

A survey of interdomain routing policies

Phillipa Gill
Citizen Lab/
Stony Brook University

Sharon Goldberg
Boston University

Michael Schapira
Hebrew University

Why research the Internet?

Study existing protocols:

- Why is IPv6 so difficult to deploy?
 - Performance of conversion? Lack of global connectivity?
- What security benefits will we get from BGPsec?
 - How many attacks will it prevent?

Design new protocols:

- **New** interdomain routing protocols
 - Need to understand performance/effectiveness



Help understand implications of policy:

- Using BGP to cut off Internet access (e.g., Egypt)

Why **model** the Internet?

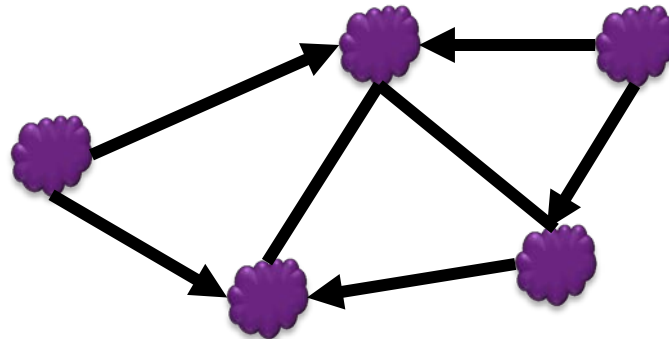
We can't always run experiments on the Internet!

For example:

- Studies of hijacking and failures
 - Cannot disrupt the Internet for the sake of research!
- Studies of unproven protocols
 - Cannot deploy a half-baked proposal and “hope it works”
 - Even deploying so-called fully-baked proposals is a daunting task!

Need to simulate behavior of the Internet to study protocols

- Models fill in gaps in empirical data to allow simulation
 - AS Topology
 - Routing policies
 - Traffic matrices



Why **model** the Internet?

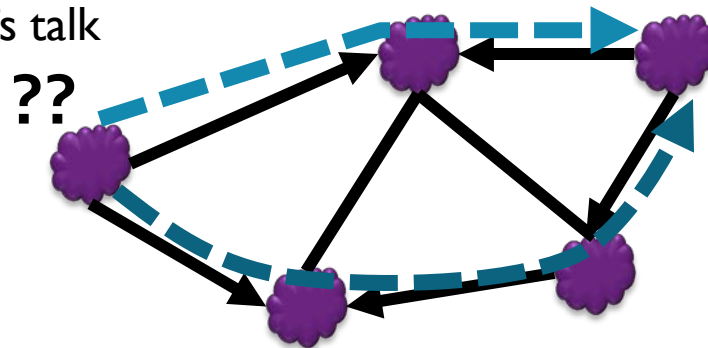
We can't always run experiments on the Internet!

For example:

- Studies of hijacking and failures
 - Cannot disrupt the Internet for the sake of research!
- Studies of unproven protocols
 - Cannot deploy a half-baked proposal and “hope it works”
 - Even deploying so-called fully-baked proposals is a daunting task!

Need to simulate **behavior** of the Internet to study protocols

- Models fill in gaps in empirical data to allow simulation
 - AS Topology
 - Routing policies ← Today's talk
 - Traffic matrices



Why **model** the Internet?


We can't always run experiments on the Internet!

For example:

- Studies of hijacking and failures
 - Cannot disrupt the Internet for the sake of research!
- Studies of unproven protocols
 - Cannot deploy a half-baked proposal and “hope it works”
 - Even deploying so-called fully-baked proposals is a daunting task!

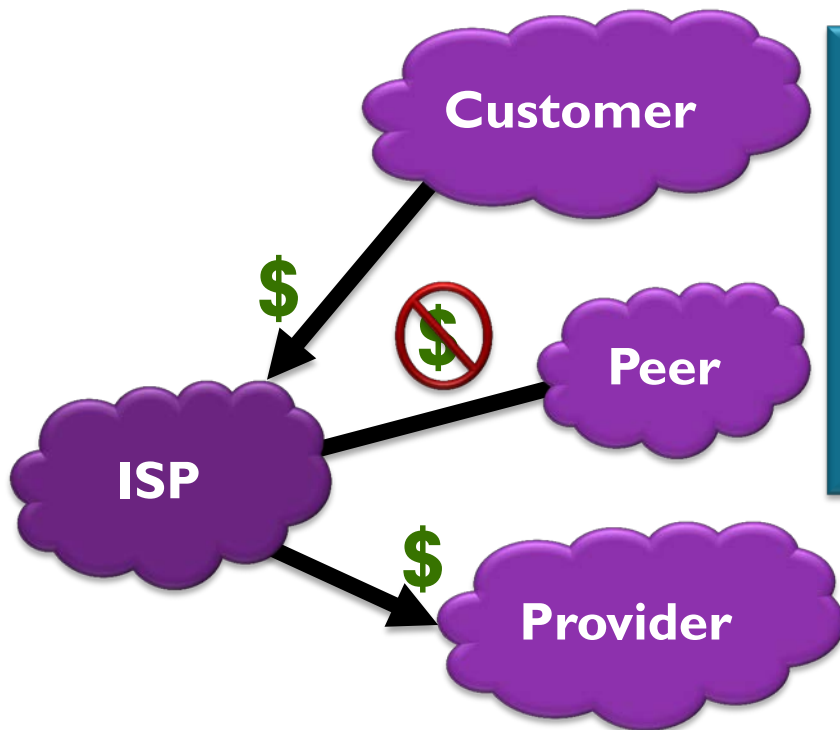
Need to simulate behavior of the Internet to study protocols

- Models fill in gaps in empirical data to allow simulation
 - AS Topology
 - Routing policies
 - Traffic matrices



Standard model of Internet routing

- Proposed by Gao & Rexford 12 years ago
- Based on practices employed by a large ISP
- Provide an intuitive model of path selection and export policy

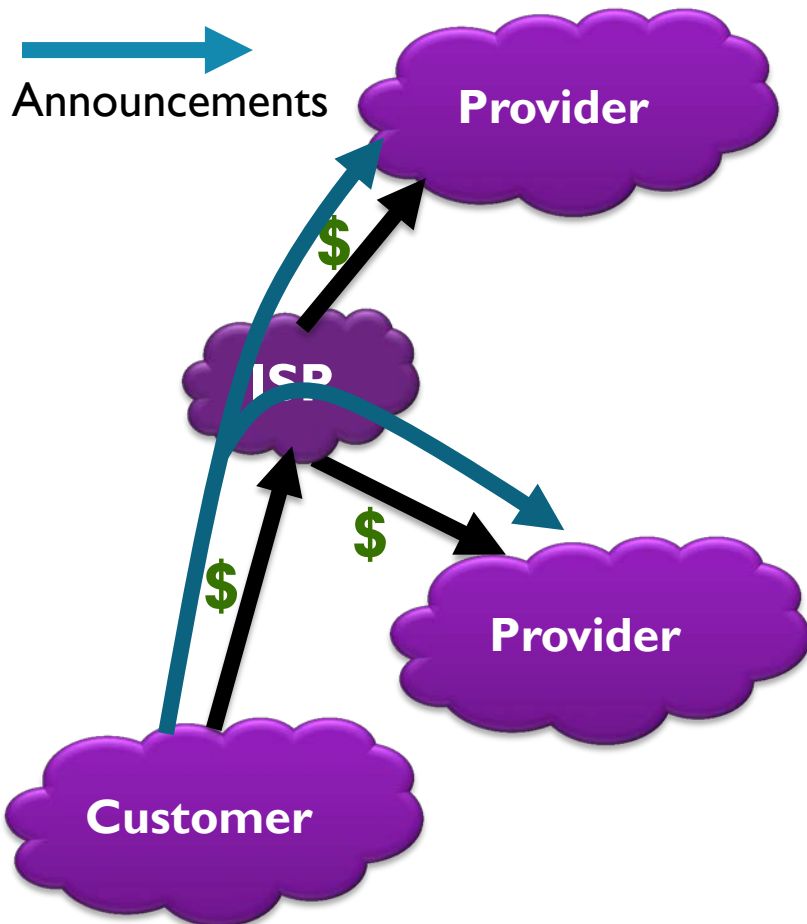


Path Selection:

1. LocalPref: Prefer customer paths over peer paths over provider paths
2. Prefer shorter paths
3. Arbitrary tiebreak

Standard model of Internet routing

- Proposed by Gao & Rexford 12 years ago
- Based on practices employed by a large ISP
- Provide an intuitive model of path selection and export policy



Path Selection:

