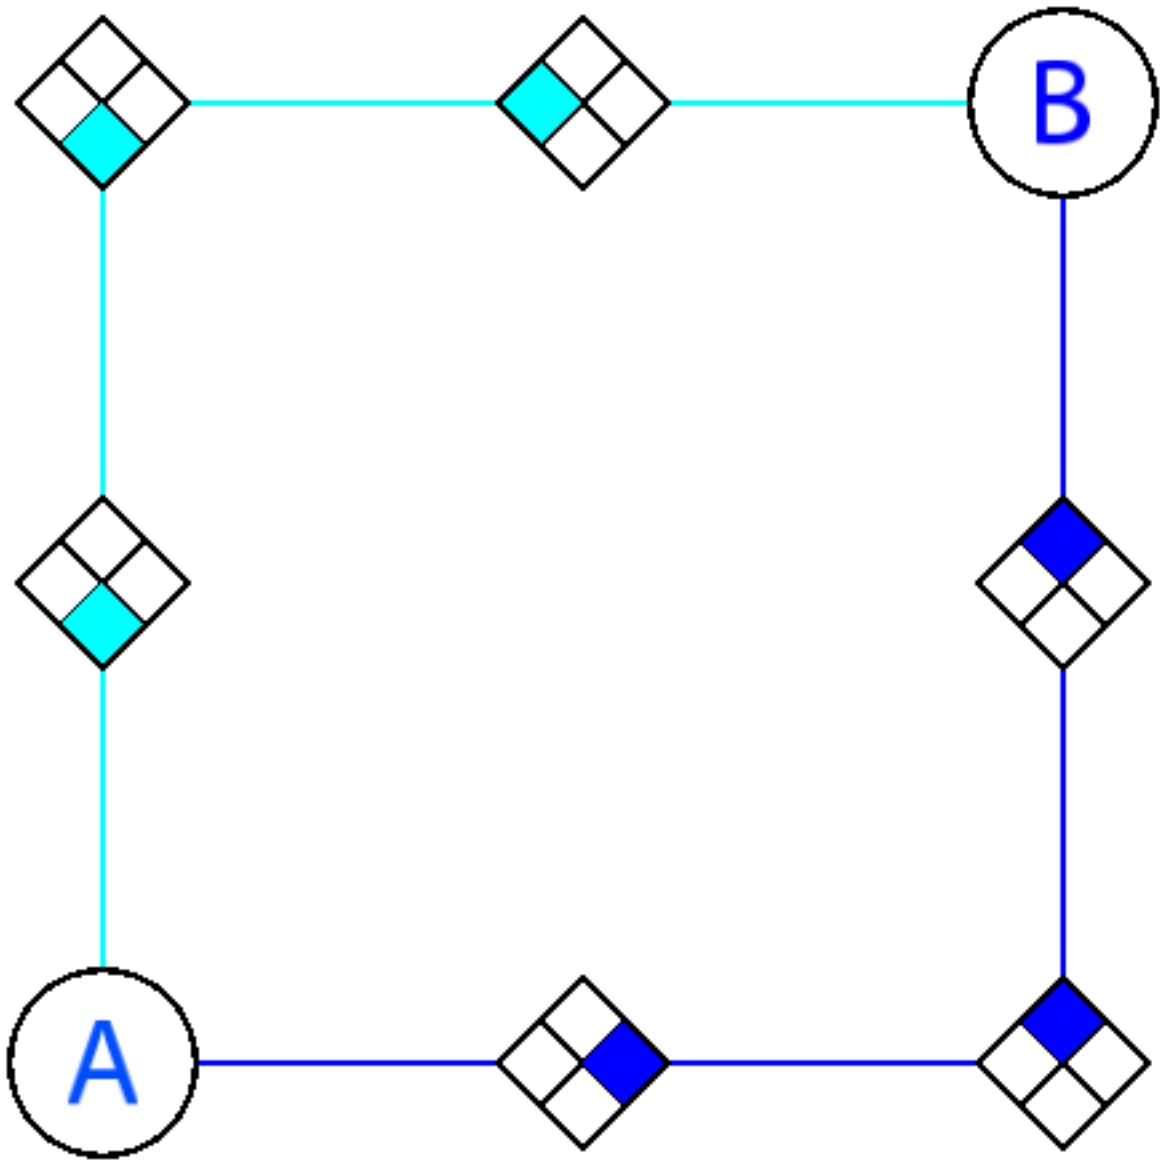

Ambiguous Topologies

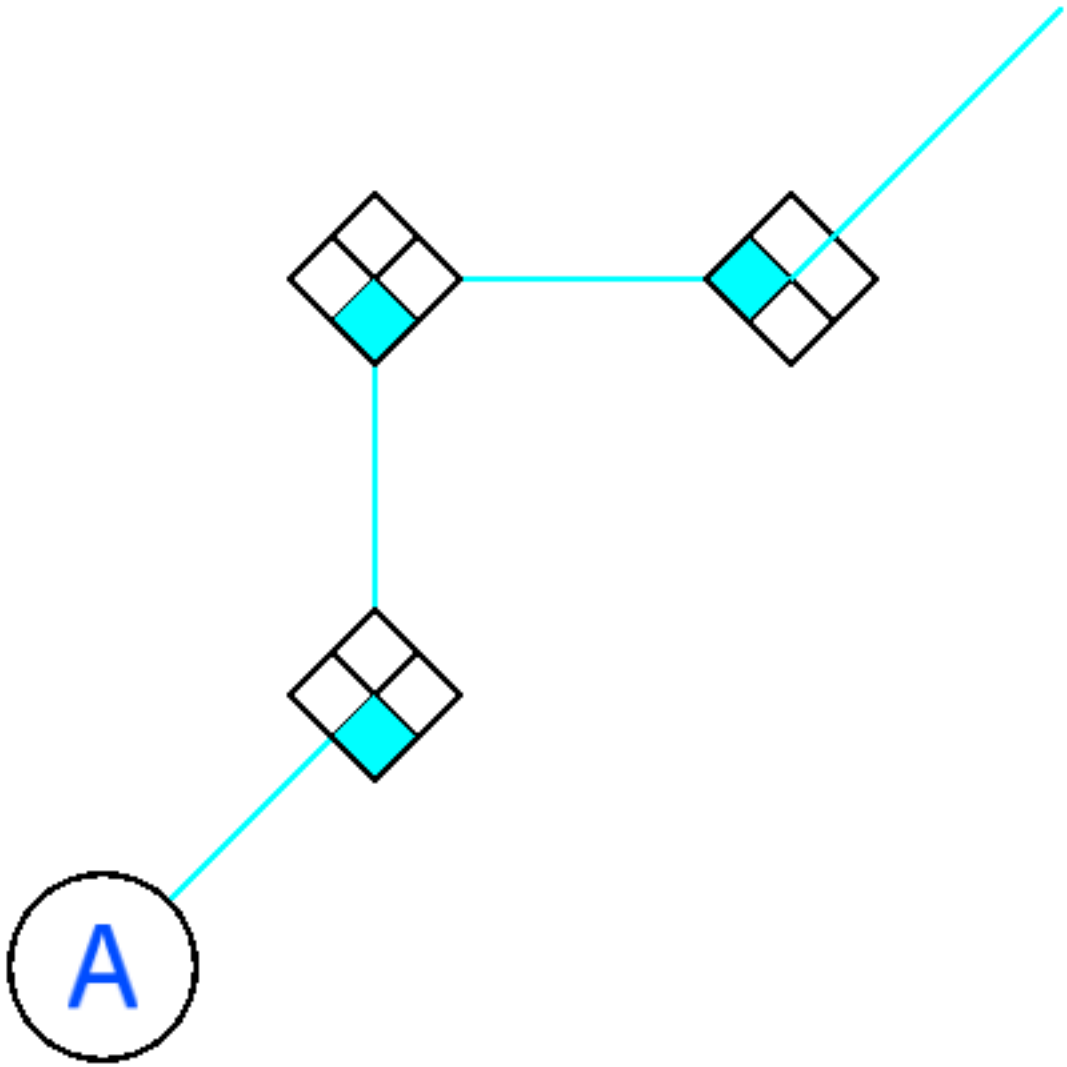

Ambiguous Topologies

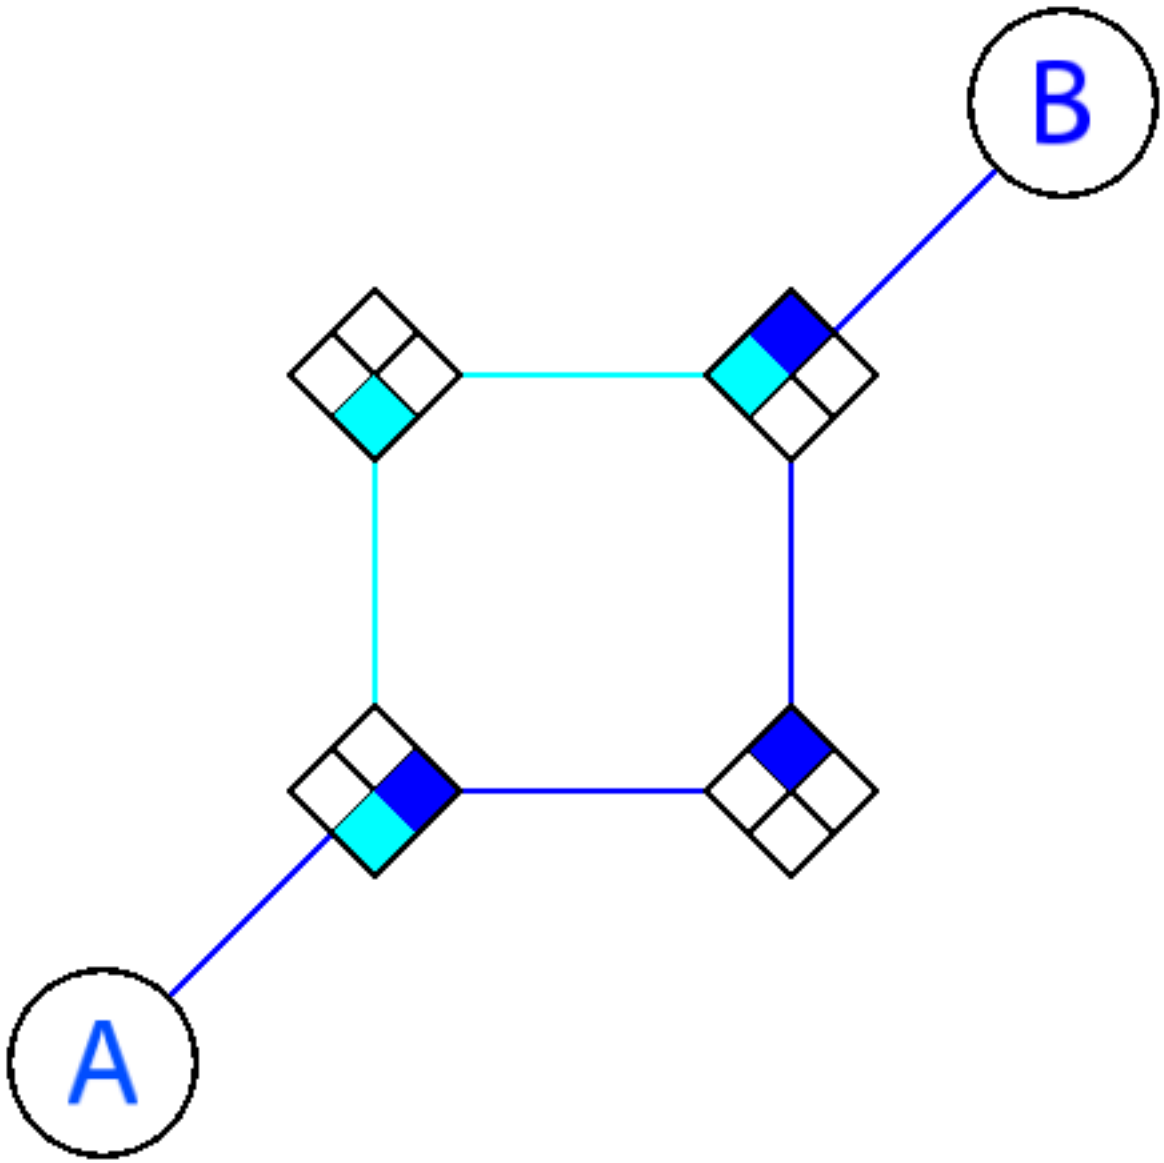

Ambiguous Topologies

Need to map interfaces to routers

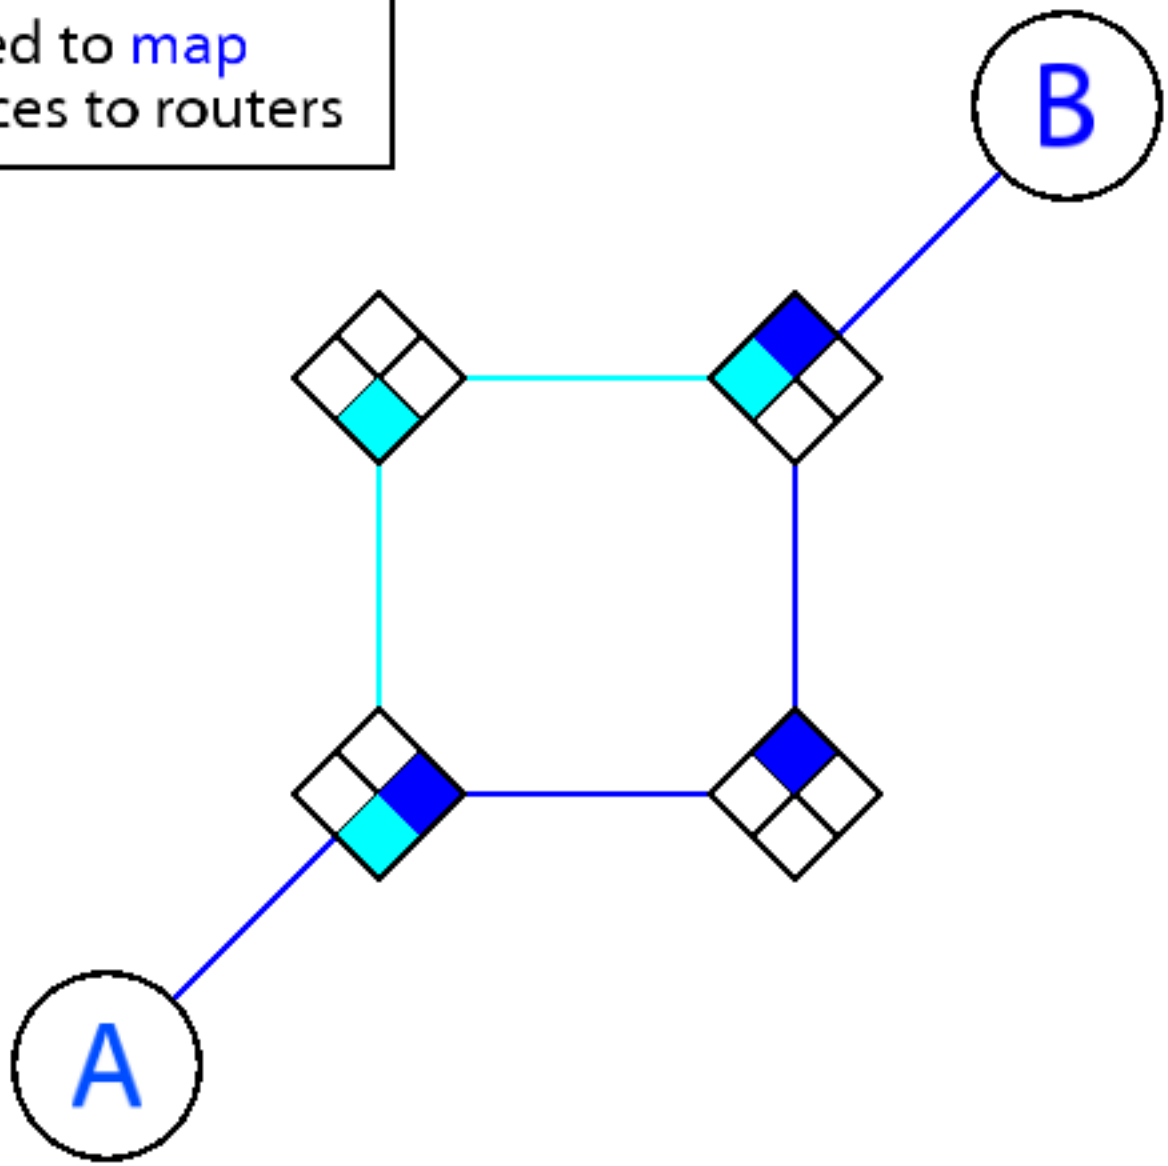

Ambiguous Topologies

Alias resolution

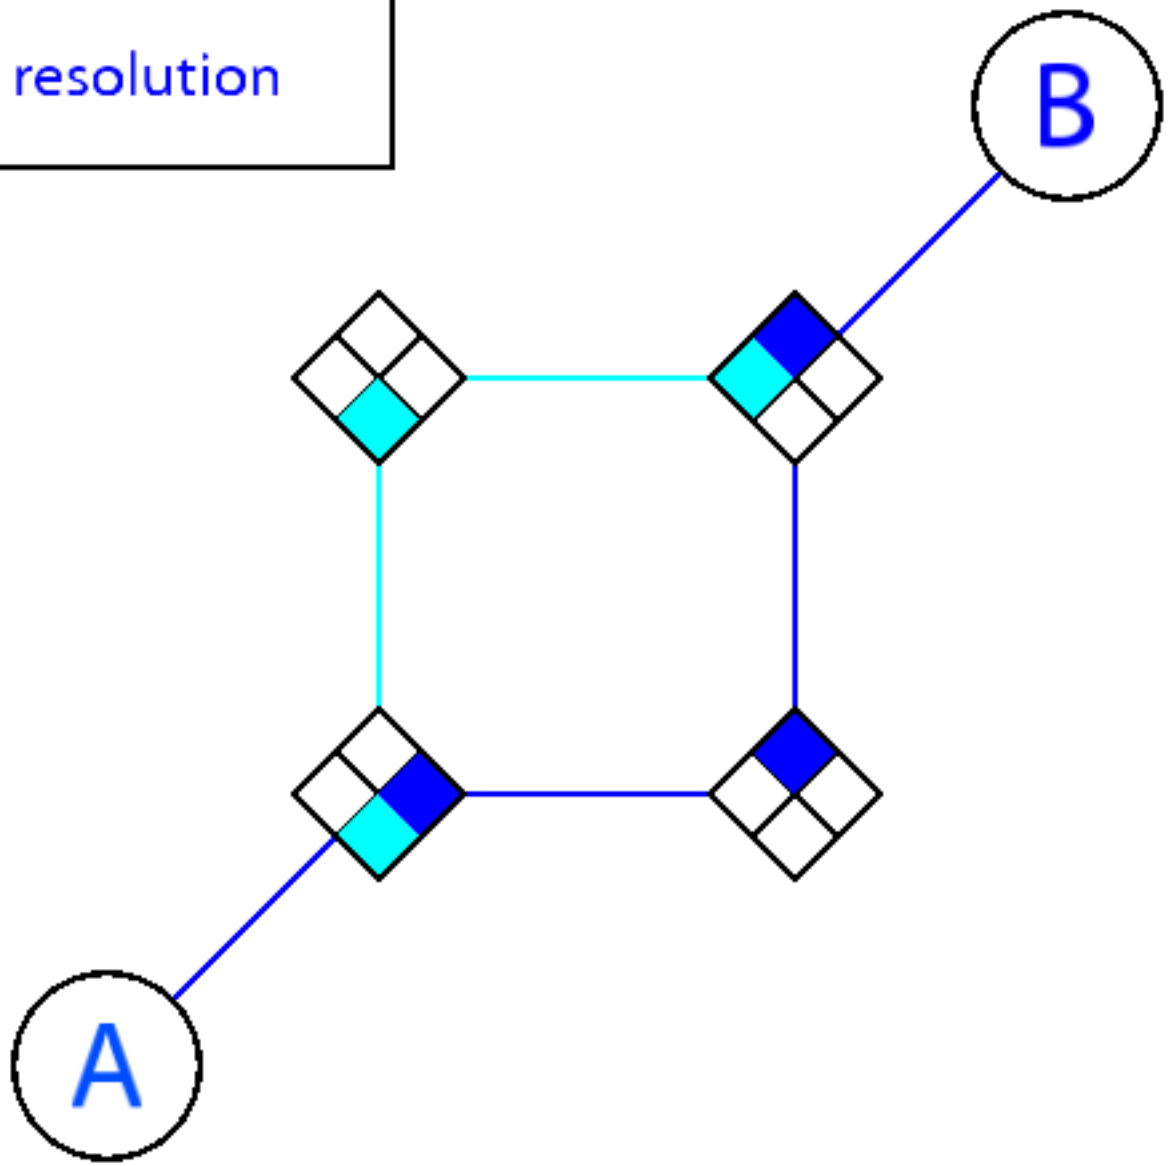

Alias Resolution

- Direct probing techniques
  - Source IP matching: Mercator
  - IP ID matching: Rocketfuel
  - 32% of IP addresses are unresponsive
- Indirect resolution techniques are incomplete and have false positives
- Bad/missing aliases affect mapping

