Exploring the Utility of BGP Location Communities in Networking Research

Thomas Krenc
(IIJ Research laboratory)

About Me

- Name: Thomas Krenc
  - Ph.D. at Technische Universitat Berlin
  - Postdoctoral Researcher
    - Naval Postgraduate School
    - CAIDA / University of California San Diego
- Research Interest
  - Improve Resilience and Security of the Internet
  - Build BGP Community repository
- Collaborations
  - CAIDA / UCSD
  - Columbia University
  - Université de Liège

Outline

- Border Gateway Protocol
- BGP Communities
- Communities in Research
- Location Inference of City Communities
- Conclusion

Chapter I:
The Border Gateway Protocol

Autonomous Systems

- Internet is a “network of networks”
- Organized in units of Autonomous Systems (short: AS)
- Each AS is identified by a number between 0 and 2^{32}-1
  - ASN of IIJ is 2497

Border Gateway Protocol

- ASes use BGP to **exchange route information**
- **De-facto standard** routing protocol (RFC1997)
- Decentralized
  - Path vector protocols (class of distance vector)
  - Information hiding protocol: best path decision
- We rely on **BGP route collectors** to study

Border Gateway Protocol

- Route announcements

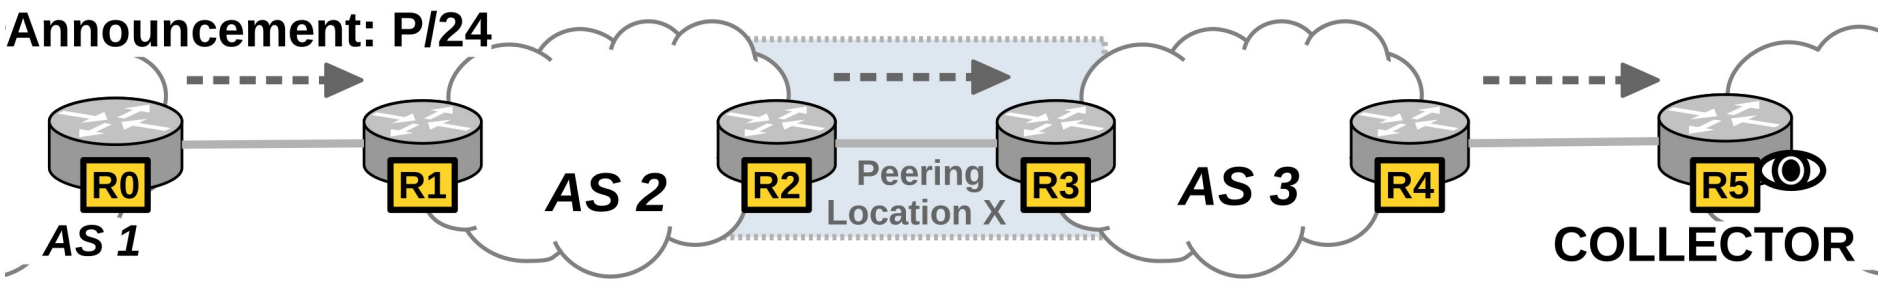

Border Gateway Protocol

- Route announcements

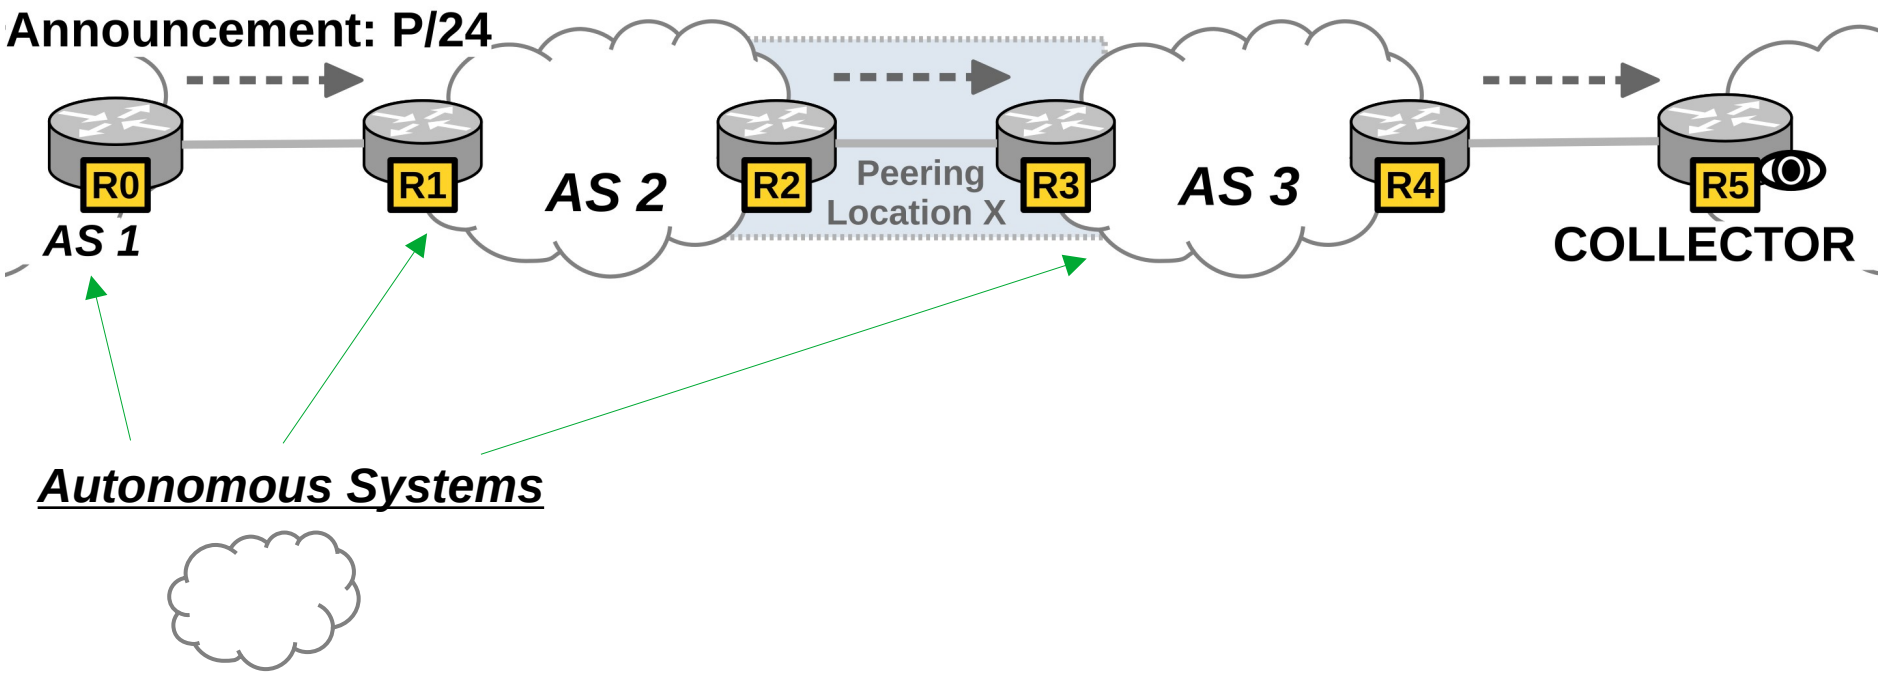

Border Gateway Protocol

- Route announcements

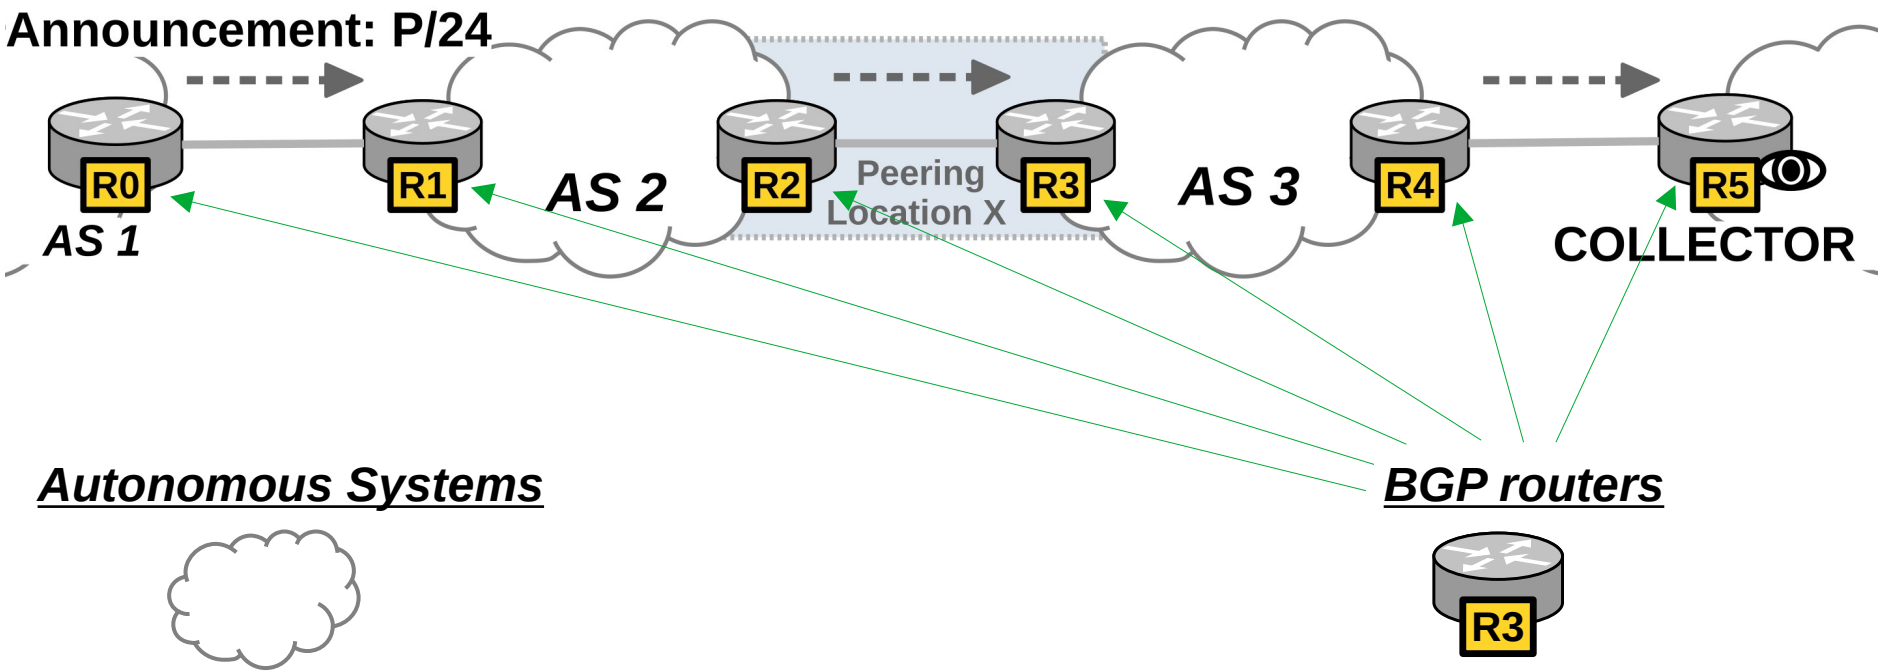

Border Gateway Protocol

- Route announcements

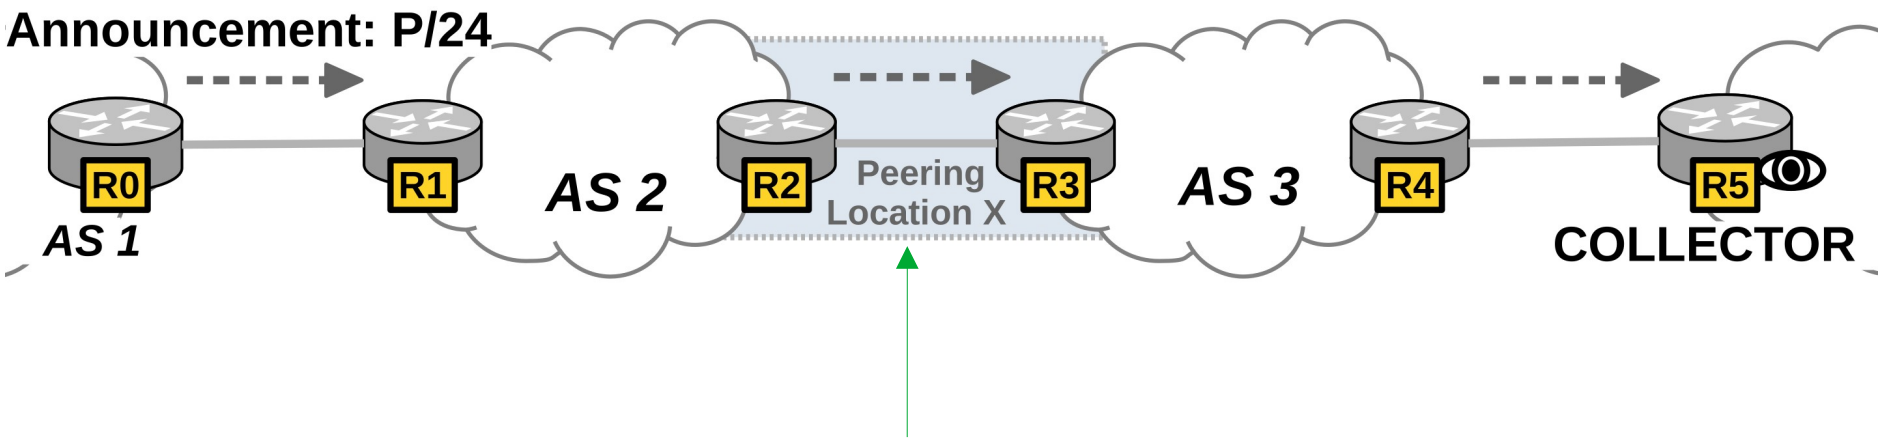

Border Gateway Protocol

- Route announcements

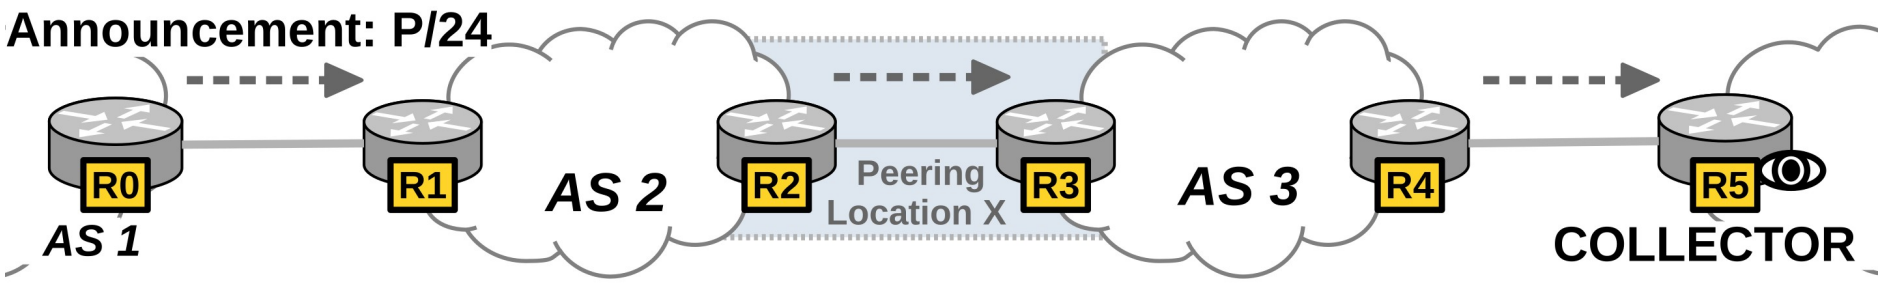

Update Message
Prefix: **[P/24]**
AS Path: **[1]**
Communities: ***[empty]***

Border Gateway Protocol

- Route announcements

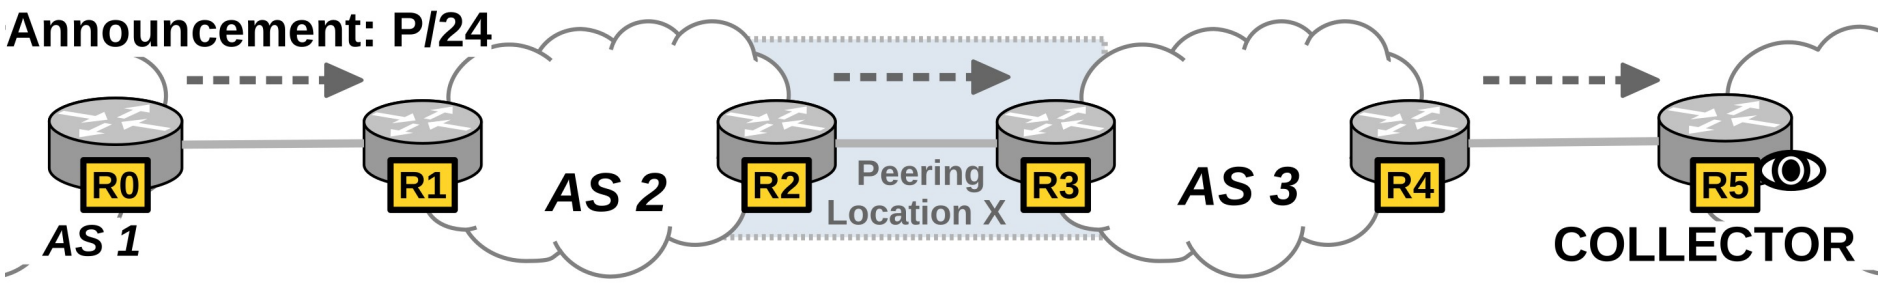

Update Message
Prefix: [P/24]
AS Path: [1]
Communities: [empty]

Update Message
Prefix: [P/24]
AS Path: [2 1]
Communities: [empty]

Border Gateway Protocol

- Route announcements

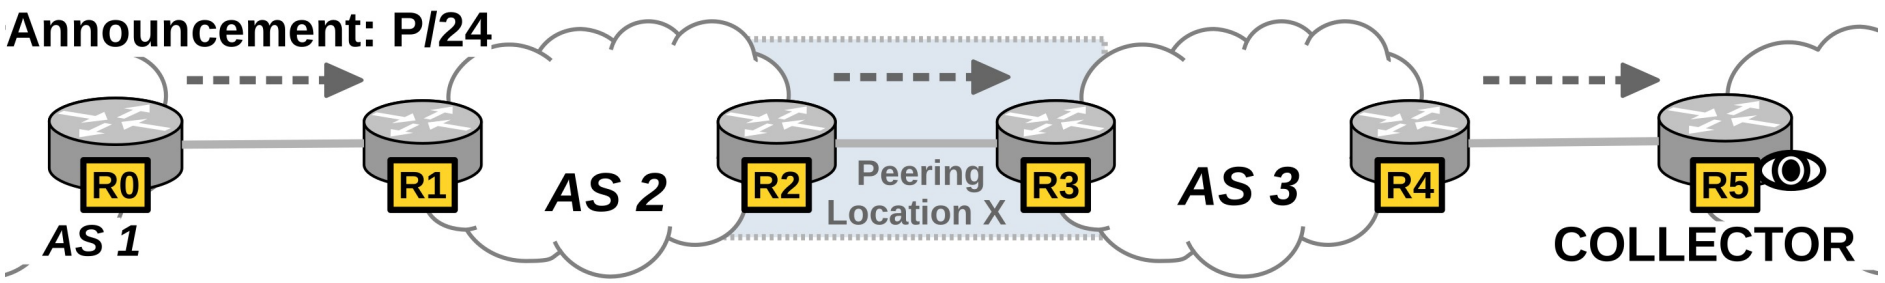

Update Message
Prefix: [P/24]
AS Path: [1]
Communities: [empty]

Update Message
Prefix: [P/24]
AS Path: [2 1]
Communities: [empty]

Update Message
Prefix: [P/24]
AS Path: [3 2 1]
Communities: [empty]

BGP Data

- Route collector projects
  - RouteViews
  - RIPE RIS
- RIBs and updates
- Tools to analyze BGP Data
  - MRT parsers: BGPKIT, BGPStream
  - bgp2go.caida.org

Chapter II:
The BGP Communities Attribute

What is a BGP community?