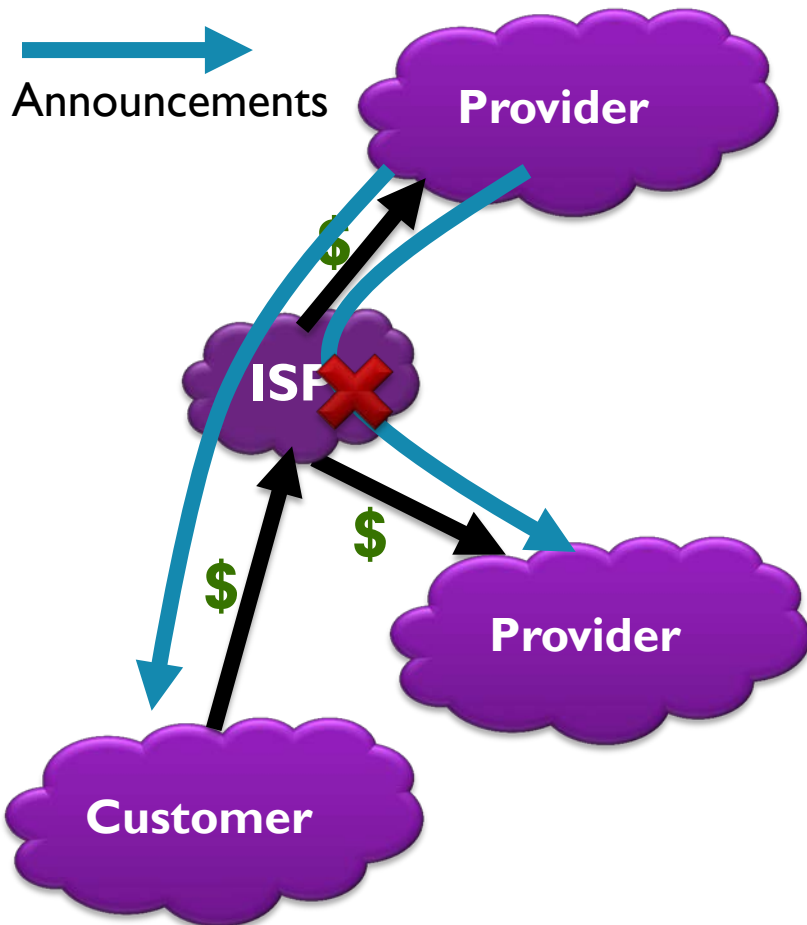
1. LocalPref: Prefer customer paths over peer paths over provider paths
2. Prefer shorter paths
3. Arbitrary tiebreak

Export Policy:

1. Export customer path to all neighbors.
2. Export peer/provider path to all customers.

Standard model of Internet routing

- Proposed by Gao & Rexford 12 years ago
- Based on practices employed by a large ISP
- Provide an intuitive model of path selection and export policy



Path Selection:

1. LocalPref: Prefer customer paths over peer paths over provider paths
2. Prefer shorter paths
3. Arbitrary tiebreak

Export Policy:

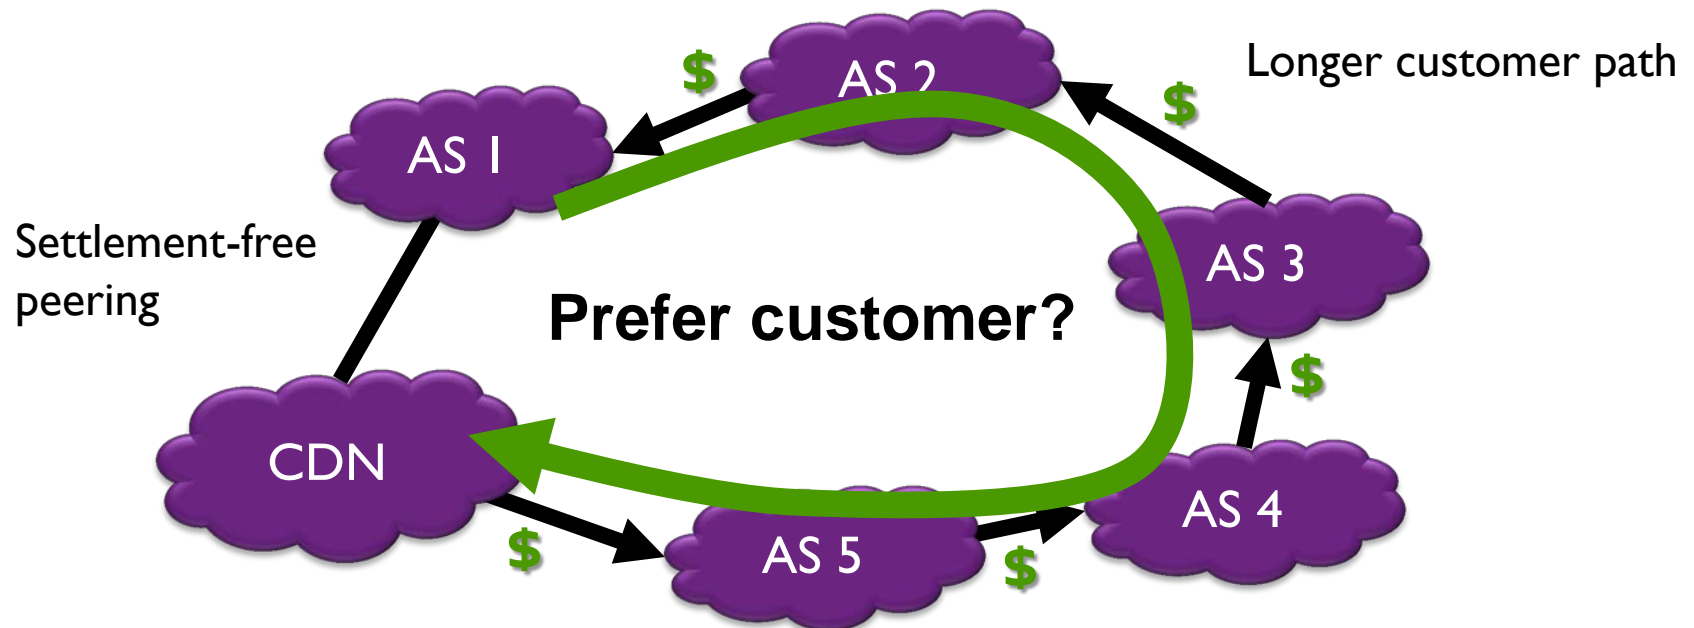
1. Export customer path to all neighbors.
2. Export peer/provider path to all customers.

Developing a new model of Internet routing

The existing model relies on routing policy assumptions...
... but how valid are these assumptions in practice?

Example: Prefer customer routes

...but what about when peer route is direct?




Developing a new model of Internet routing

The existing model relies on routing policy assumptions...
... but how valid are these assumptions in practice?

Example: Prefer customer routes

...but what about when peer route is direct?



Key questions:

How often does the model hold?

What exceptions arise and how frequent are they?

How to understand Internet routing?

Challenges

- Policies can vary from network to network!
 - Tier I vs. Large Content Providers
- Understanding **exception vs. rule**

Survey network operators about their routing policies

- **Today:** Preliminary survey results
- ... still a long way to go!

How you can help?

- Fill out our survey! [<http://bit.ly/routingsurvey>]
- Come talk to me in the break!

The survey

Routing Policy Survey

This is an informal survey to better understand how routing policies look in practice at a variety of networks. Please answer as many questions as possible (leaving some blank is ok too).

When answering the survey questions, please refer to default configurations, and leave out corner cases (e.g., if a customer asks you, via communities or otherwise, to do something different.)

Please refer to the partial list of steps in the BGP path selection algorithm to your right for some of the questions.

E-mail routingsurvey@cs.toronto.edu with questions or concerns.