The Importance of Aliases

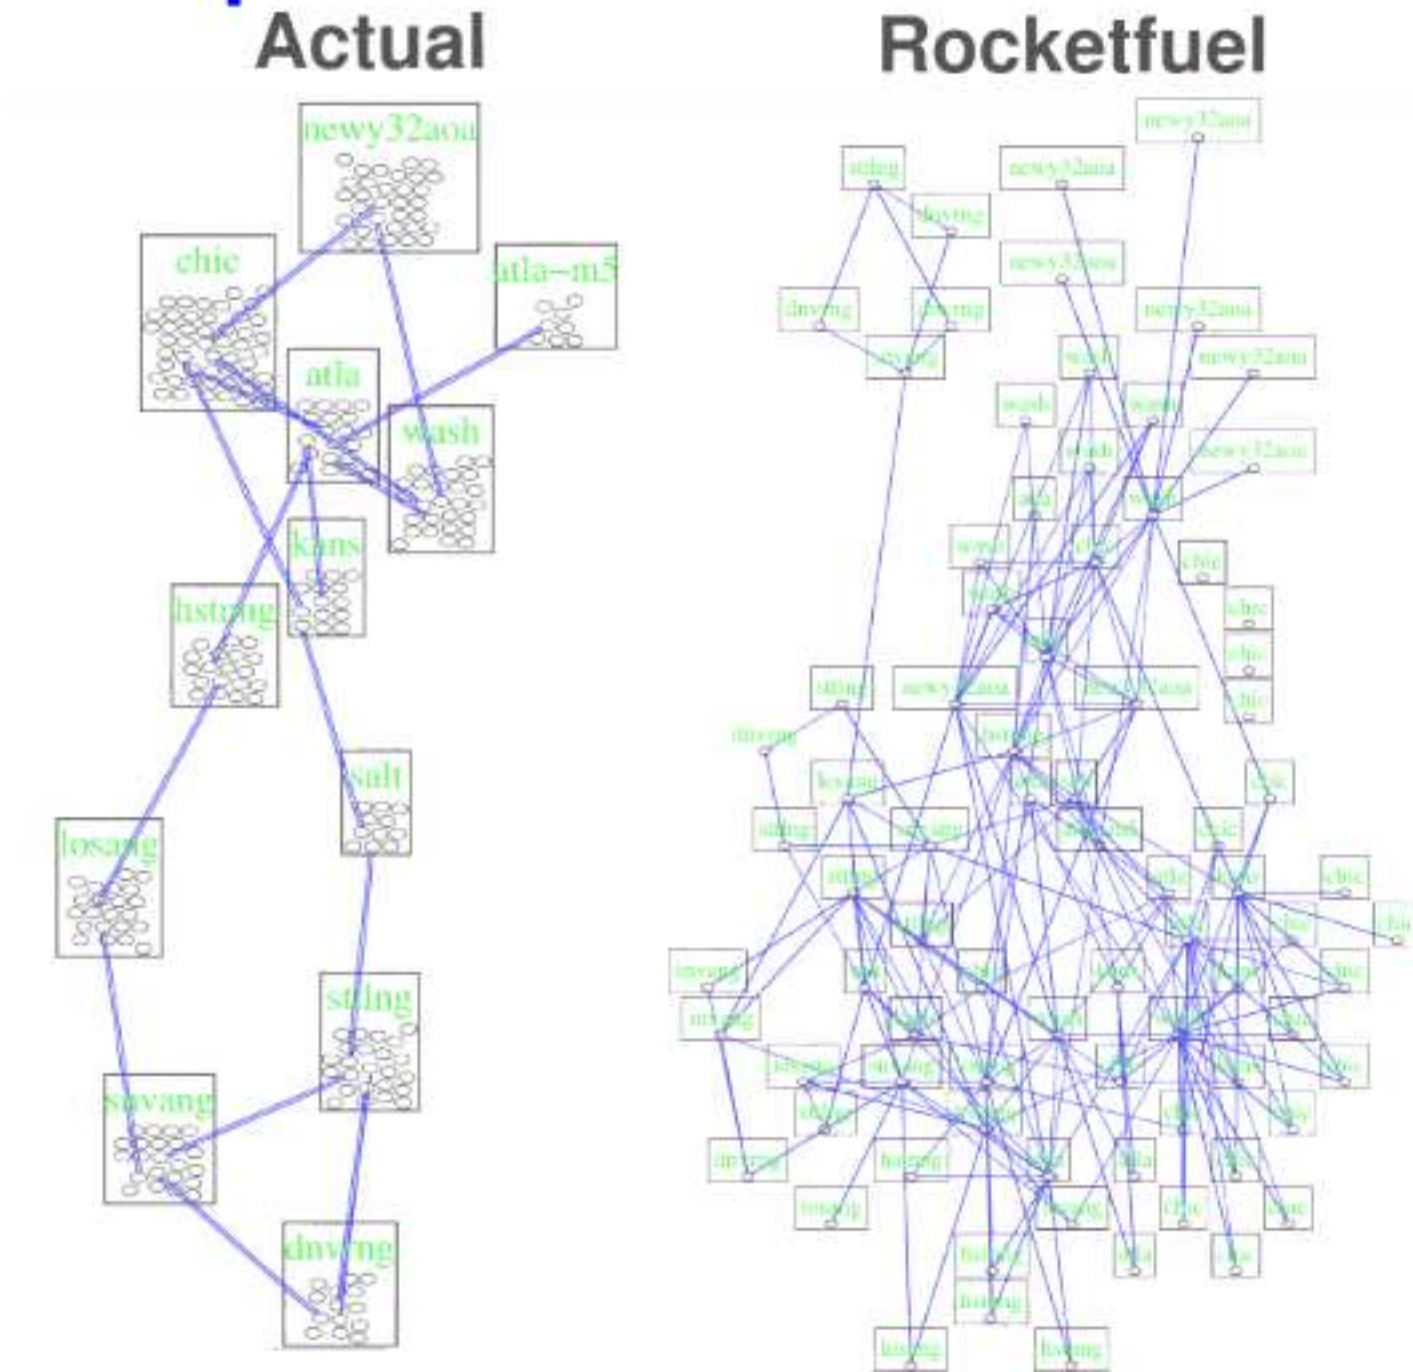

The Importance of Aliases

Actual

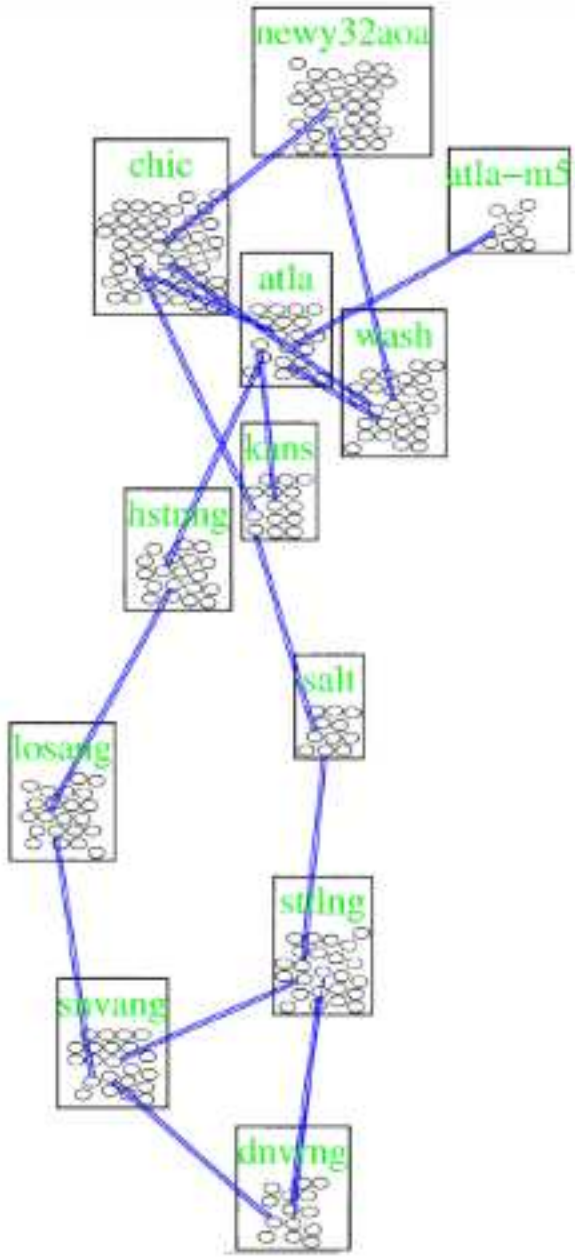

Rocketfuel

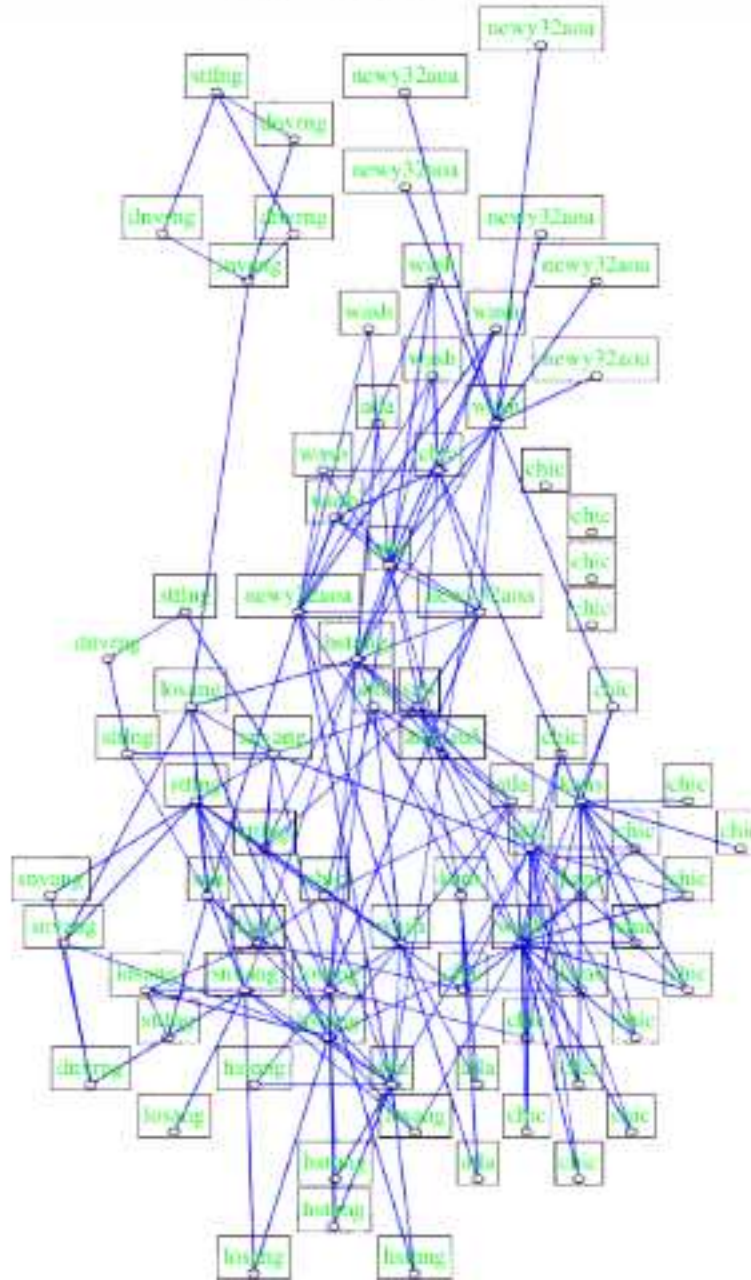

The Importance of Aliases

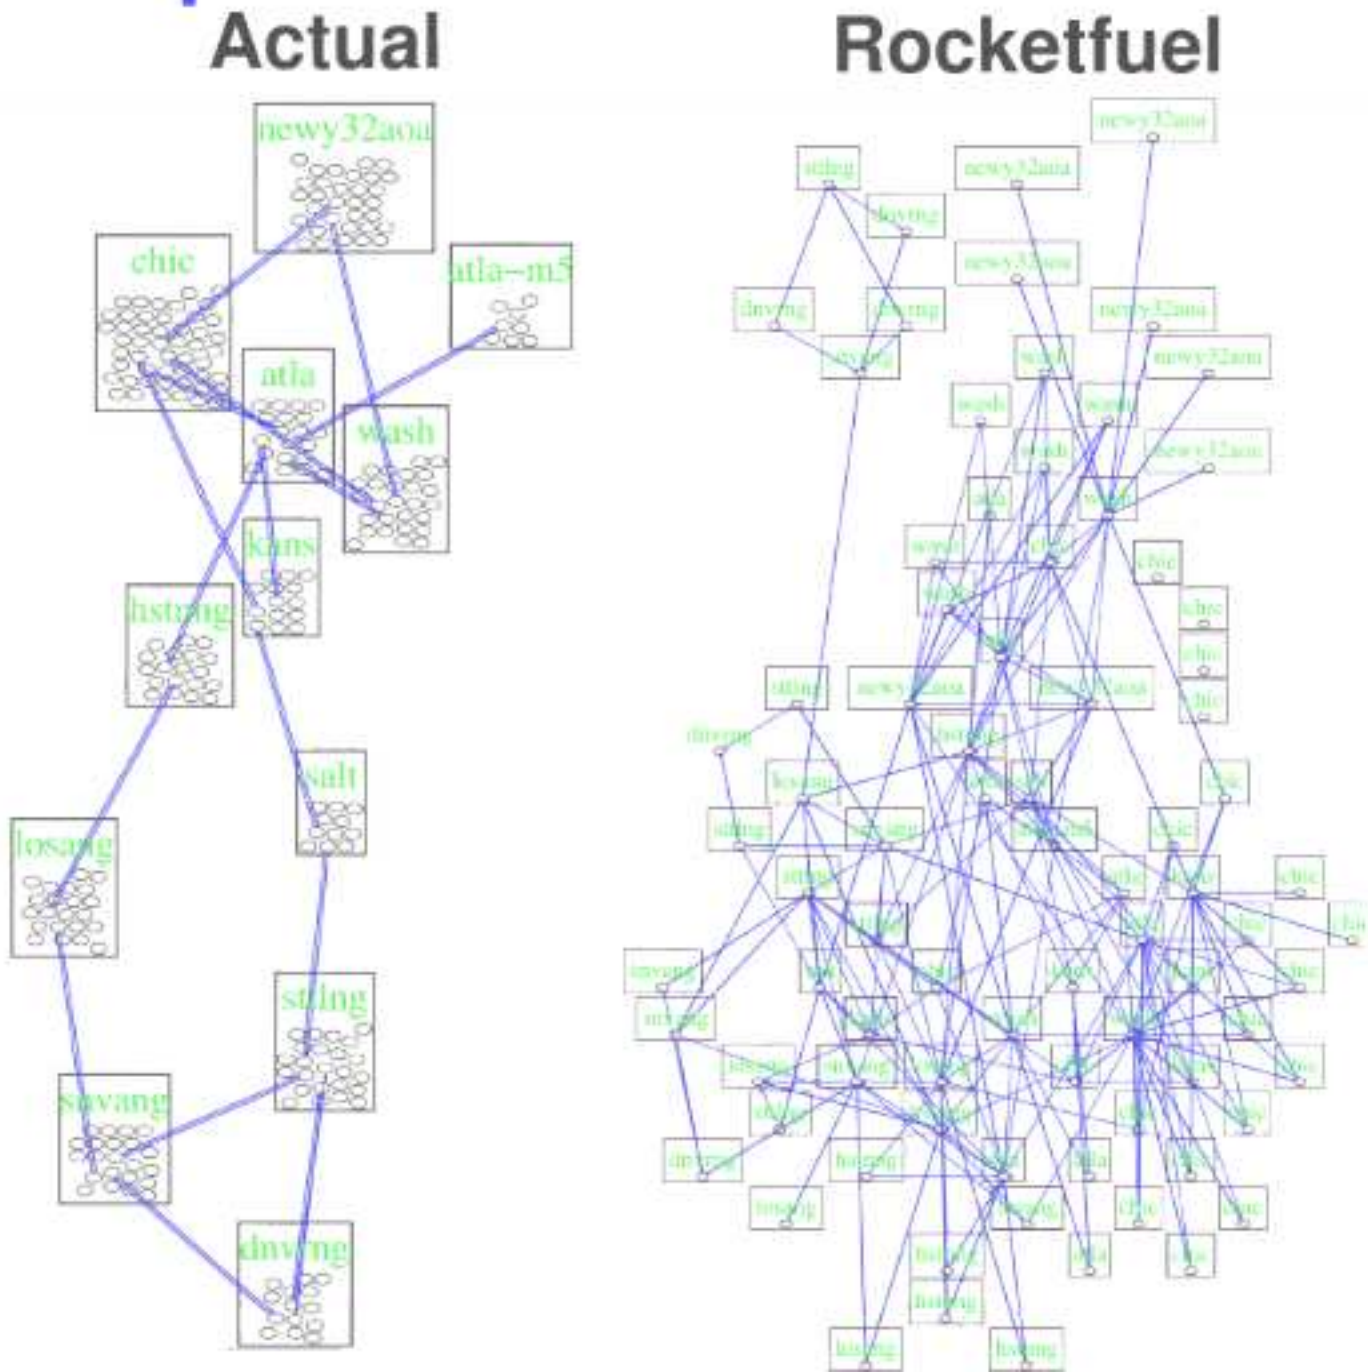

Record Route Finds Aliases

Record Route Finds Aliases

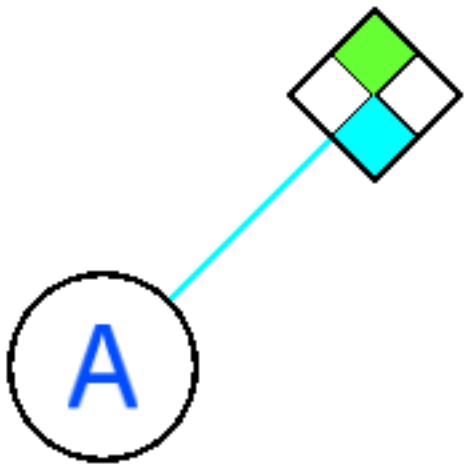

Record Route Finds Aliases

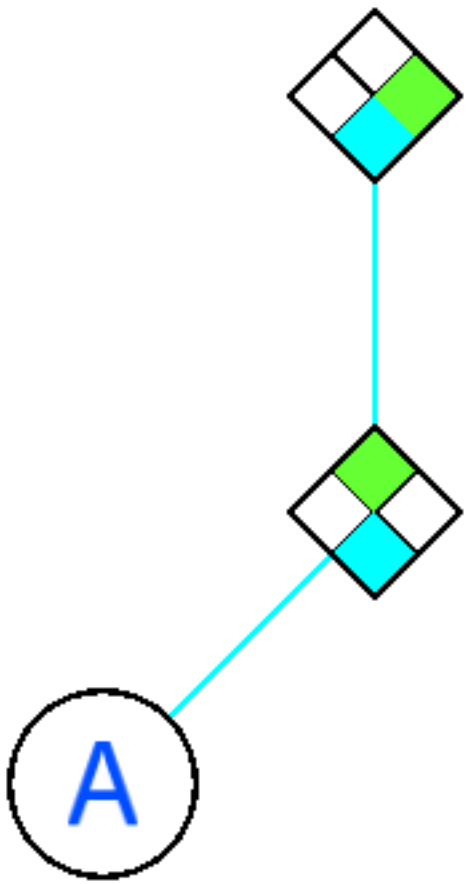

Record Route Finds Aliases

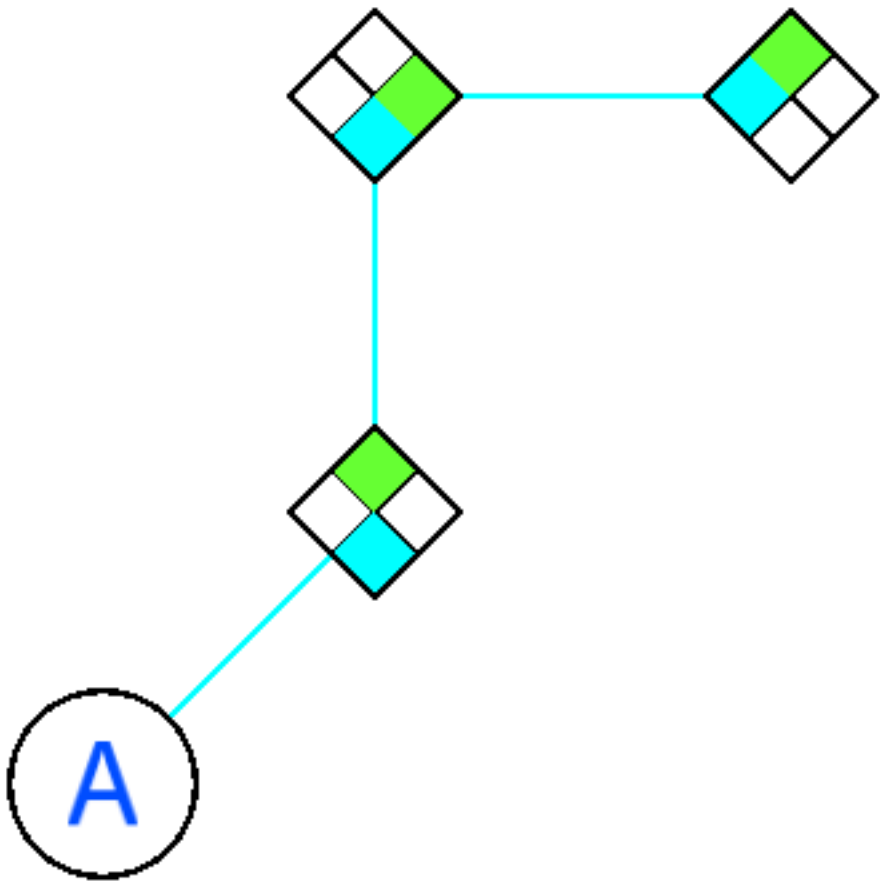

Record Route Finds Aliases

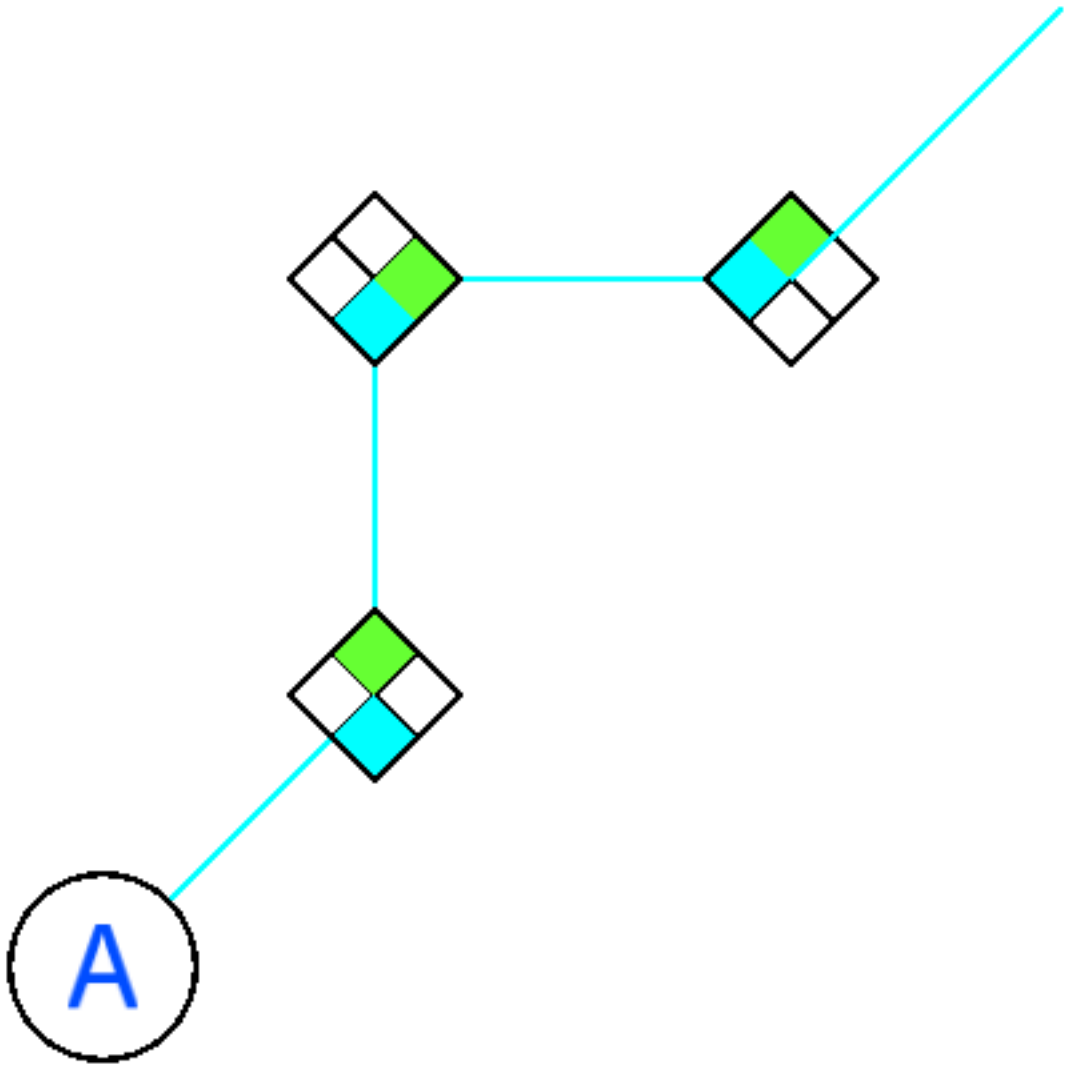

Record Route Finds Aliases

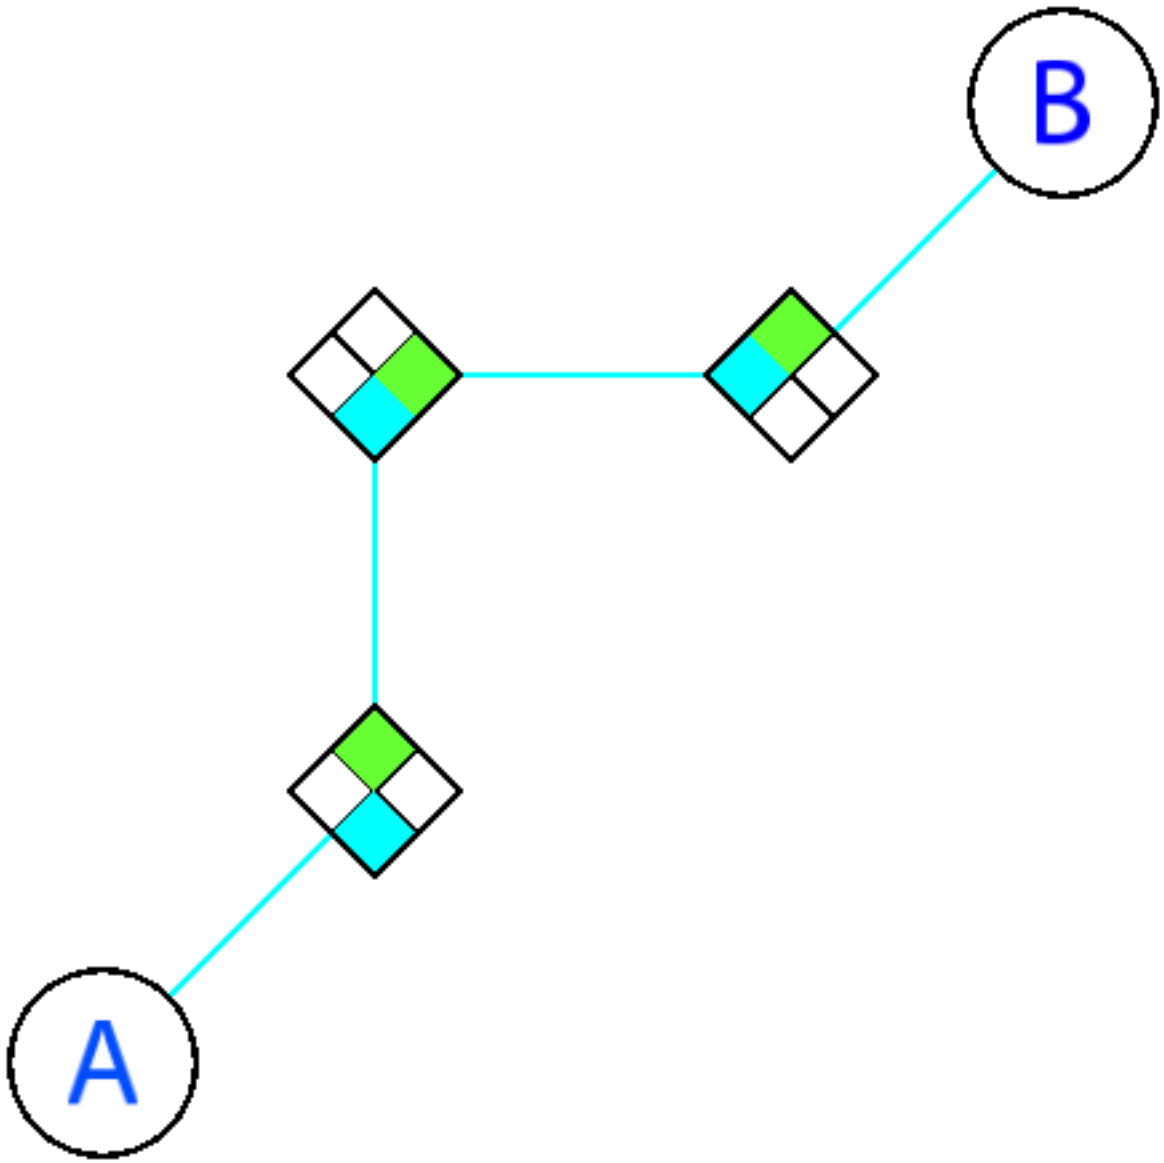

Record Route Finds Aliases

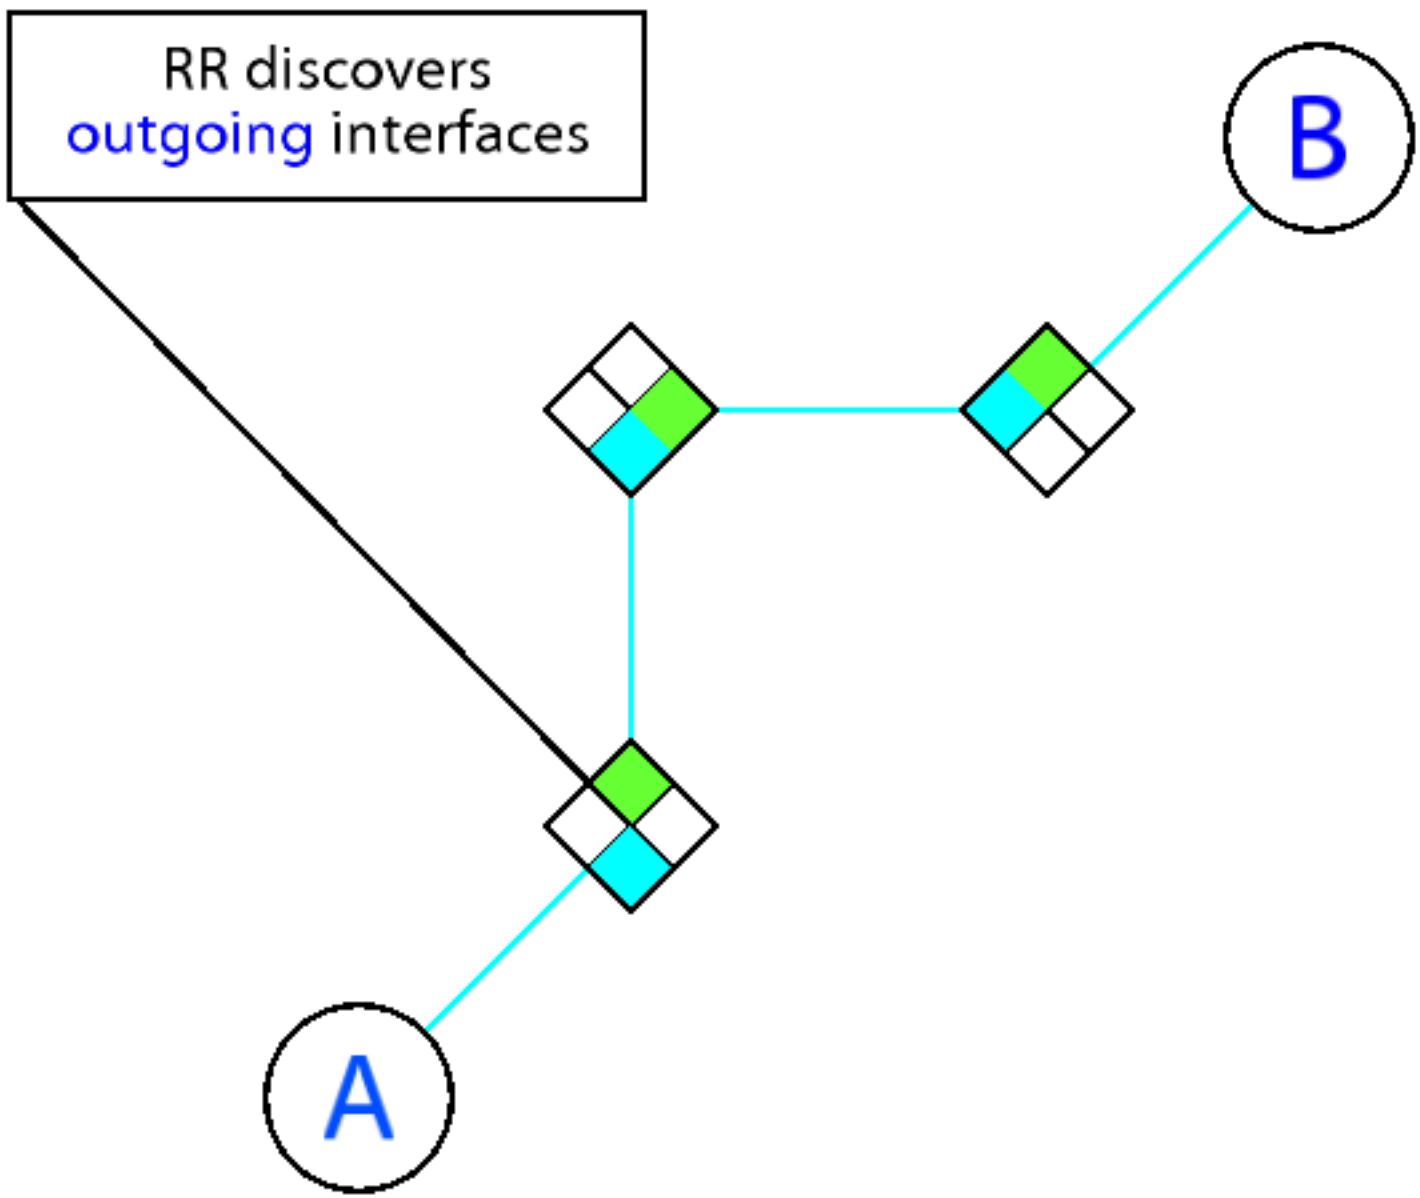

Record Route Finds Aliases

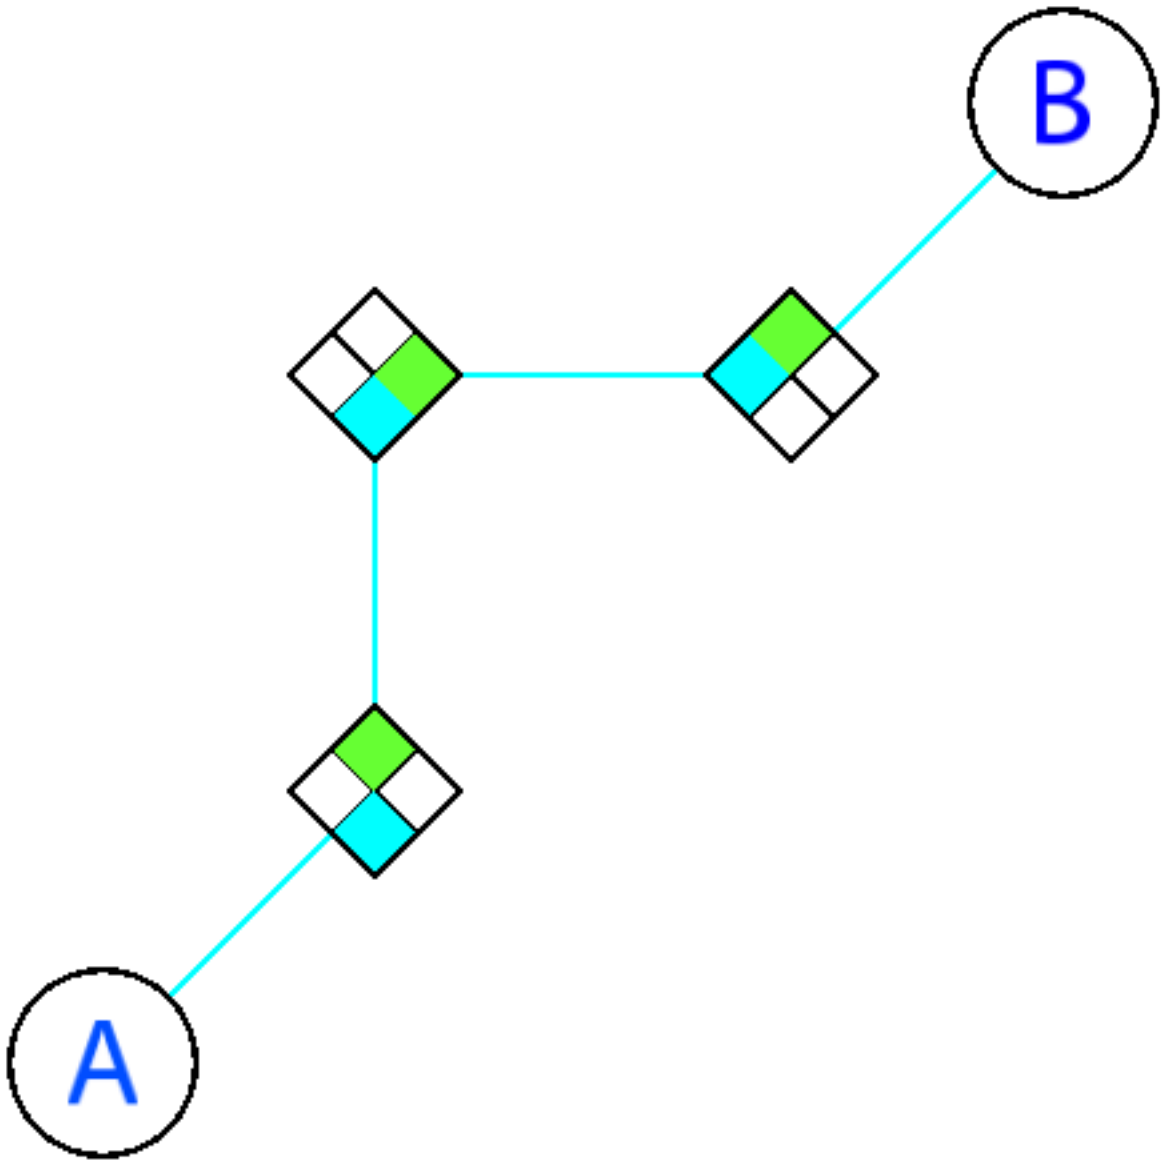

Record Route Finds Aliases

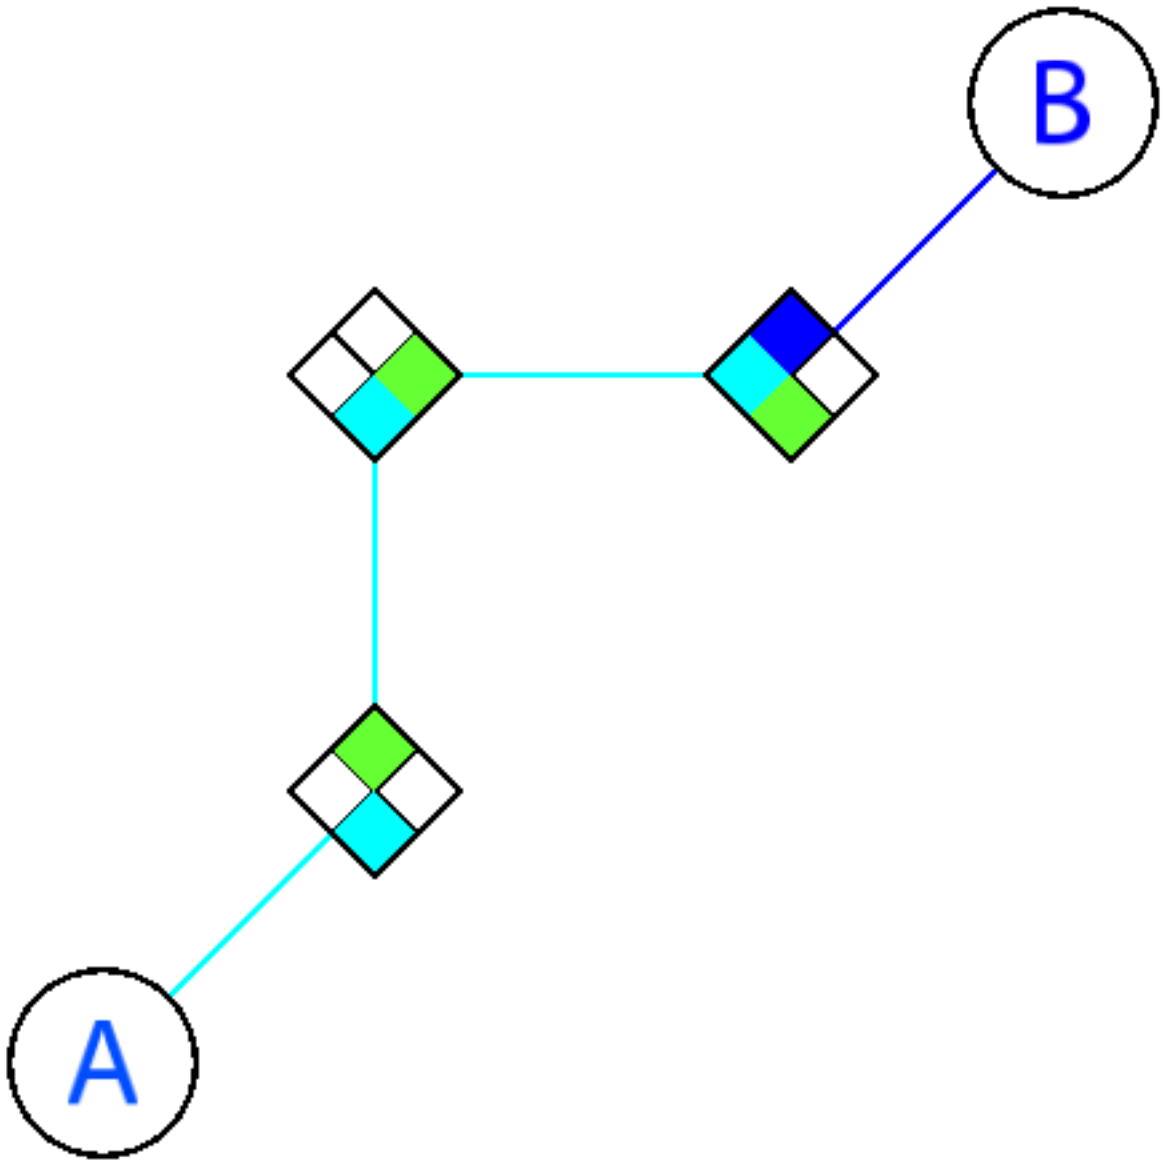

Record Route Finds Aliases

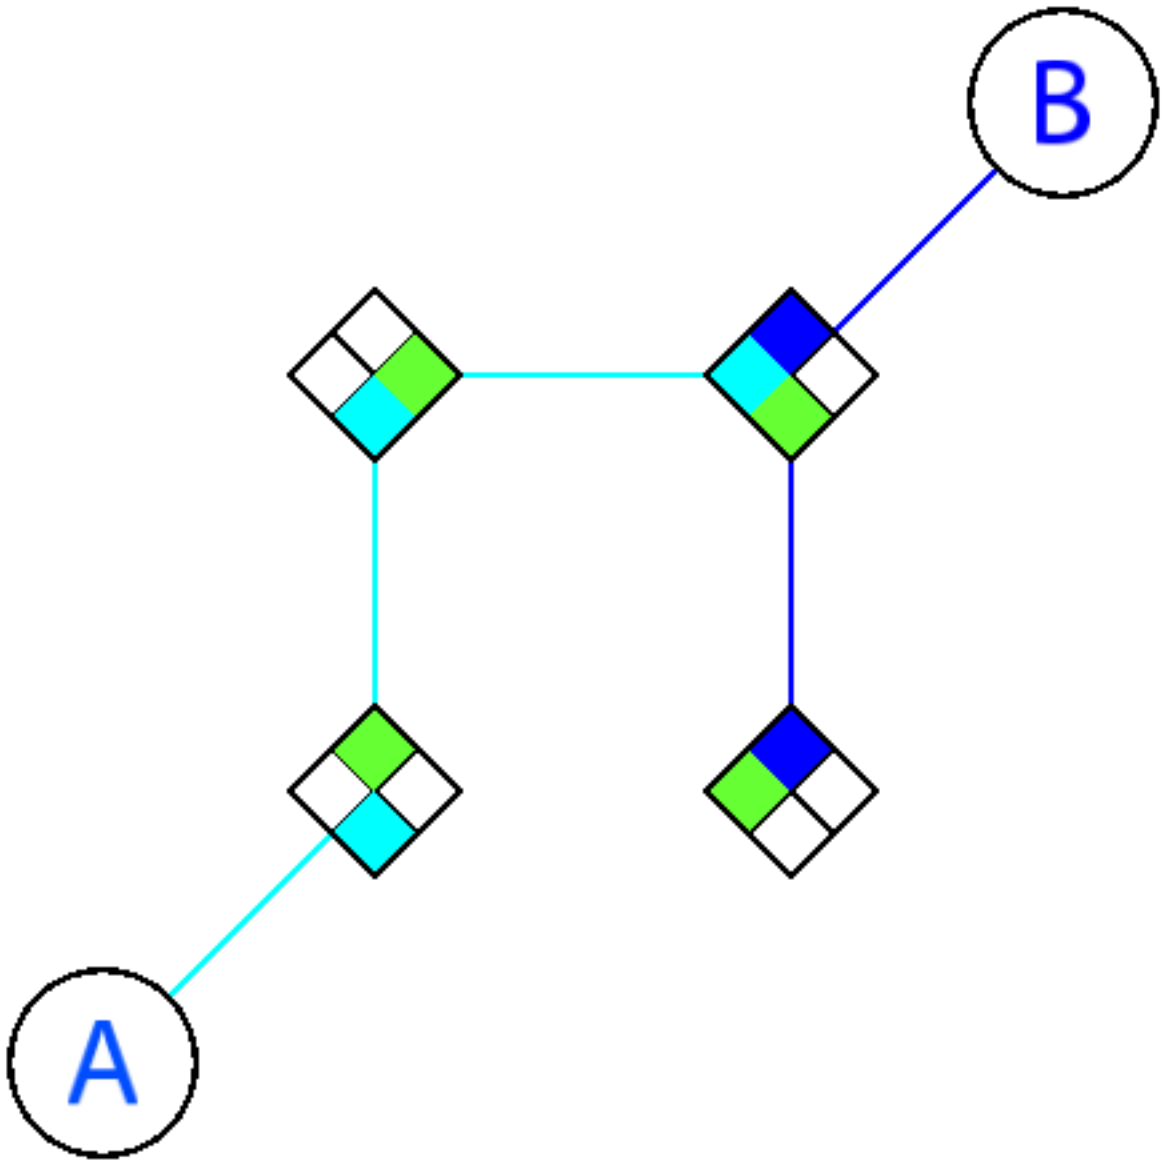

Record Route Finds Aliases

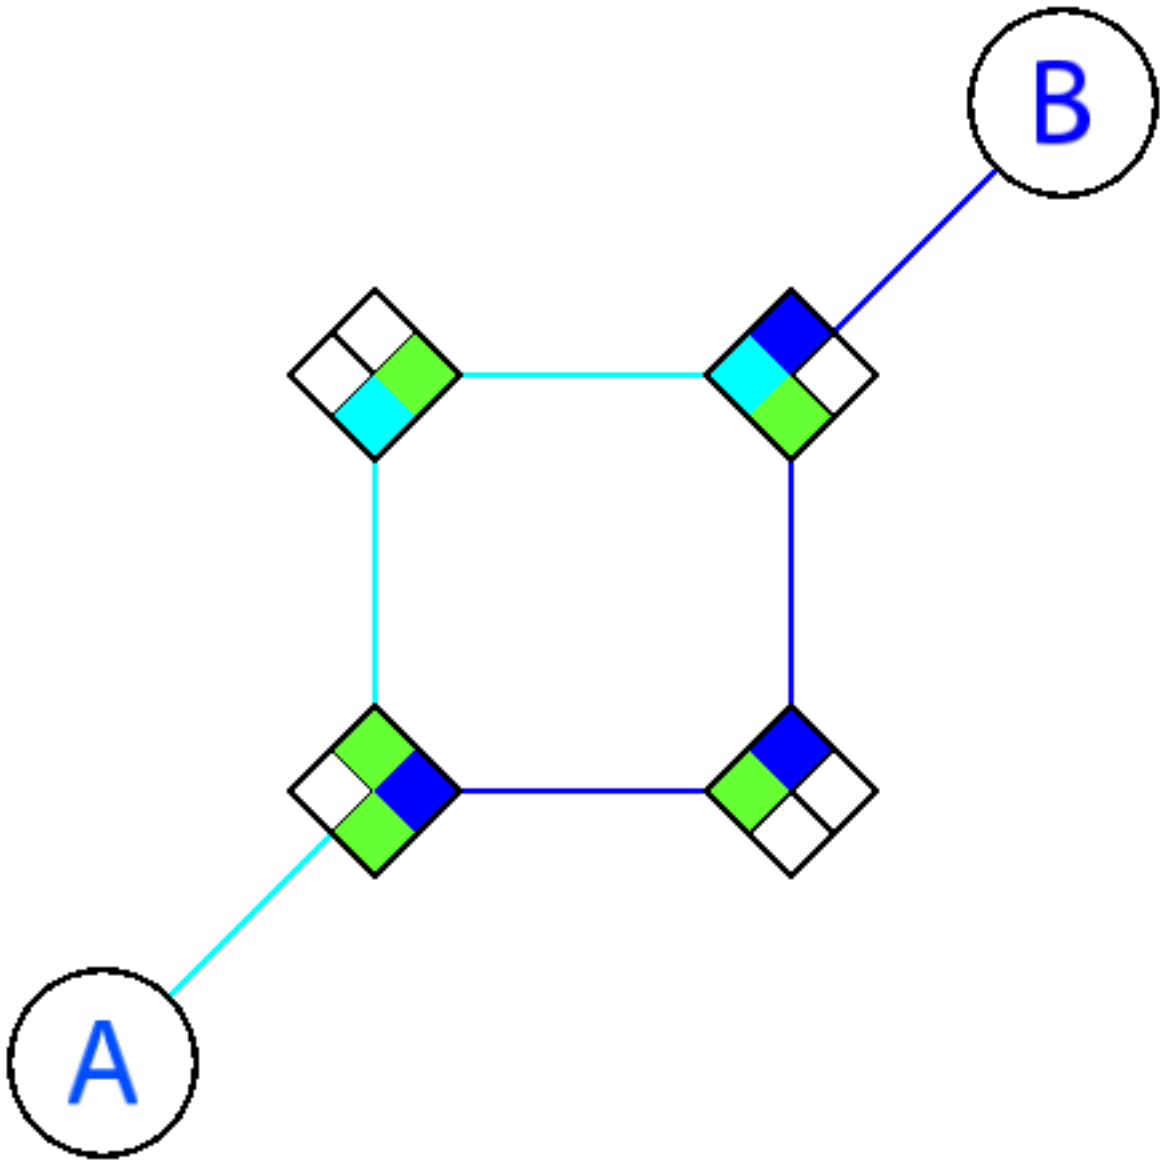

Record Route Finds Aliases

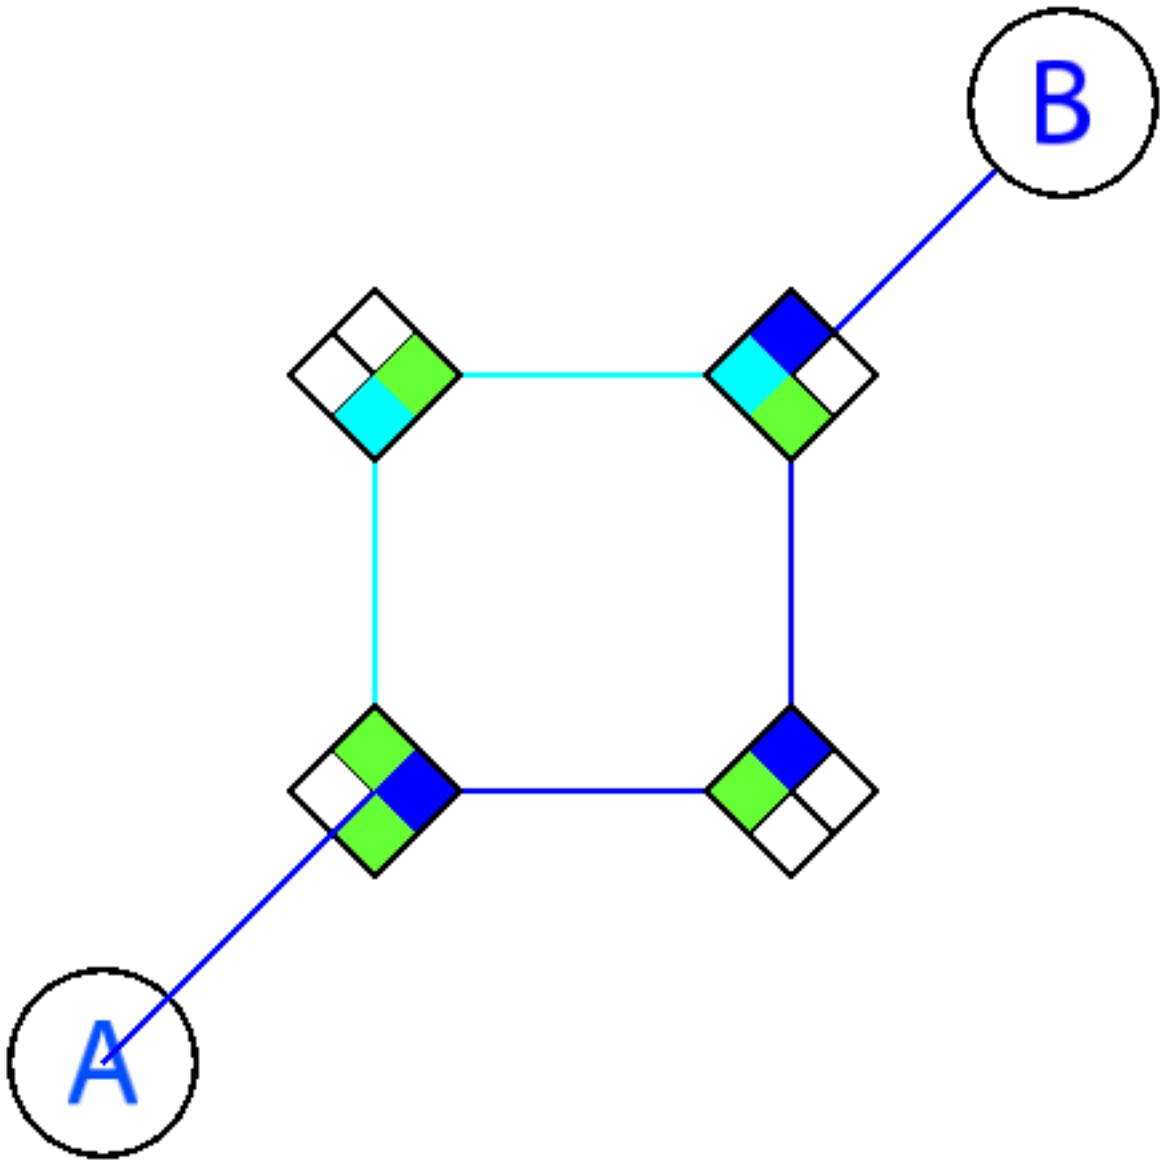

Record Route Finds Aliases