- "*A community is a group of destinations which share some common property.*"

  RFC1997 (25+ years old)

- Signaling mechanism between BGP routers
- Simple integer (32 bits) – opaque value

BGP Communities

- Route announcements ***with communities***

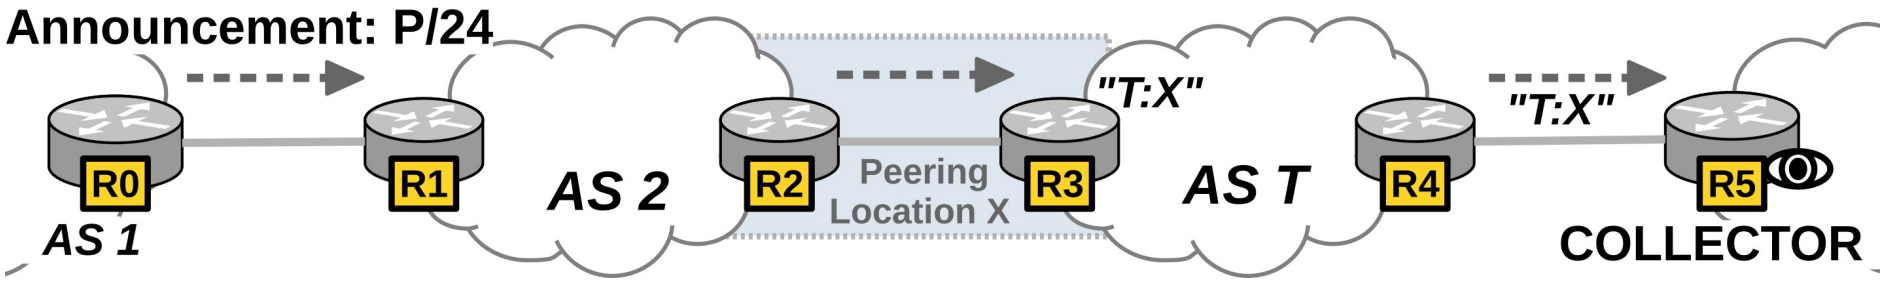

BGP Communities

- Route announcements ***with communities***

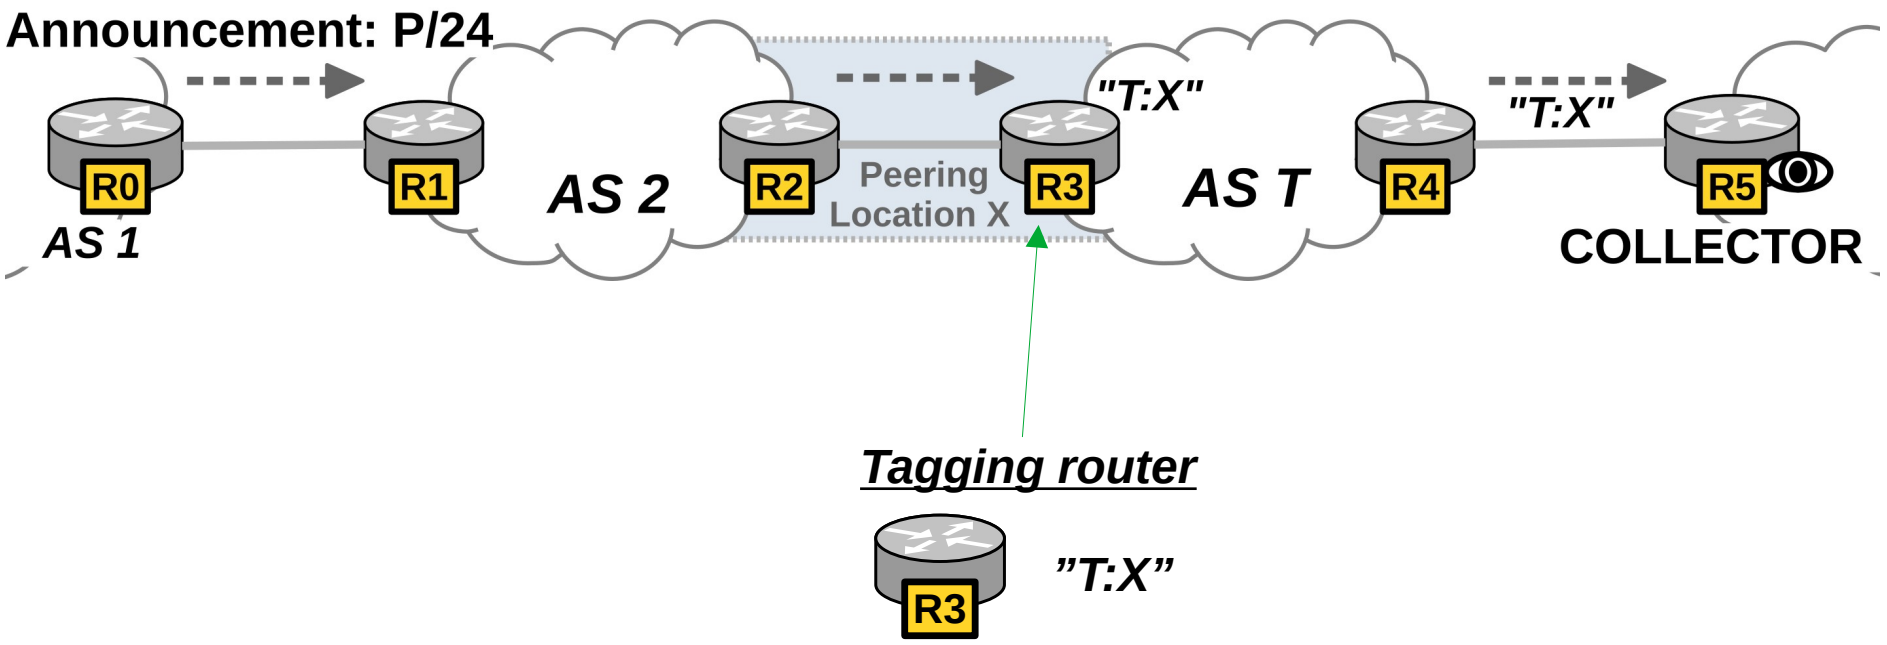

BGP Communities

- Route announcements ***with communities***

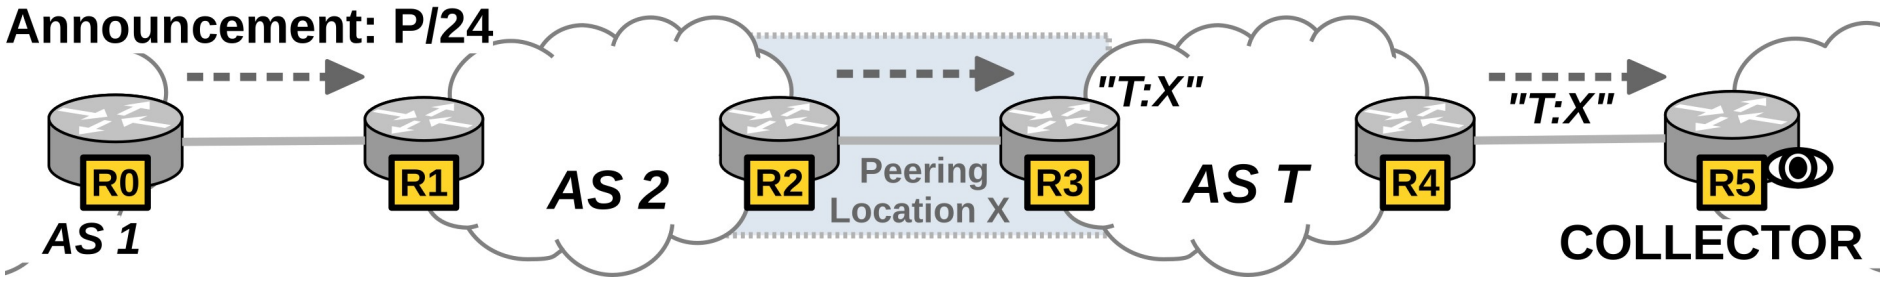

Update Message
Prefix: **[P/24]**
AS Path: **[1]**
Communities: **[*empty*]**

BGP Communities

- Route announcements ***with communities***

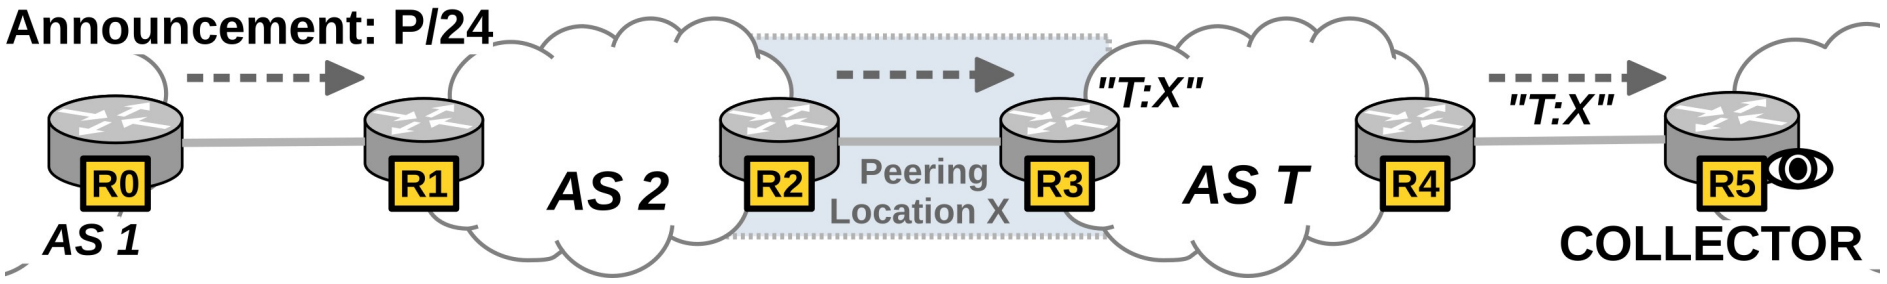

Update Message
Prefix: **[P/24]**
AS Path: **[1]**
Communities: **[*empty*]**

Update Message
Prefix: **[P/24]**
AS Path: **[2 1]**
Communities: **[*empty*]**

BGP Communities

- Route announcements ***with communities***

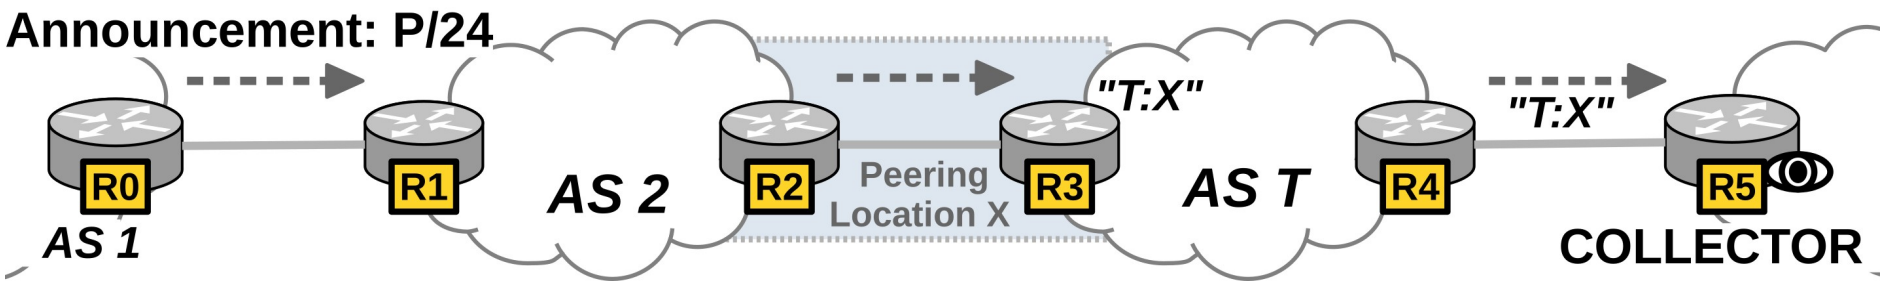

Update Message
Prefix: [P/24]
AS Path: [1]
Communities: [*empty*]

Update Message
Prefix: [P/24]
AS Path: [2 1]
Communities: [*empty*]

Update Message
Prefix: [P/24]
AS Path: [T 2 1]
Communities: [T:X]

BGP Community ***Examples***

1299:2569

BGP Community ***Examples***

1299:2569

Autonomous System Number

BGP Community ***Examples***

1299:2569

Autonomous System Number

“Arelion”

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

(Action)

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

(Information)

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

3356:100

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

3356:100 – “Set local preference to 100”

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

3356:100 – “Set local preference to 100”

3356:2073

BGP Community ***Examples***

1299:2569 – “Do not export this route to AS3356 in Europe”

1299:35130 – “Route was learned in Boston”

3356:100 – “Set local preference to 100”

3356:2073 – “Route was learned in London”

BGP Community ***Categories***

Information Categories:

- **Location**: *City, Facility, Router, Session*
- **Region**: *Continent, Country, State/Province*
- **Non-geo**: *ASN, Relationship, RPKI status*

Action Categories:

- **Action**: *(No-)export, LocalPref, Prepend*
- **Target**: *Location, Region, ASN*

BGP Community ***Categories***

Information Categories:

Coarse-grained

- **Location**: *City, Facility, Router, Session*
- **Region**: *Continent, Country, State/Province*
- **Non-geo**: *ASN, Relationship, RPKI status*

Action Categories:

- **Action**: *(No-)export, LocalPref, Prepend*
- **Target**: *Location, Region, ASN*

BGP Community ***Categories***

Fine-grained

Information Categories:

- **Location**: *City, Facility, Router, Session*
- **Region**: *Continent, Country, State/Province*
- **Non-geo**: *ASN, Relationship, RPKI status*

Action Categories:

- **Action**: *(No-)export, LocalPref, Prepend*
- **Target**: *Location, Region, ASN*

BGP Community ***Use Cases***

- Scaling of large networks
  - Differentiate customer routers from transit/peer route: *Prevent route leaks*
  - Tagging of anycast instances
    - *“blind” without communities; most peering sessions w/ route collectors*

BGP Community ***Use Cases***

- Scaling of large networks
  - Differentiate customer routers from transit/peer route: *Prevent route leaks*
  - Tagging of anycast instances
    - *“blind” without communities; most peering sessions w/ route collectors*
- Implement routing policies
  - Upstream preference: *influence local preference, AS path prepending behavior*
  - Cold potato routing

BGP Community ***Use Cases***

- Scaling of large networks
  - Differentiate customer routers from transit/peer route: *Prevent route leaks*
  - Tagging of anycast instances
    - *“blind” without communities; most peering sessions w/ route collectors*
- Implement routing policies
  - Upstream preference: *influence local preference, AS path prepending behavior*
  - Cold potato routing
- Security
  - DDoS traffic Blackholing
  - Filtering RPKI invalids

BGP Community ***Use Cases***

- Scaling of large networks
  - Differentiate customer routers from transit/peer route: *Prevent route leaks*
  - Tagging of anycast instances
    - *“blind” without communities; most peering sessions w/ route collectors*
- Implement routing policies
  - Upstream preference: *influence local preference, AS path prepending behavior*
  - Cold potato routing
- Security
  - DDoS traffic Blackholing
  - Filtering RPKI invalids
- *… anything that can be configured in BGP routers*

Use Case: Cold Potato Routing

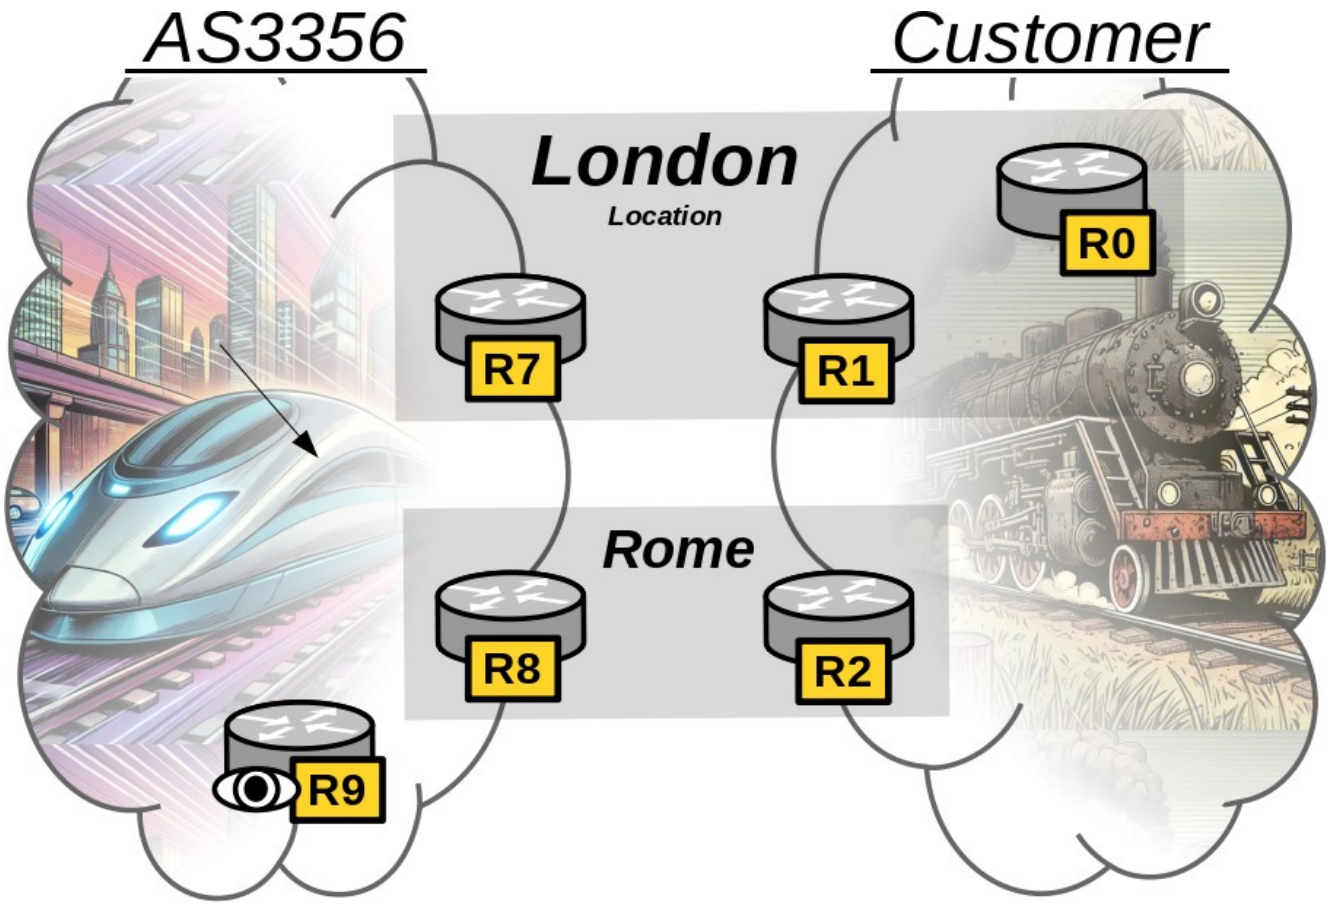

Use Case: Cold Potato Routing

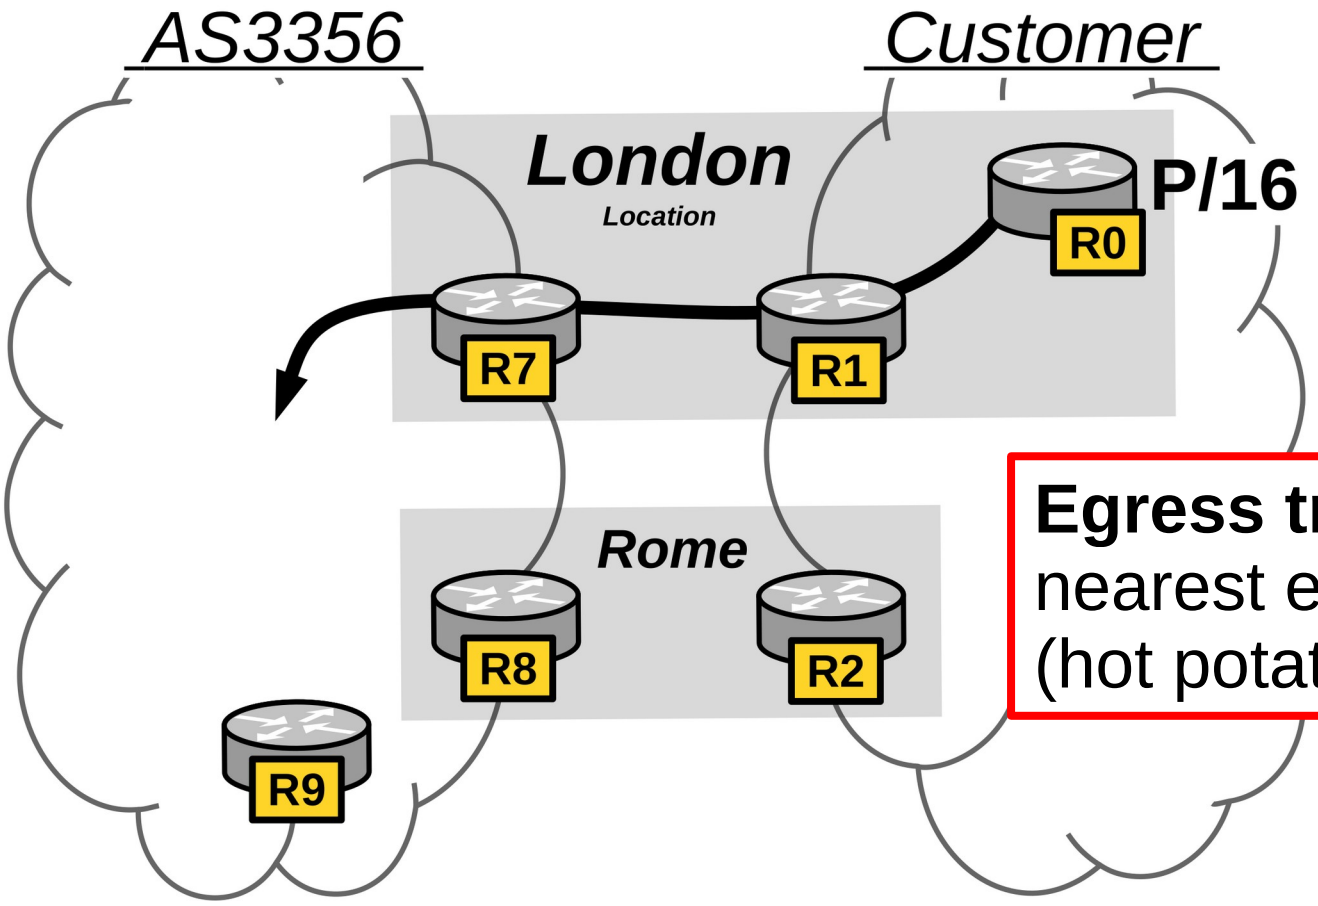

Egress traffic takes nearest exit point ✅ (hot potato routing)