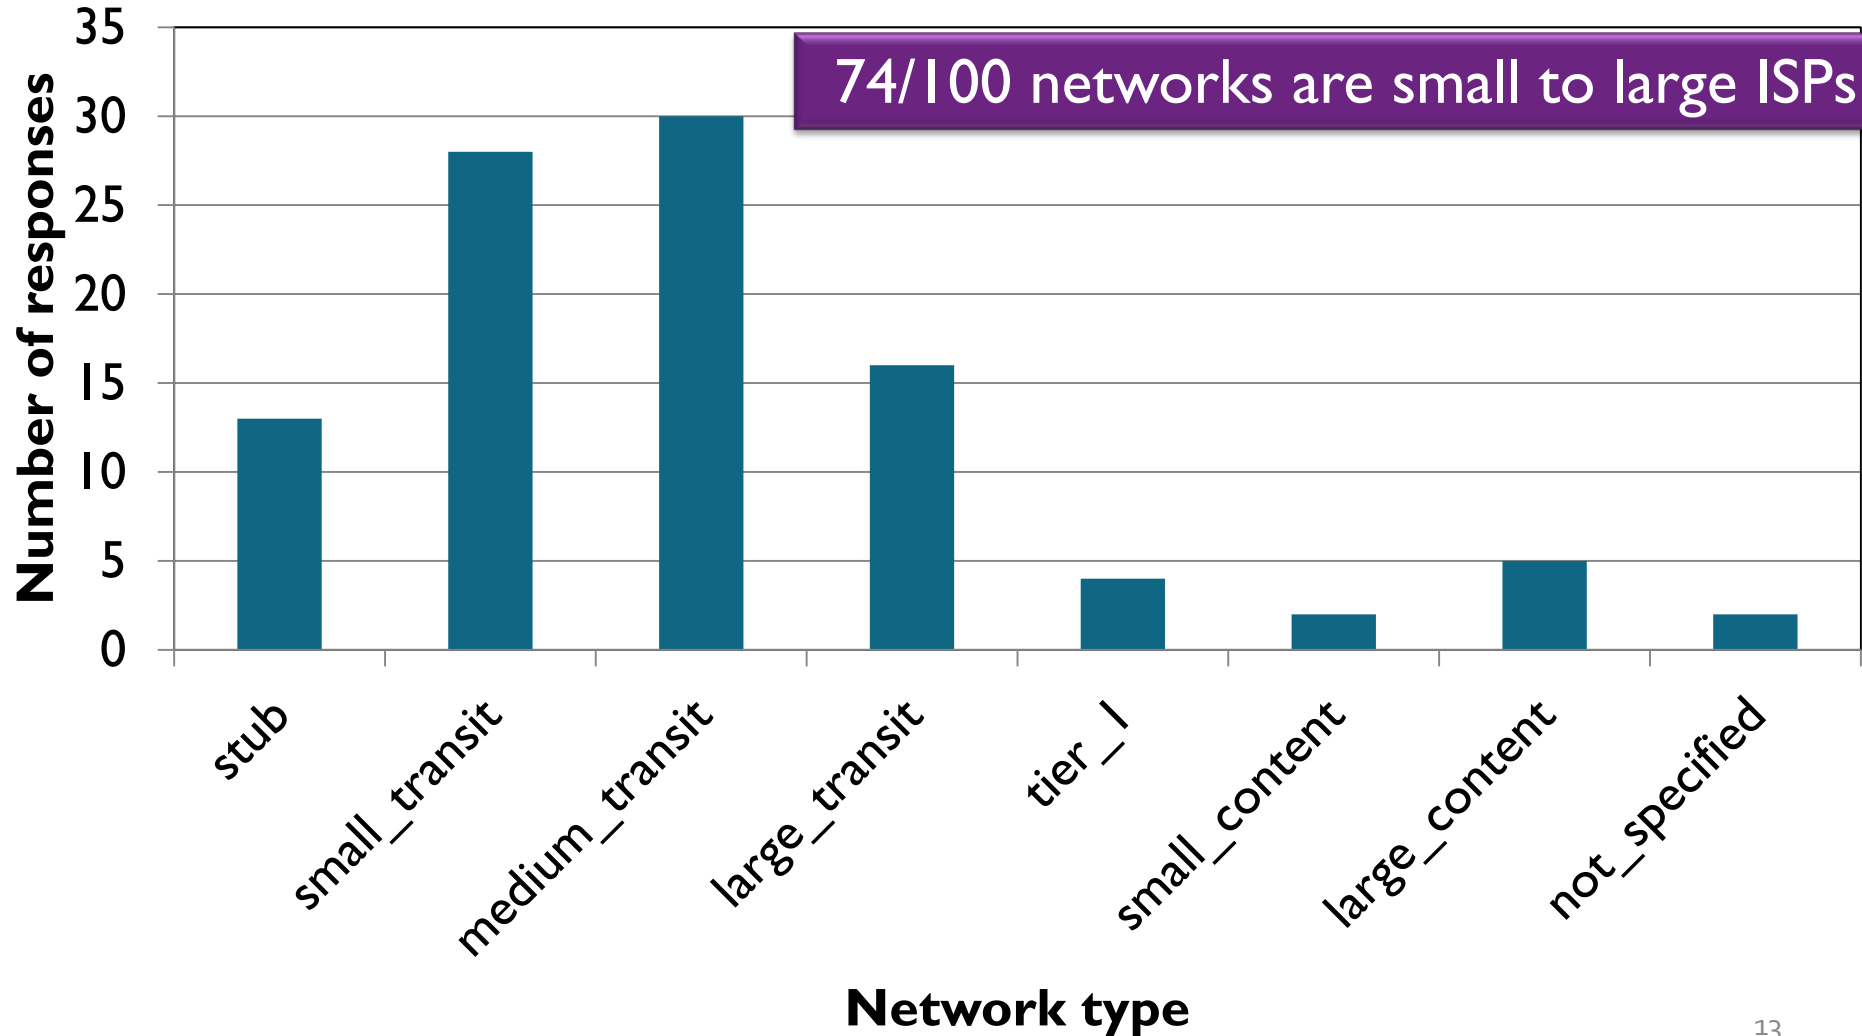
BGP Path Selection Algorithm [\[1,2\]](#)

1. Highest LocalPref
2. Lowest AS path length
3. Lowest origin type
4. Lowest MED
5. eBGP-learned over iBGP-learned
6. Lowest IGP cost to border router (hot-potato routing)
7. If both paths are external, prefer the path that was received first (i.e., the oldest path) [\[1\]](#)
8. Lowest router ID (to break ties)

1. What kind of network do you operate? Other:
2. On what continent is your network?
3. Do you always assign a higher LocalPref (see Step 1 in the table) to a path through your customer than to a path through your peer or transit provider? (Note: exclude cases where routes through customers are tagged as backup.) Why? (optional)
4. Does your LocalPref configuration depend only on the next-hop AS (and not on other ASes on the path)? Why? (optional)
5. Do you use the same LocalPref configuration across all BGP-speaking routers in your network? Why? (optional)
6. Is the "prefer oldest path" step (see Step 7 in the table) enabled on your BGP-speaking routers? (Note: this step is enabled by default on Cisco routers in the last few years.) Why? (optional)
7. If path validation (eg BGPsec) was deployed in your network, **before** what step (1-8) in the table would you place the following step: "Prefer secure paths (validated paths) over insecure paths"? Select a number from 1-8. Why? (optional)

The survey

- Initial survey circulated on NANOG mailing list + others
- Breakdown of responses: 100 responses in total



Preliminary results

Configuring LocalPref:

A1: Assign higher LocalPref to a path through customer

A2: LocalPref only depends on next-hop AS

Export Policies:

A3: Do **not** export paths from non-customers to non-customers

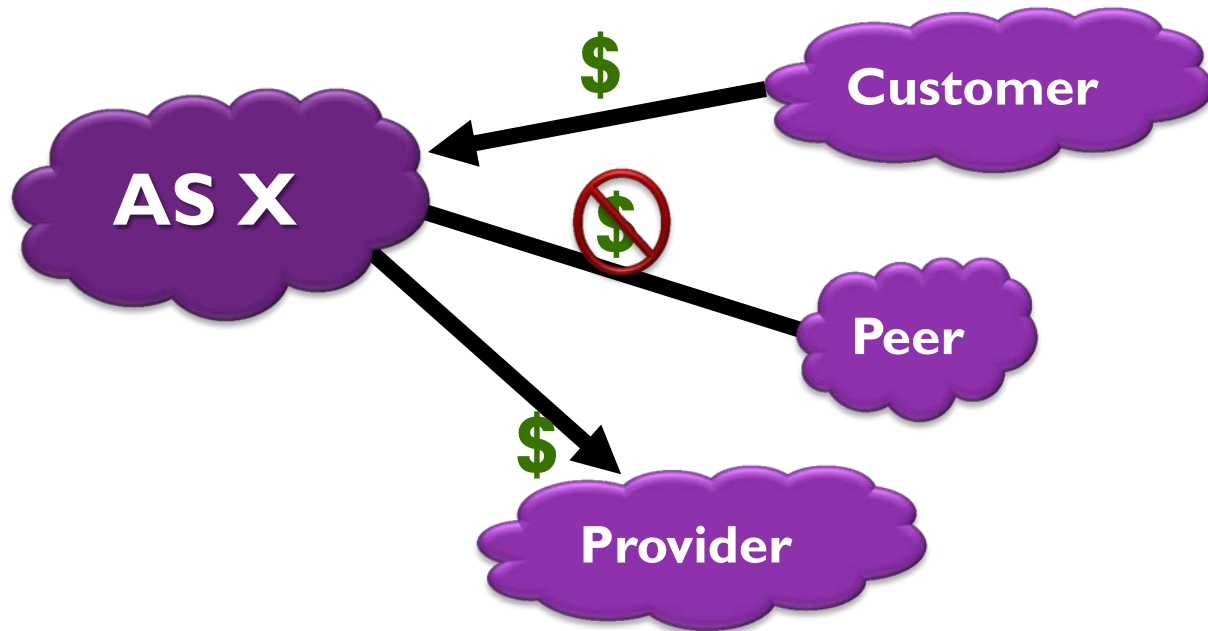
A4: Export the same path to neighbors of the same type

Other topics: (not today...)