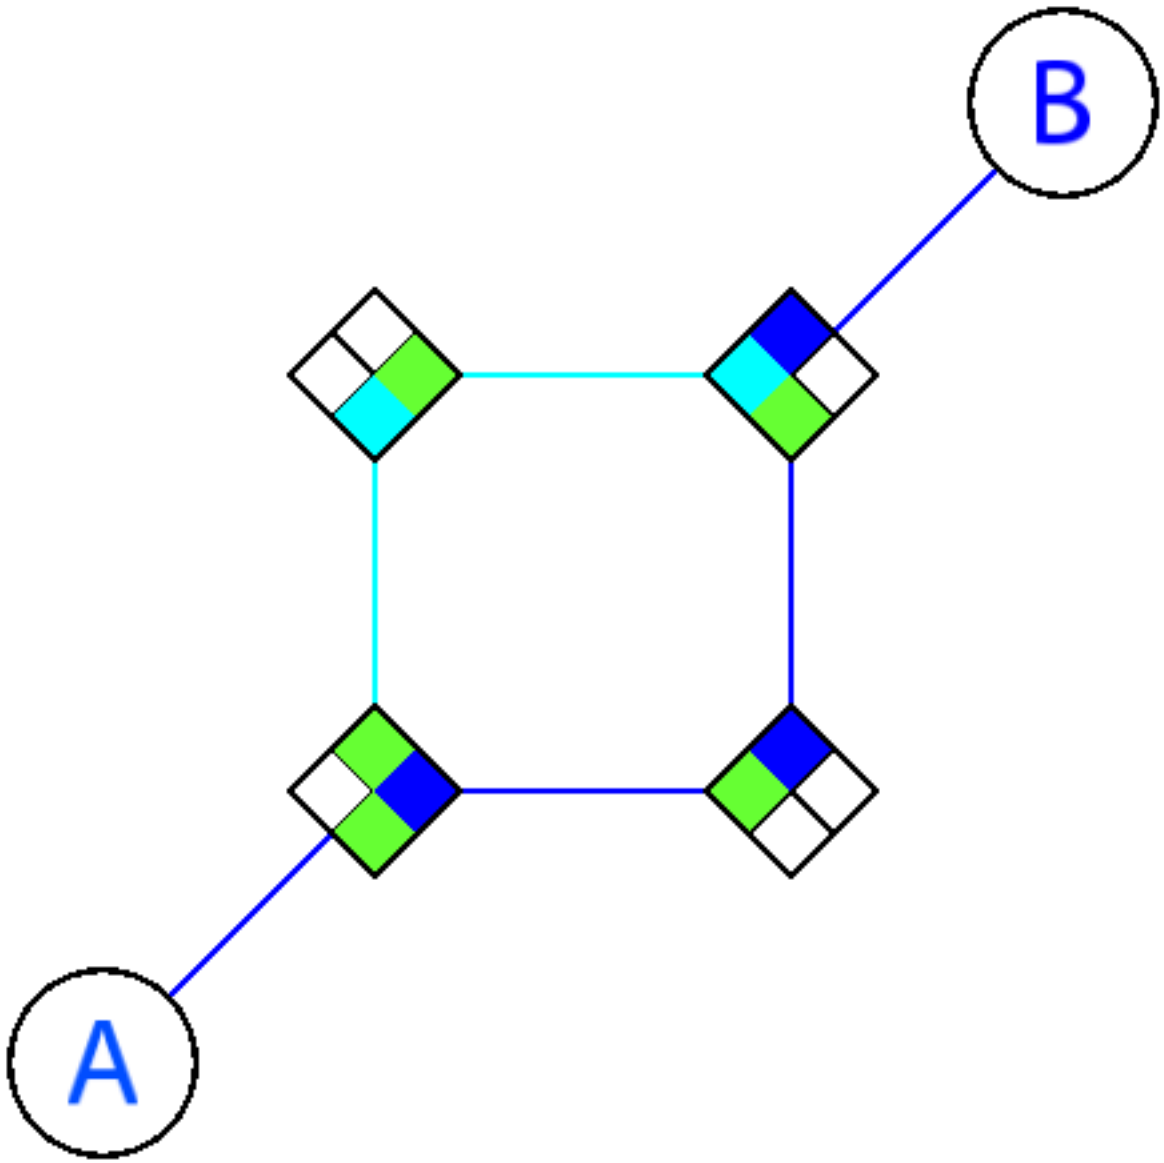

Record Route Finds Aliases

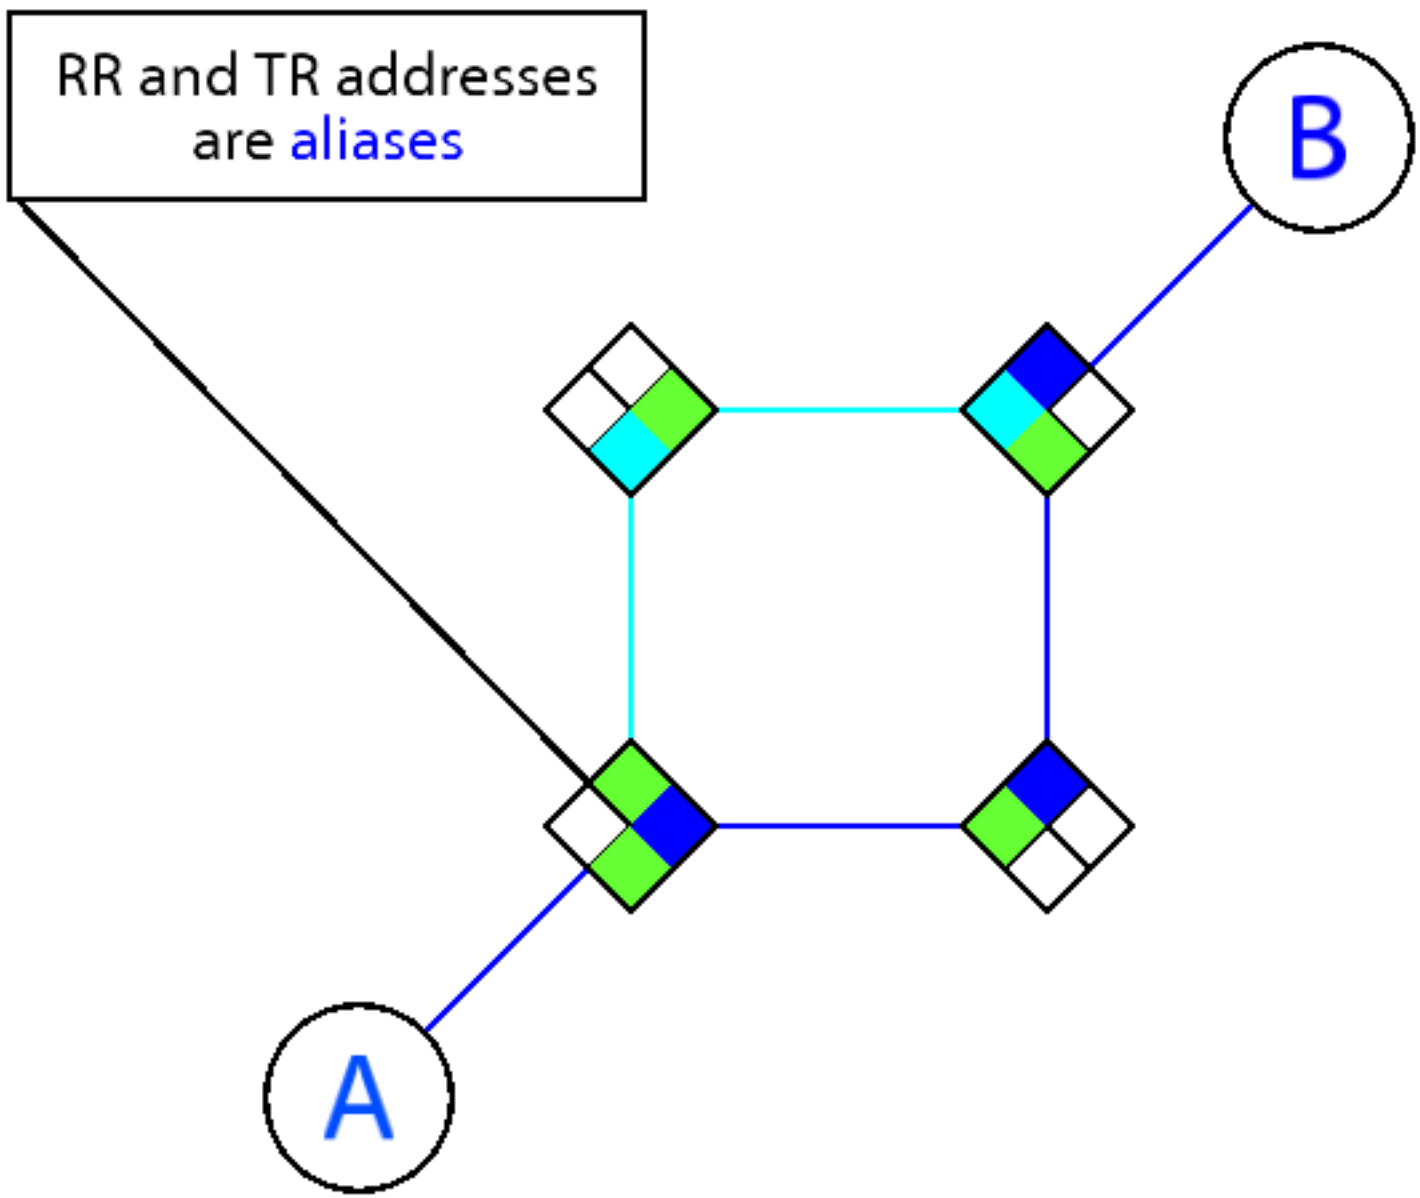

Record Route Finds Aliases

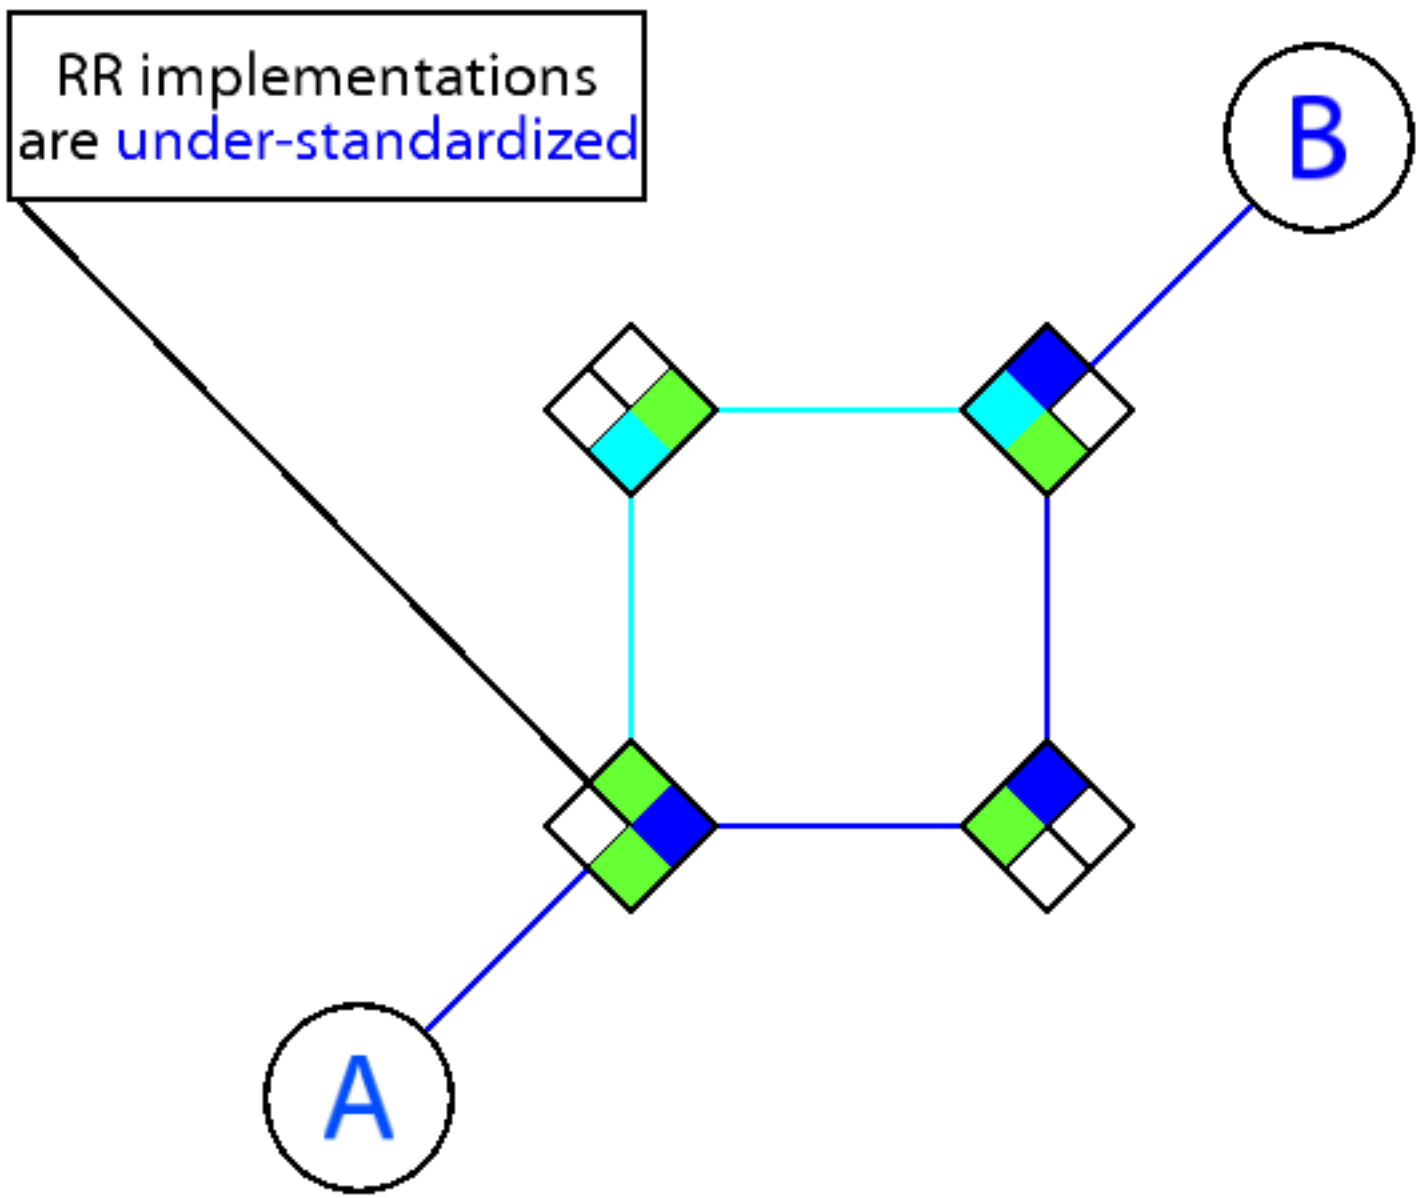

RR Implementations

	TTL decremented before RR field updated	Interface	Decrement TTL
Departing (61.9)	Yes	Outgoing	Yes
MPLS (13.3%)	Like Departing, but does not implement RR on MPLS interfaces		
NotImpl (9.1%)	--	--	Yes
Arriving (7.1%)	No	Outgoing or loopback	Yes
Lazy (5.8%)	?	?	No if RR is set
Mixed (2.7%)	No	Outgoing, but incoming if TTL=1	Yes
Hidden (<1%)	?	--	No

RR Diversity

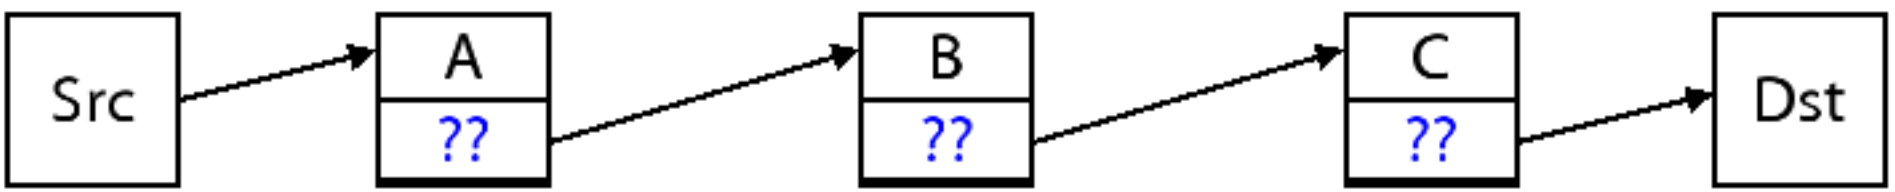

RR Diversity

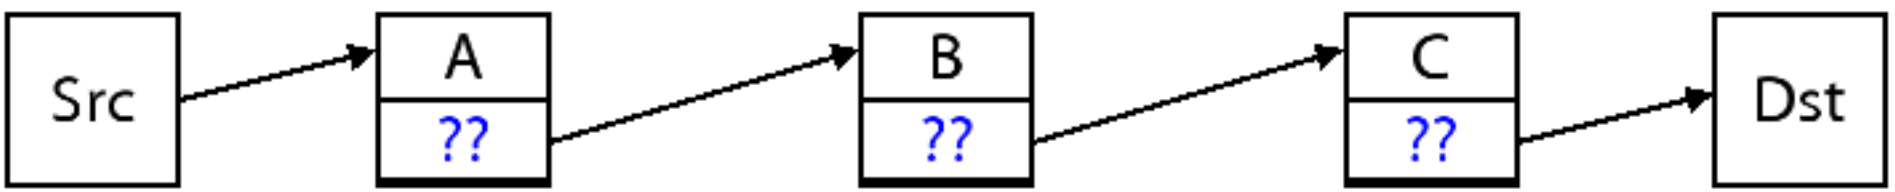

RR Diversity

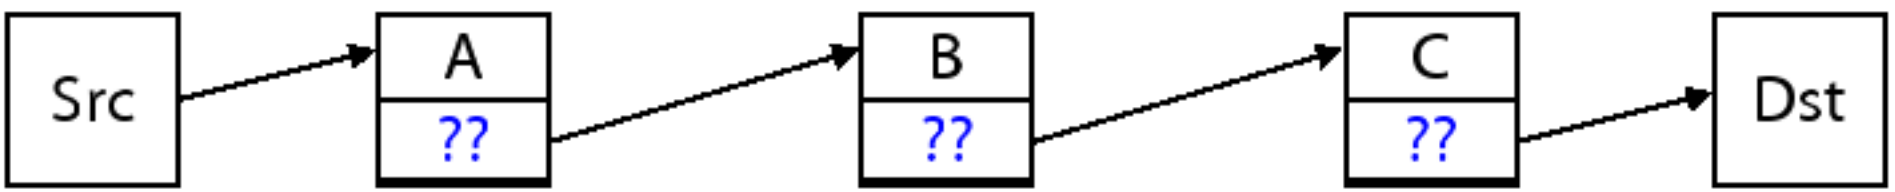

RR Diversity

Src

Record route

Traceroute

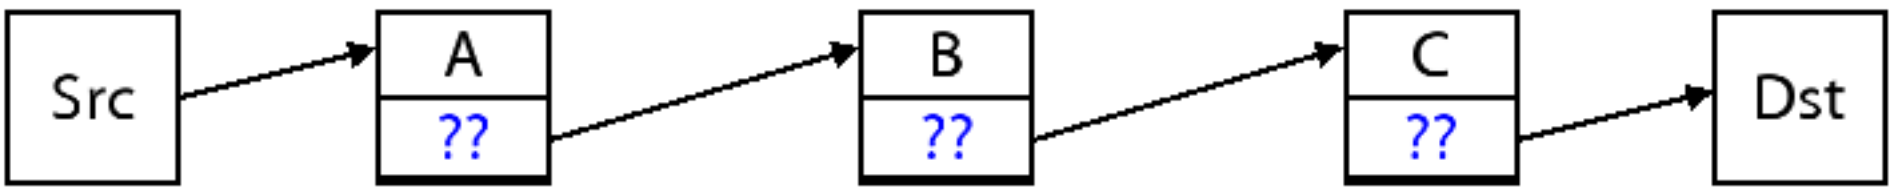

RR Diversity

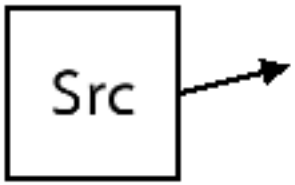

Record route

Traceroute

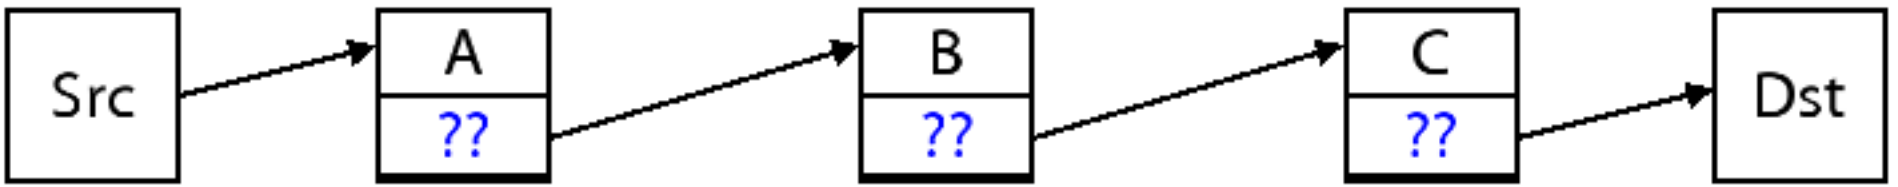

RR Diversity

Record route

Traceroute

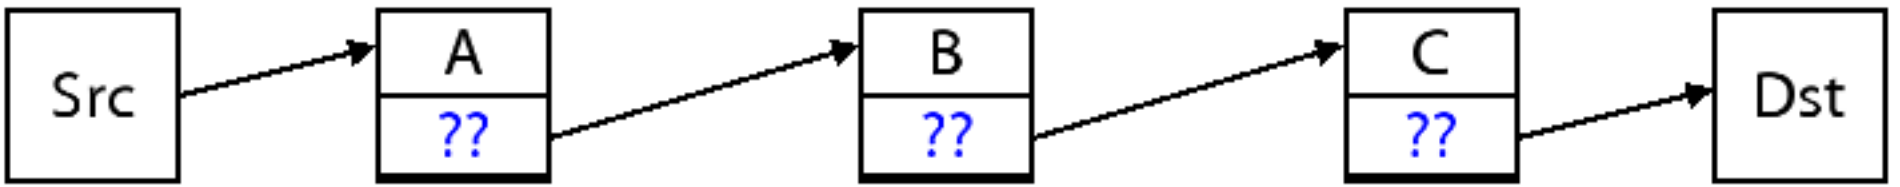

RR Diversity

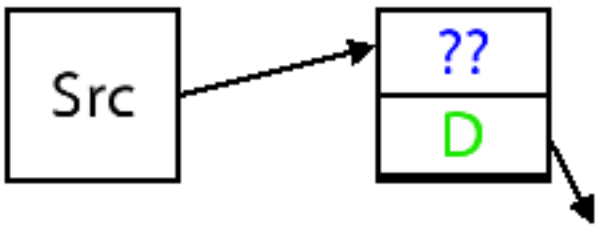

Record route

Traceroute