Use Case: Cold Potato Routing

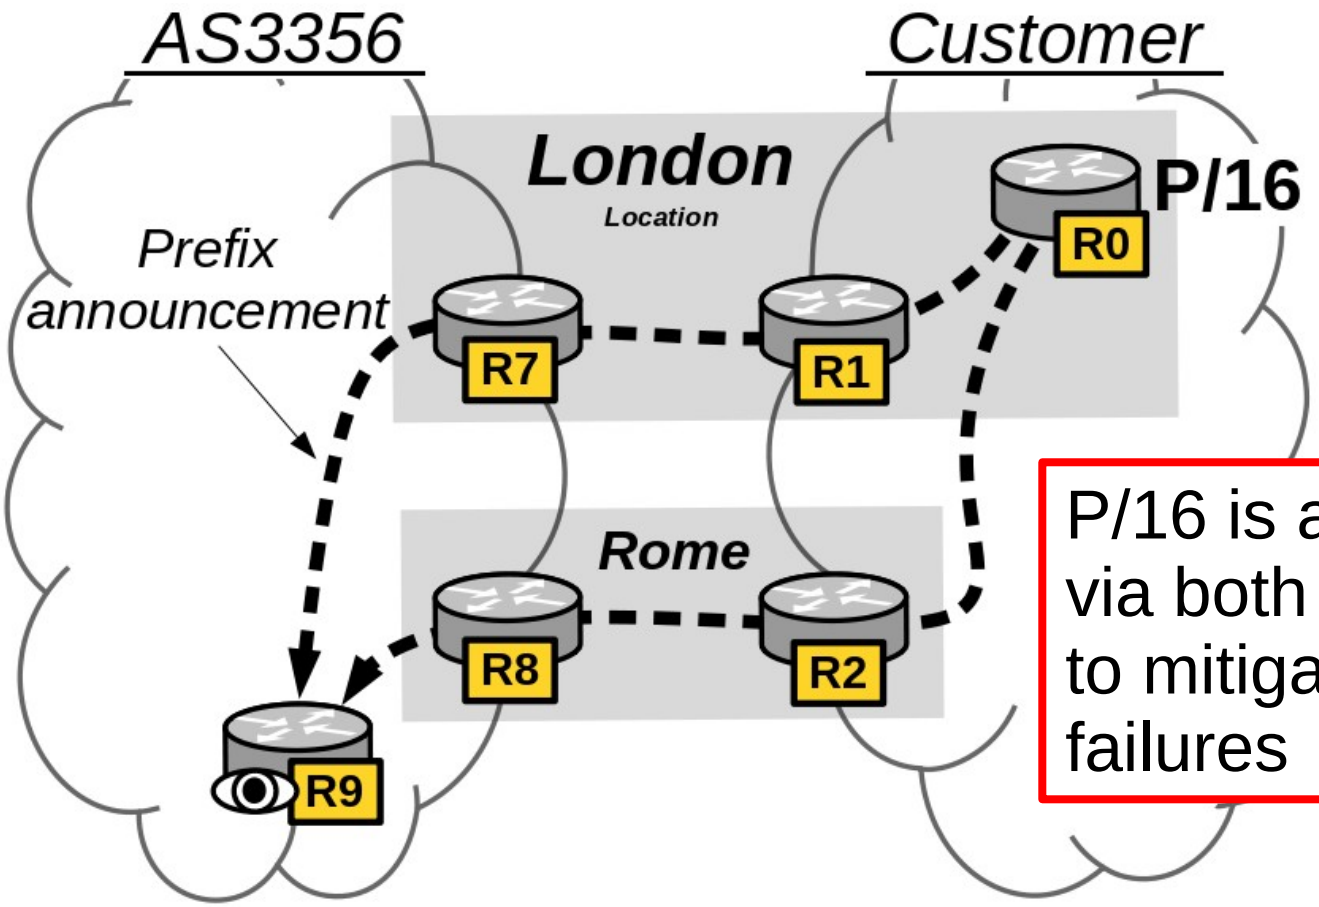

P/16 is announced via both locations to mitigate link failures

Use Case: Cold Potato Routing

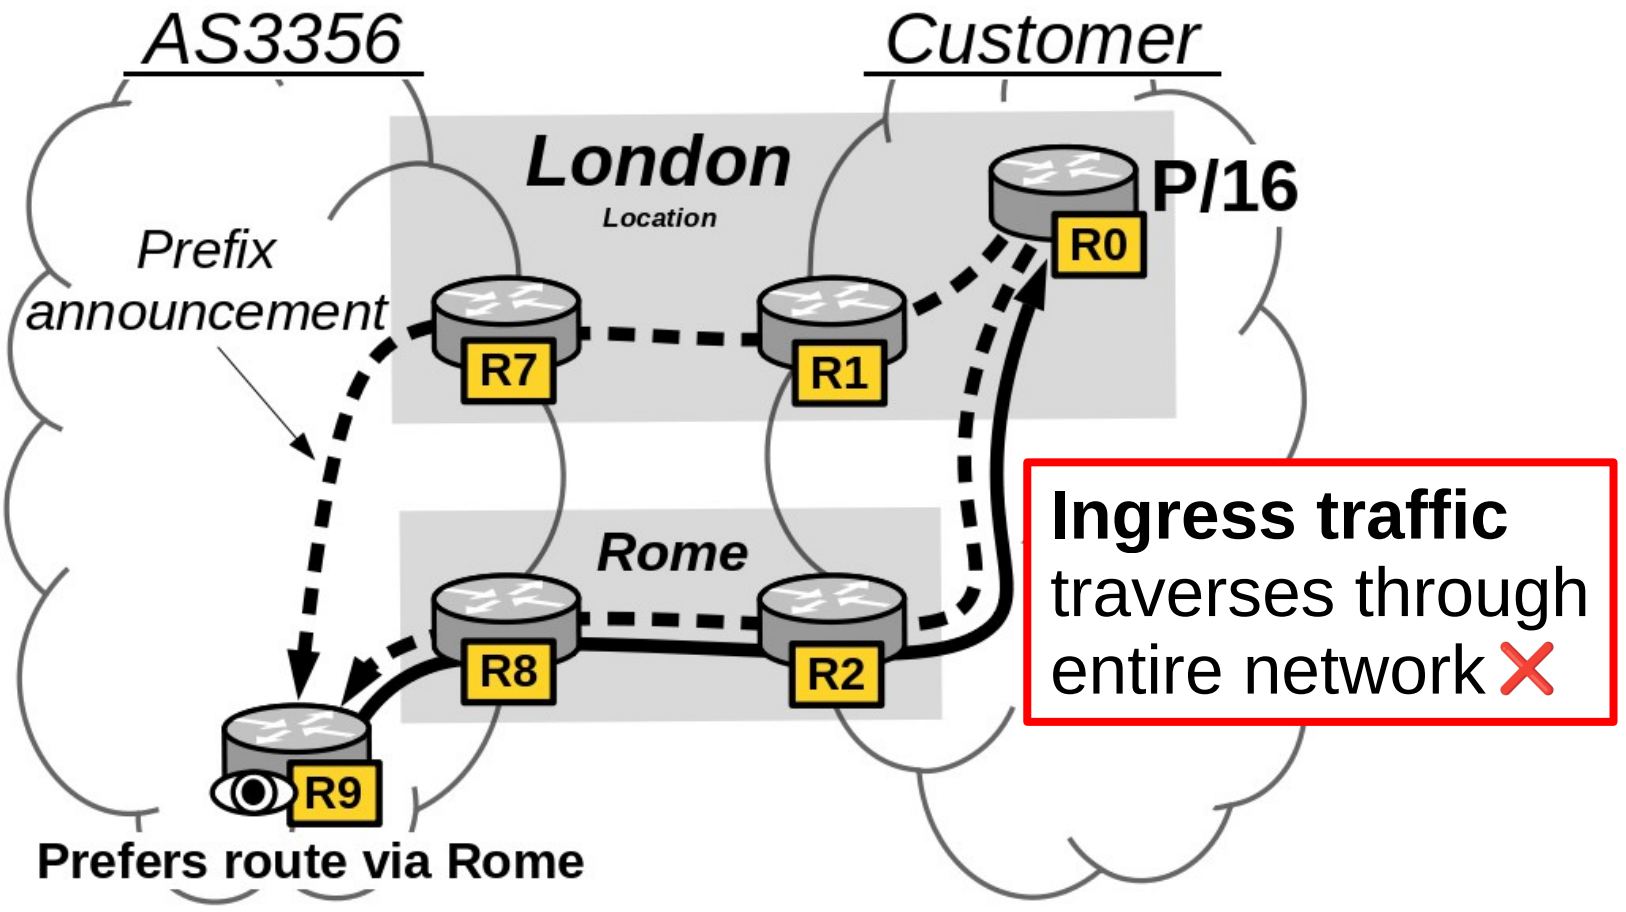

Use Case: Cold Potato Routing

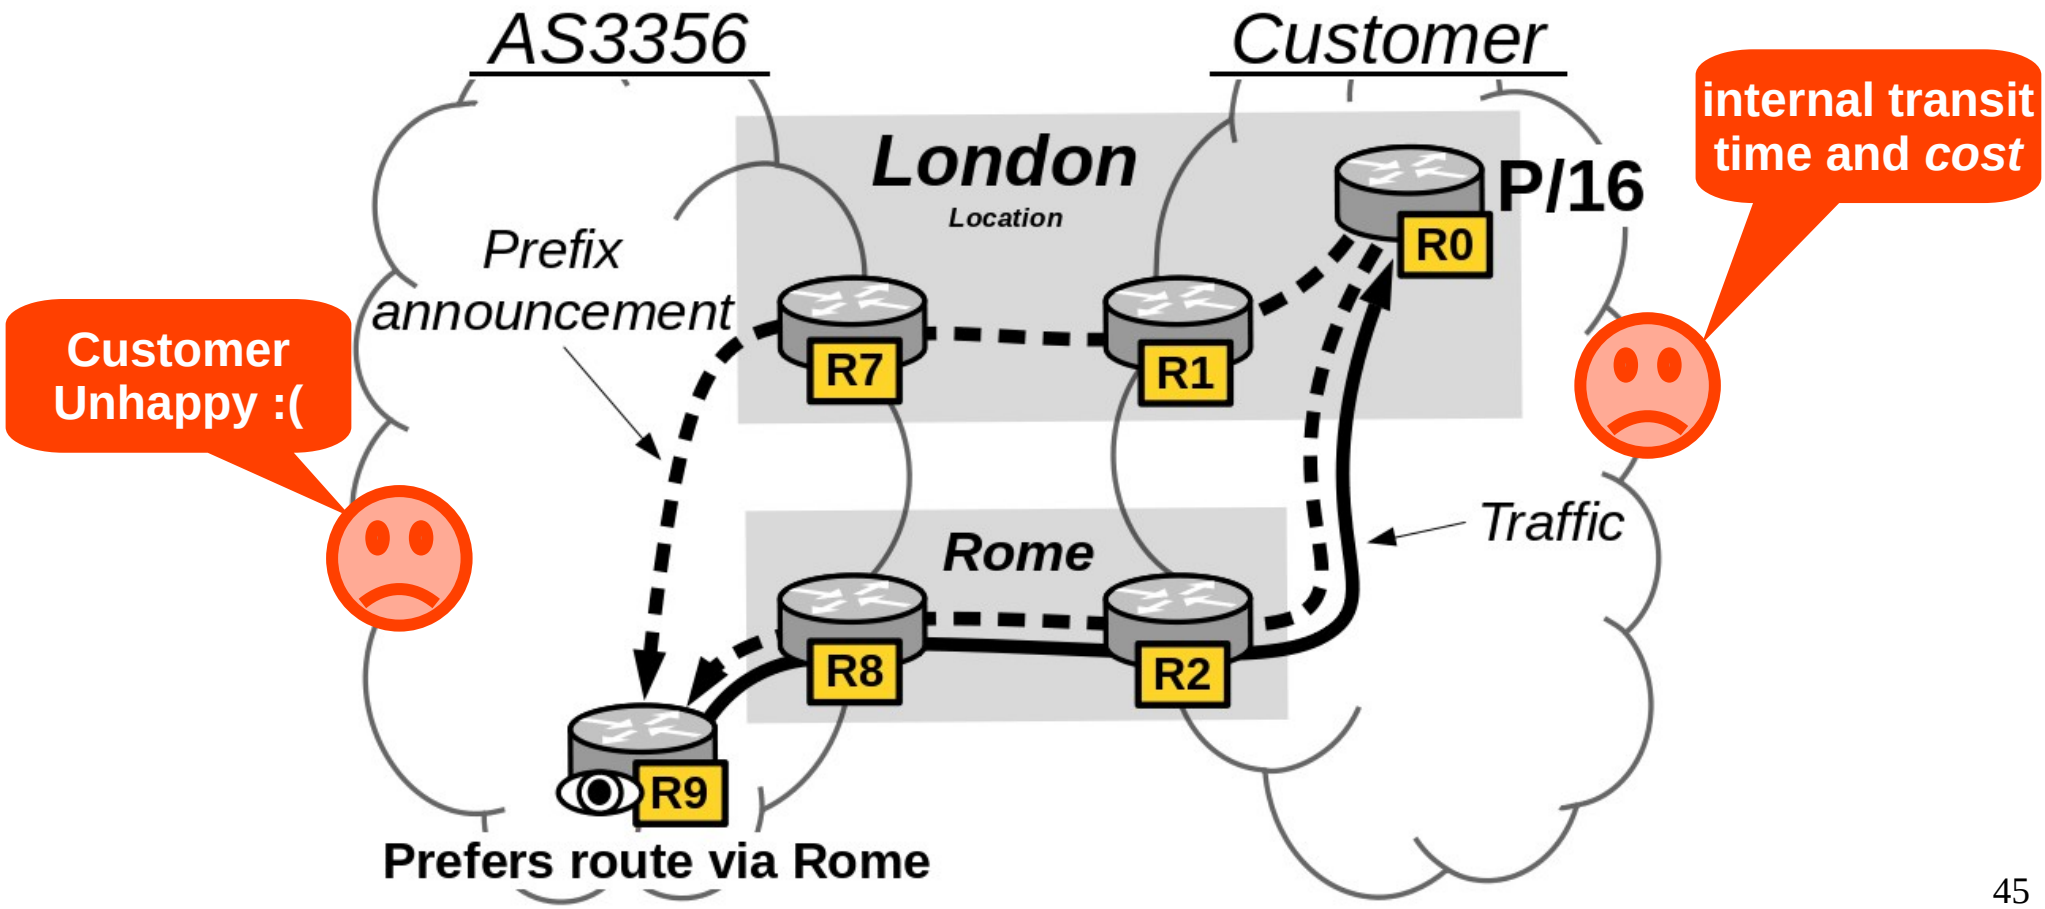

Use Case: Cold Potato Routing

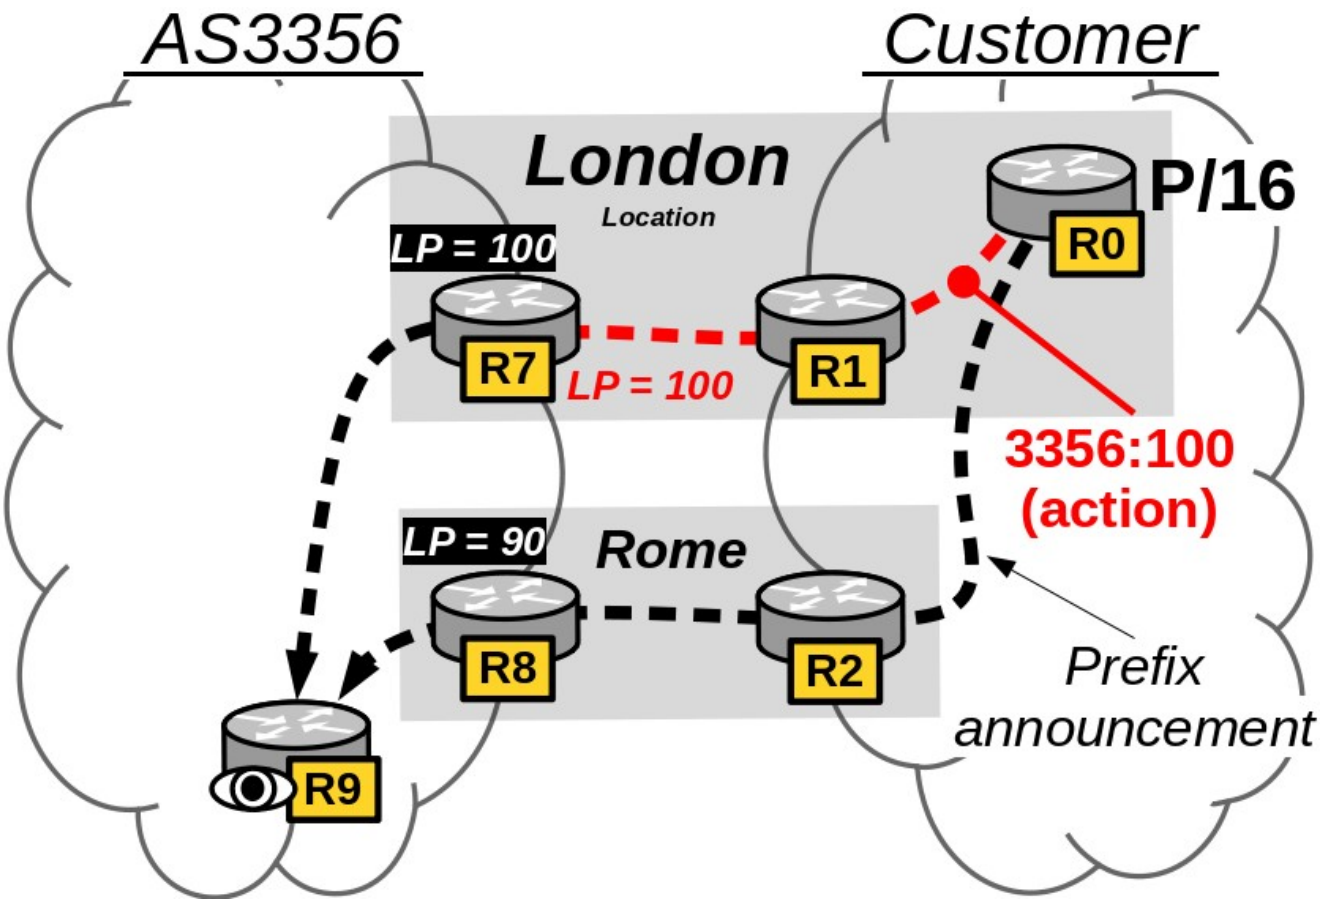

Use Case: Cold Potato Routing

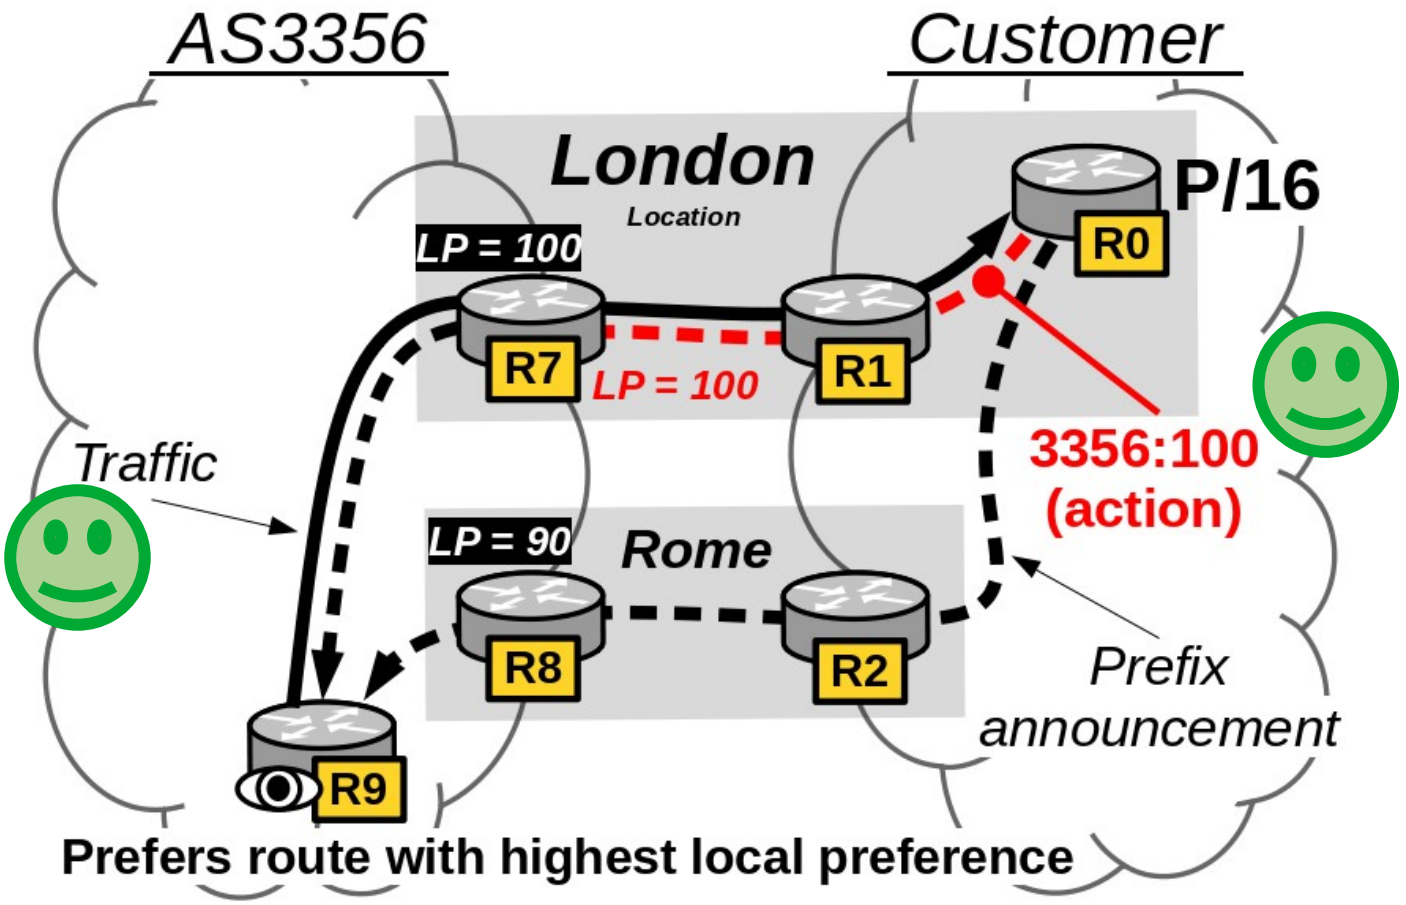

Use Case: Cold Potato Routing

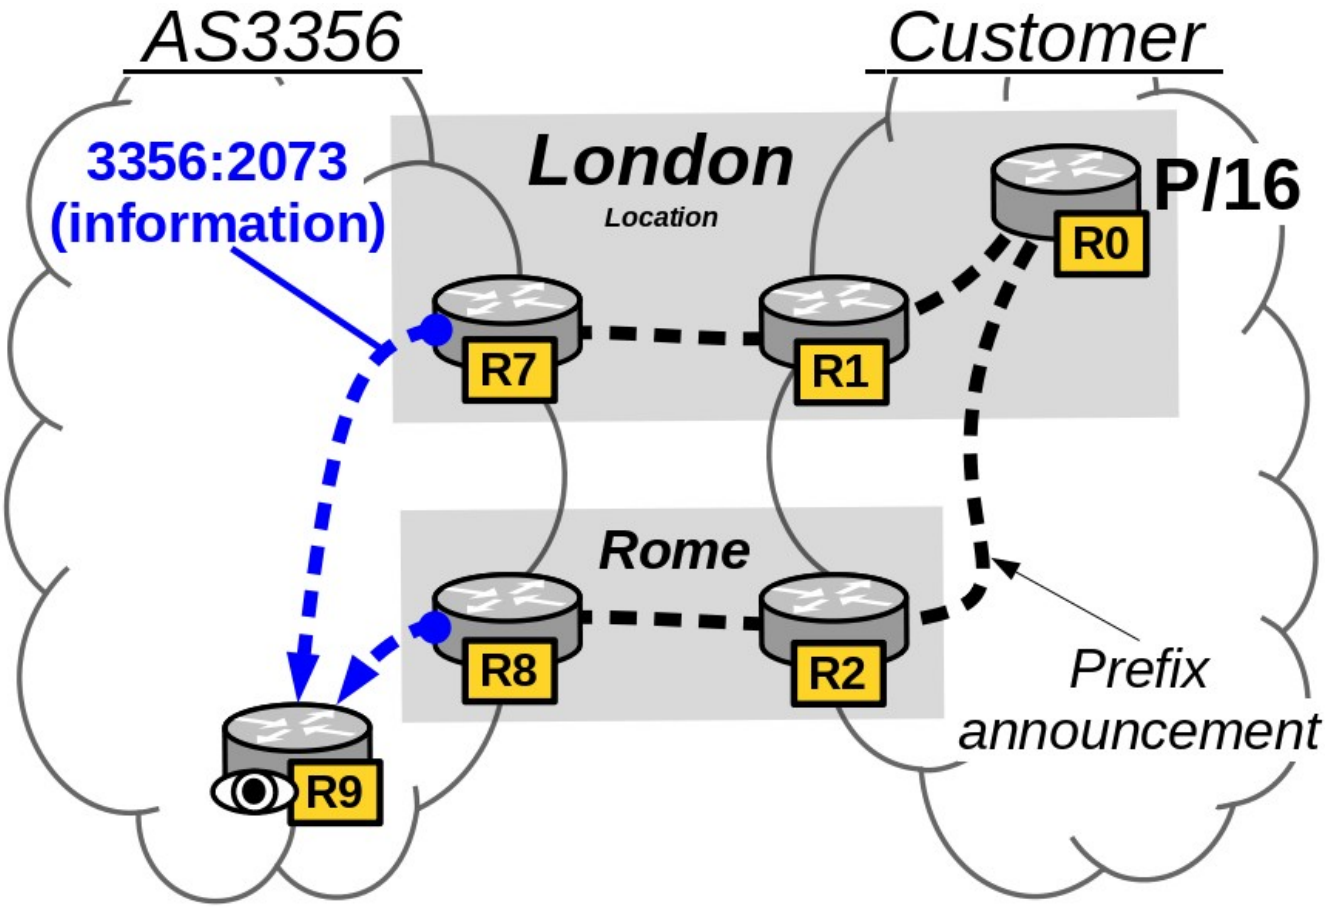

Use Case: Cold Potato Routing

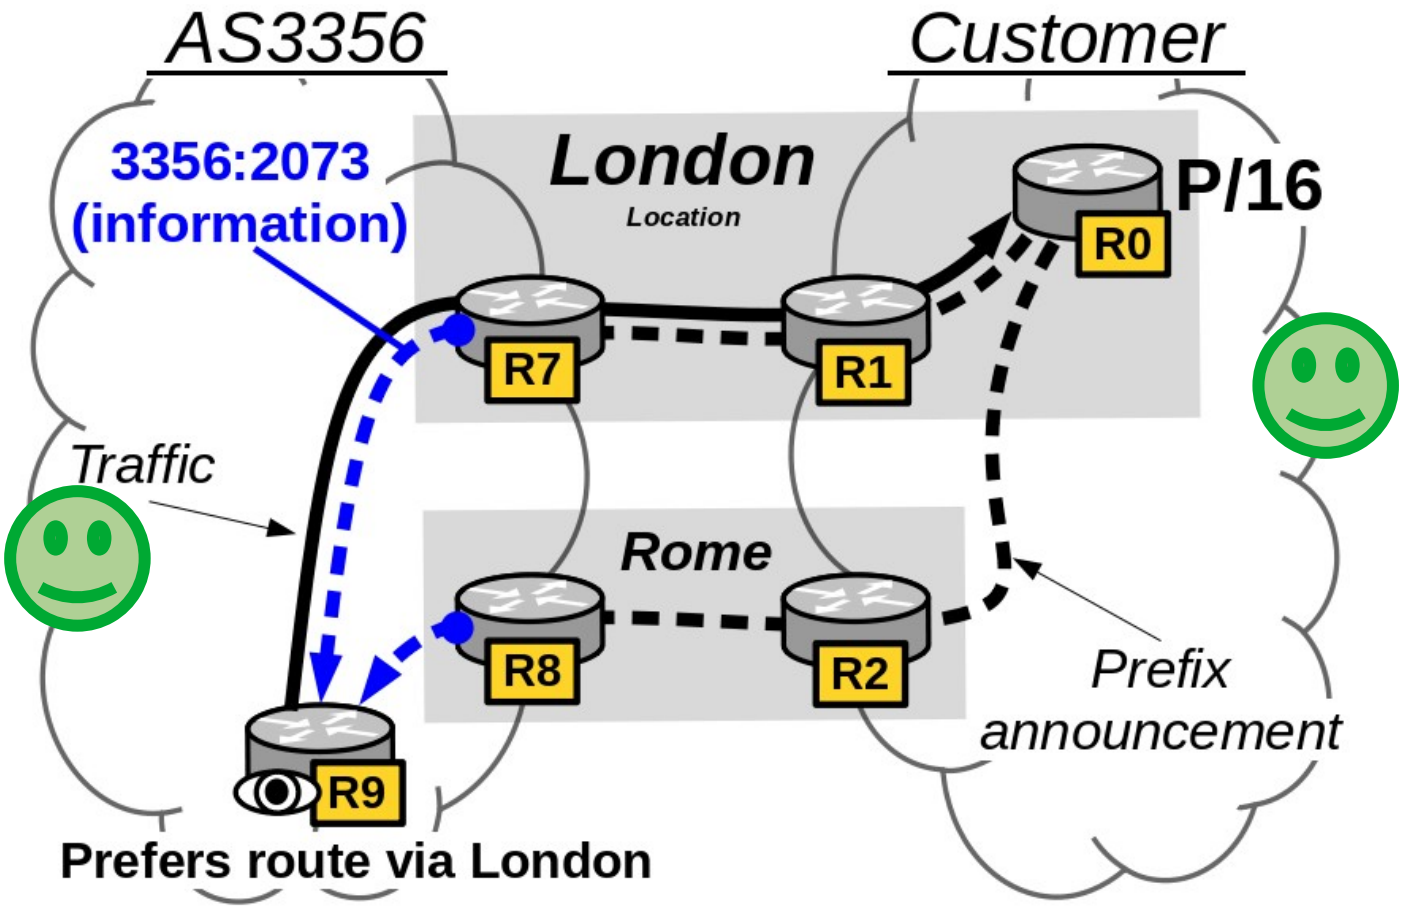

Use Case: Cold Potato Routing

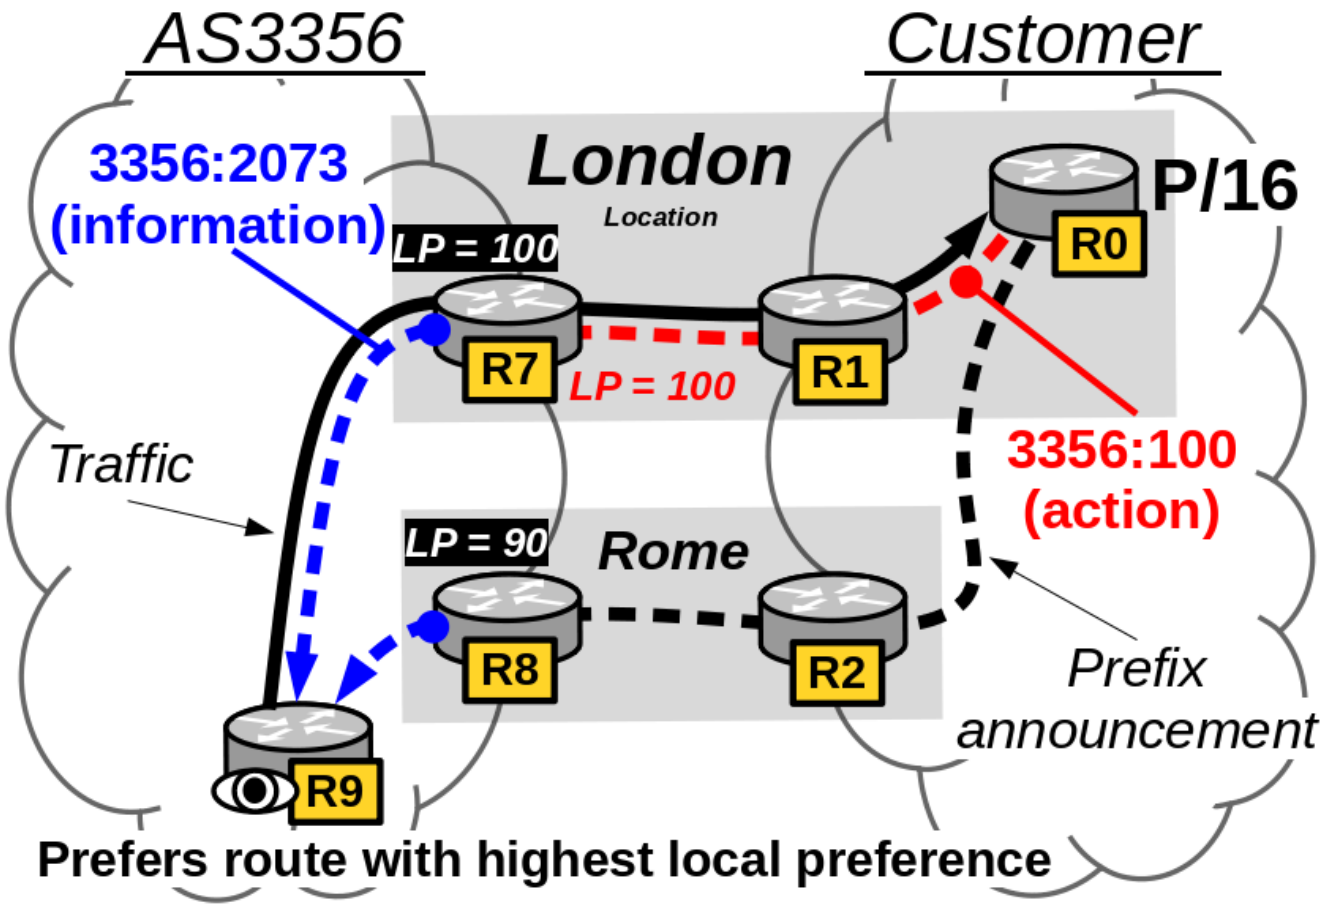

Use Case: Cold Potato Routing

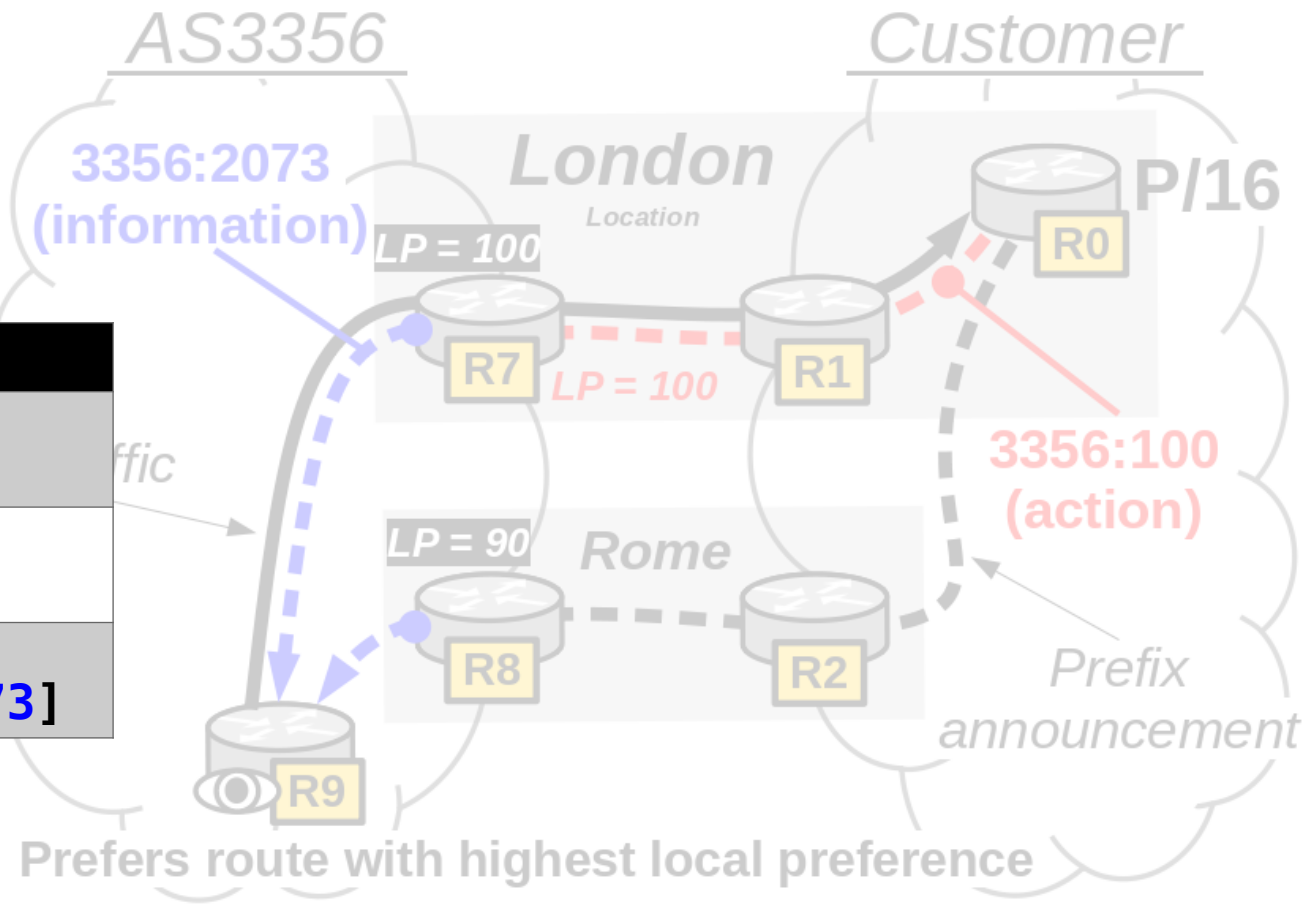

Route
Prefix: **[P/16]**
AS Path: [3356 X]
Communities: **[3356:2073]**

Chapter III:
BGP Communities in Research

Research utilizing BGP communities

Topology

- Mapping Peering Interconnections to a Facility, CoNEXT’15
- Improving the Discovery of IXP Peering Links. INFOCOM’13
- Valley-Free Violation in Internet Routing. ICC’12

Usage

- Collecting Self-reported Semantics of BGP Communities. IMC’24
- Usage of IXPs' Action BGP Communities. CoNEXT’22
- AS-Level BGP Community Usage Classification. IMC’21
- On BGP Communities. CCR’08

Research utilizing BGP communities

Events

- Large Scale Outage Visibility on the Control Plane. CoNEXT-SW’21
- DoS Attacks and BGP Blackholing in the Wild. IMC’18
- Inferring BGP Blackholing Activity in the Internet. IMC’17
- Detecting Peering Infrastructure Outages in the Wild. SIGCOMM’17

Update rate

- Exploring the Routing Message Impact of BGP Communities. CoNEXT’20
- What do parrots and BGP routers have in common?. CCR’16

Research utilizing BGP communities

Security

- Surgical Interception Attacks by Manipulating BGP Communities. CCS’19