MRAl, prefer oldest path, pricing models, BGP security

Configuring LocalPref

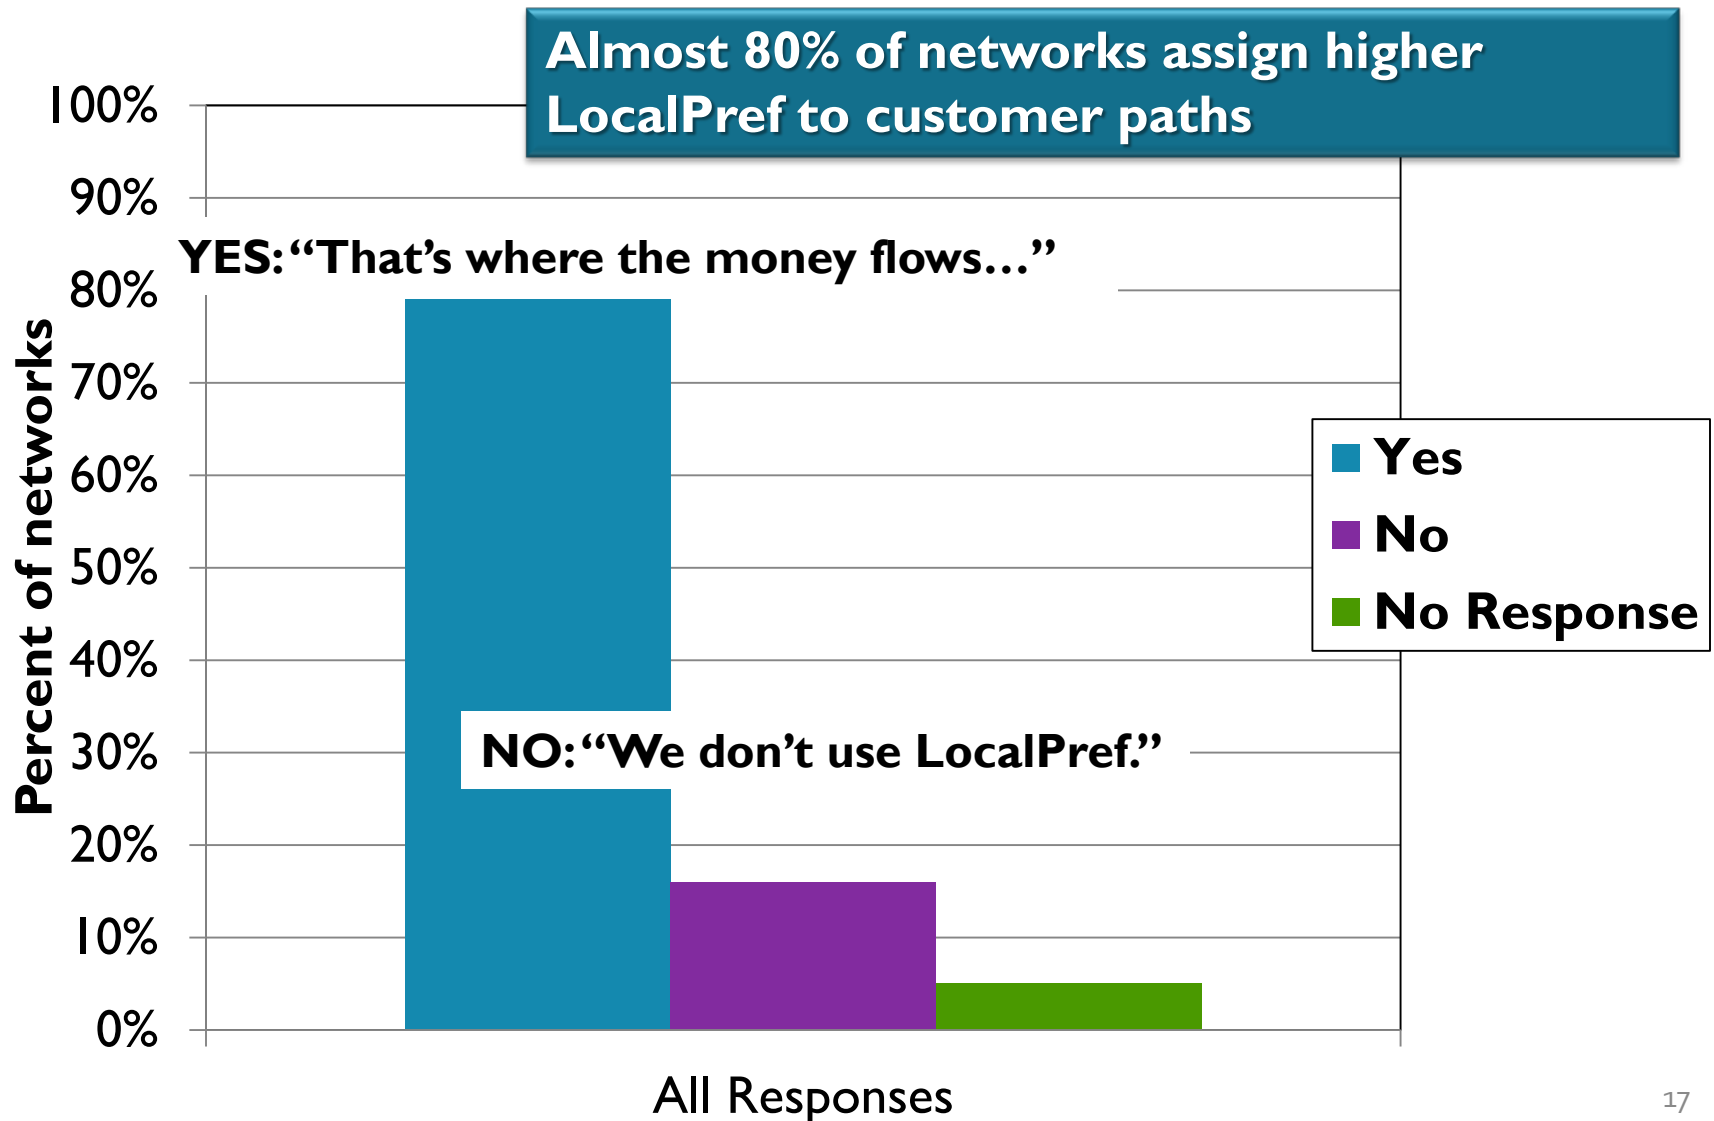
AI: Assign higher LocalPref to a path through customer (than to peer or provider)



A1: Assign higher LocalPref to a path through customer (than to peer or provider)

Do you always assign a higher LocalPref to a path through your customer than to a path through your peer or transit provider?

A1: Assign higher LocalPref to a path through customer (than to peer or provider)



Preliminary results

Configuring LocalPref:

 **A1:** Assign higher LocalPref to a path through customer

A2: LocalPref only depends on next-hop AS

Export Policies:

A3: Do **not** export paths from non-customers to non-customers

A4: Export the same path to neighbors of the same type

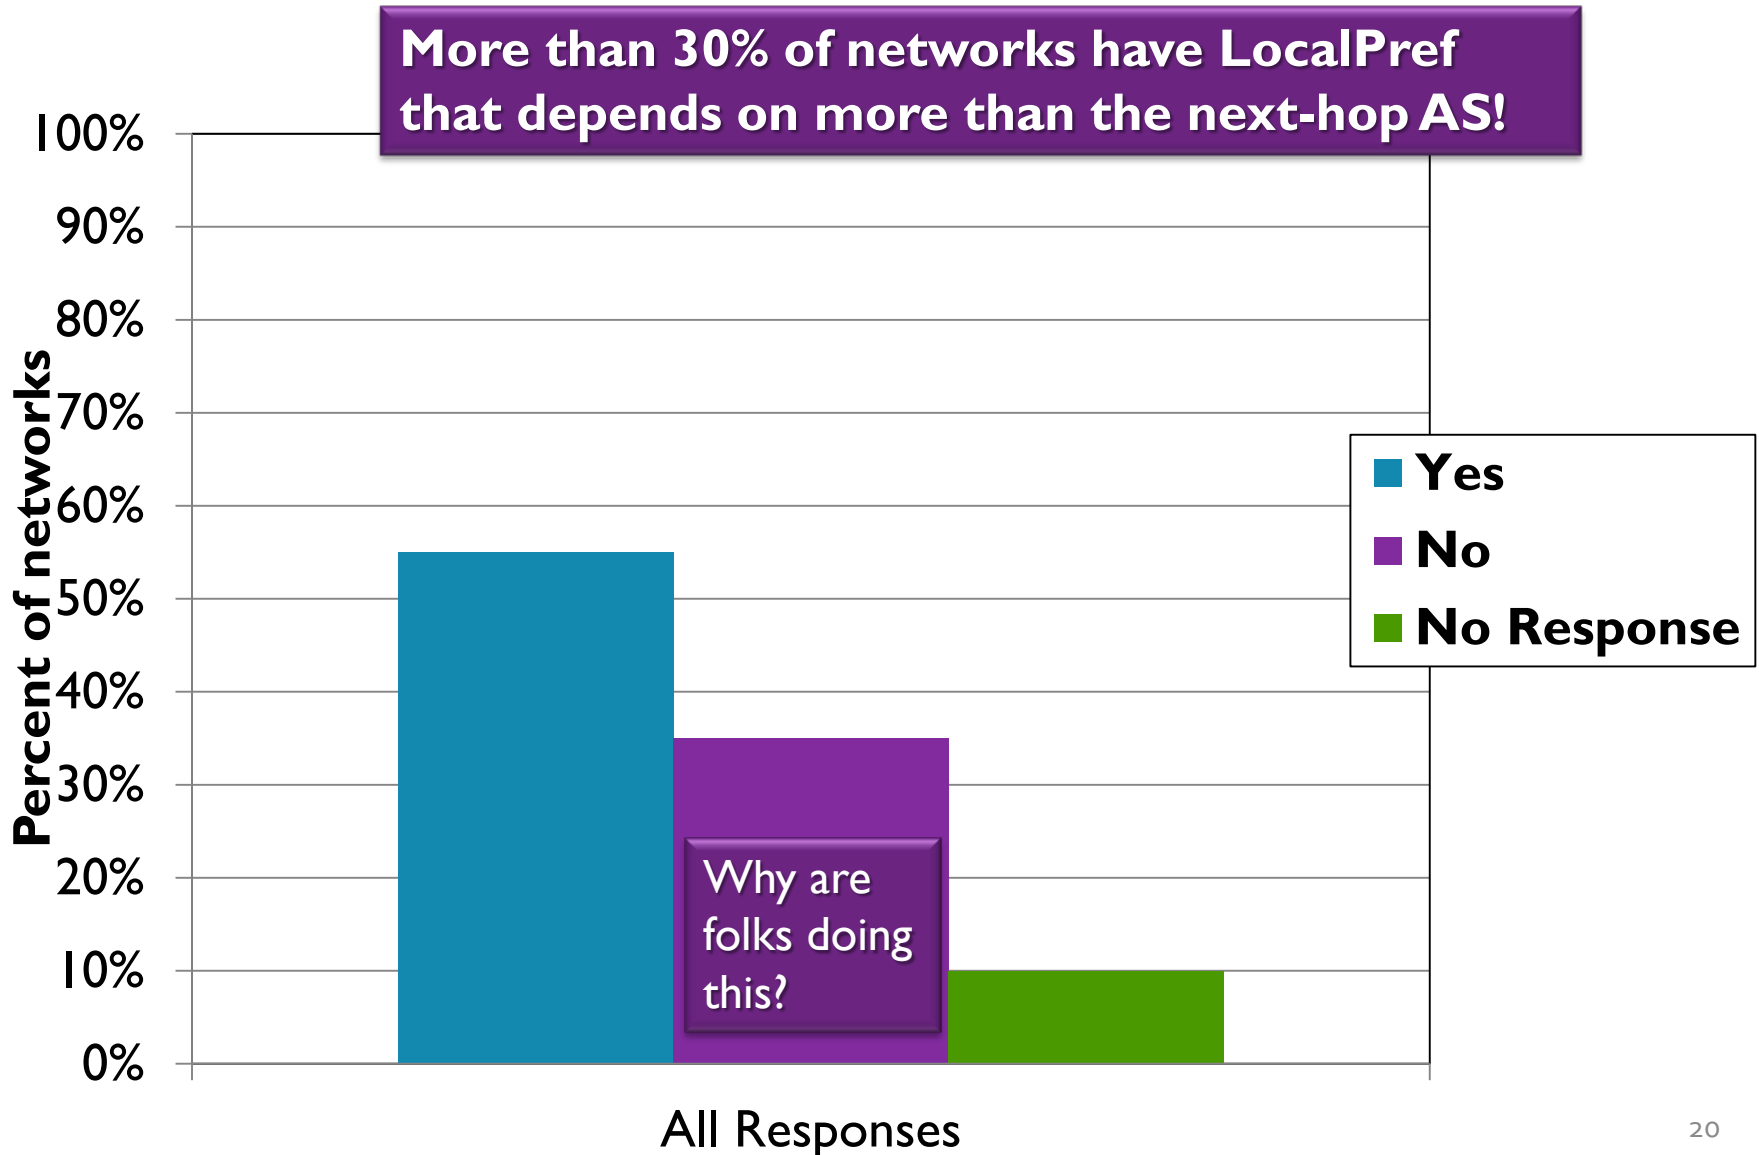
Other topics: (not today...)