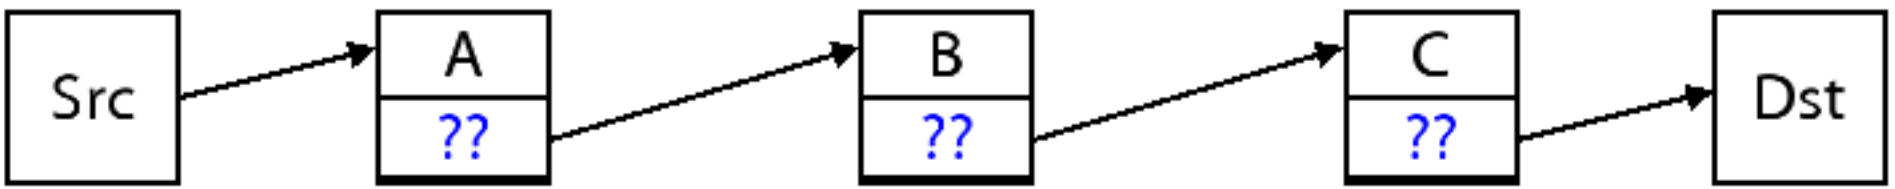

RR Diversity

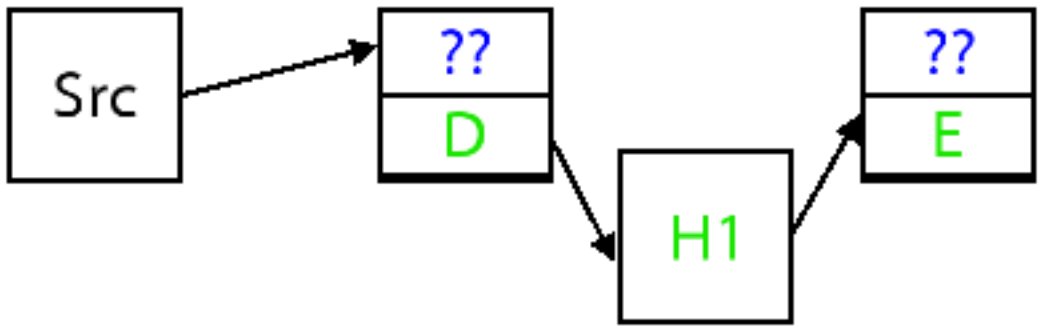

Record route

Traceroute

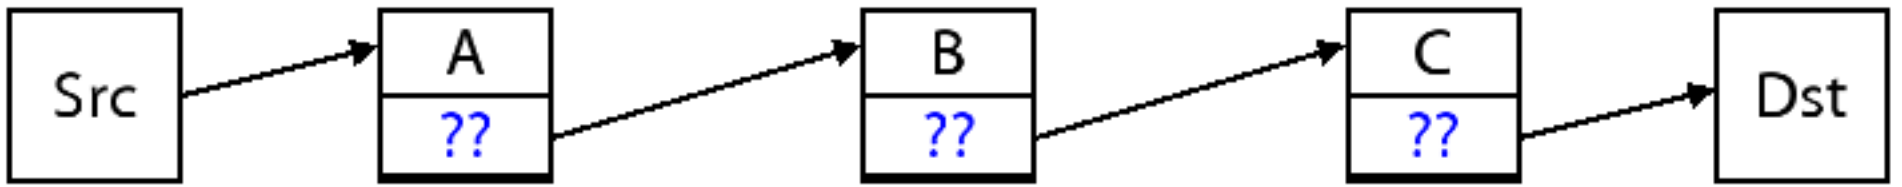

RR Diversity

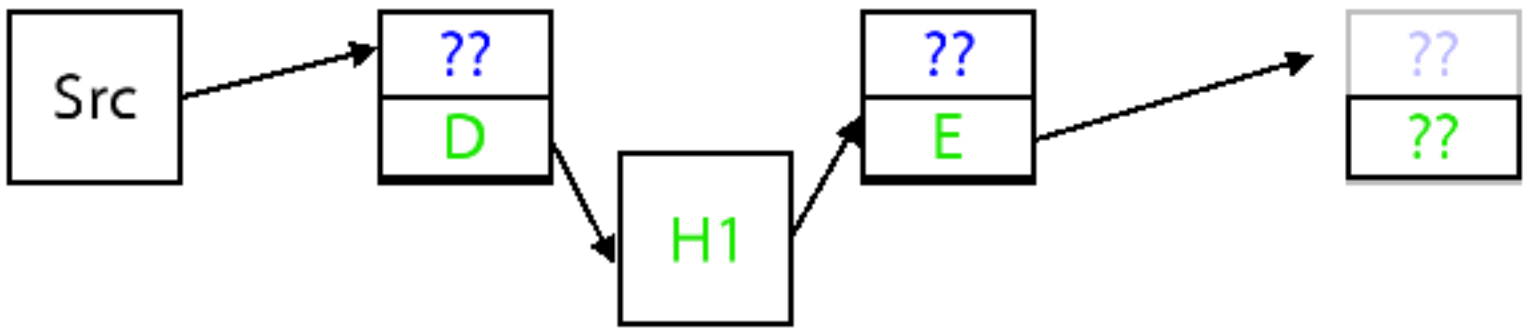

Record route

Traceroute

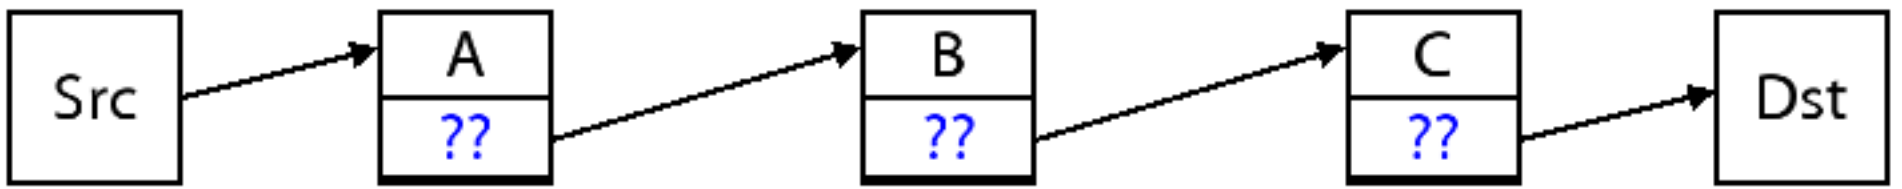

RR Diversity

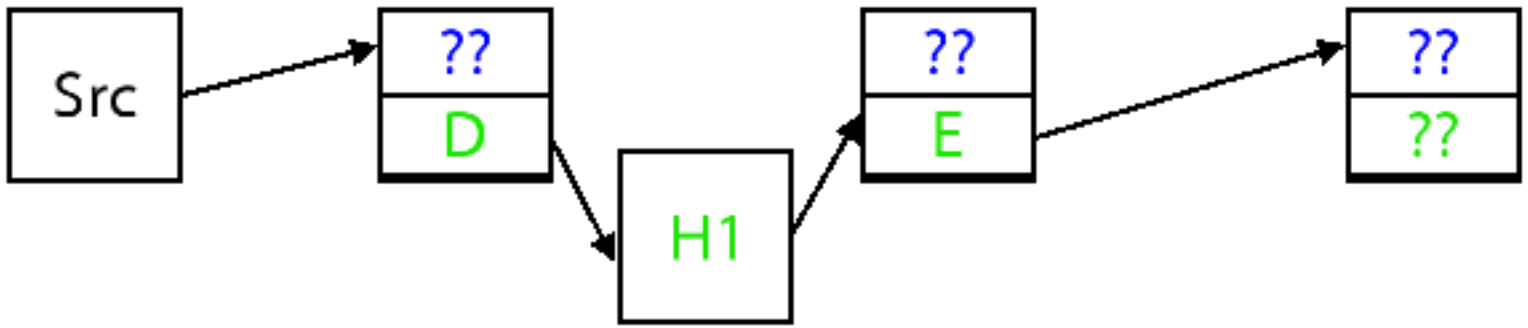

Record route

Traceroute

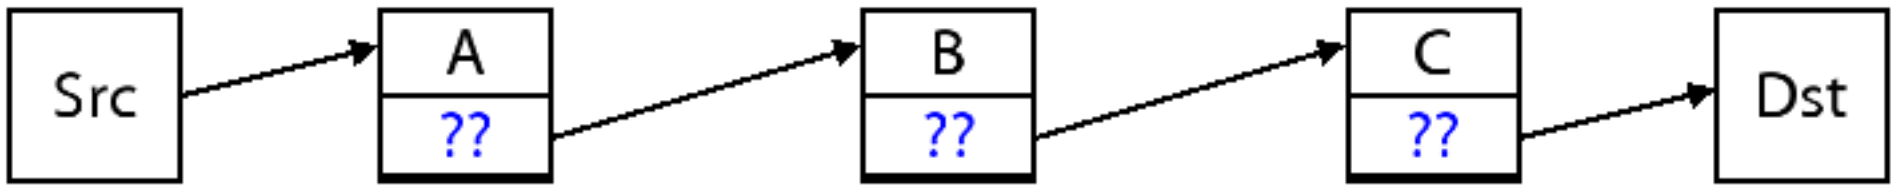

RR Diversity

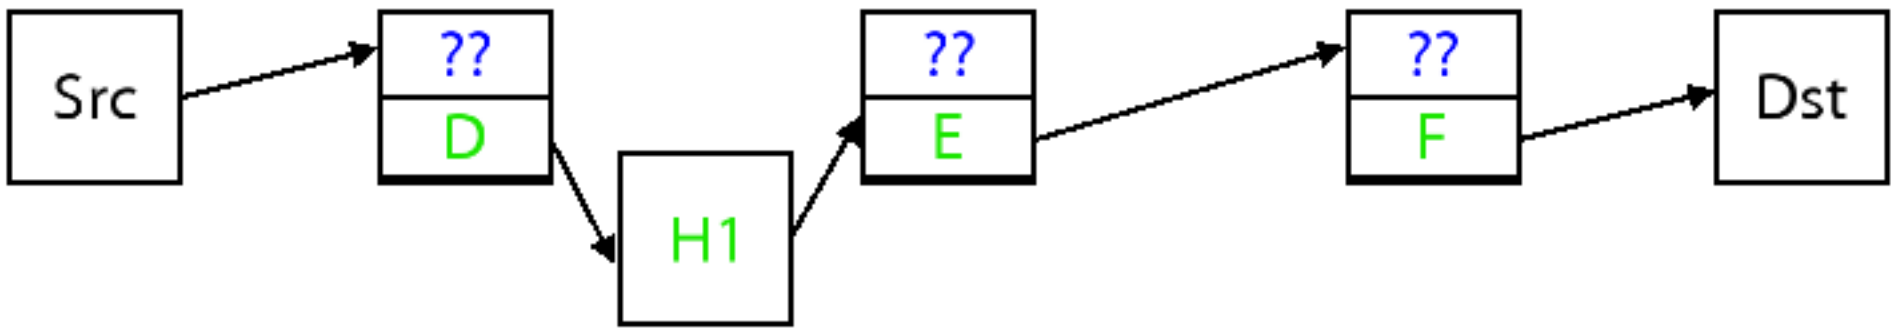

Record route

Traceroute

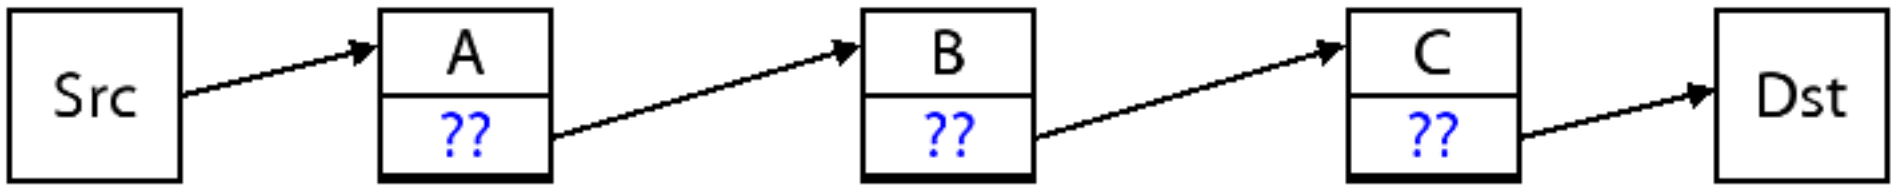

RR Diversity

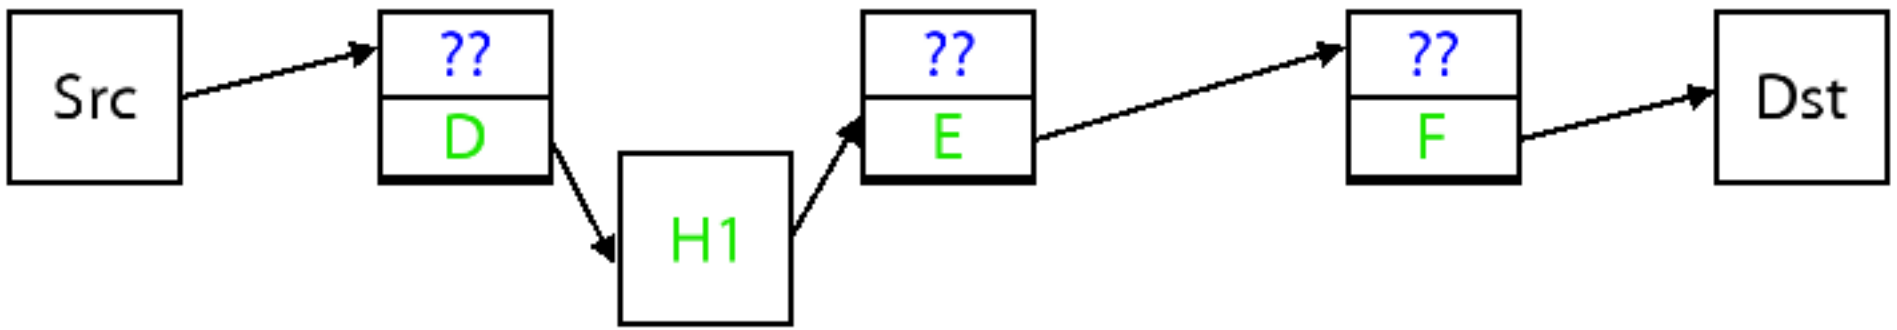

Record route

Traceroute

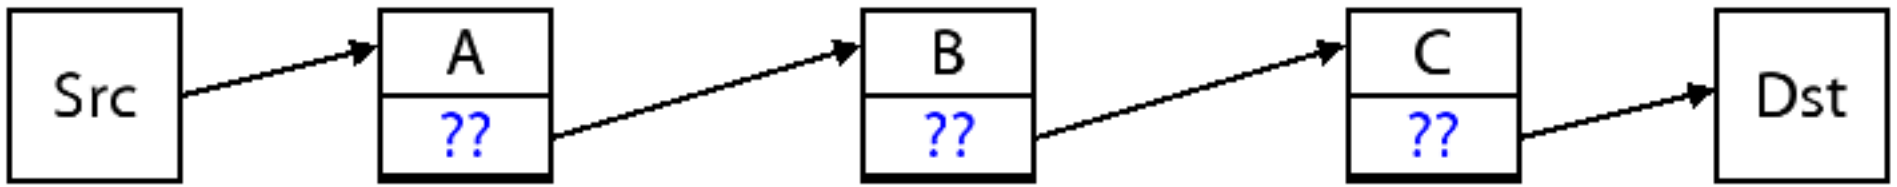

Merged TR and RR

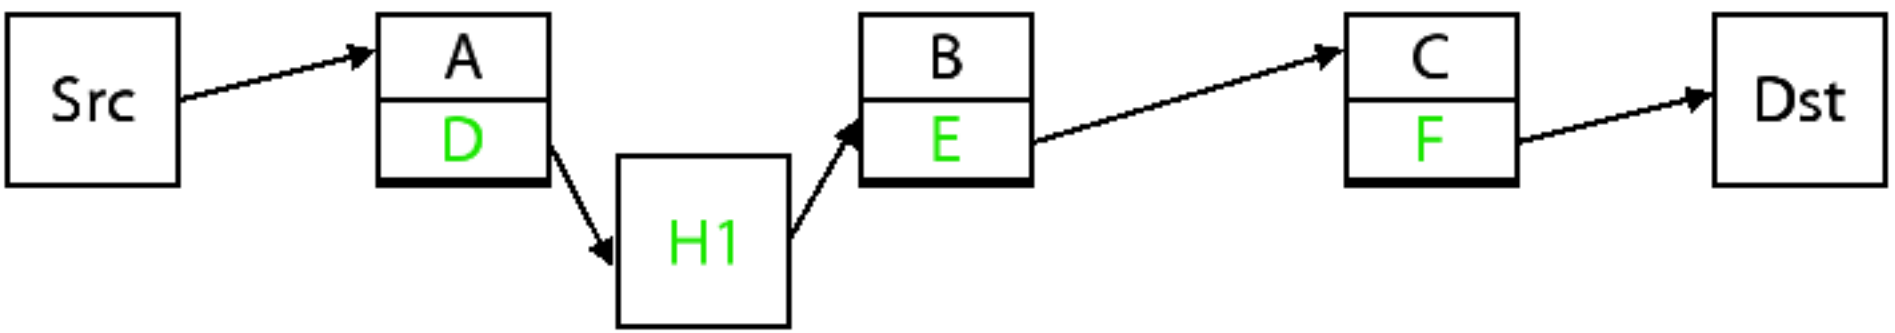

Merged TR and RR

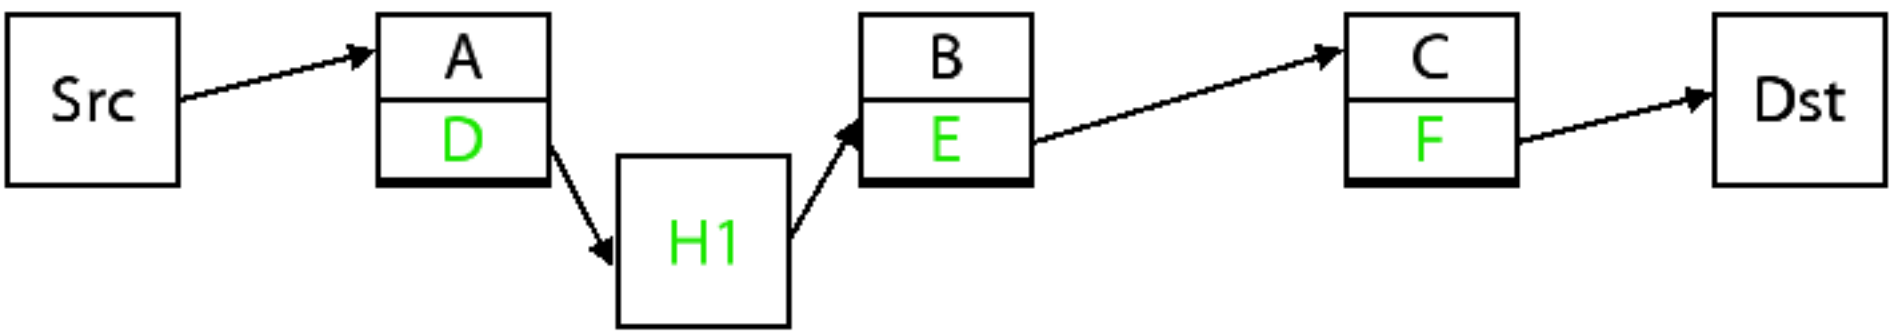

RR is Hard to Use

- Implementation diversity adds ambiguity to RR and TR matching
  - Each traceroute hop adds 0-4 RR addresses
  - 7 implementations times 7 implementations
  - Alias resolution requires implementation classification
  - Bad classifications lead to bad aliases
- RR probes take different paths