- BGP Communities: Even more Worms in the Routing Can. IMC’18

Classification

- Uncovering BGP Action Communities. ACM MAC’24
- Coarse-Grained Inference of BGP Community Intent. IMC’23
- Automatic Inference of BGP Location Communities. SIGMETRICS’22

Research utilizing BGP communities

Collecting Self-reported Semantics of BGP Communities. IMC'24
Usage of IXPs' Action BGP Communities. CoNEXT'22
AS-Level BGP Community Usage Classification. IMC'21
On BGP Communities. CCR'08
Mapping Peering Interconnections to a Facility, CoNEXT'15
Improving the Discovery of IXP Peering Links. INFOCOM'13
Valley-Free Violation in Internet Routing. ICC'12
Large Scale Outage Visibility on the Control Plane. CoNEXT-SW'21
DoS Attacks and BGP Blackholing in the Wild. IMC'18
Inferring BGP Blackholing Activity in the Internet. IMC'17
Detecting Peering Infrastructure Outages in the Wild. SIGCOMM'17
What do parrots and BGP routers have in common?. CCR'16
Exploring the Routing Message Impact of BGP Communities. CoNEXT'20
Surgical Interception Attacks by Manipulating BGP Communities. CCS'19
BGP Communities: Even more Worms in the Routing Can. IMC'18
Automatic Inference of BGP Location Communities. SIGMETRICS'22
Coarse-Grained Inference of BGP Community Intent. IMC'23

Usage
Topology
Events
Update rate
Security
Classification

Community Documentation

- Because community values are opaque, **dictionaries are needed**
  - websites
  - NLNOG repository
  - Bgp.tools
- Recent study shows **about 90%** of routed communities are **not documented** [1]
- We need to infer communities ourselves
  - Goal: **<u>Create BGP community dictionary for Research</u>**

[1] “Collecting Self-reported Semantics of BGP Communities” IMC’24

level Origin Communities.