MRAl, prefer oldest path, pricing models, BGP security

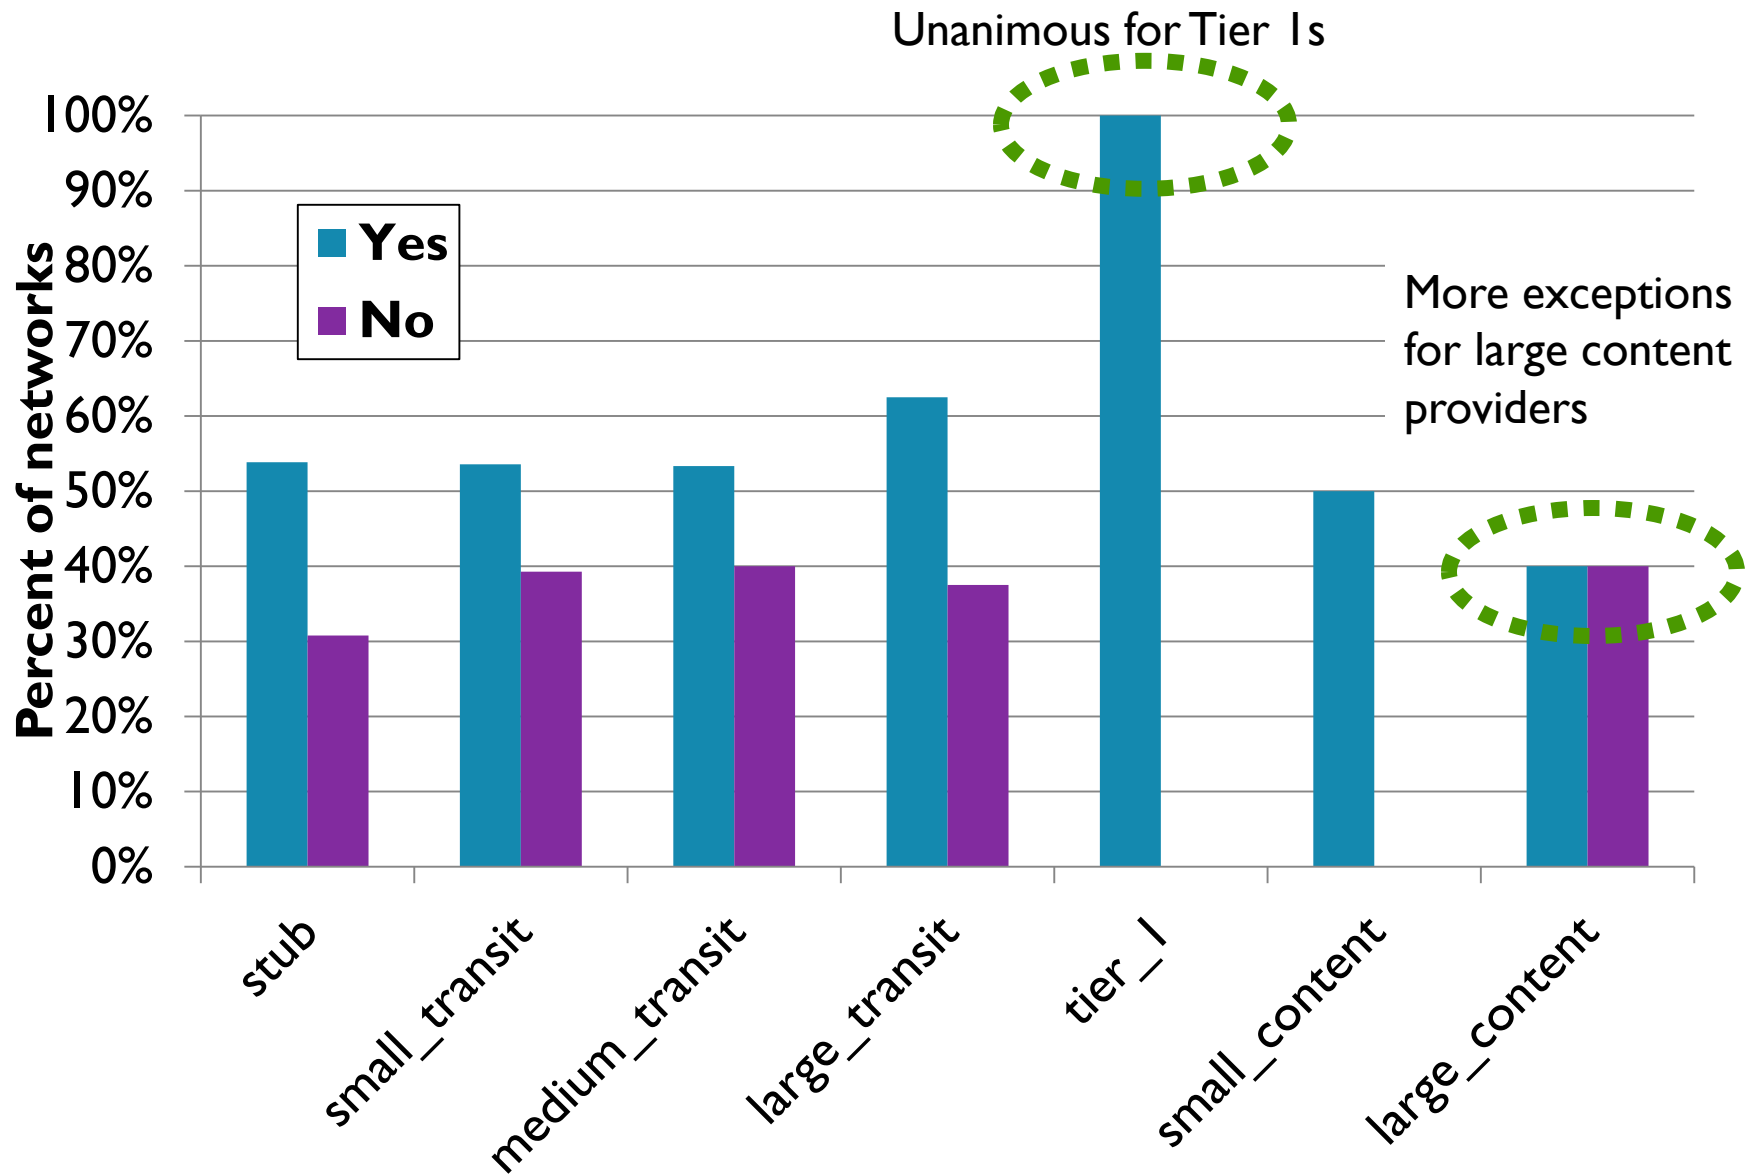
A2: LocalPref only depends on next-hop AS

Does your LocalPref configuration depend only on the next-hop AS (and not on other ASes on the path)?