Example Merge

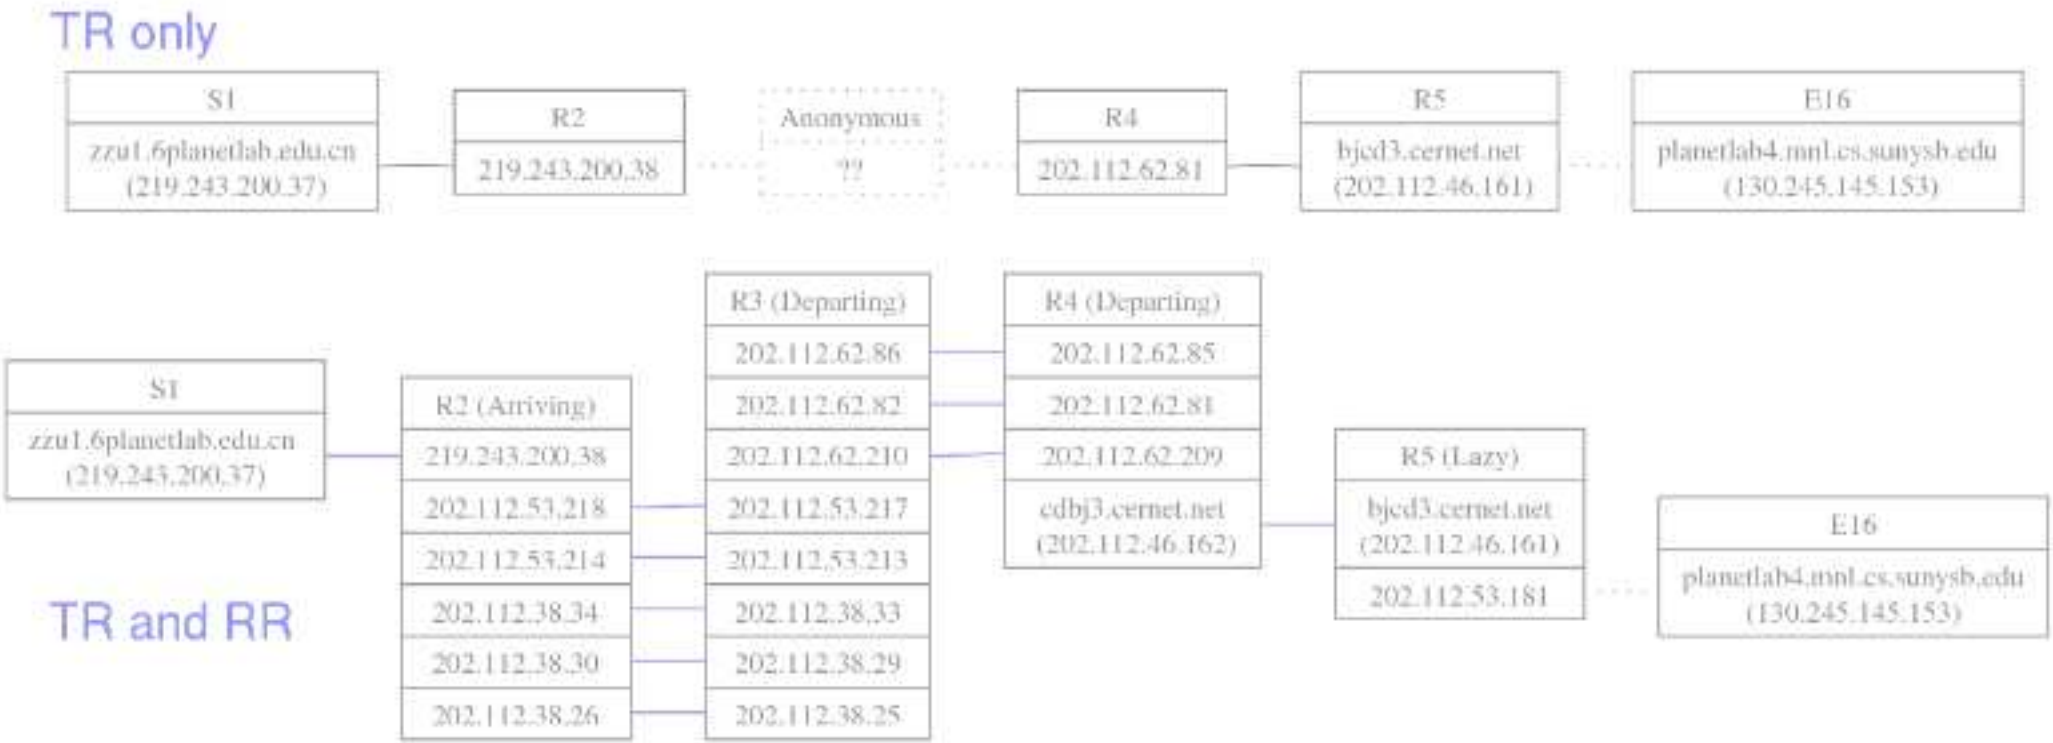

Example Merge

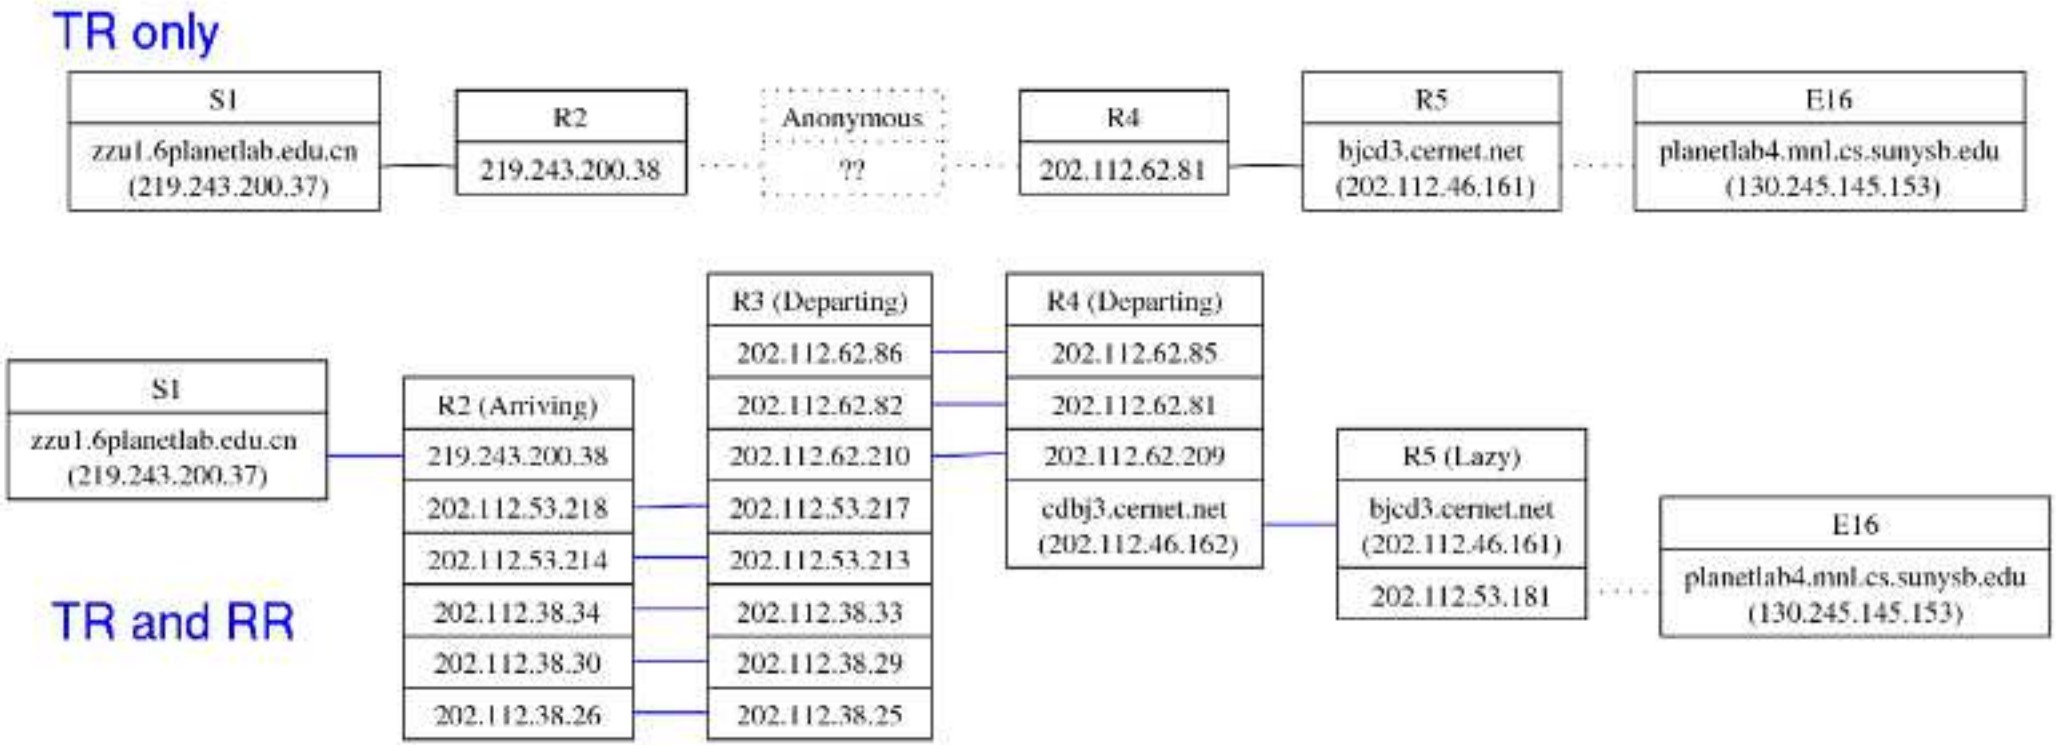

TR only

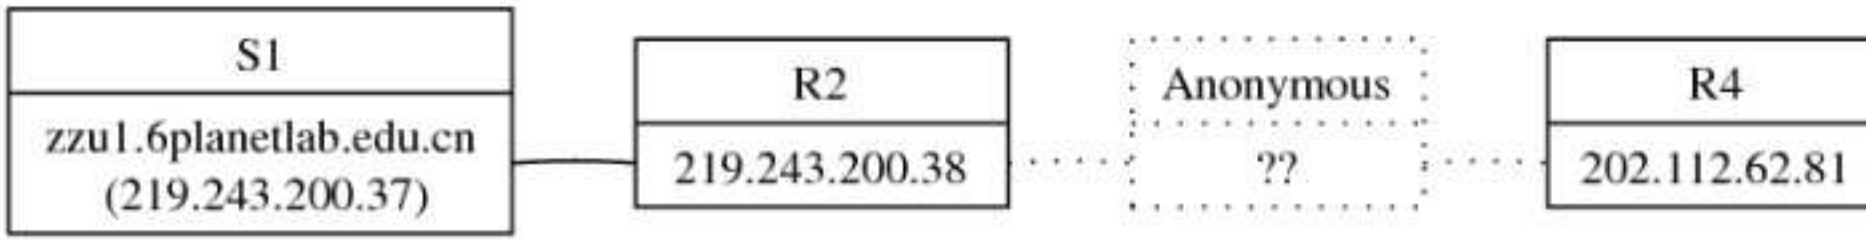

TR and RR

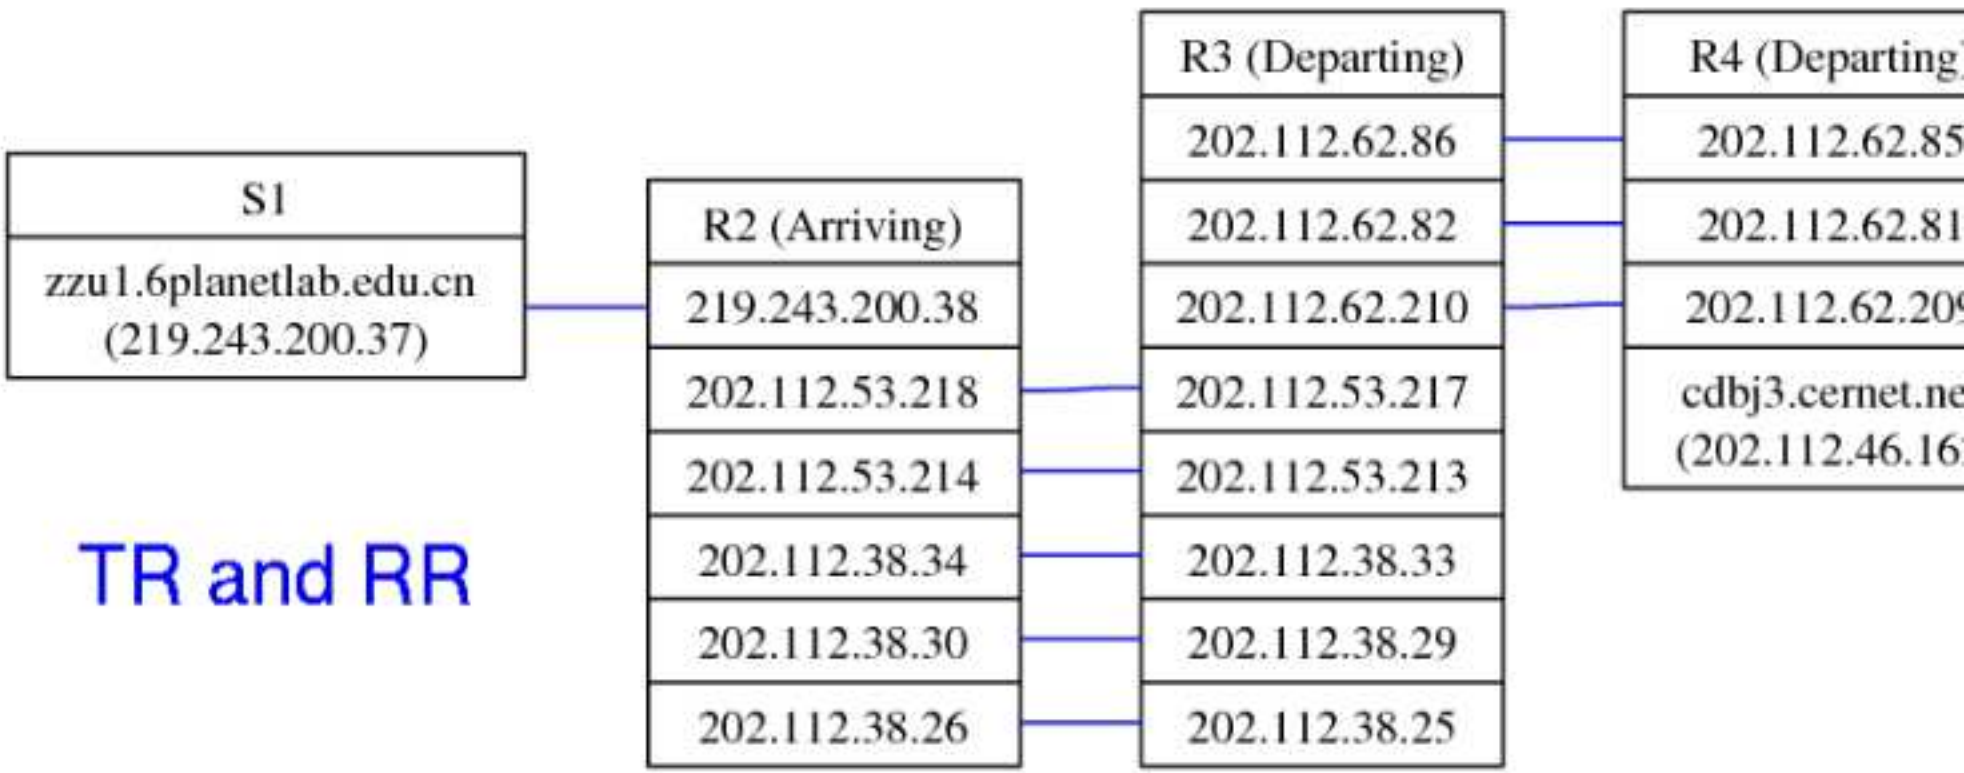

TR only

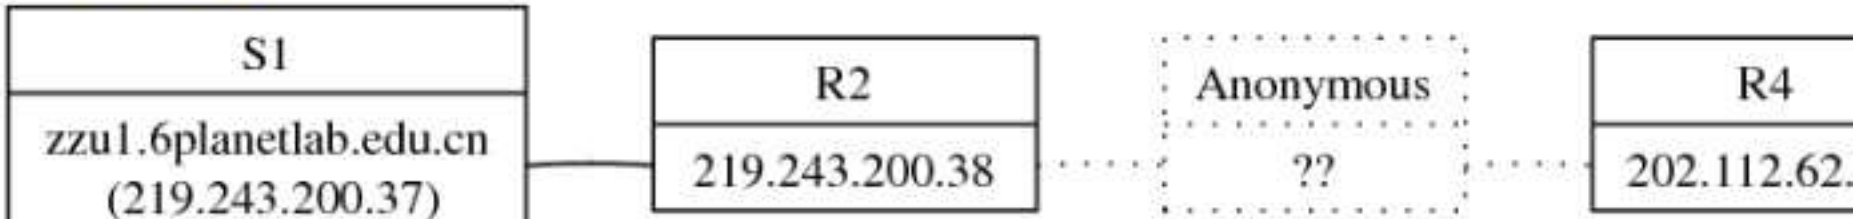

TR and RR

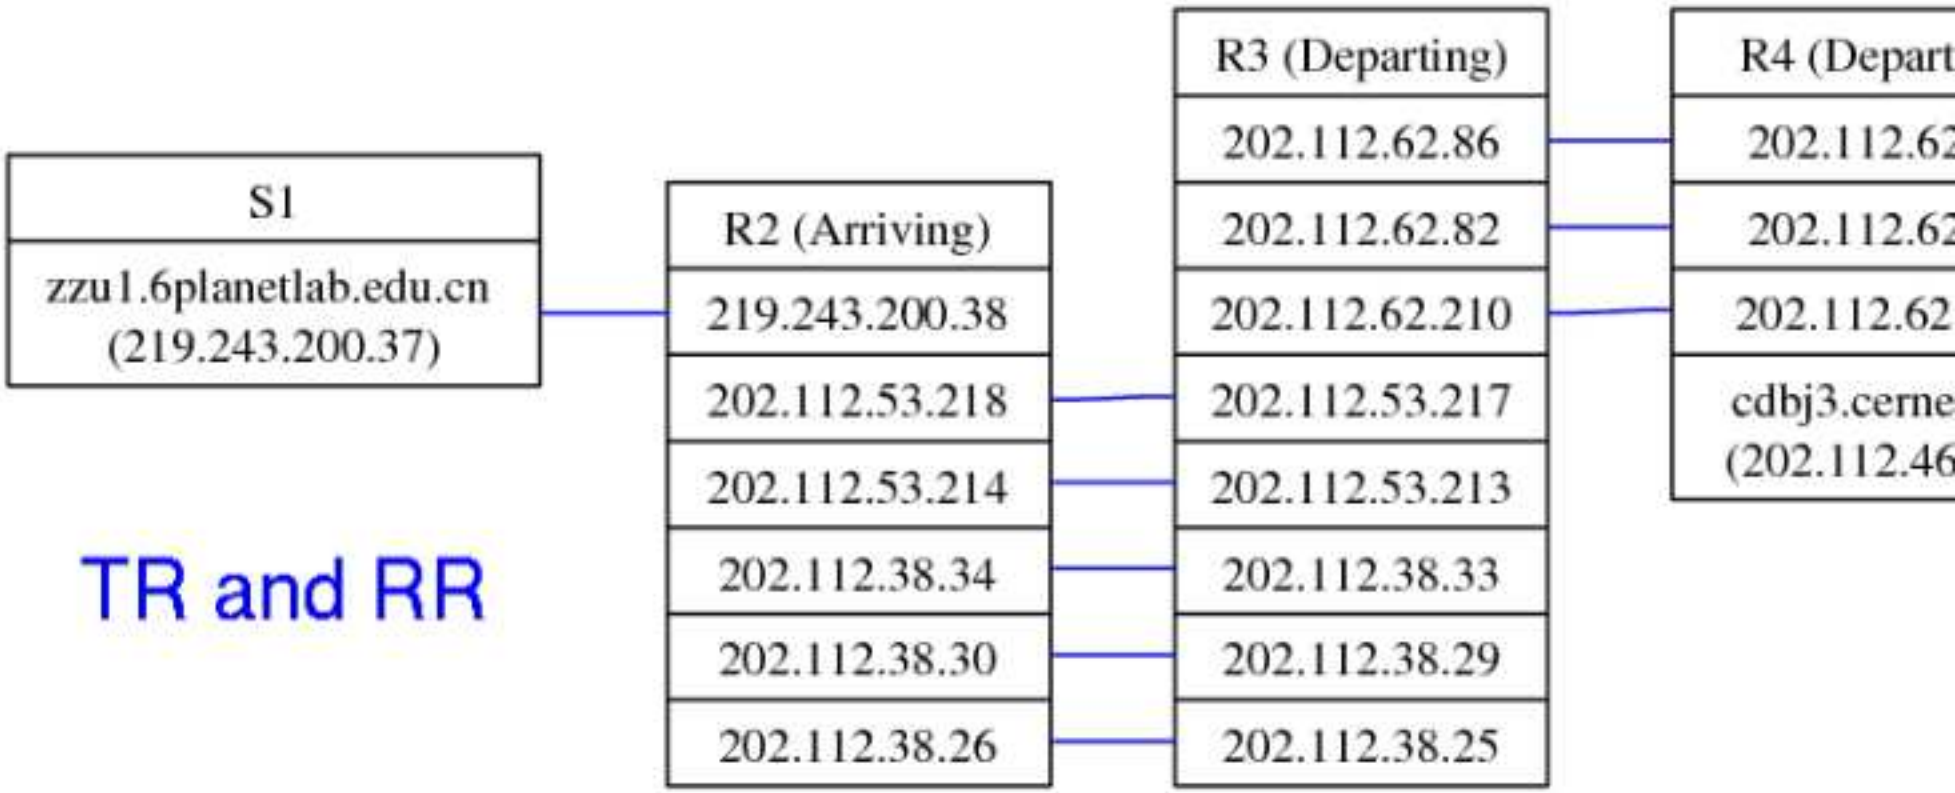

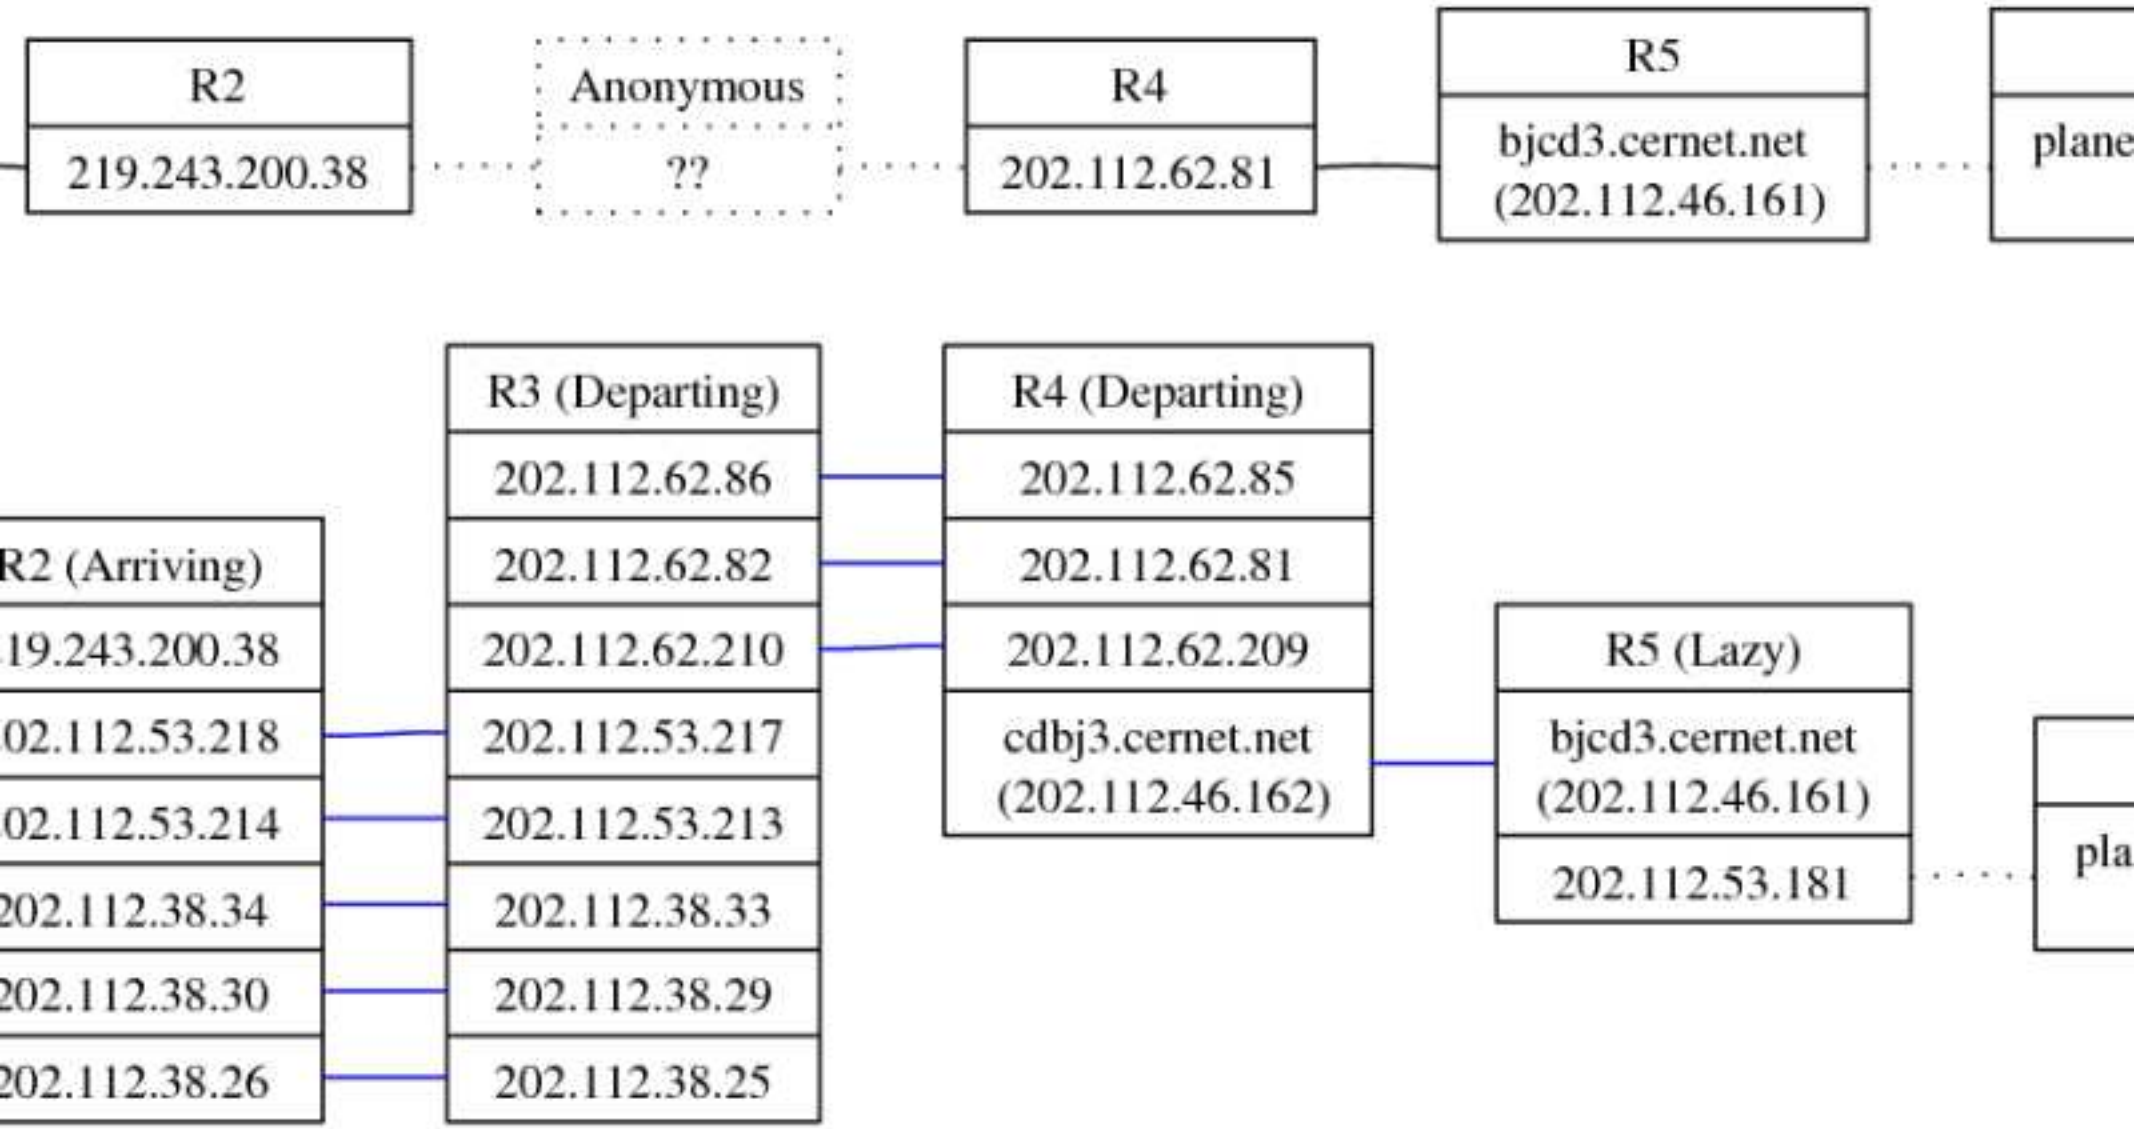
R2
219.243.200.38
Anonymous
??
R4
202.112.62.81
R5
bjcd3.cernet.net
(202.112.46.161)
plane
R3 (Departing)
202.112.62.86
202.112.62.82
202.112.62.210
202.112.53.217
202.112.53.213
202.112.38.33
202.112.38.29
202.112.38.25
R4 (Departing)
202.112.62.85
202.112.62.81
202.112.62.209
cdbj3.cernet.net
(202.112.46.162)
R5 (Lazy)
bjcd3.cernet.net
(202.112.46.161)
202.112.53.181
pla
R2 (Arriving)
19.243.200.38
02.112.53.218
02.112.53.214
202.112.38.34
202.112.38.30
202.112.38.26

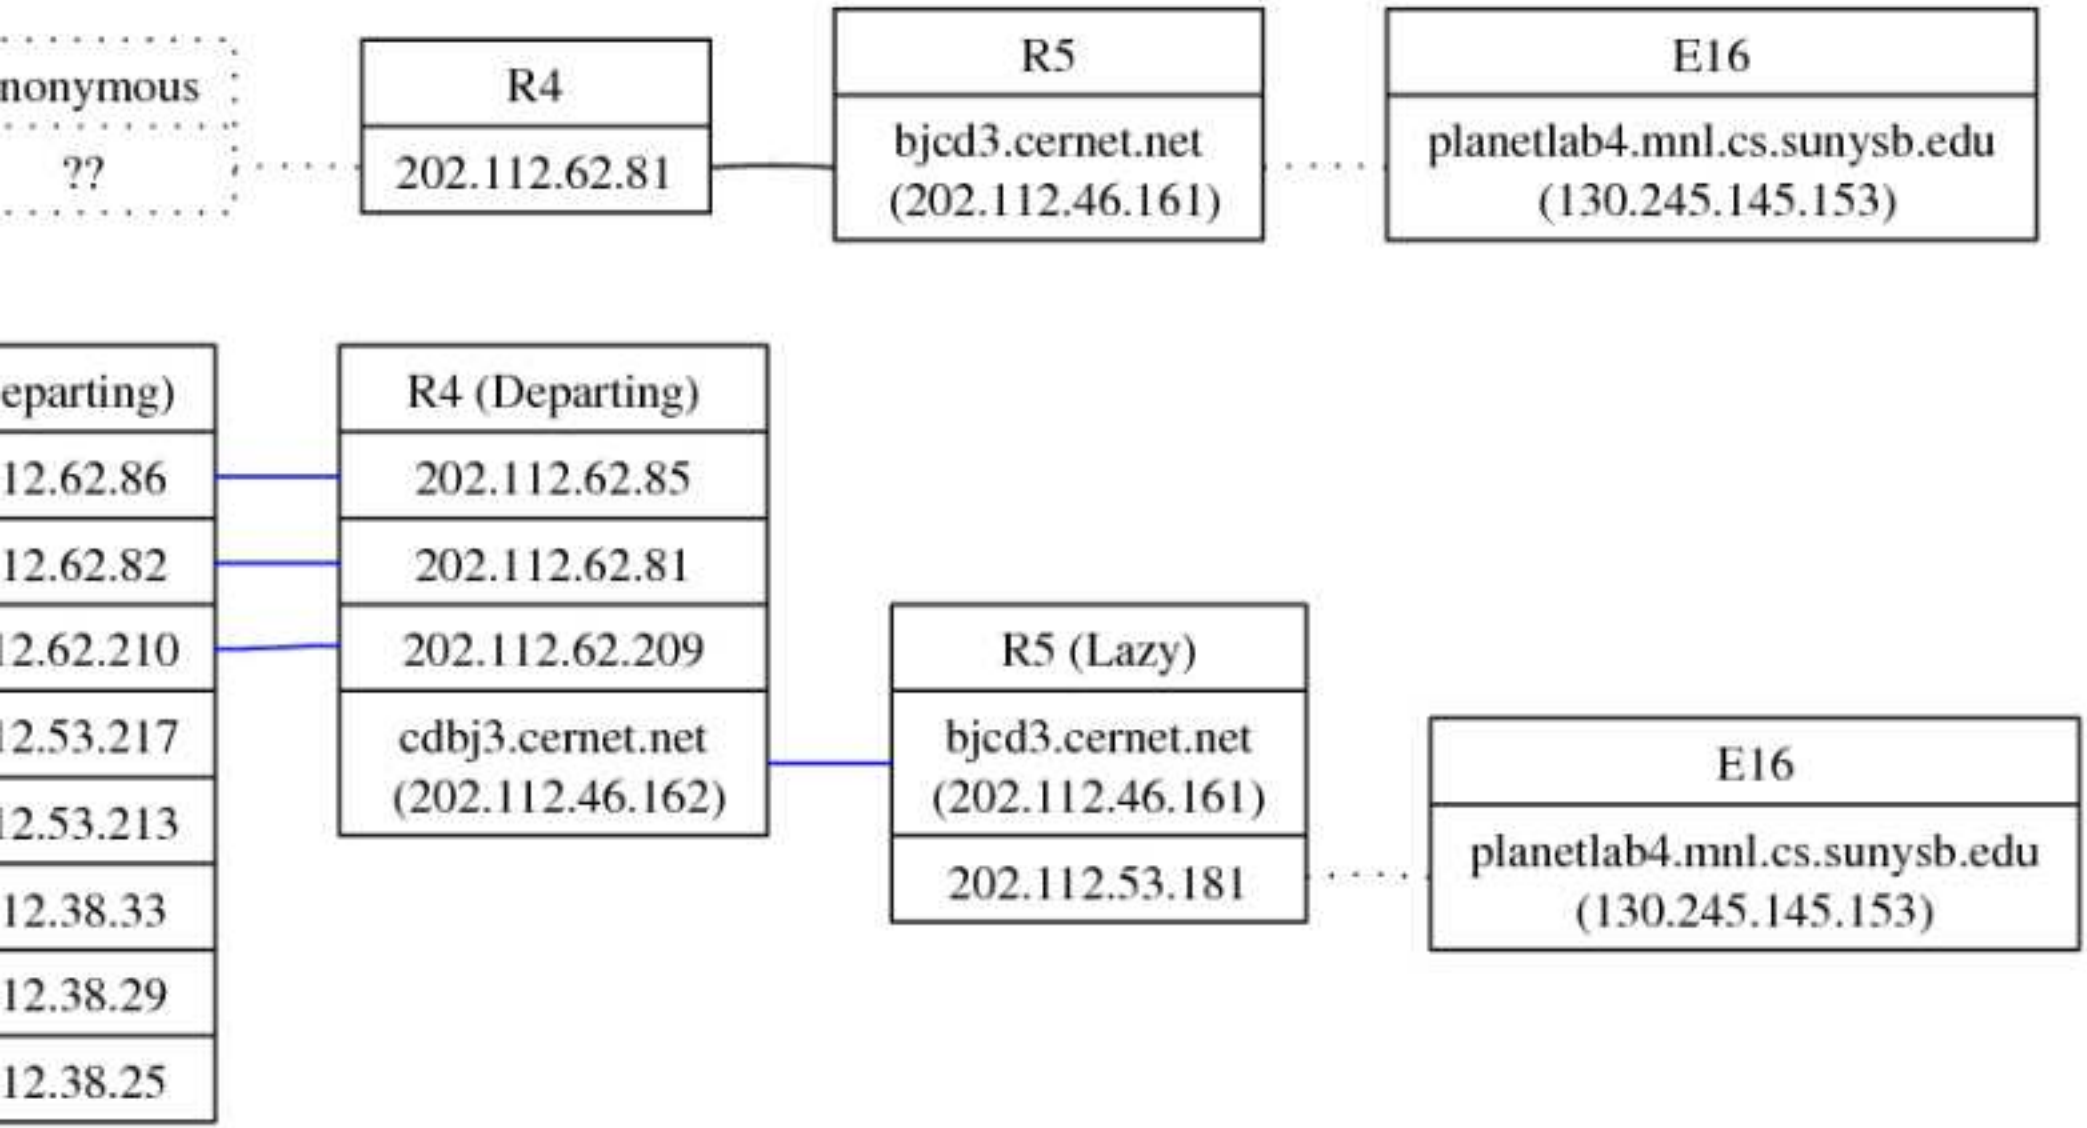
nonymous
??
R4
202.112.62.81
R5
bjcd3.cernet.net
(202.112.46.161)