Numbering Structure

Community numbering uses the following structure:

- 1299:xyzzz
- Where:
  - x is BGP Neighbour type; 2 for Peers or 3 for Customers
  - y is Region; 0 for Europe, 5 for North America or 7 for Asia & Pacific
  - zzz is City; see below

Currently available Customer Origin + Communities are listed below:

Europe

Community	Country	IP city prefix	Description
1299:30000			**European Customers**
1299:30100	DK	kbn	Copenhagen
1299:30110	SE	got	Gothenburg
1299:30200	SE	sto	Stockholm
1299:30210	LT	kau	Kaunas
		vls	Vilnius
1299:30220	NO	oso	Oslo

https://www.arelion.com/our-network/bgp-routing/bgp-communities

Chapter IV:
City Communities

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*

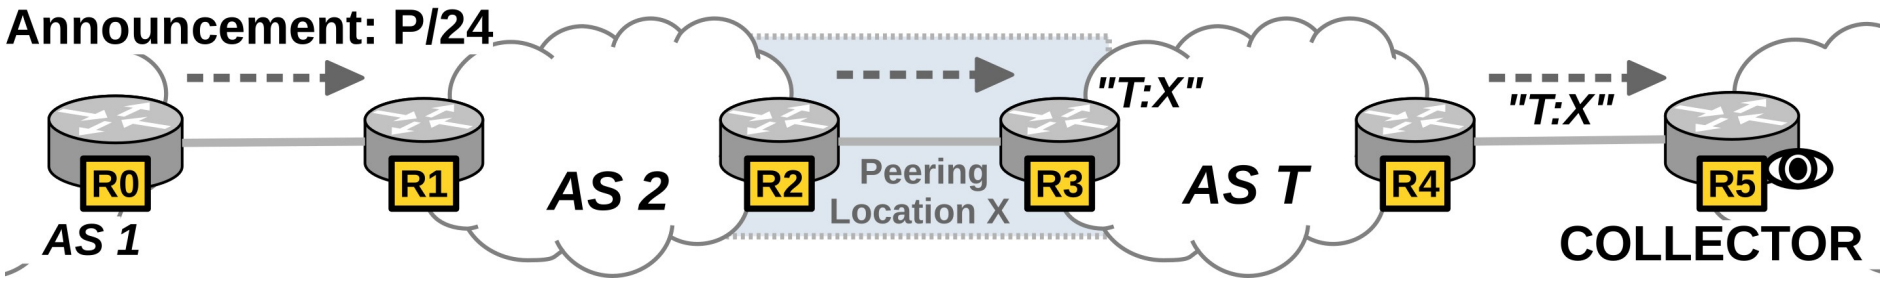

***Challenge*: In BGP data, there is no notion of tagging routers**

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*

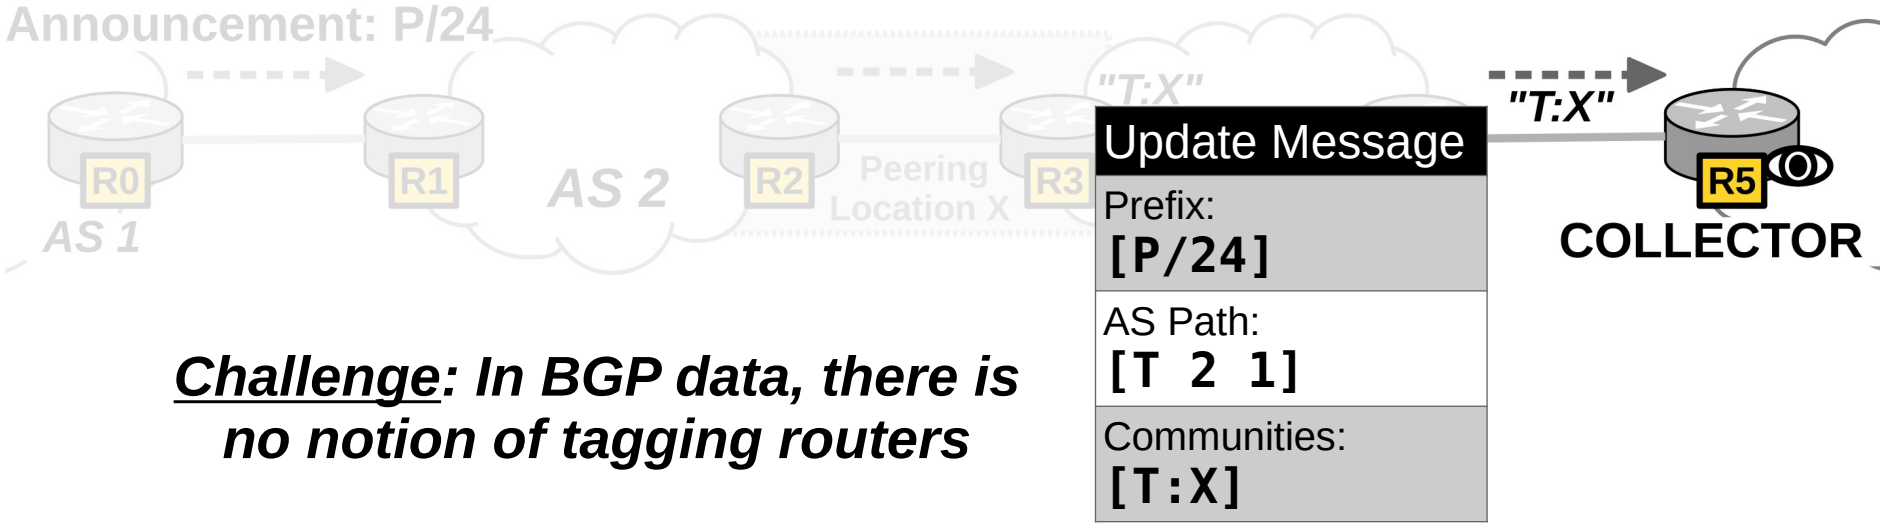

<u>Challenge</u>: In BGP data, there is no notion of tagging routers

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- **Two basic approaches** ***(to determine location of tagging router)***

  A) Traceroute + router geolocation (active)

  B) BGP + prefix geolocation (passive)

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- Two basic approaches *(to determine location of tagging router)*

  A) Traceroute + router geolocation (active)

  B) BGP + prefix geolocation (passive)

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- Two basic approaches *(to determine location of tagging router)*

  A) Traceroute + router geolocation (active)

  B) BGP + prefix geolocation (passive)

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- Two basic approaches *(to determine location of tagging router)*

A) Traceroute + router geolocation (active)

B) BGP + prefix geolocation (passive)

Cold potato routing

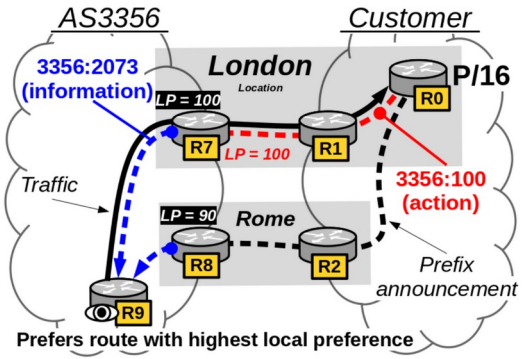

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- Two basic approaches *(to determine location of tagging router)*

  A) Traceroute + router geolocation (active)

  B) BGP + prefix geolocation (passive)
- **Validation using ground truth**
  - We manually collect coordinates for ~1,500 city communities