A2: LocalPref only depends on next-hop AS



A2: LocalPref only depends on next-hop AS



Preliminary results

Configuring LocalPref:

 **A1:** Assign higher LocalPref to a path through customer

A2: LocalPref only depends on next-hop AS

...exceptions for large content providers

Export Policies:

A3: Do **not** export paths from non-customers to non-customers

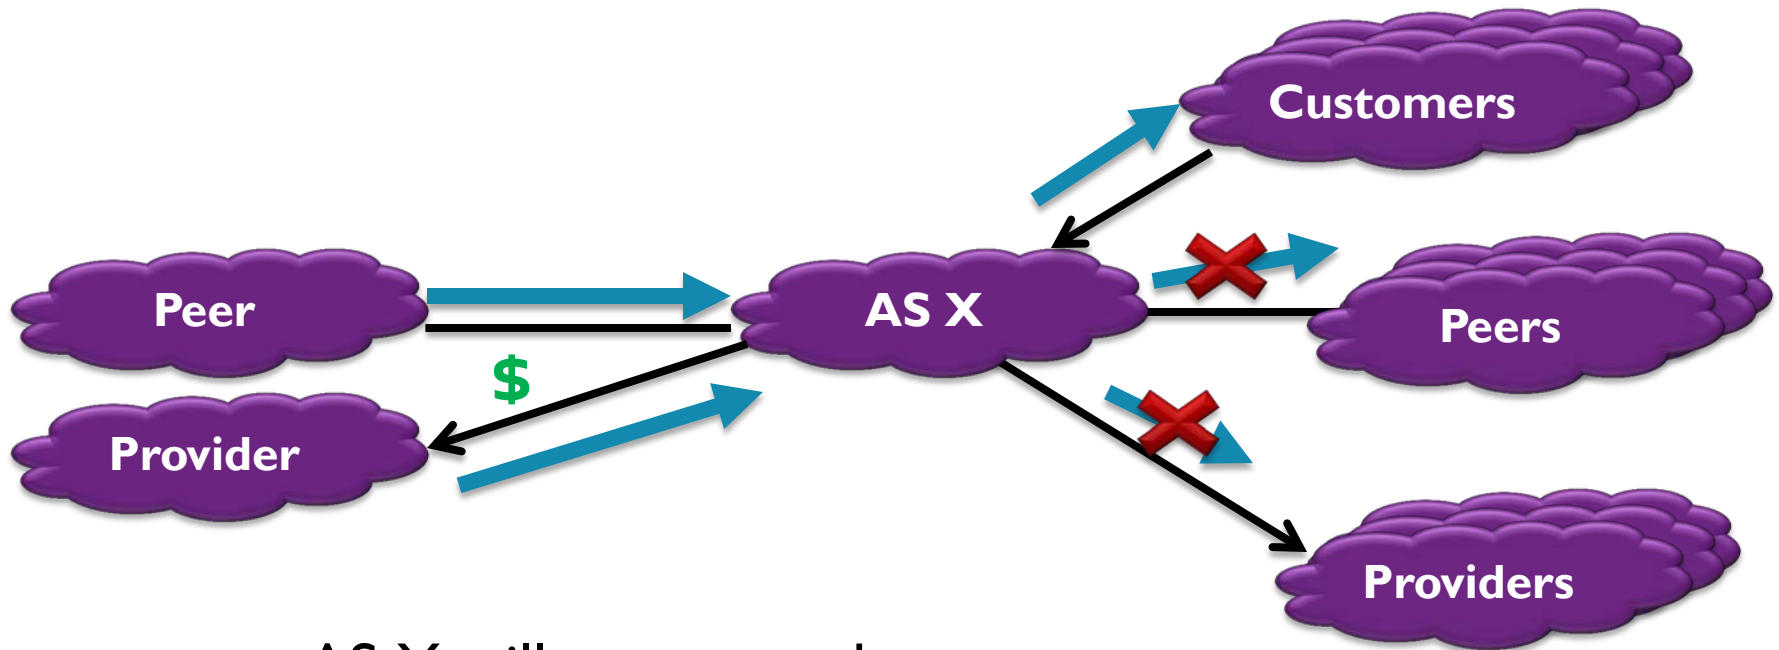
A4: Export the same path to neighbors of the same type

Other topics: (not today...)

MRAl, prefer oldest path, pricing models, BGP security

Export Policies

A3: Do **not** export paths from non-customers to non-customers



AS X will announce the non-customer routes only to his customers!