E16
planetlab4.mnl.cs.sunysb.edu
(130.245.145.153)
eparting)
12.62.86
12.62.82
12.62.210
12.53.217
12.53.213
12.38.33
12.38.29
12.38.25
R4 (Departing)
202.112.62.85
202.112.62.81
202.112.62.209
cdbj3.cernet.net
(202.112.46.162)
R5 (Lazy)
bjcd3.cernet.net
(202.112.46.161)
202.112.53.181
E16
planetlab4.mnl.cs.sunysb.edu
(130.245.145.153)

DisCarte System

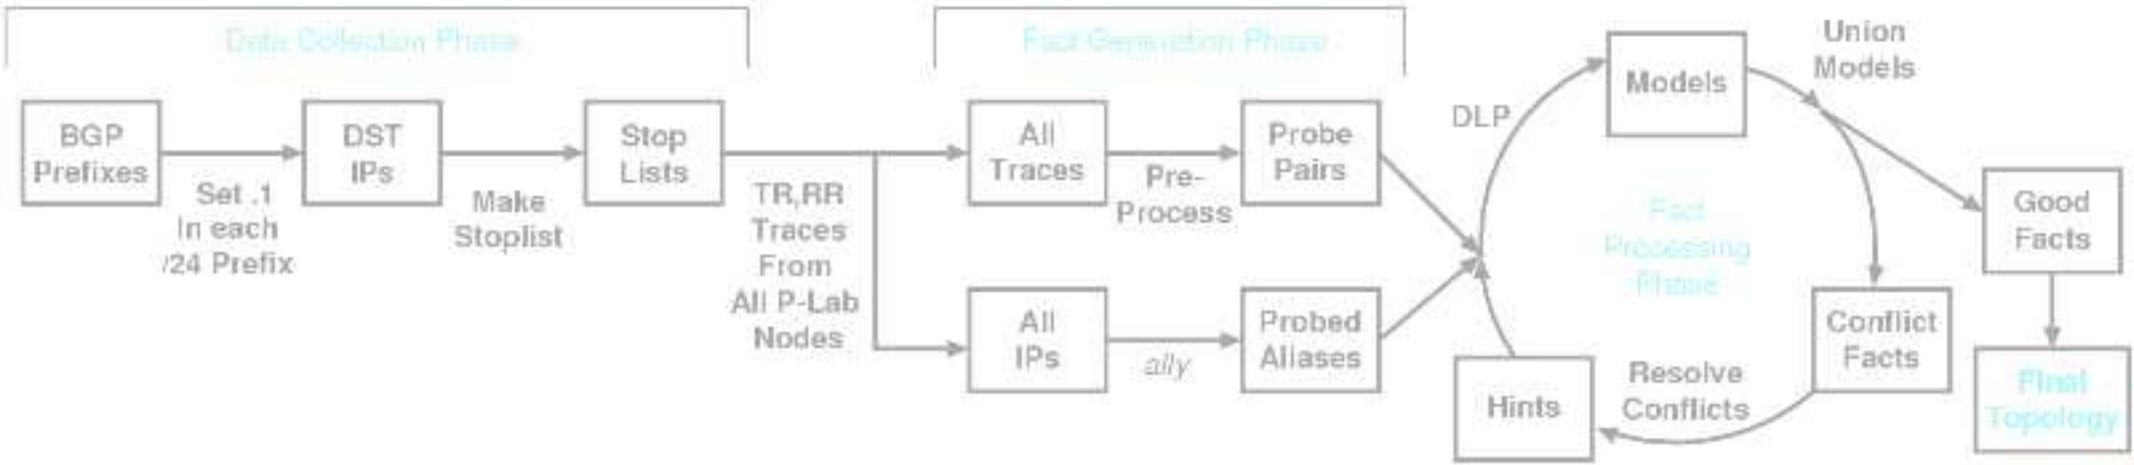

DisCarte System

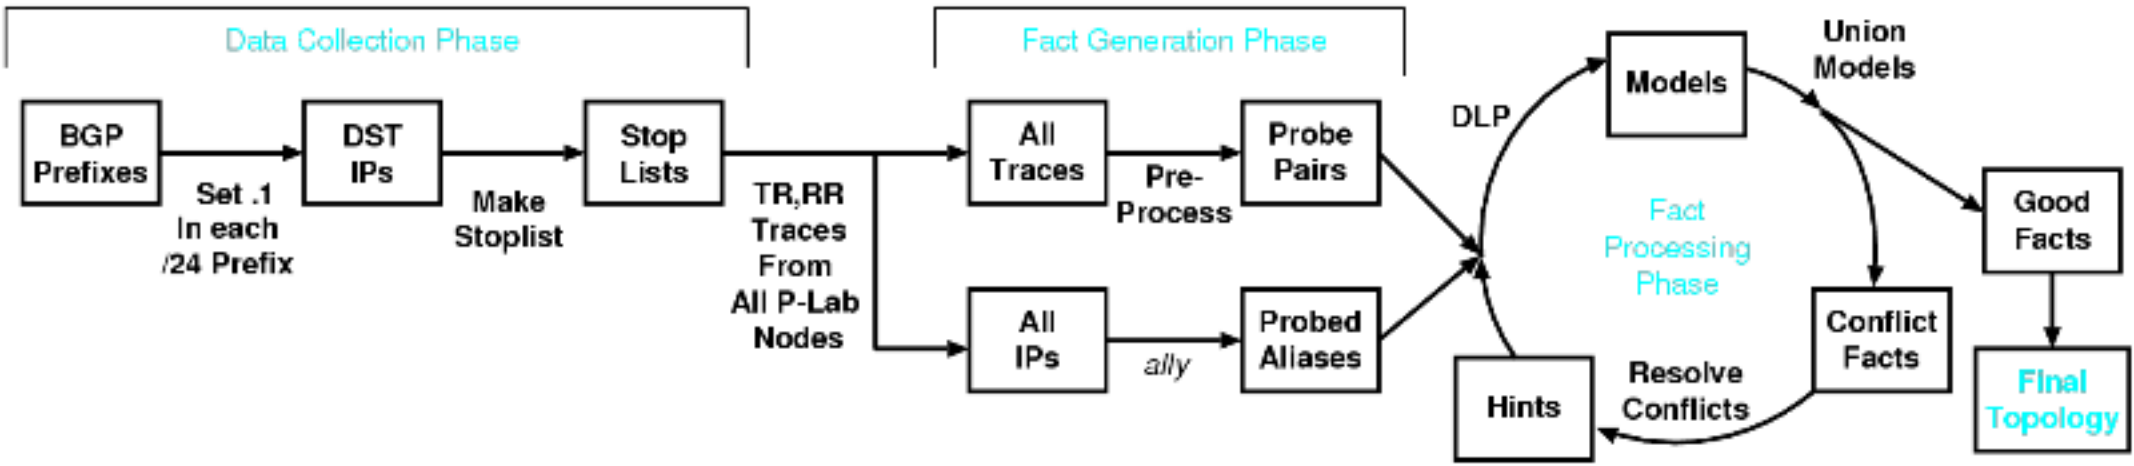

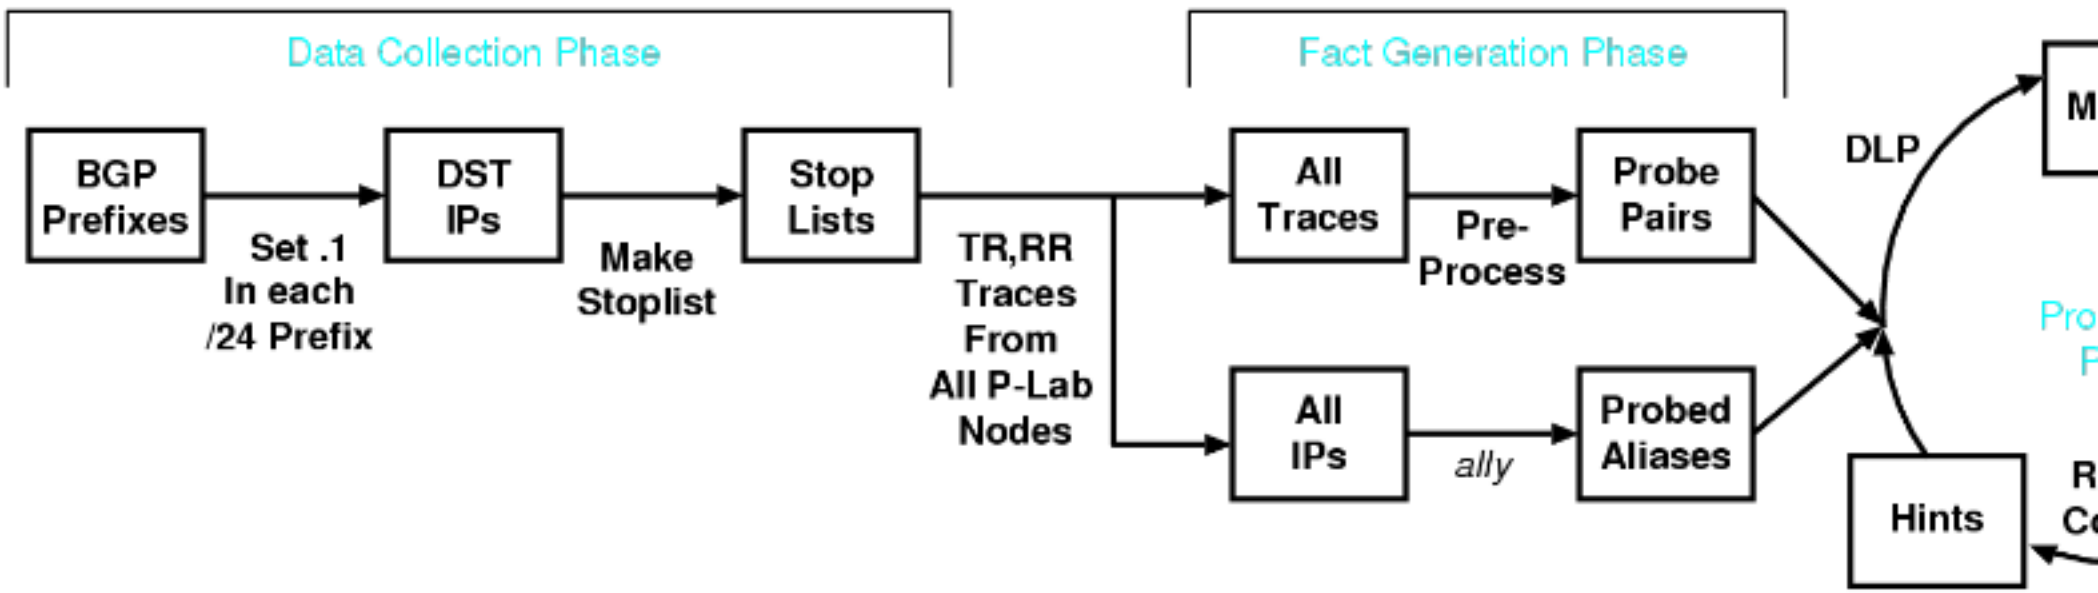
Data Collection Phase
Fact Generation Phase
BGP Prefixes
Set .1 In each /24 Prefix
DST IPs
Make Stoplist
Stop Lists
TR,RR Traces From All P-Lab Nodes
All Traces
Pre-Process
Probe Pairs
All IPs
ally
Probed Aliases
DLP
Hints

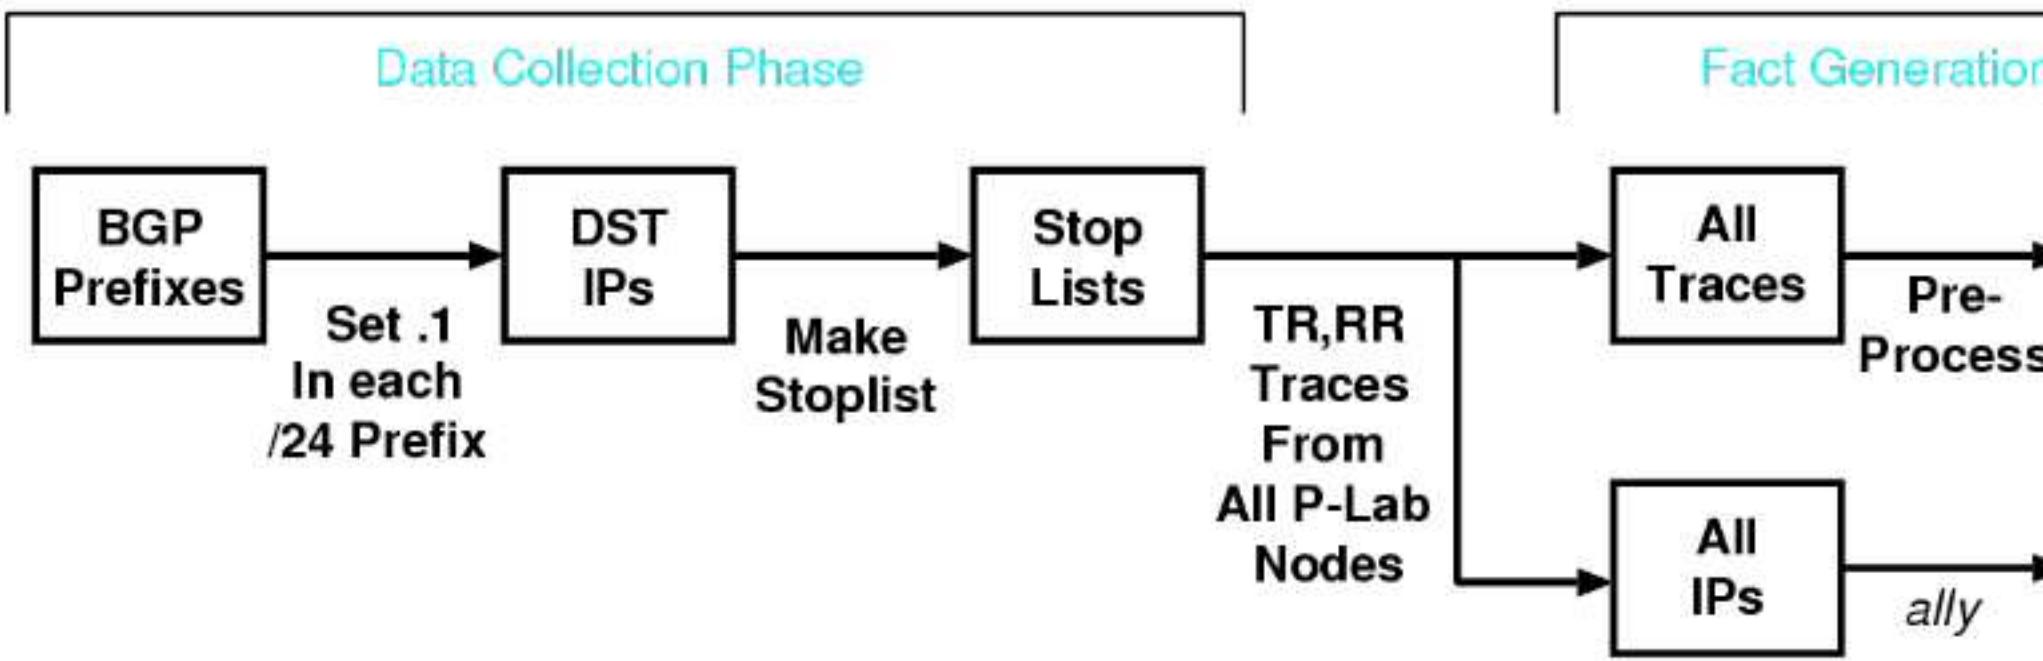
Data Collection Phase
Fact Generation
BGP Prefixes
Set .1 In each /24 Prefix
DST IPs
Make Stoplist
Stop Lists
TR,RR Traces From All P-Lab Nodes
All Traces
Pre-Process
All IPs
ally