A) Traceroute + router geolocation

- Active approach to identify location ***X***

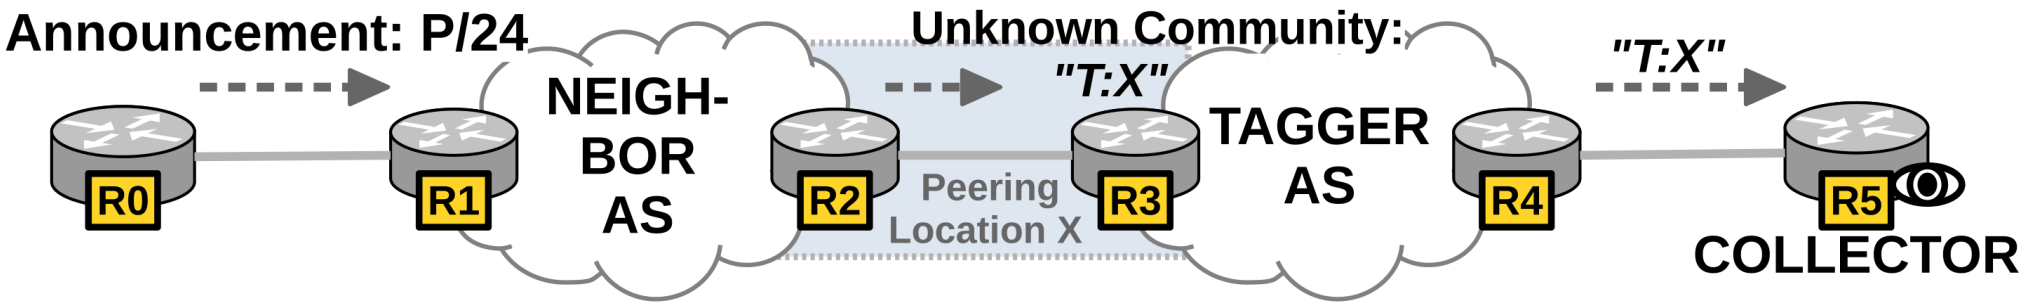

A) Traceroute + router geolocation

- Active approach to identify location ***X***

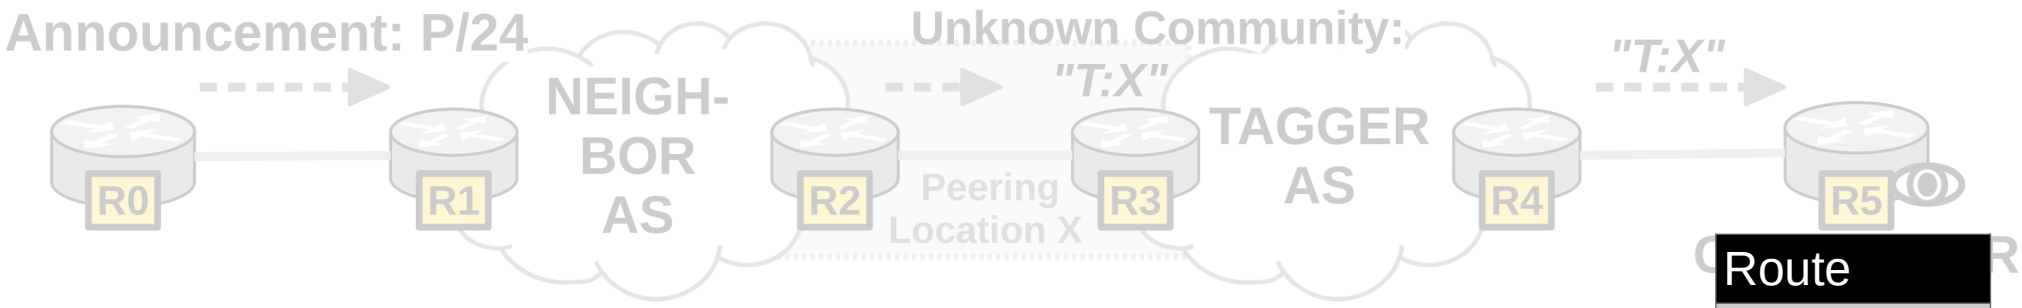

Route
Prefix: **[P/24]**
AS Path: **[T 2 1]**
Communities: **[T:X]**

A) Traceroute + router geolocation

- Active approach to identify location ***X***

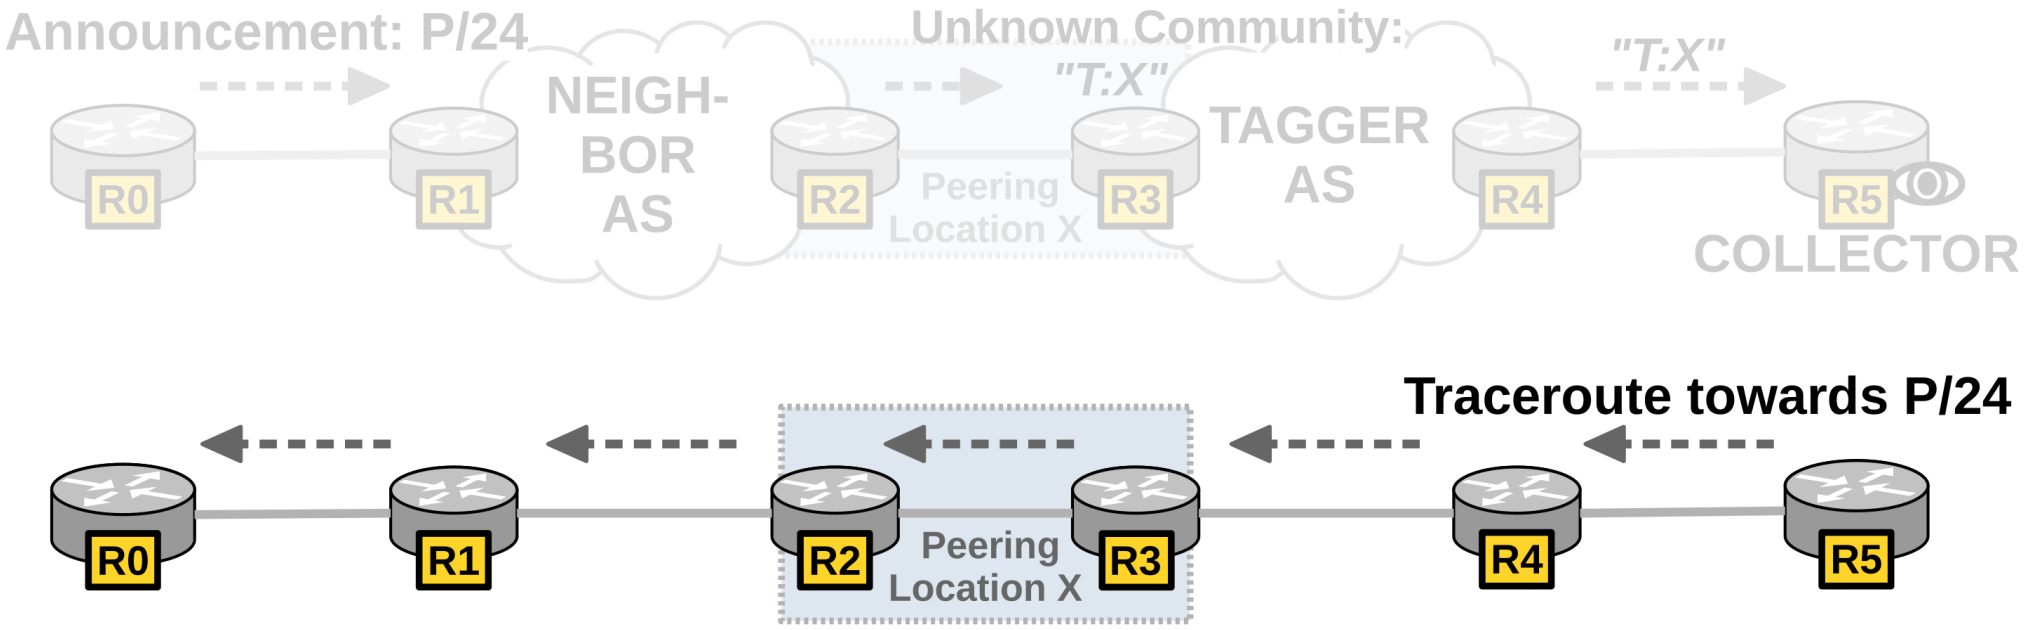

A) Traceroute + router geolocation

- Active approach to identify location ***X***

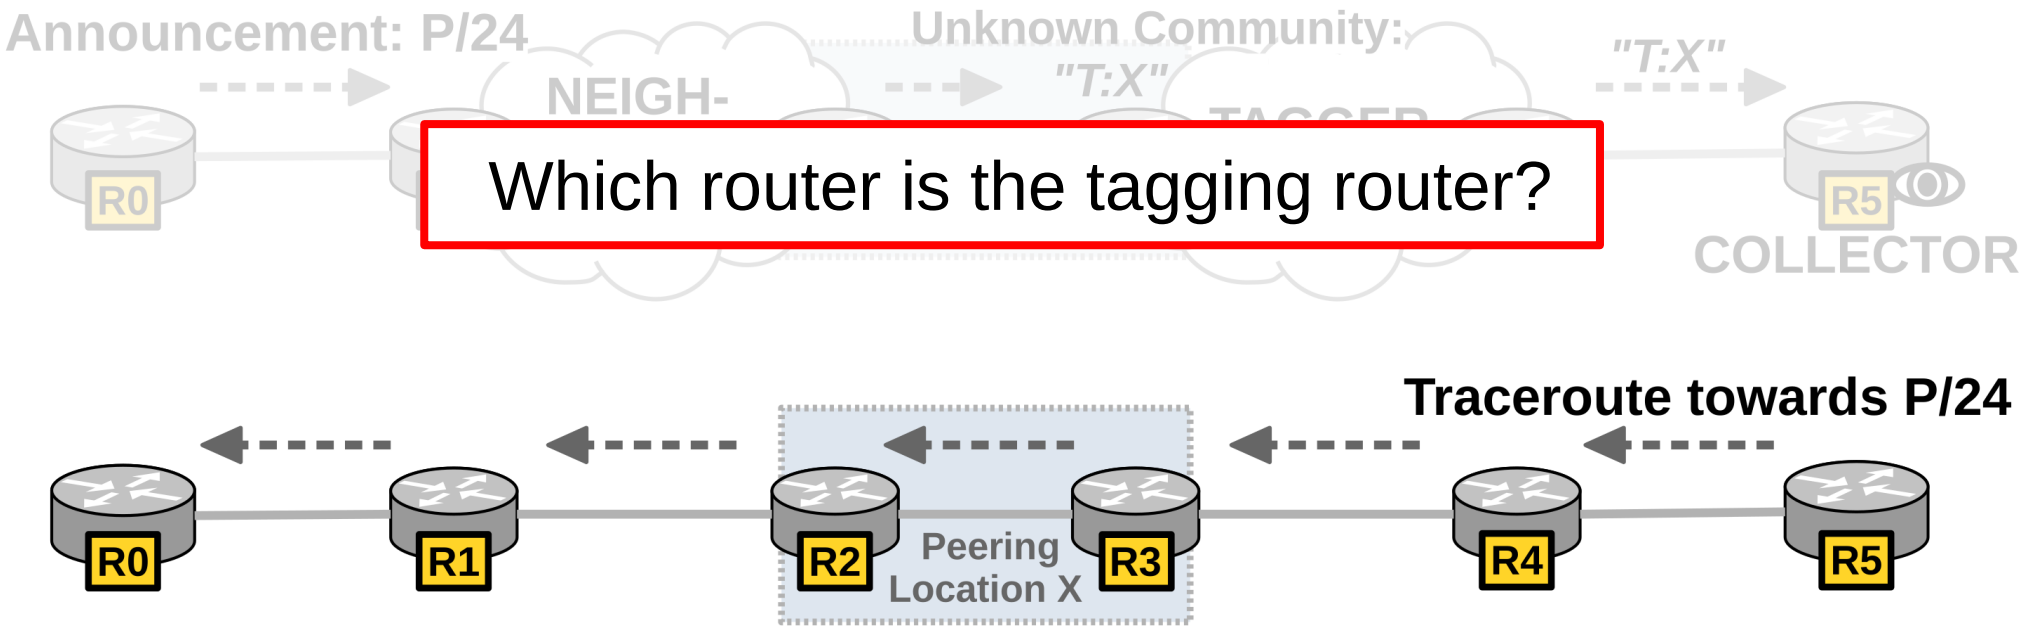

A) Traceroute + router geolocation

- Active approach to identify location ***X***

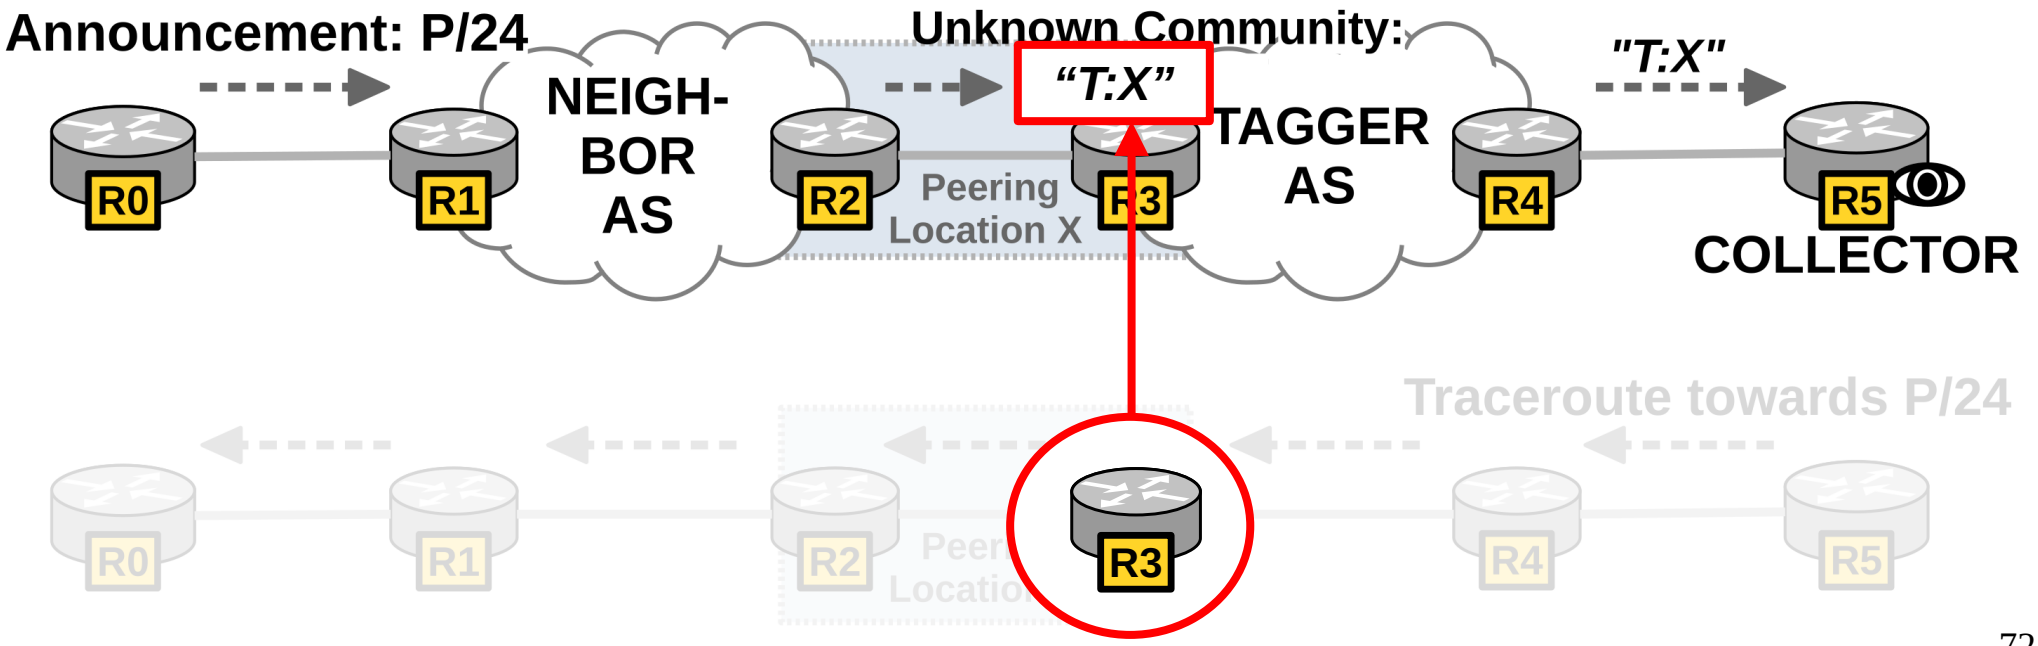

A) Traceroute + router geolocation

- Active approach to identify location ***X***

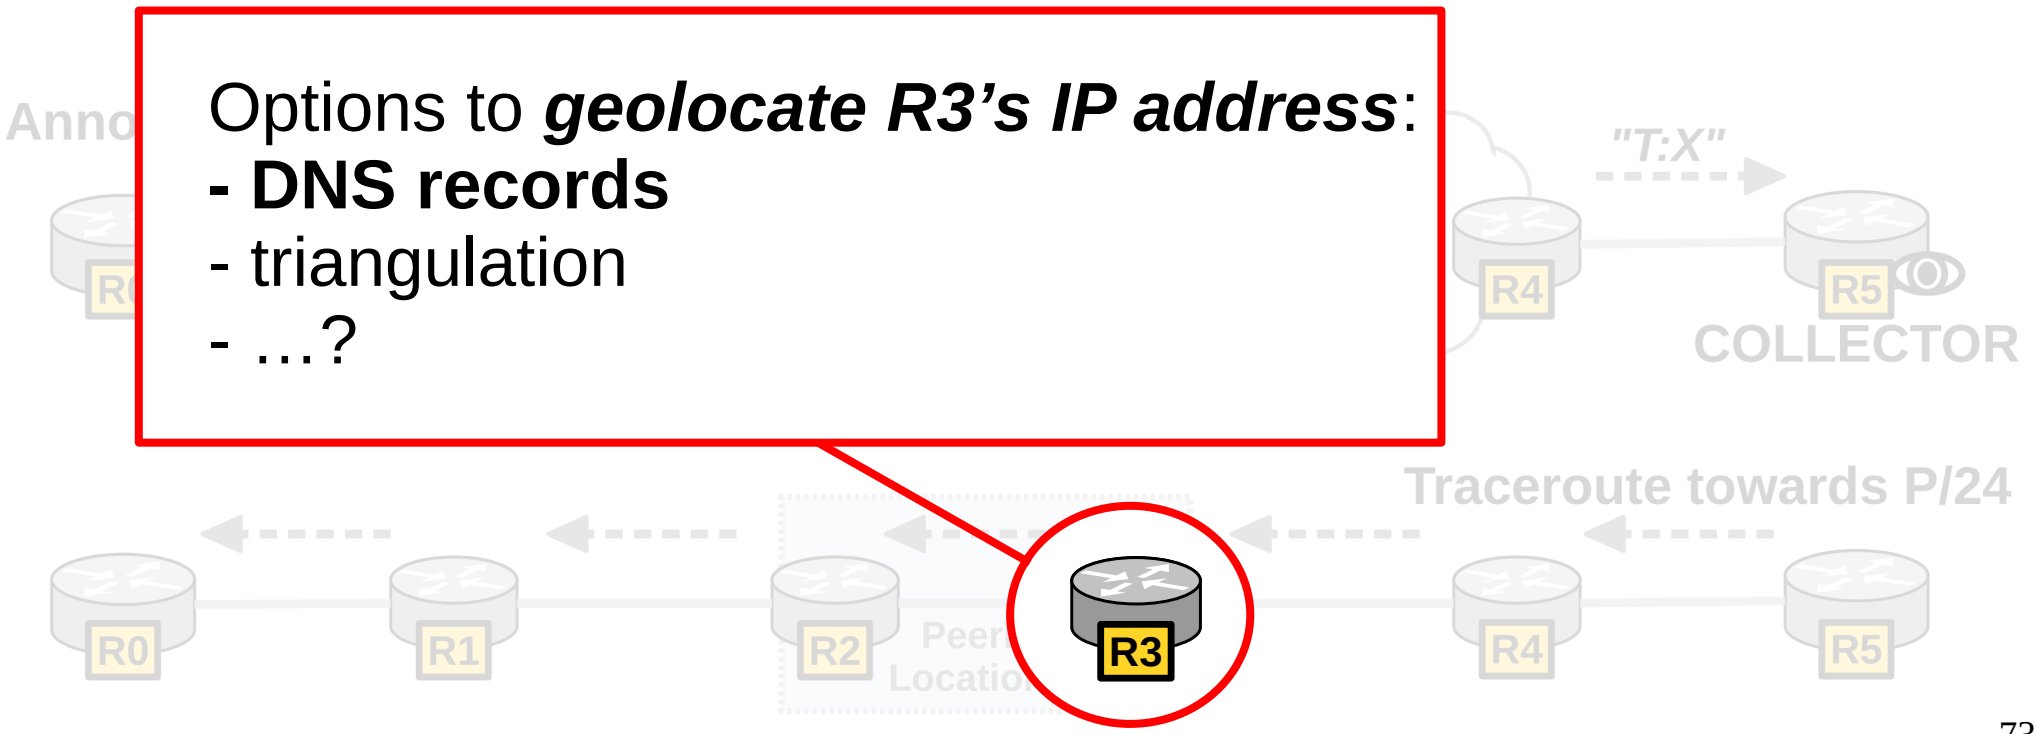

A) Traceroute + router geolocation

- General approach
  1) Perform traceroute
  2) Identify tagging router
  3) Geolocate tagging router
  4) Assign coordinates to community

A) Traceroute + router geolocation

- Results for AS2914 (NTT) using DNS records

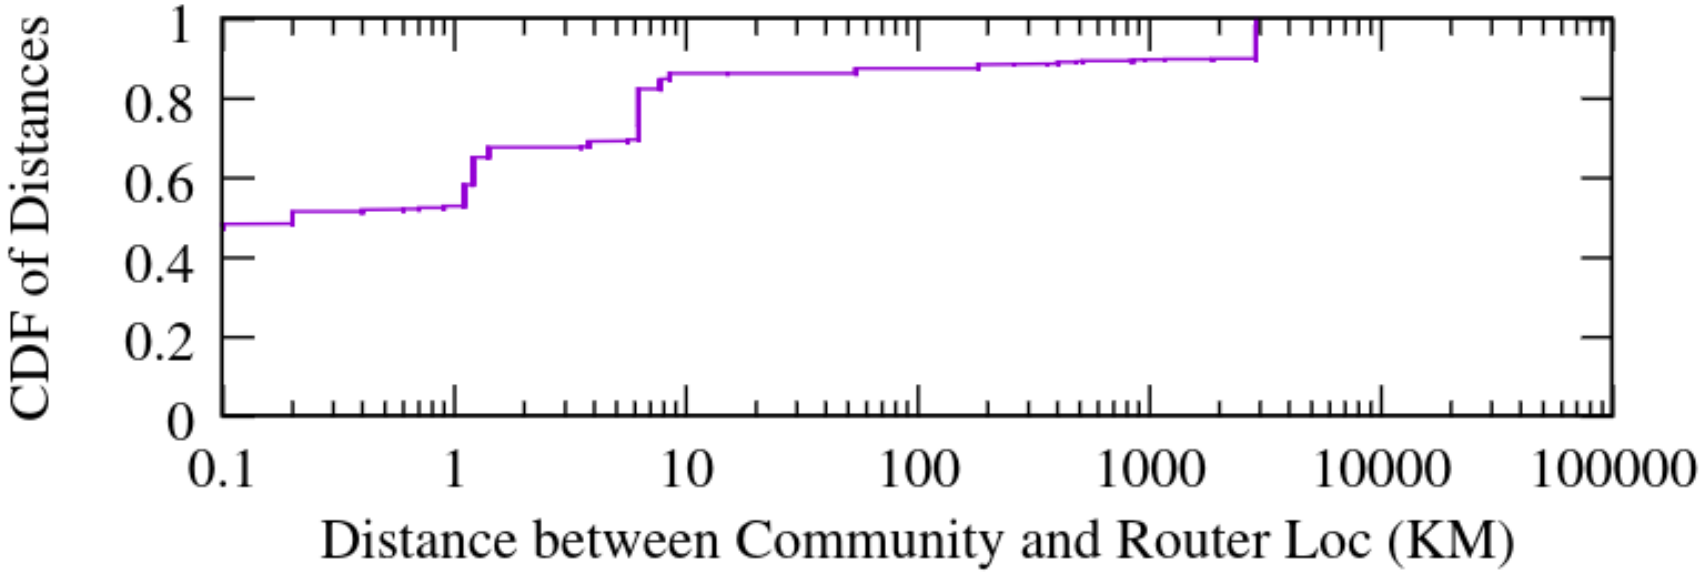

A) Traceroute + router geolocation

- Results for AS2914 (NTT) using DNS records

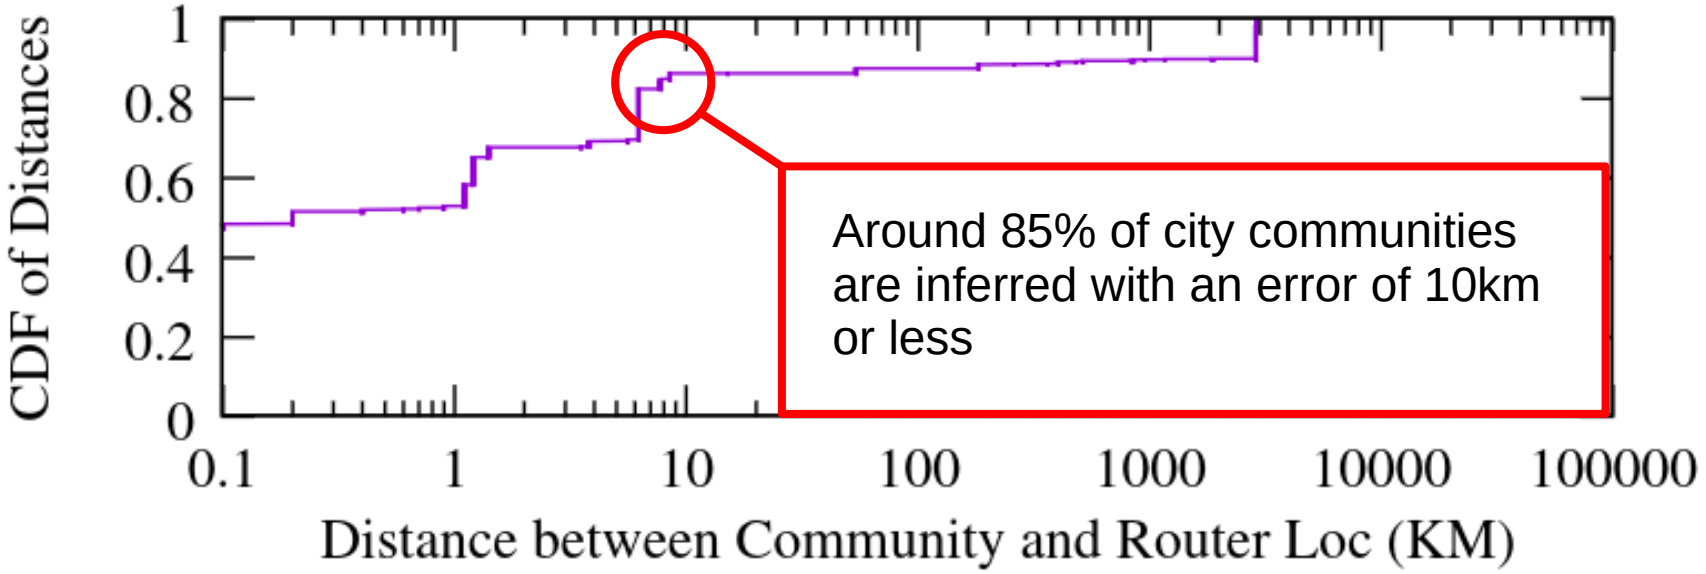

A) Traceroute + router geolocation

- Limitations
  - Vantage point for traceroute
  - IP Aliasing
  - DNS records
  - Triangulation

Inference of unknown city communities

- Given unknown city community T:X
  - what city does it signal?
  - *or: where is the tagging router located?*
- Two basic approaches *(to determine location of tagging router)*

  A) Traceroute + router geolocation (active)

  B) BGP + prefix geolocation (passive)

B) BGP + prefix geolocation

- Passive approach to identify location ***X***

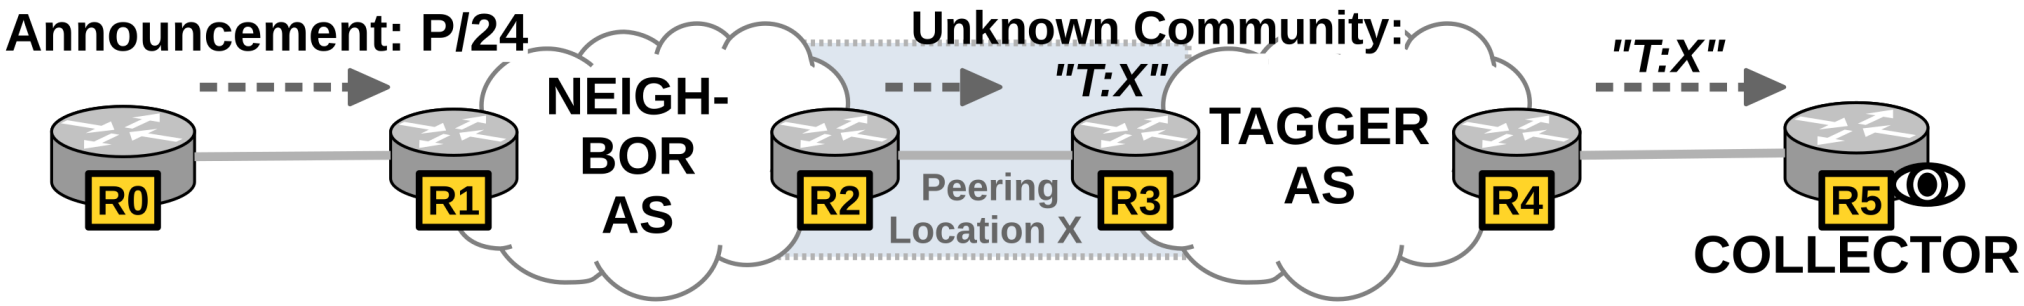

B) BGP + prefix geolocation

- Passive approach to identify location ***X***

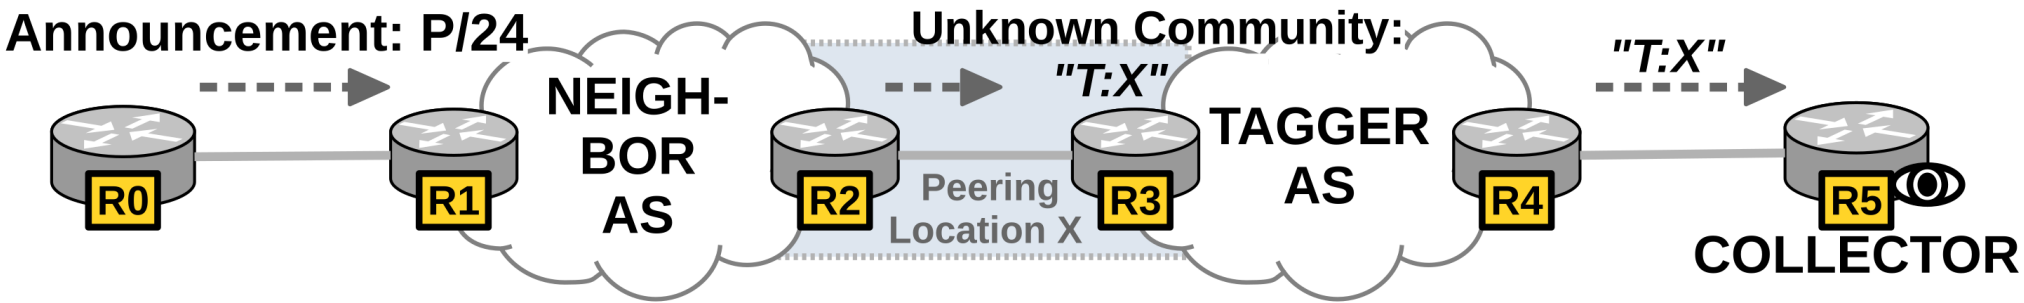

Research question:

Do city-tagged prefixes typically originate near the tagging router?