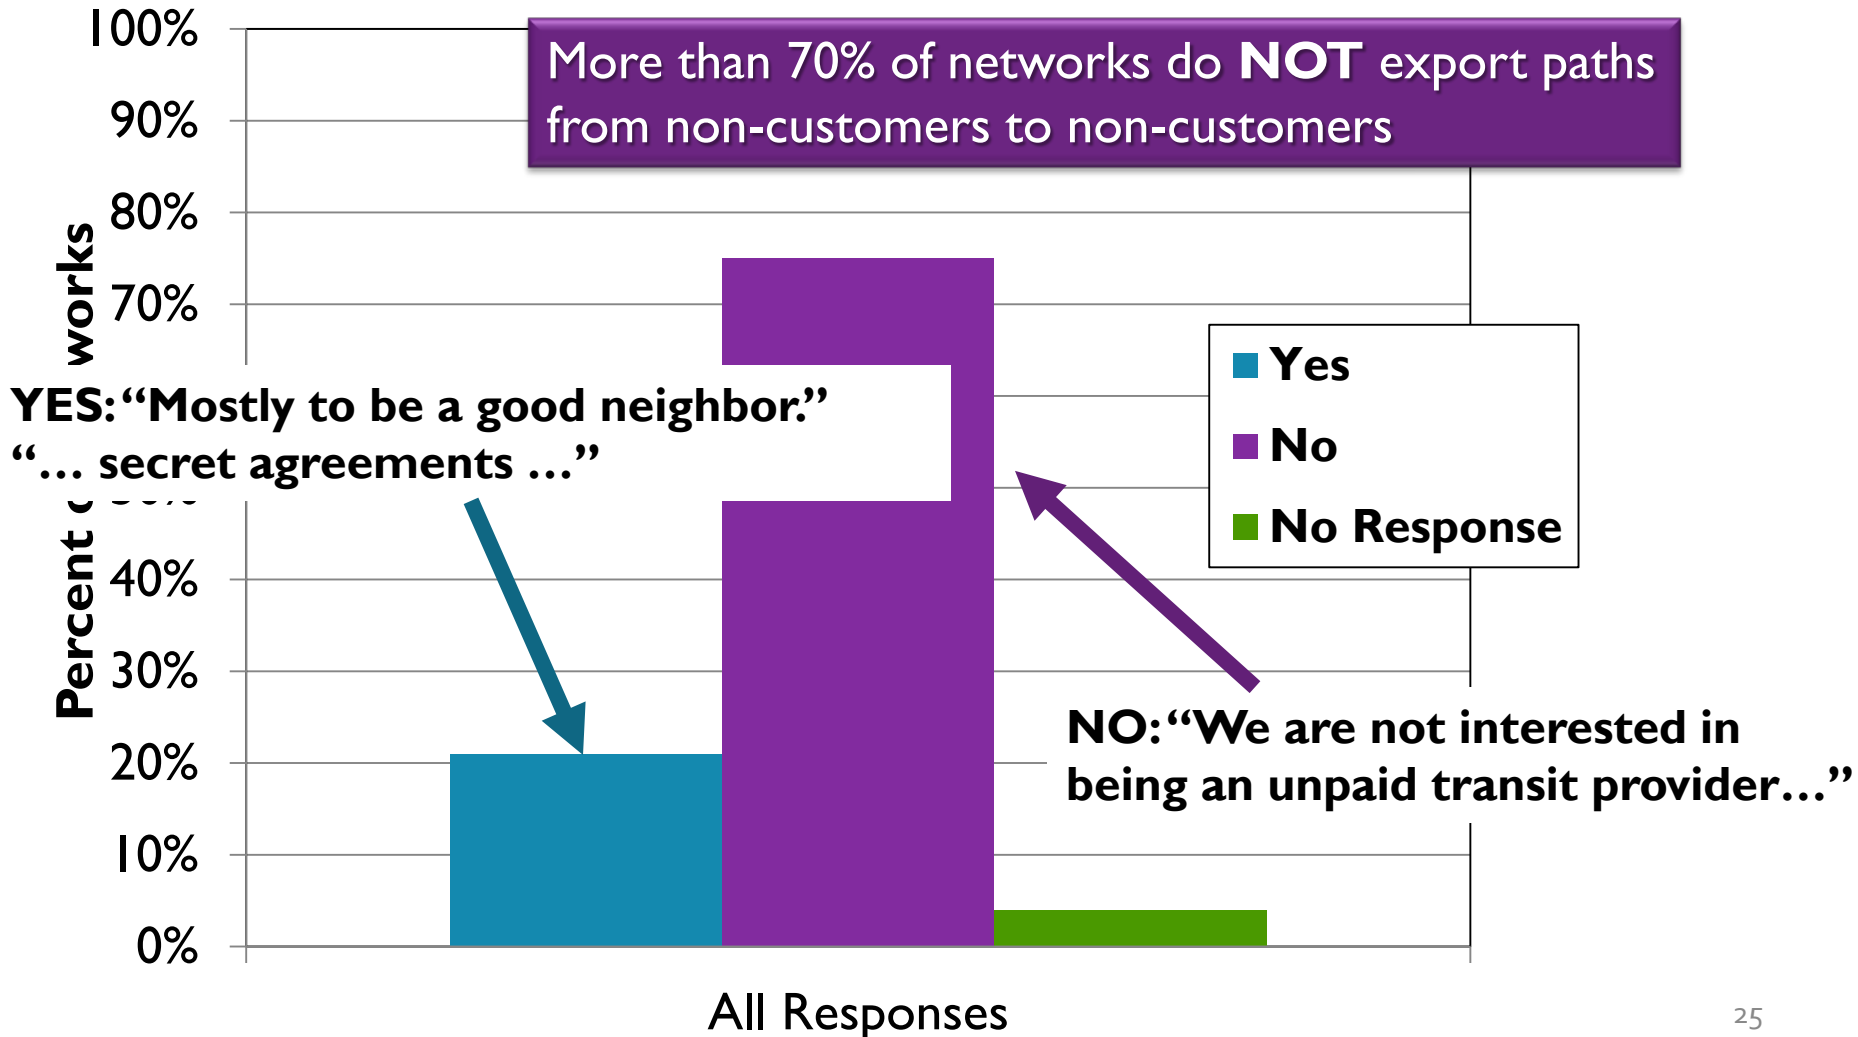

Announcements

A3: Do not export paths from non-customers to non-customers

Do you announce paths from peers and providers to other peers and providers?

A3: Do **not** export paths from non-customers to non-customers

Q: Do you announce paths through peers/providers to **other** peers/providers?



Preliminary results

Configuring LocalPref:

 **A1:** Assign higher LocalPref to a path through customer

A2: LocalPref only depends on next-hop AS

...exceptions for large content providers

Export Policies:

 **A3:** Do **not** export paths from non-customers to non-customers

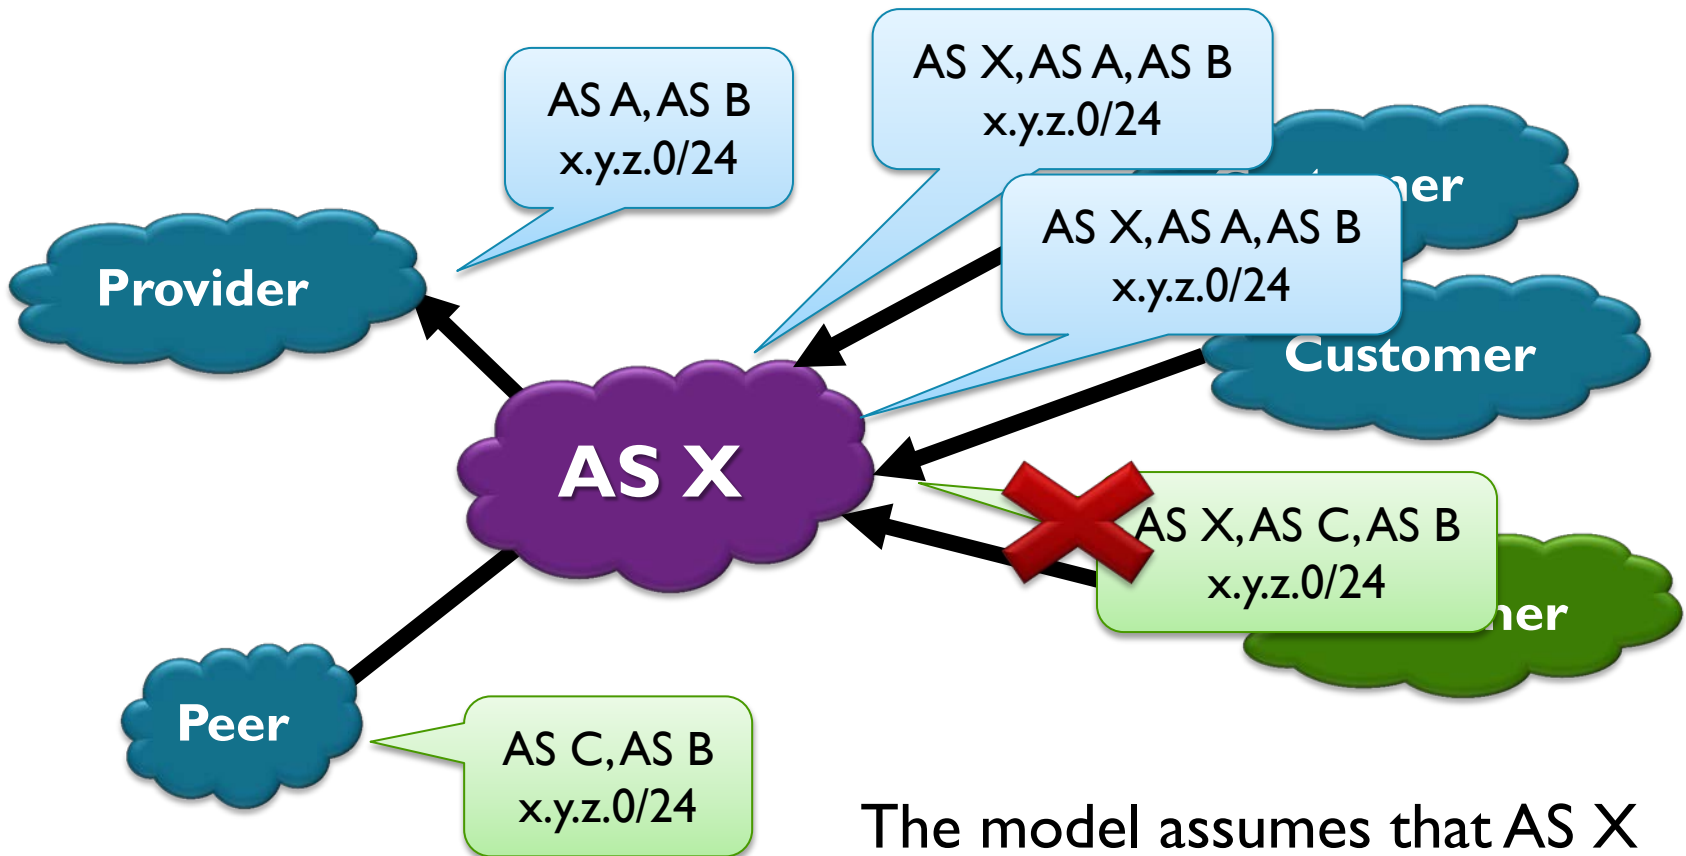
A4: Export the same path to neighbors of the same type

Other topics: (not today...)

MRAl, prefer oldest path, pricing models, BGP security

Export Policies (2)