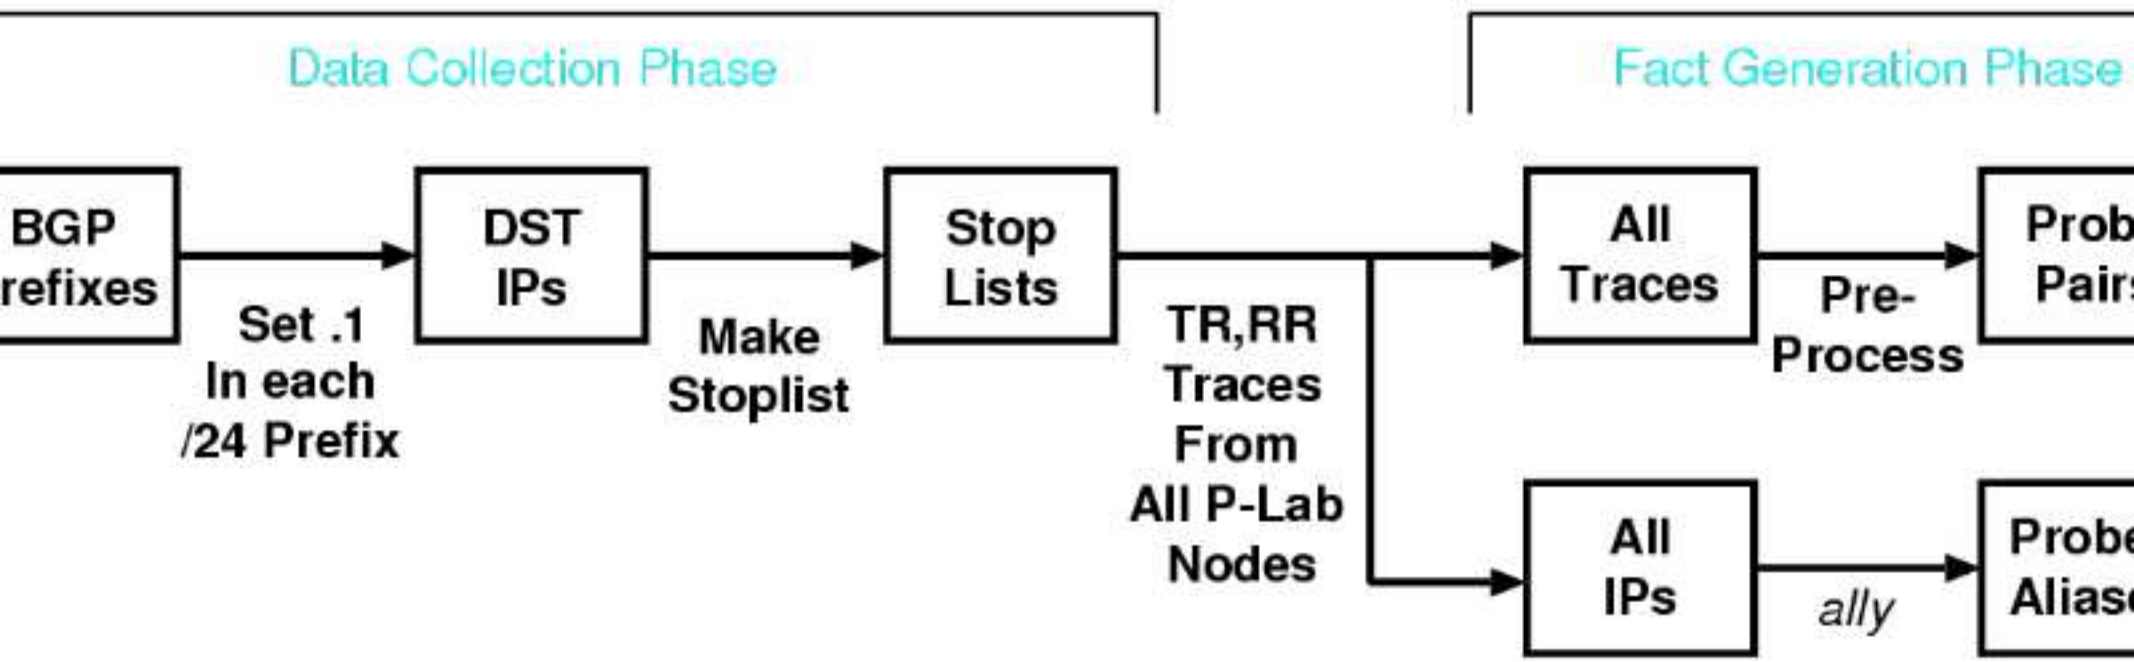
Data Collection Phase
Fact Generation Phase
BGP
refixes
Set .1
In each
/24 Prefix
DST
IPs
Make
Stoplist
Stop
Lists
TR,RR
Traces
From
All P-Lab
Nodes
All
Traces
Pre-
Process
Prob
Pair
All
IPs
ally
Probe
Alias

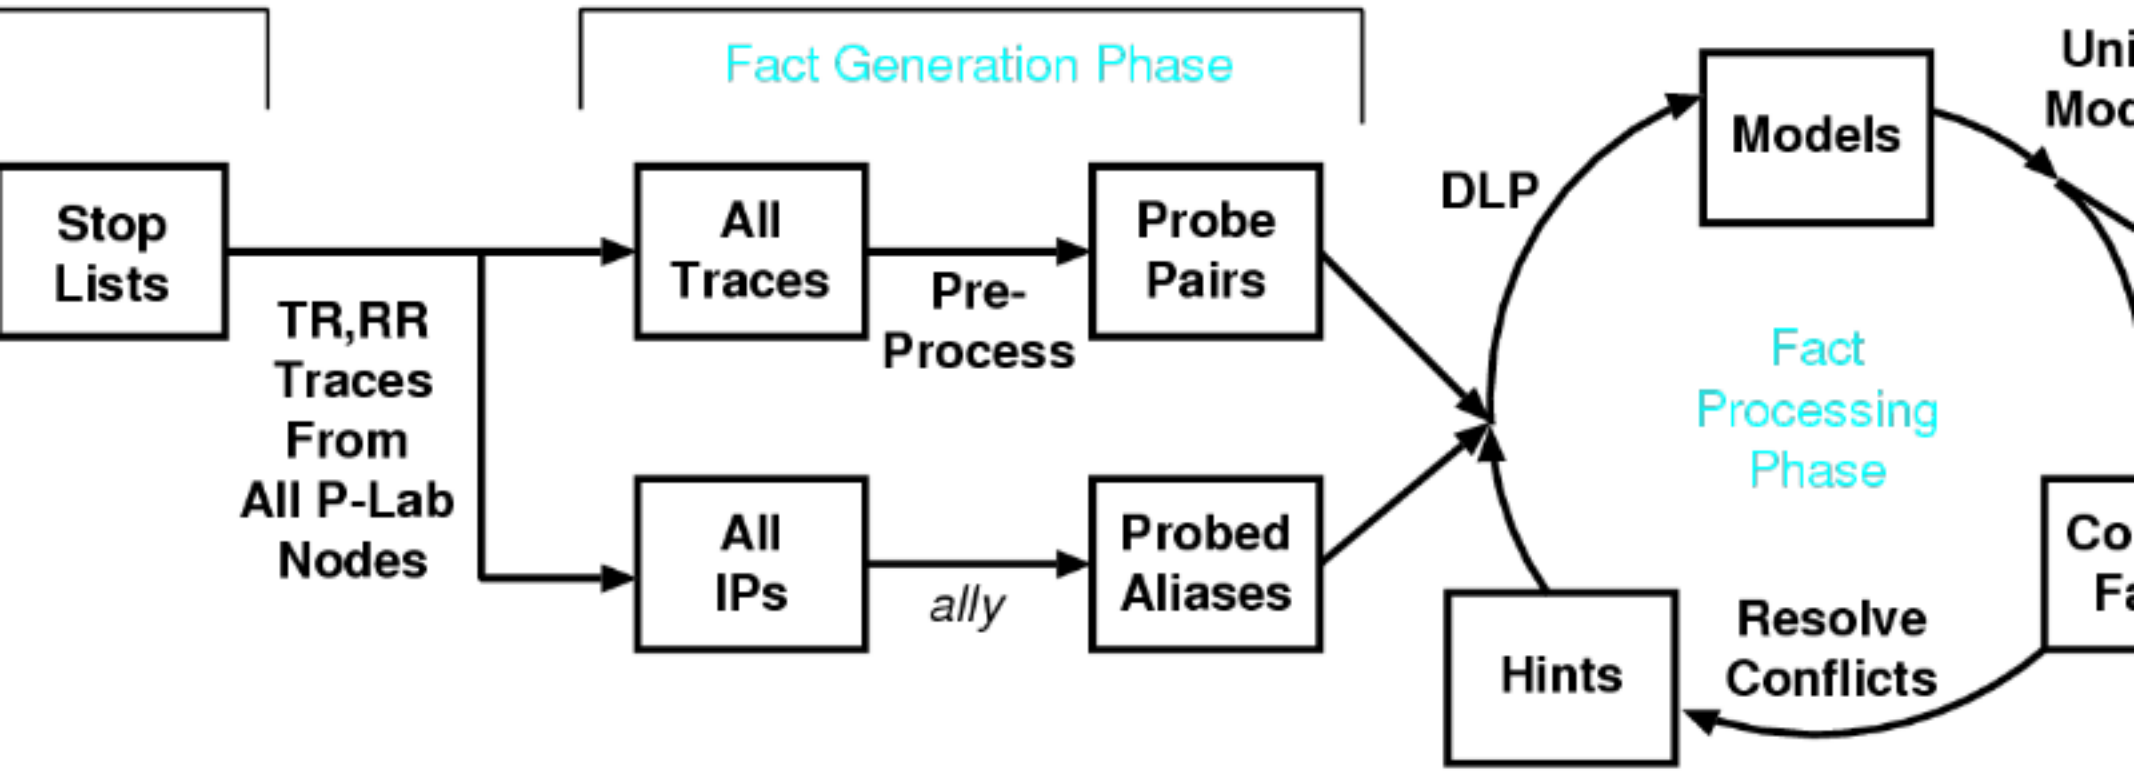
Fact Generation Phase
Stop Lists
TR,RR Traces From All P-Lab Nodes
All Traces
Pre-Process
Probe Pairs
All IPs
ally
Probed Aliases
DLP
Models
Fact Processing Phase
Hints
Resolve Conflicts

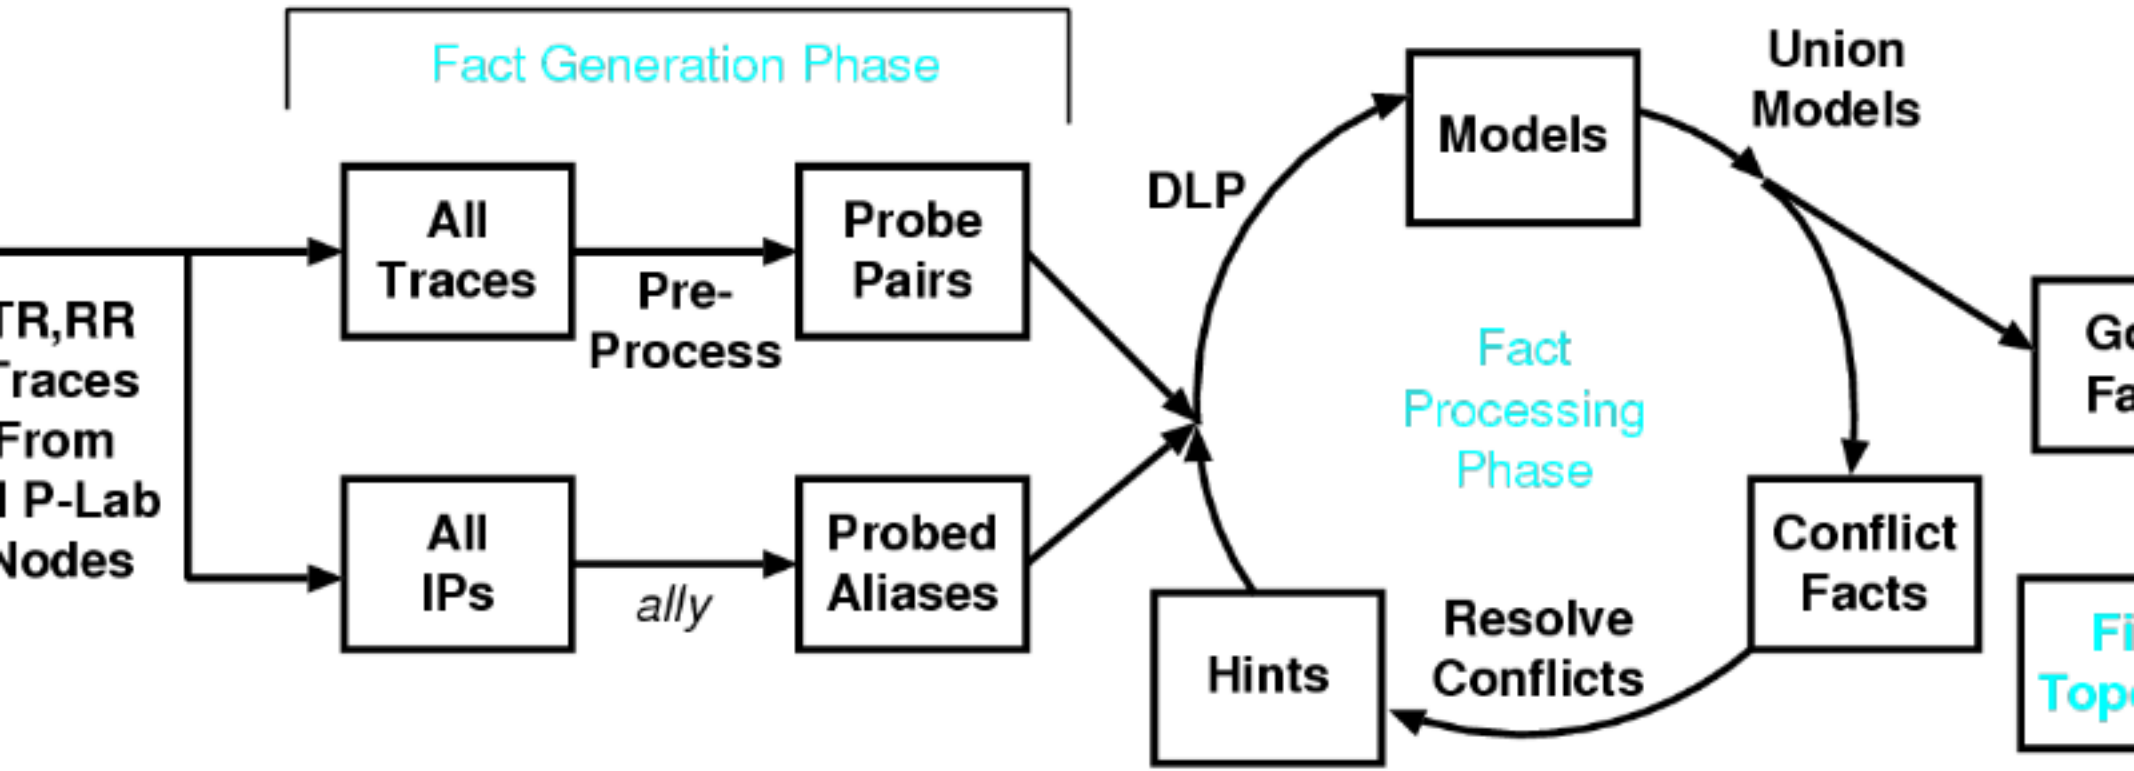
Fact Generation Phase
TR,RR
Traces
From
l P-Lab
Nodes
All
Traces
Pre-
Process
Probe
Pairs
All
IPs
ally
Probed
Aliases
DLP
Models
Union
Models
Fact
Processing
Phase
Conflict
Facts
Resolve
Conflicts
Hints
Go
Fa
Fi
Top

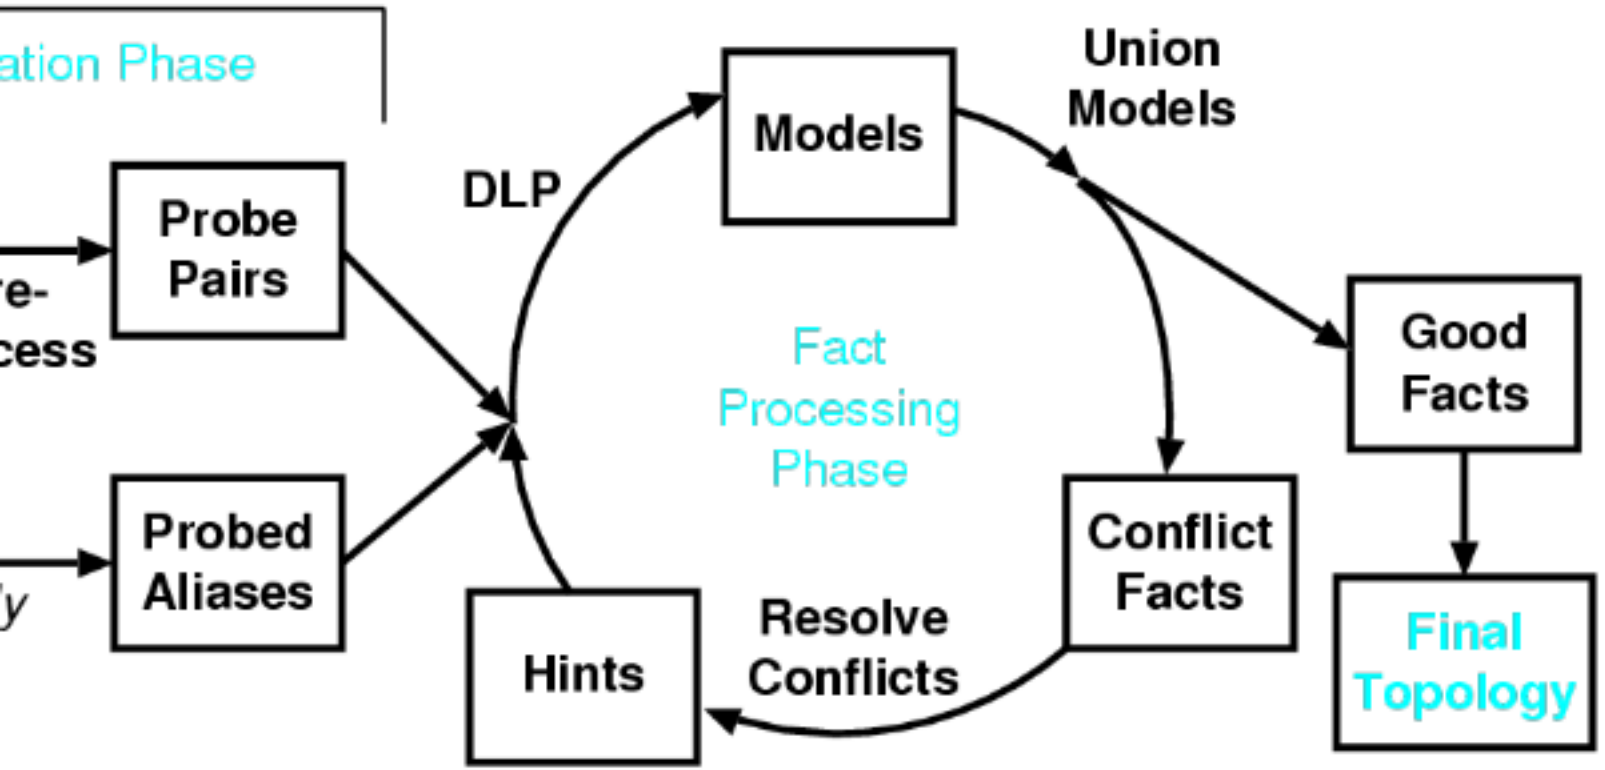
ation Phase
Probe Pairs
Probed Aliases
DLP
Models
Union Models
Good Facts
Fact Processing Phase
Conflict Facts
Final Topology
Hints
Resolve Conflicts

Disjunctive Logic Programs

- DLP: Inference engine and constraint solver
- Facts: TR and RR measurement
  - probe(*src=128.8.128.118, dst=128.208.4.198, ttl=9, resp=206.196.178.90*)
- Inference rules
  - 0 new RR entries→ Arriving to Departing, NotImpl to NotImpl, etc...
- Constraints:
  - Observed engineering practices

Engineering Practices

Prefer models that violate fewest beliefs

- Routers rarely have self-loops
- Linked IPs are typically off-by-one
- Prior techniques are most often correct
- Hidden routers are uncommon

Scaling Challenges

- Data collections
  - 387 sources, 376,408 destinations
  - 100,256 routers found
- Fact generation
  - 1.3 billion trace facts
- 1475 lines of DLP code
- DLP has exponential runtime

Divide and Conquer

- Collect multiple views of the same topology features
  - Single view results in too many conflicts
- Keep inputs facts sizes manageable
- Aggregate output facts
- Work out potential disagreements

Divide and Conquer

- Collect multiple views of the same topology features
  - Single view results in too many conflicts
- Keep inputs facts sizes manageable
- Aggregate output facts
- Work out potential disagreements

Divide and Conquer

- Collect multiple views of the same topology features
  - Single view results in too many conflicts
- Keep inputs facts sizes manageable
- Aggregate output facts
- Work out potential disagreements