B) BGP + prefix geolocation

- Passive approach to identify location ***X***

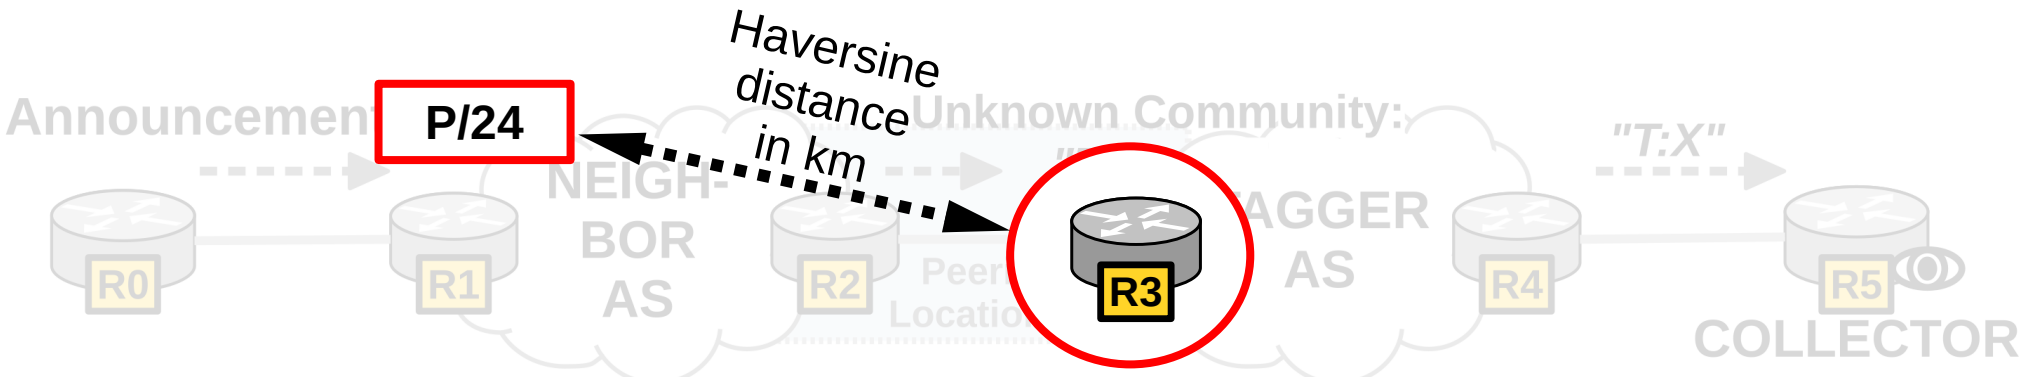

Research question:

Do city-tagged prefixes typically originate near the tagging router?

B) BGP + prefix geolocation

- Passive approach to identify location ***X***

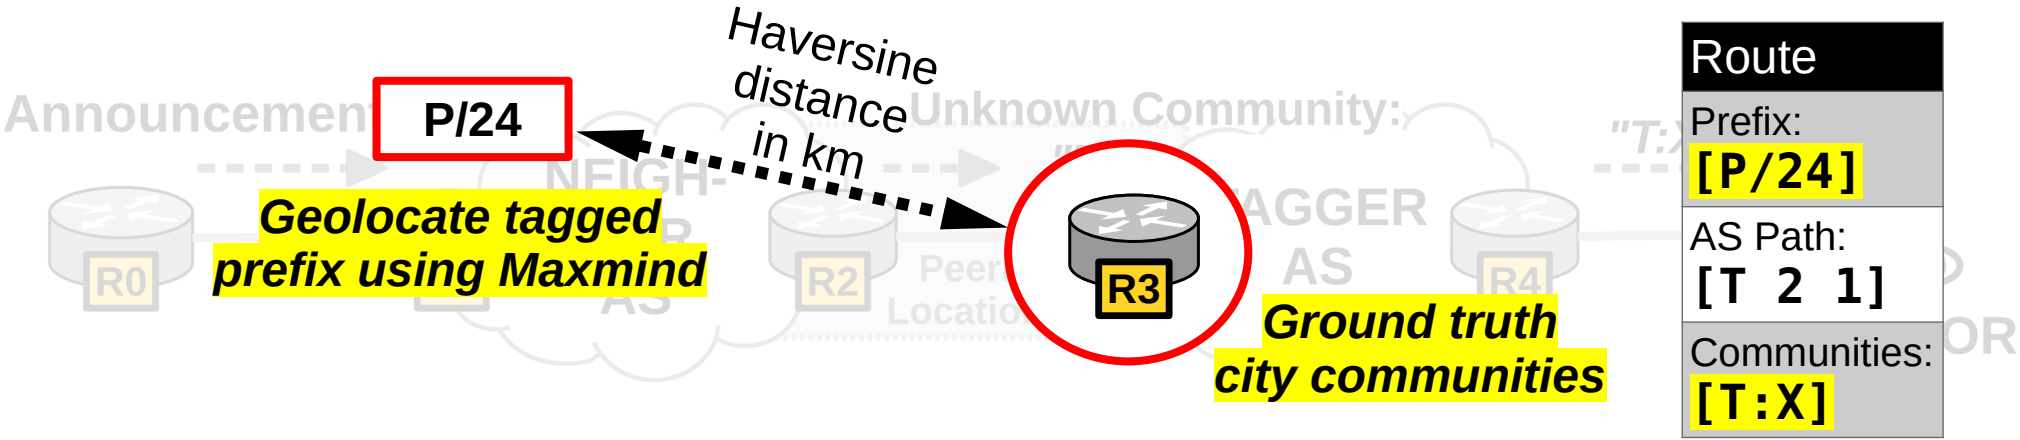

Research question:

Do city-tagged prefixes typically originate near the tagging router?

B) BGP + prefix geolocation

- Distances between tagged prefixes and ground truth city communities

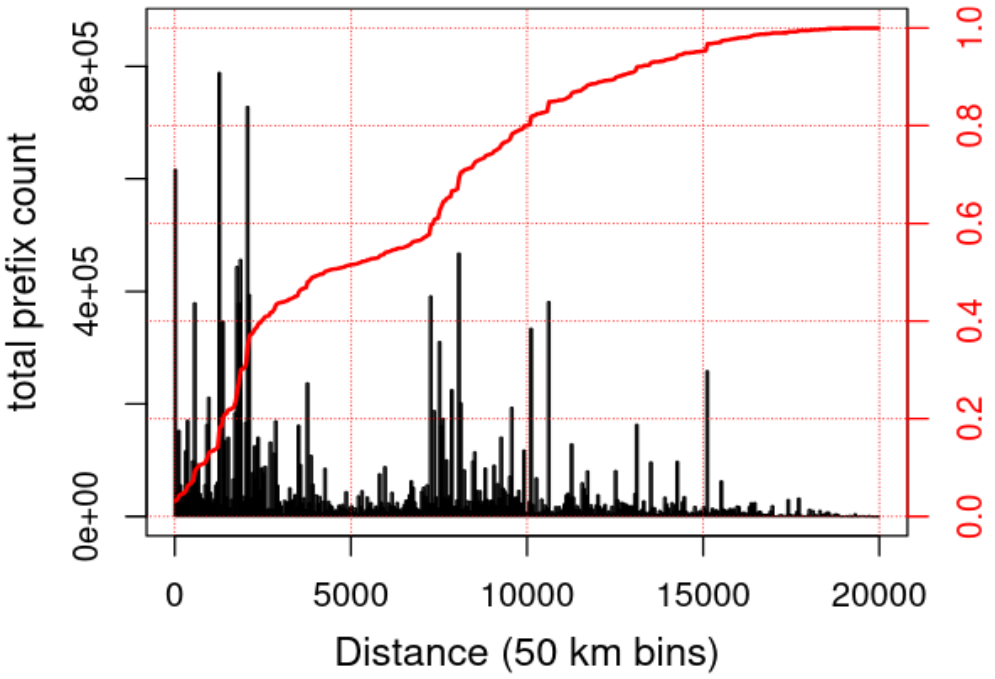

B) BGP + prefix geolocation

- Distances between tagged prefixes and ground truth city communities
- Only Around 600K (<5%) originate near the tagging router → local prefixes

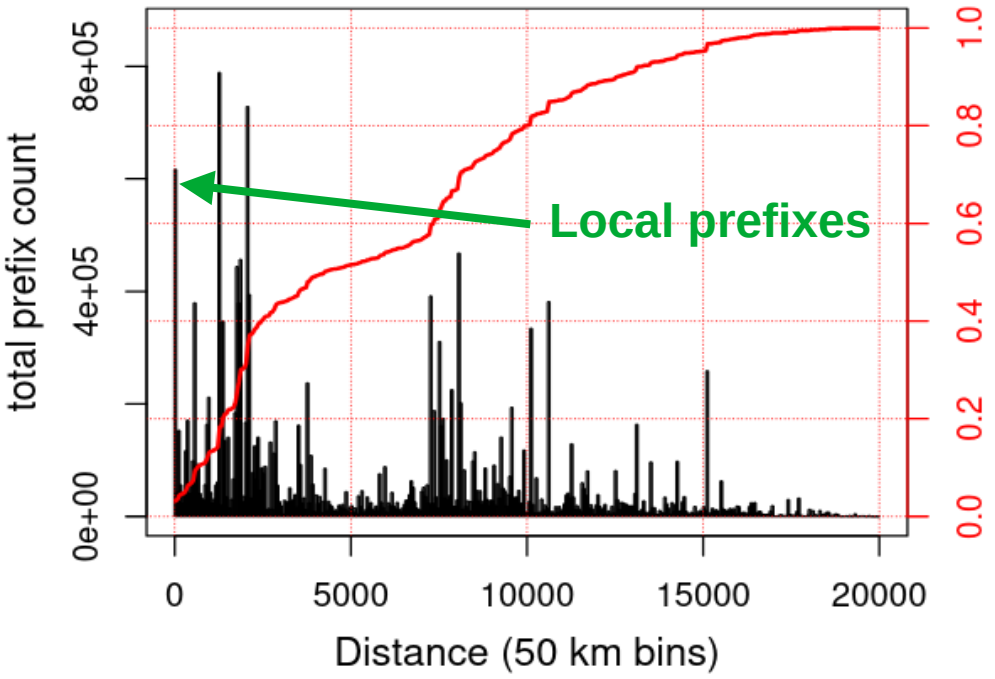

B) BGP + prefix geolocation

- Distances between tagged prefixes and ground truth city communities
- Only Around 600K (<5%) originate near the tagging router → local prefixes
- **Challenge: isolate local prefixes and use to infer location of tagging router**

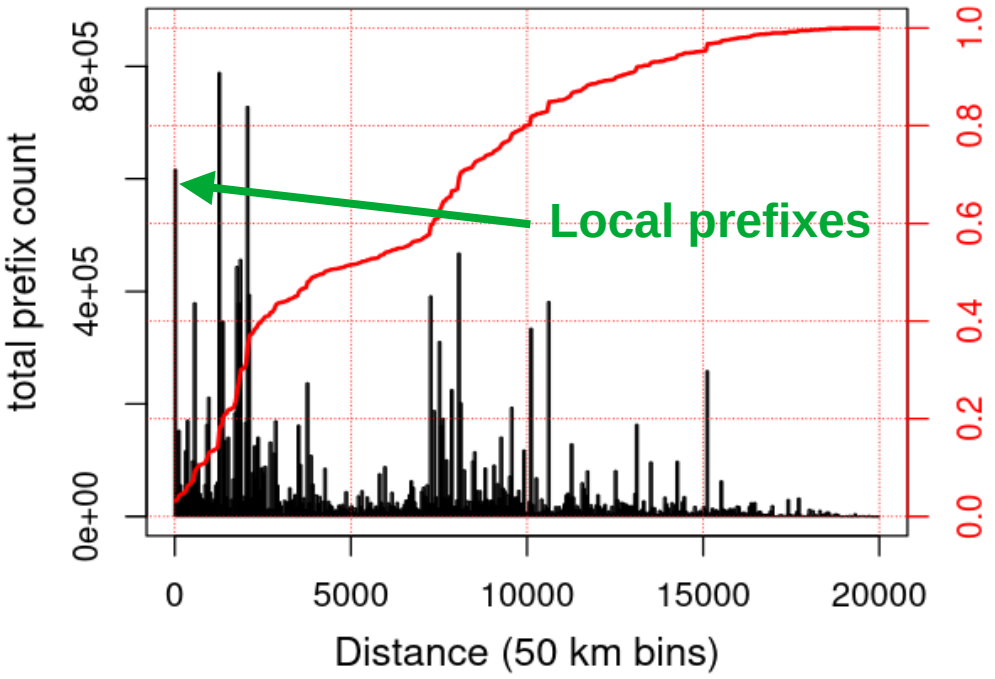

B) BGP + prefix geolocation

- General approach
  1) Obtain tagged prefixes
  2) Maximize share of local prefixes
  3) Cluster geographic locations
  4) Assign densest cluster to community

B) BGP + prefix geolocation

- General approach
  1) Obtain tagged prefixes
  2) Maximize share of local prefixes

B) BGP + prefix geolocation

2) Maximize share of local prefixes

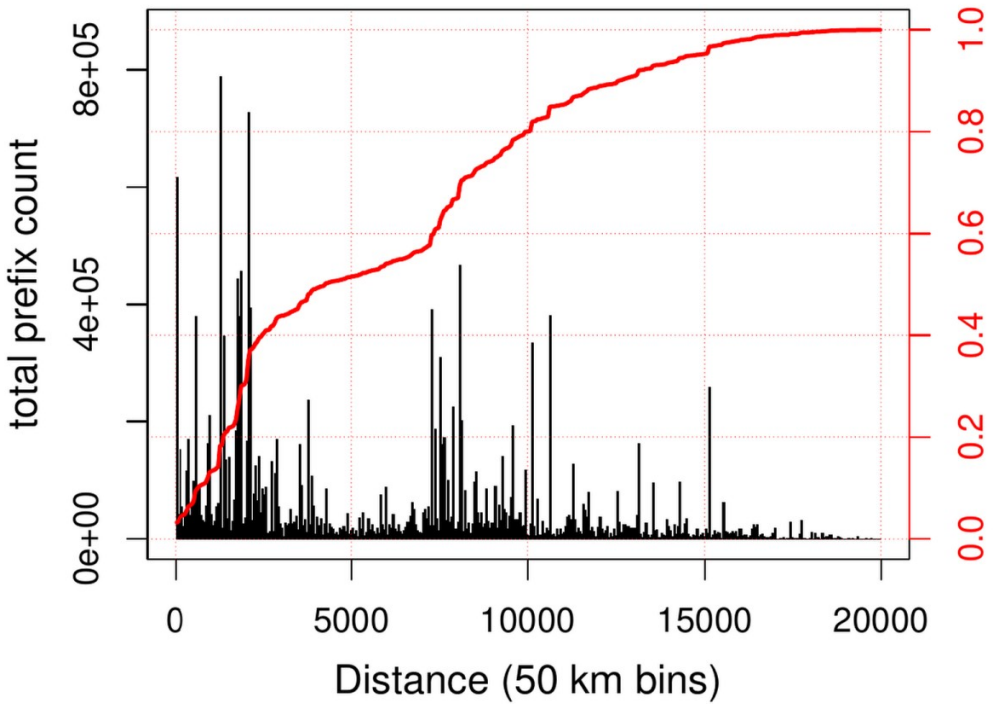

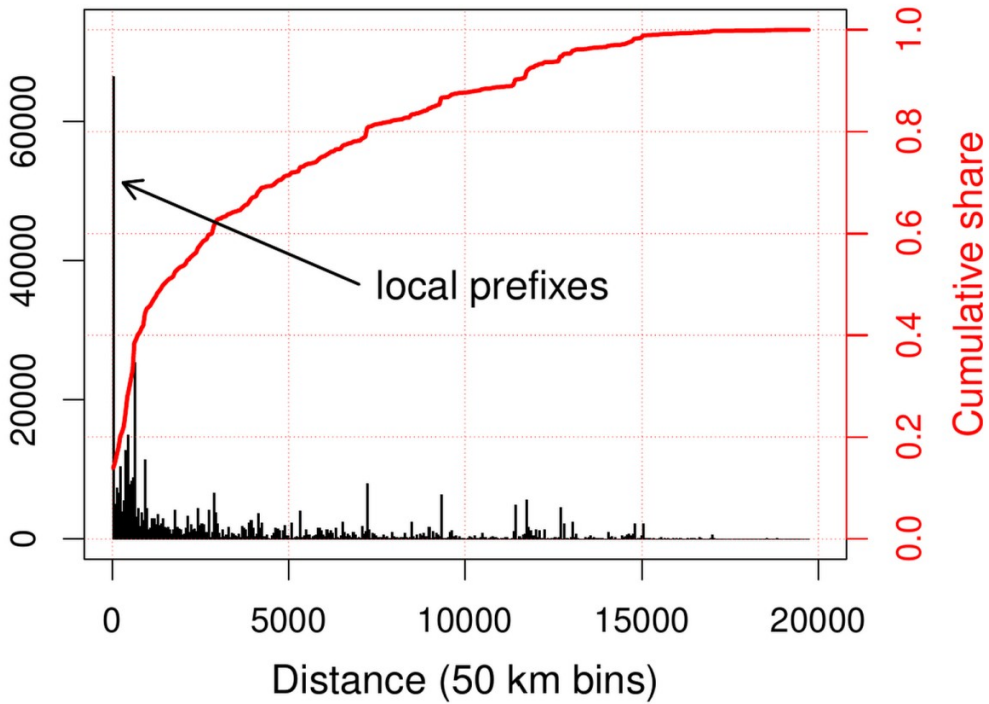

B) BGP + prefix geolocation

- General approach
  1) Obtain tagged prefixes
  2) Maximize share of local prefixes
  3) Cluster geographic locations