A4: Export the same path to neighbors of the same type



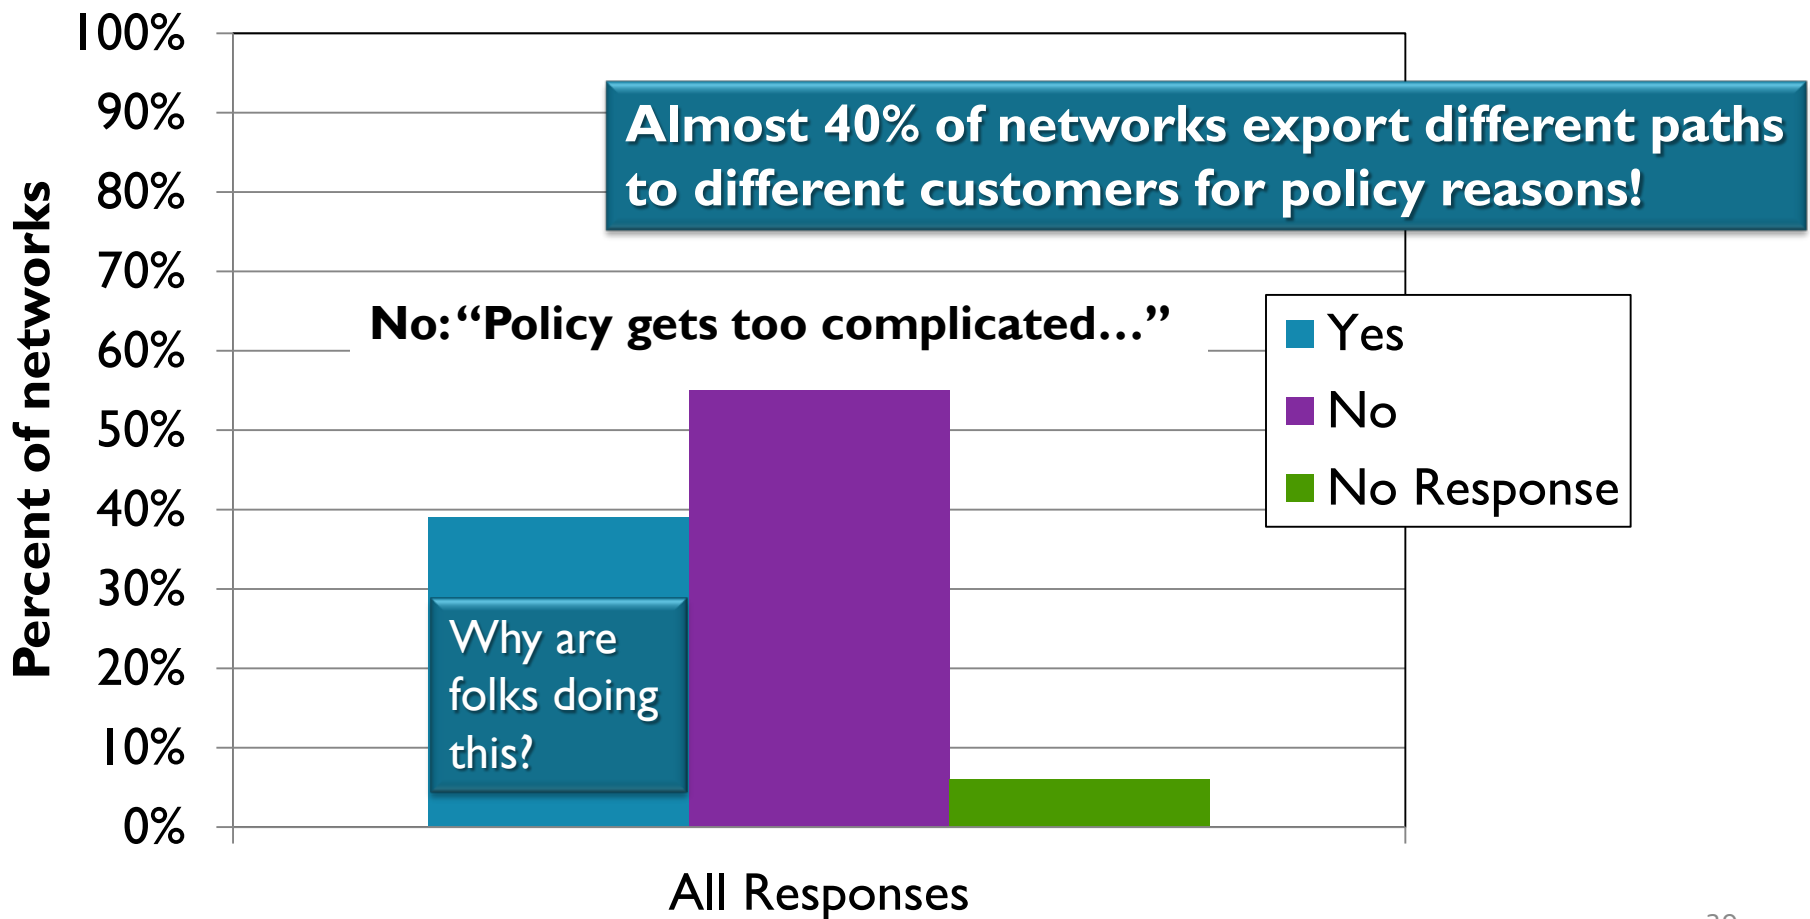
The model assumes that AS X will announce the same path to the destination to all customers!

A4: Export the same path to neighbors of the same type

Do you do neighbor-specific path selection, *e.g.*, select a different path for different customers for **policy reasons (and **not due to hot-potato routing** *etc.*)**

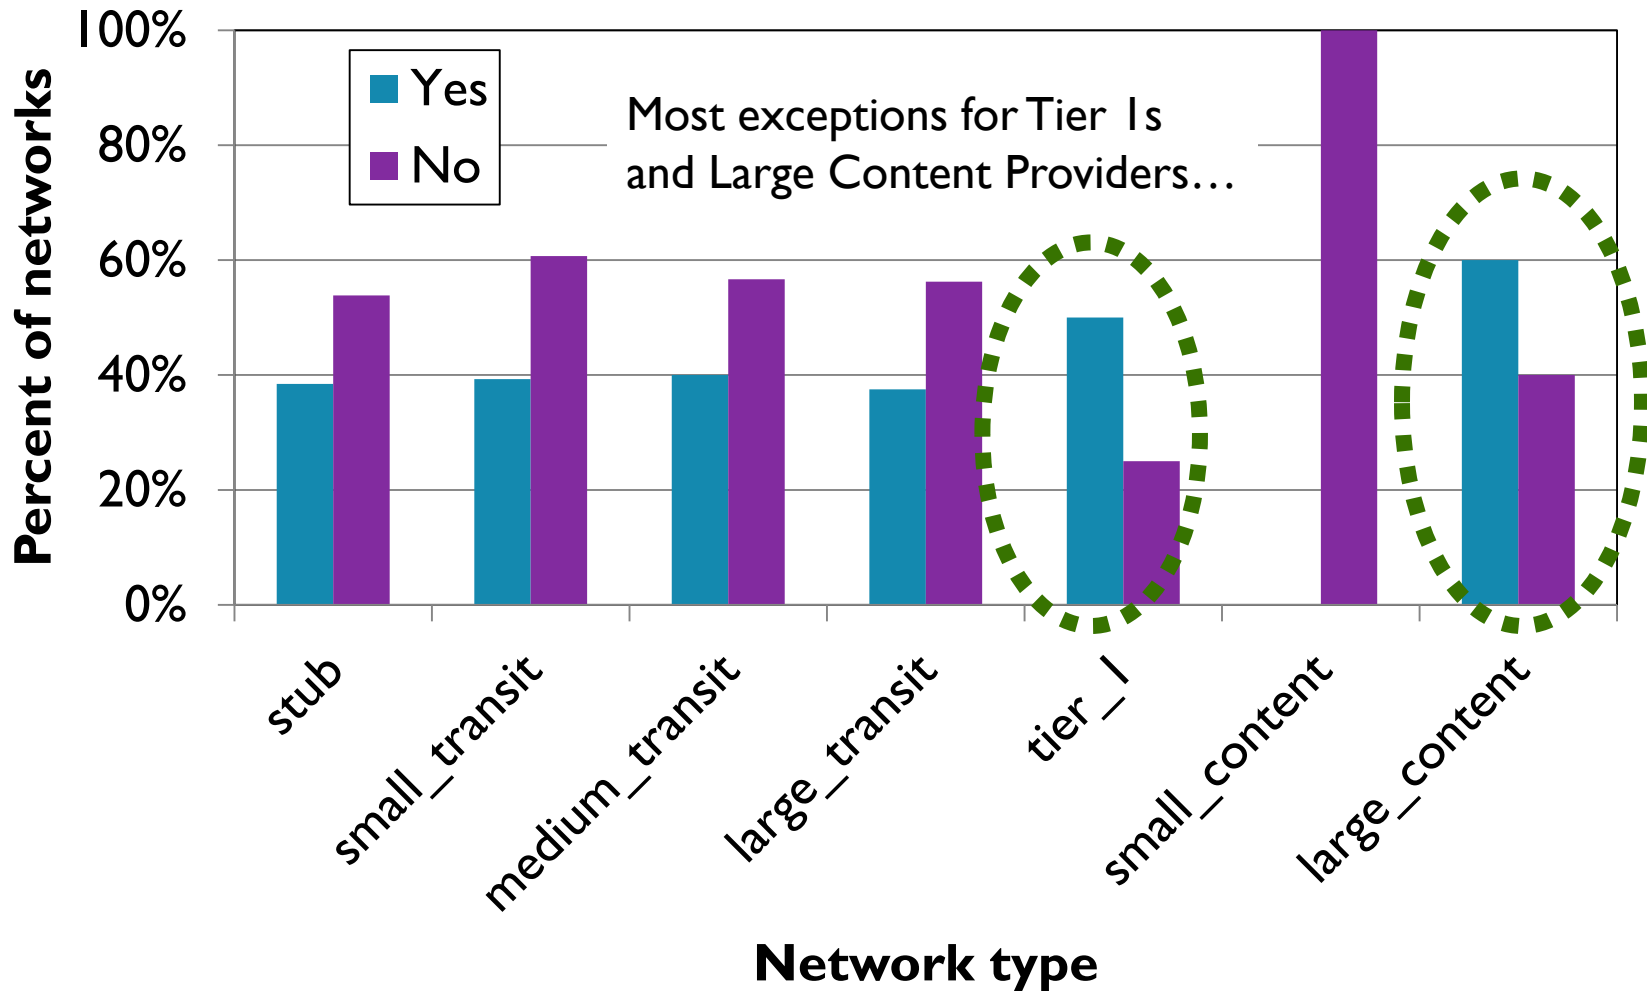
A4: Export the same path to neighbors of the same type

Q: Do you do neighbor-specific path selection, e.g., select a different path for different customers for **policy reasons** (and **not due to hot-potato routing** etc.)



A4: Export the same path to neighbors of the same type

Q: Do you do neighbor-specific path selection, e.g., select a different path for different customers for **policy reasons** (and **not due to hot-potato routing etc.**)



Preliminary results

Configuring LocalPref:

 **A1:** Assign higher LocalPref to a path through customer

A2: LocalPref only depends on next-hop AS

...exceptions for large content providers

Export Policies:

 **A3:** Do **not** export paths from non-customers to non-customers

A4: Export the same path to neighbors of the same type

...exceptions for tier 1s and large content providers

In all cases exceptions exist!

Going forward...

Need to better understand corner cases:

How often do these things happen?

Why do they happen?

When do they happen?

Who is doing them?

Come tell us about your experiences!

<http://bit.ly/routingsurvey>

What questions would you like answered about routing policies?

Contact us:

phillipa@cs.toronto.edu, goldbe@cs.bu.edu, schapiram@huji.ac.il

Fin.