Divide and Conquer Scheme

- Conflicts reduced from 1,547 to 28

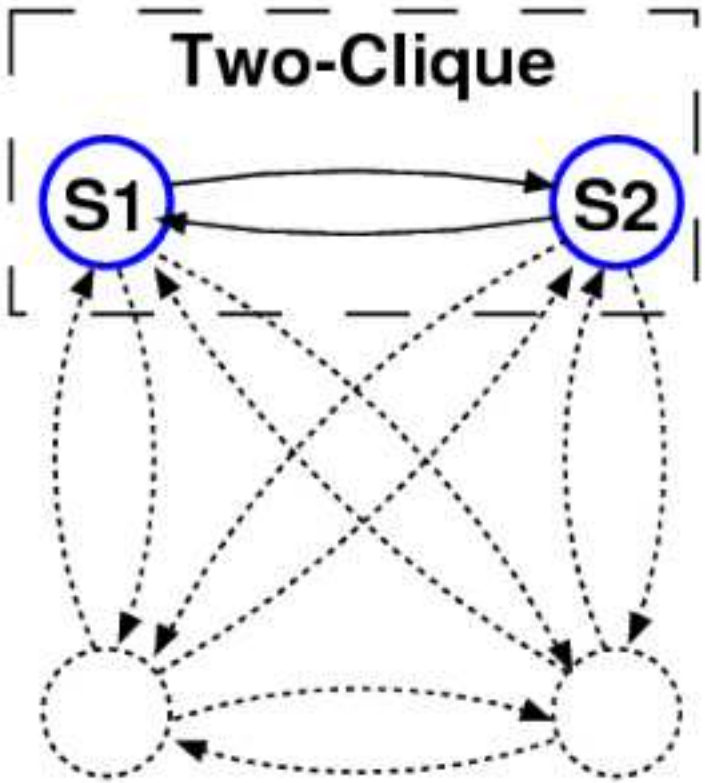

Step #1

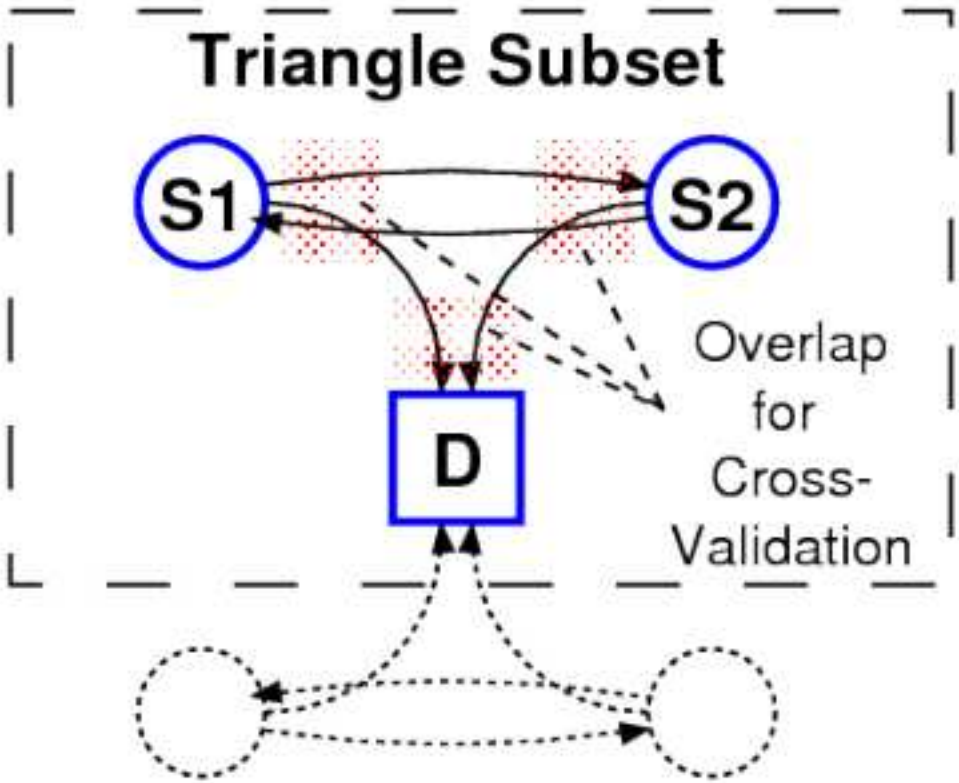

Step #2

Abilene Network

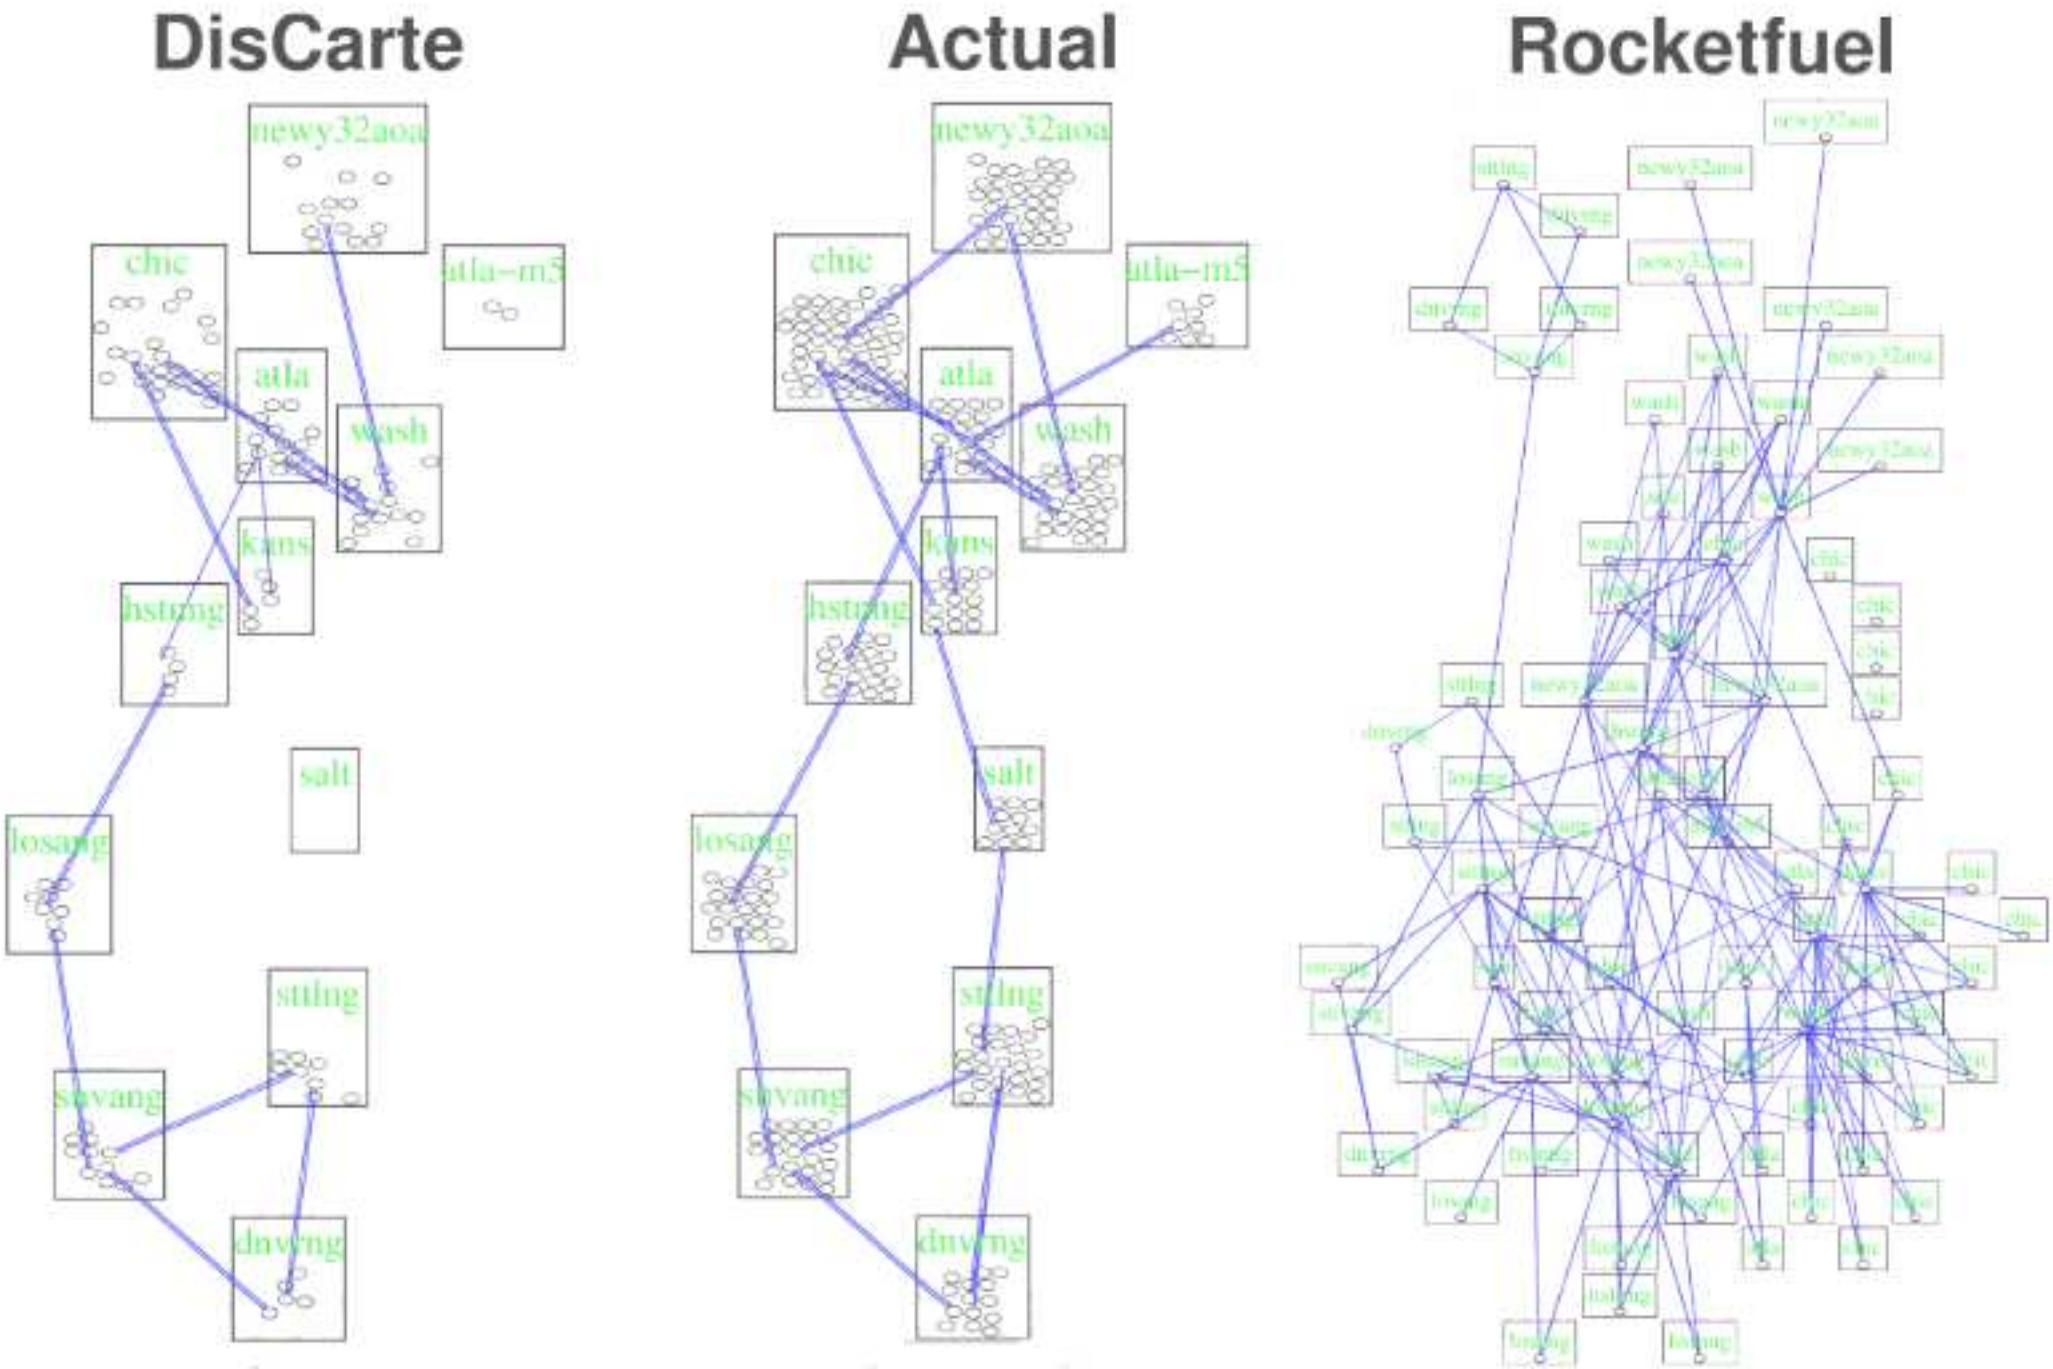

Abilene Network

DisCarte

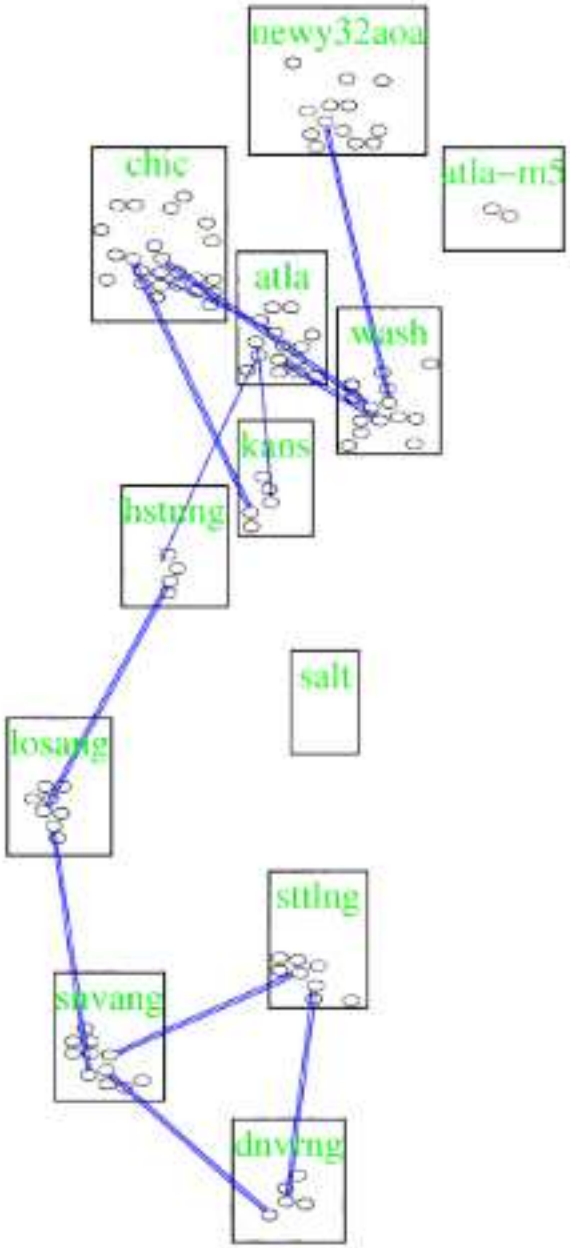

Actual

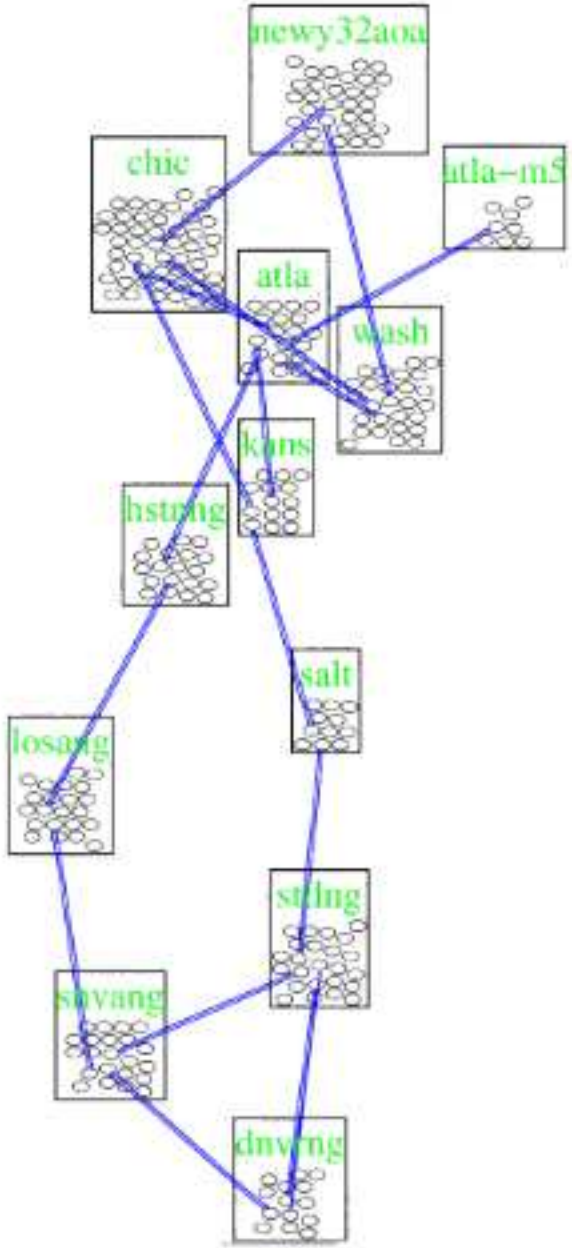

Rocketfuel

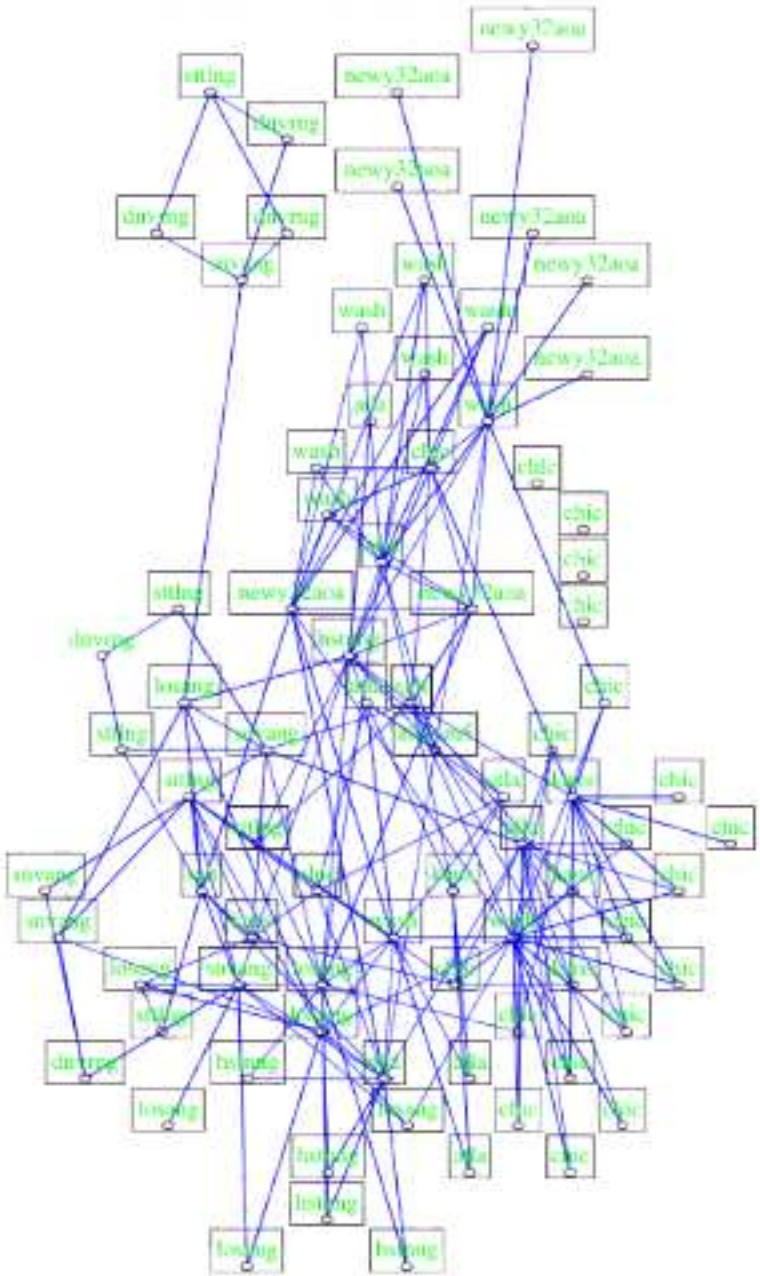

Abilene Network

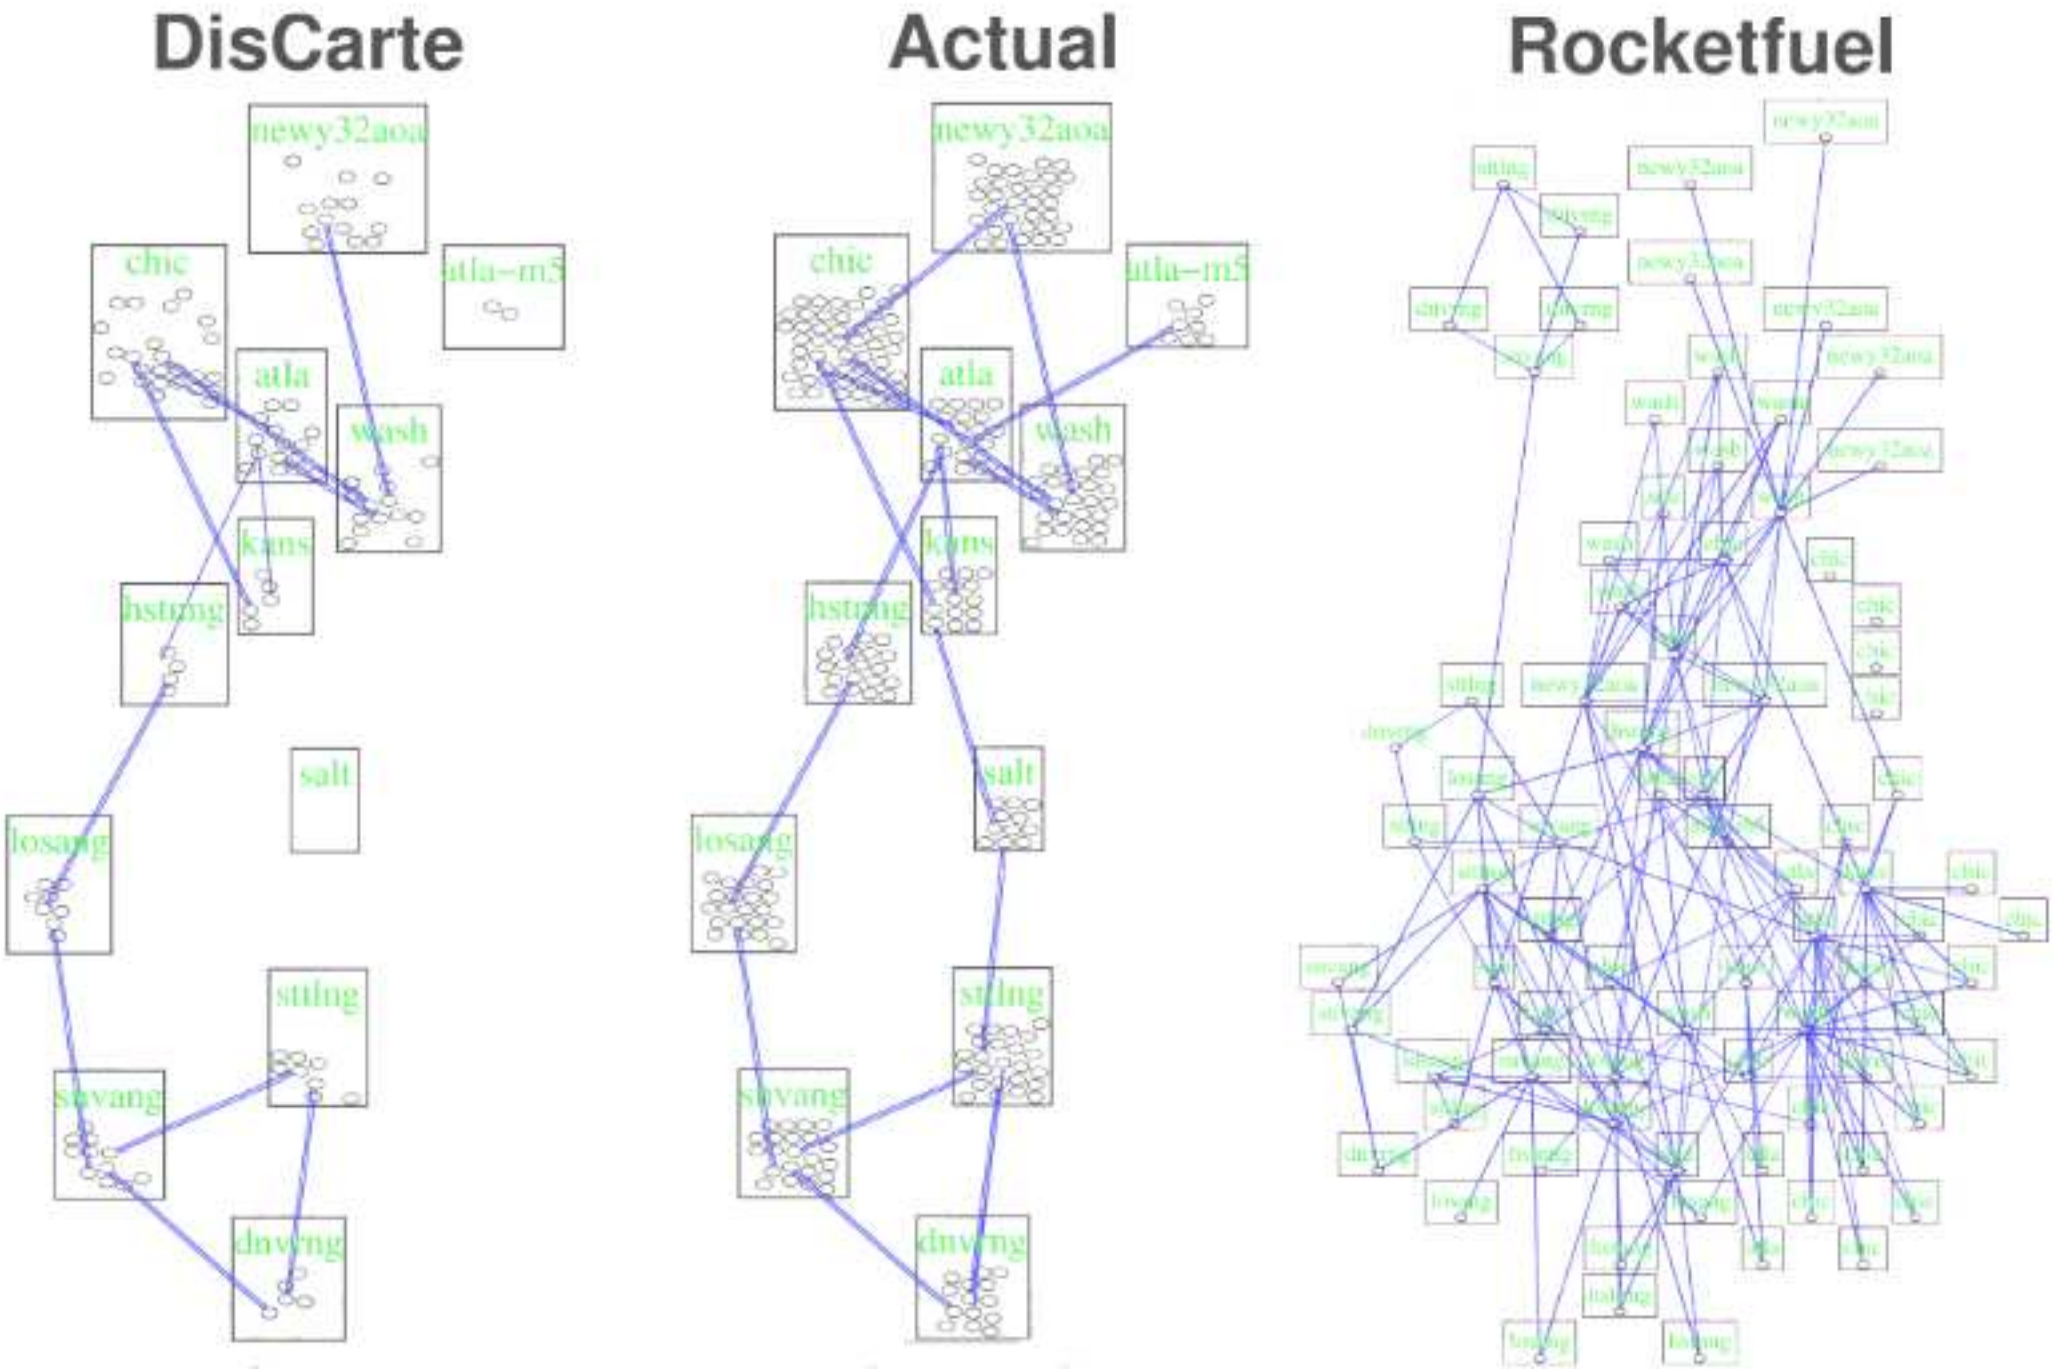

Conclusions

- Shown how to use Record Route
- Described DisCarte
  - Cross-validates measurements to reduce error
- Introduced disjunctive logic programming (DLP)
  - Developed divide and conquer scheme to scale DLP to Internet-sized topologies
- Shown DisCarte was more accurate than previous techniques

Conclusions

- Shown how to use Record Route
- Described DisCarte
  - Cross-validates measurements to reduce error
- Introduced disjunctive logic programming (DLP)
  - Developed divide and conquer scheme to scale DLP to Internet-sized topologies
- Shown DisCarte was more accurate than previous techniques

Published Networks

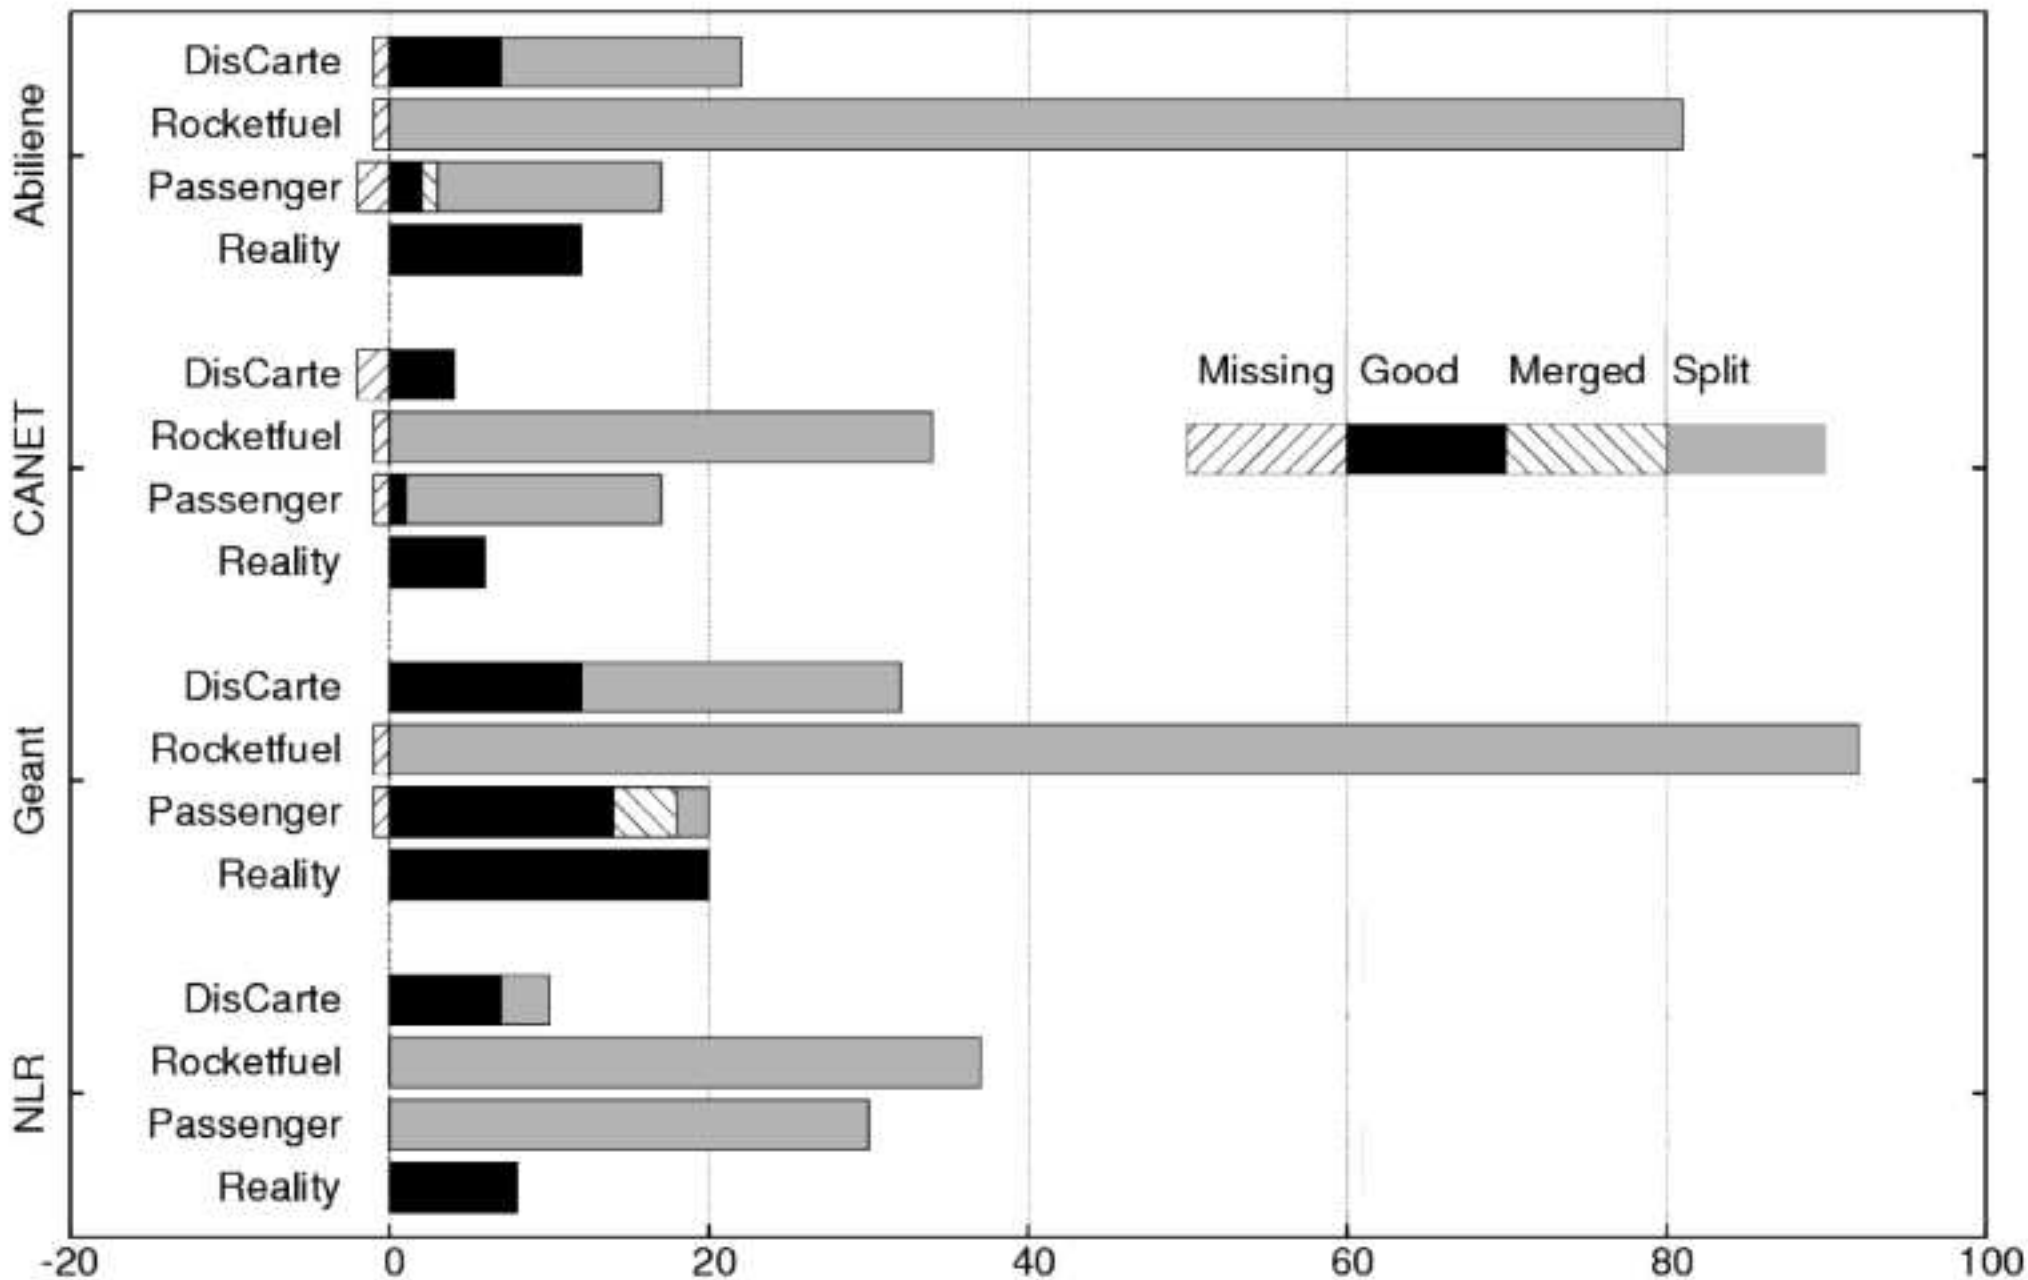

Published Networks