B) BGP + prefix geolocation

3) Clustering geolocations of tagged prefixes

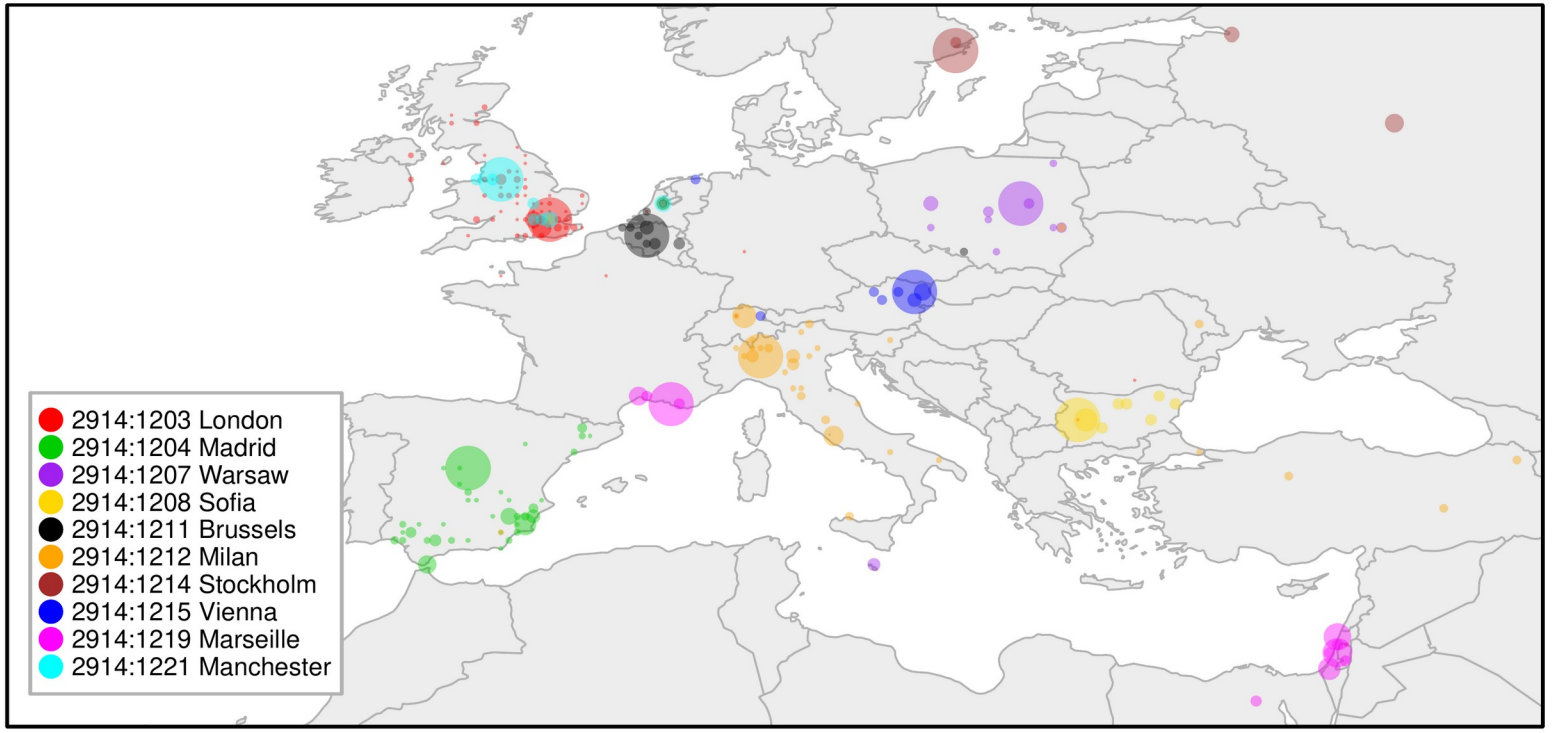

B) BGP + prefix geolocation

3) Clustering geolocations of tagged prefixes

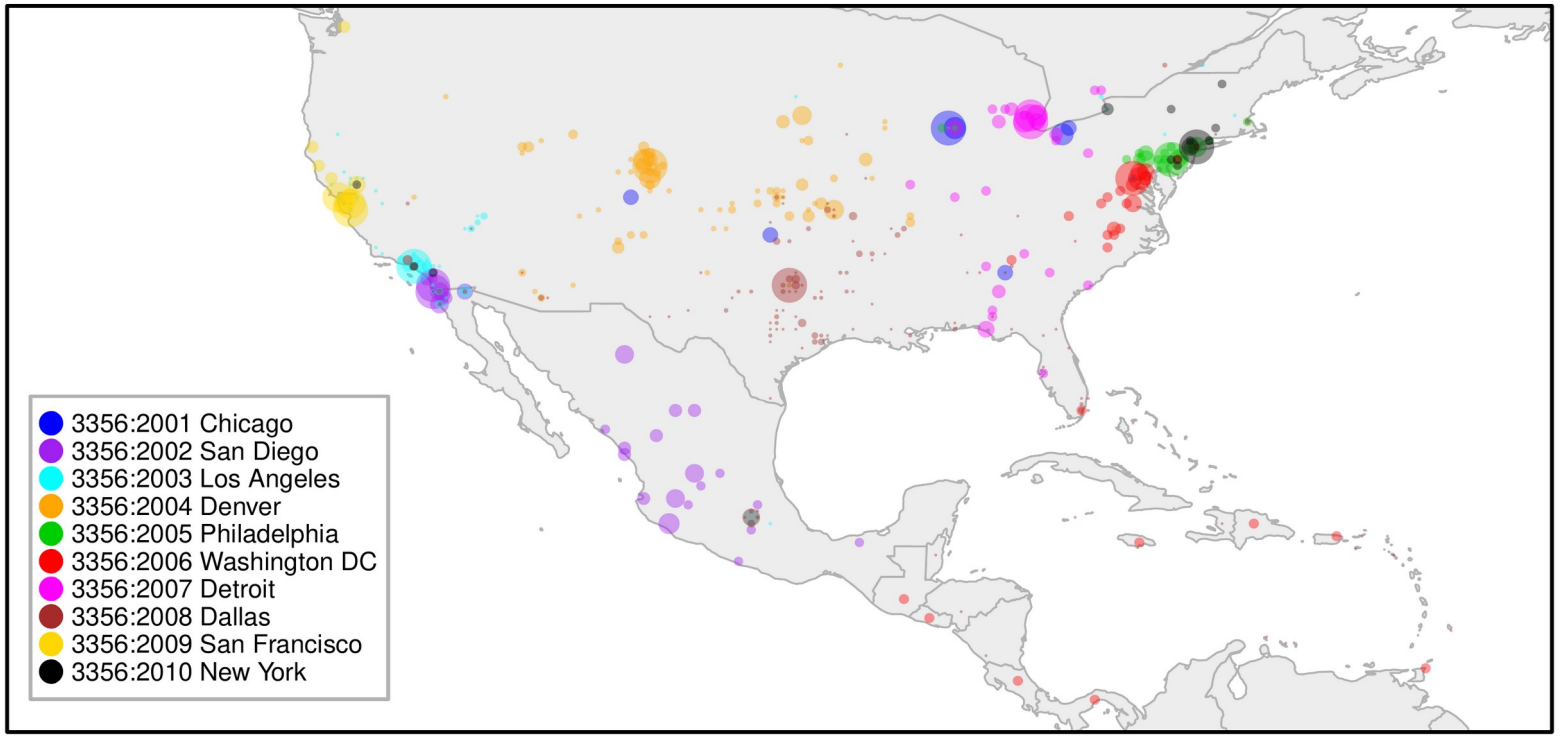

B) BGP + prefix geolocation

- General approach
  1) Obtain tagged prefixes
  2) Maximize share of local prefixes
  3) Cluster geographic locations
  4) Assign densest cluster to community

B) BGP + prefix geolocation

- 4) Assign densest cluster to community

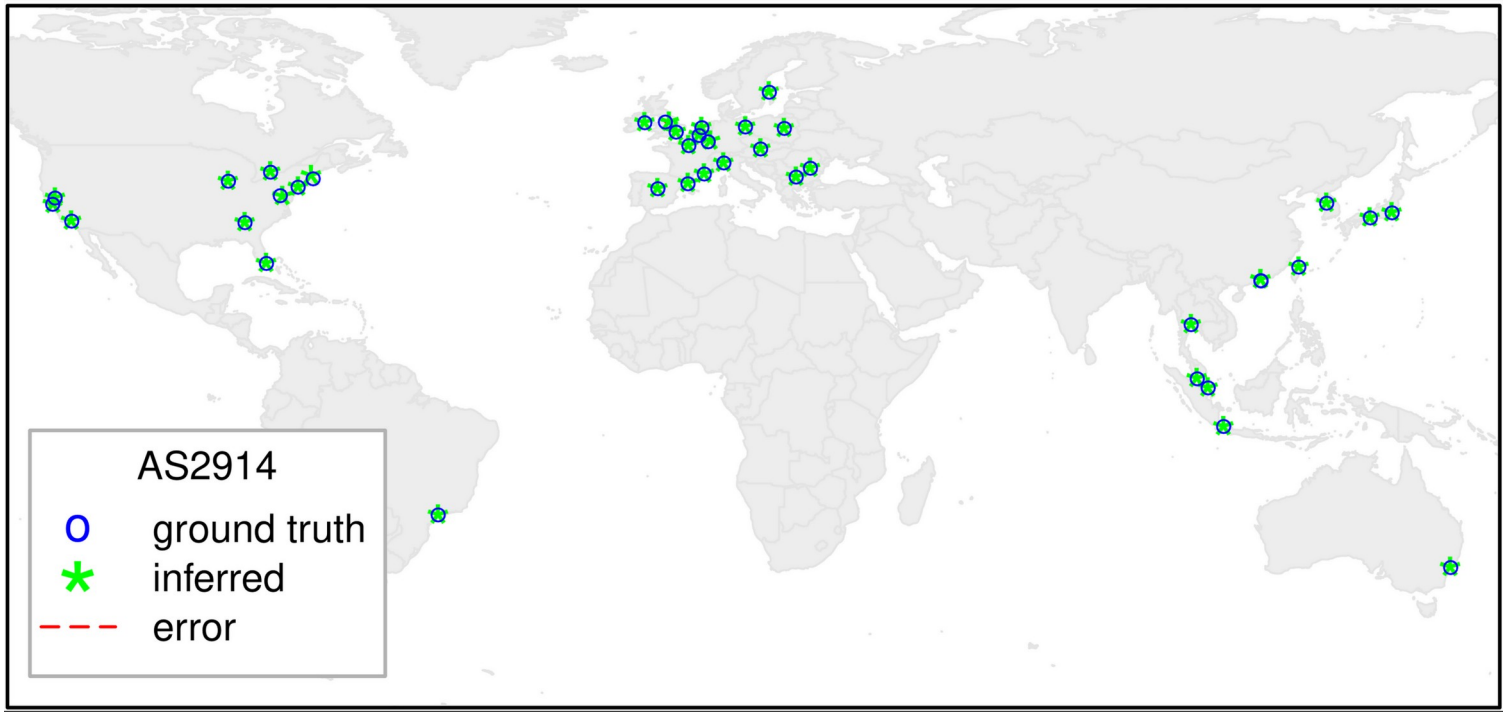

B) BGP + prefix geolocation

- Results for AS2914 (NTT)

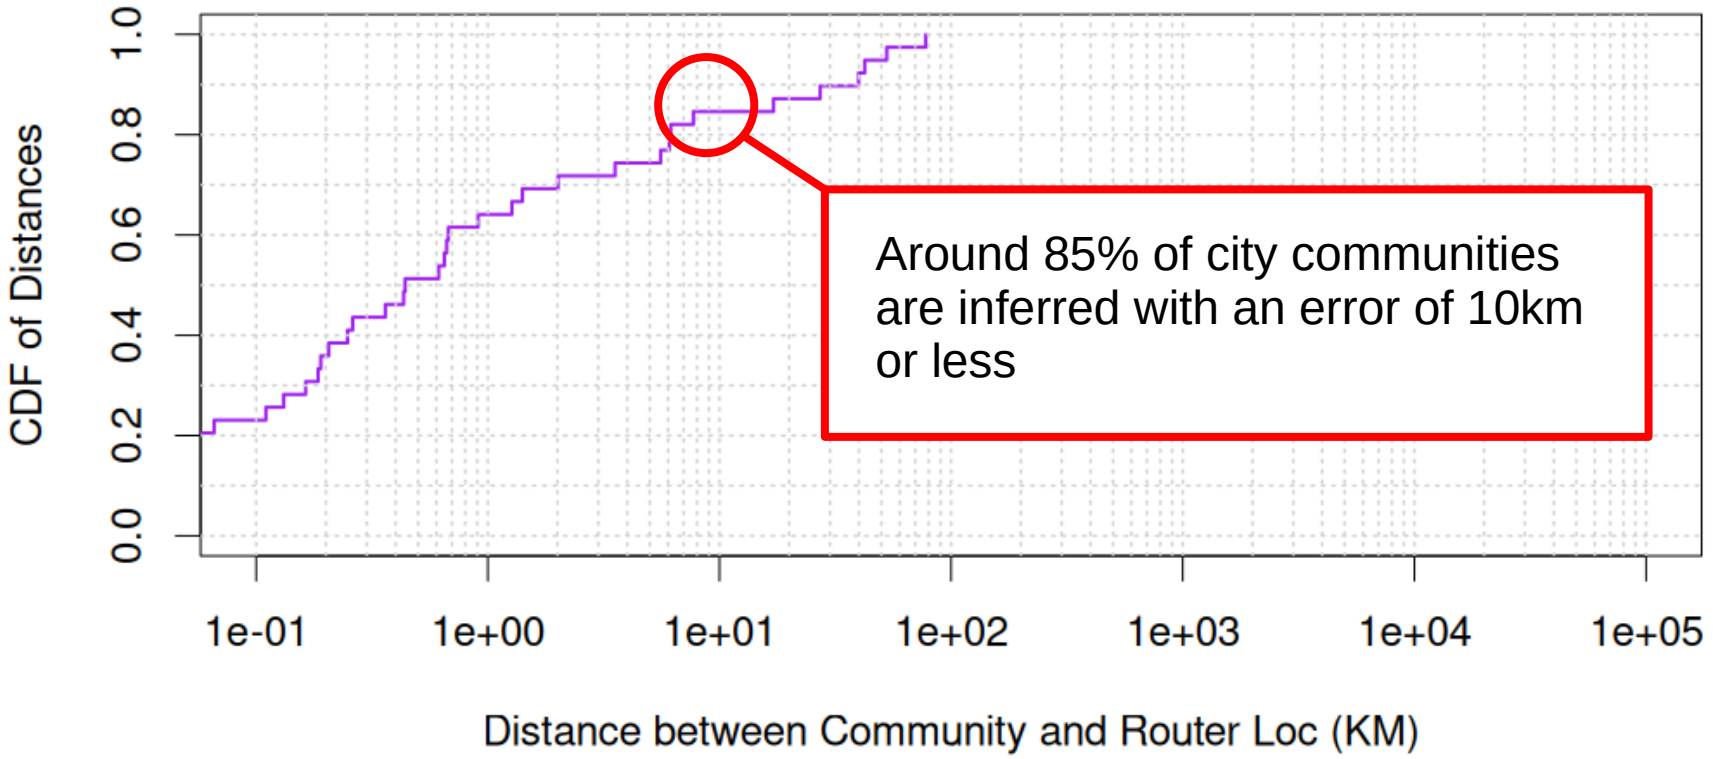

Chapter V:
Conclusions

Conclusions

- BGP + Communities

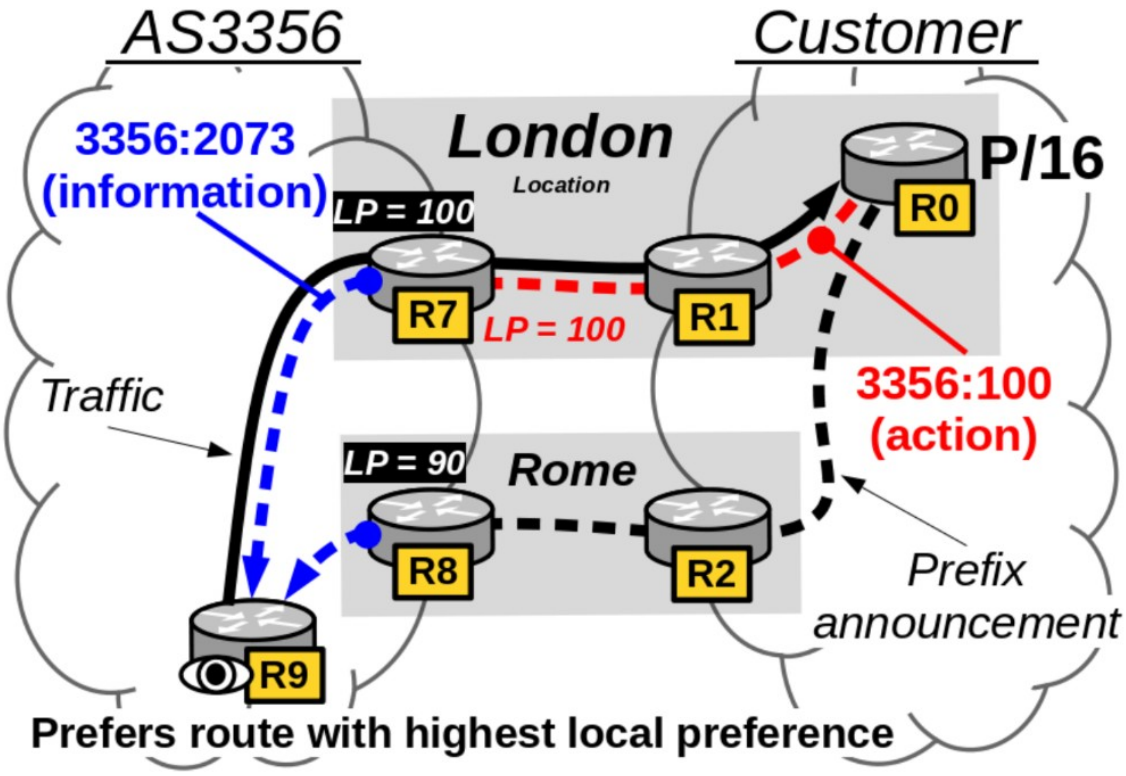

Conclusions

- BGP + Communities
- Inferring city communities:
  - Two approaches that show similar performance

A) Traceroute + router geolocation

1) Perform traceroute
2) Identify tagging router
3) Geolocate tagging router
4) Assign coordinates to community

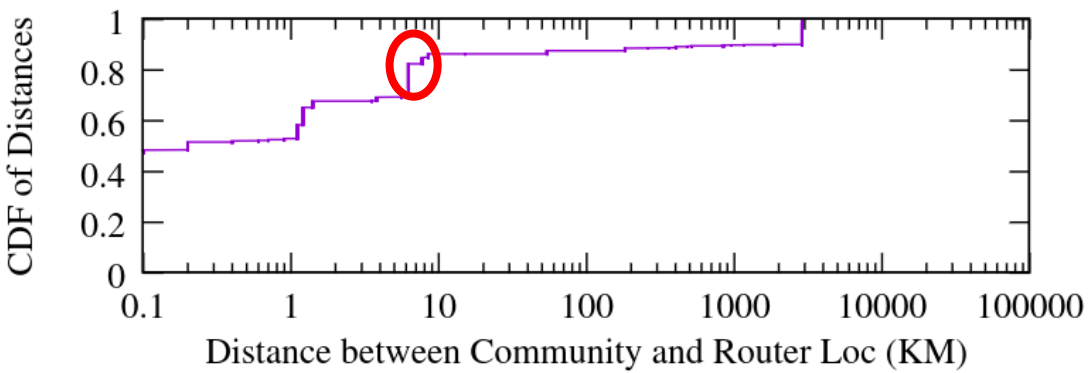

B) BGP + prefix geolocation

1) Obtain tagged prefixes
2) Maximize share of local prefixes
3) Cluster geographic locations
4) Assign densest cluster to community

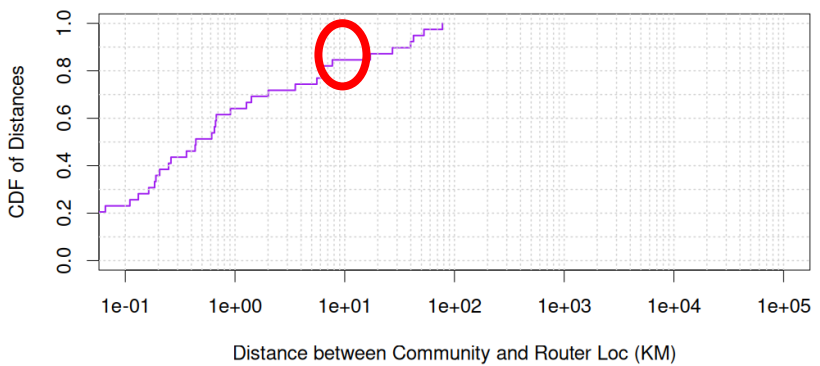

Conclusions

- BGP + Communities
- Inferring city communities:
  - Two approaches that show similar performance
  - Overall using ~1,500 city comms: 80% with error <70km

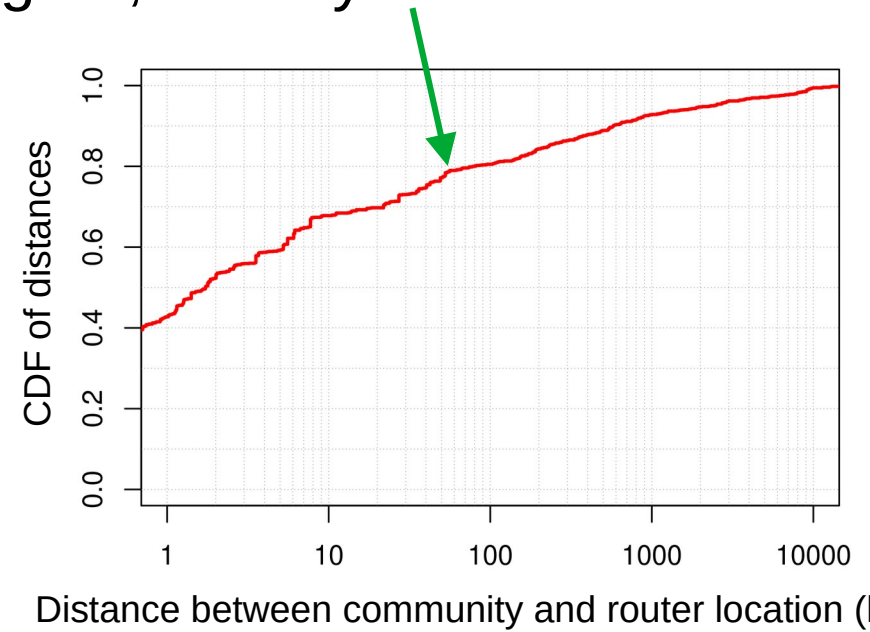

Conclusions

- BGP + Communities
- Inferring city communities:
  - Two approaches that show similar performance
  - Overall using ~1,500 city comms: 80% with error <70km
  - Outliers can help understanding network configuration