NIRA: A New Inter-Domain Routing Architecture

Xiaowei Yang, David Clark, Arthur W. Berger

Rachit Agarwal

(Results are by others, any errors are by me)

(‘Animated’ slides shamelessly stolen from Prasad’s slides (CS495, Northwestern University), Thanks Google!)

What this paper talks about!

- Routing at domain level
  - Giving more control to the user over the route
    - Fosters competition among ISPs
    - Routes chosen by BGP not the most efficient
    - Only users know whether a path suits his/her application

What this paper talks about!

- Claims to answer the questions:
  - Supporting user choice
  - provider compensation
  - scalable route discovery
  - efficient route representation
  - fast route fail-over
  - security

What this paper does not talk about!

- Acknowledged Issues:
  - Autonomy Issues (why would an ISP allow that?)
  - Potential performance problems
- Issues not acknowledged:
  - Where is “design for tussle”? (stronger users means stronger attacks?)

NIRA

Core

- tier-I ISPs: ISPs that have no providers
- Core: Region where tier-I ISPs interconnect
- Up-graph (of an user): network of user's providers, provider's providers (and peers) until the core is reached

Example: Core

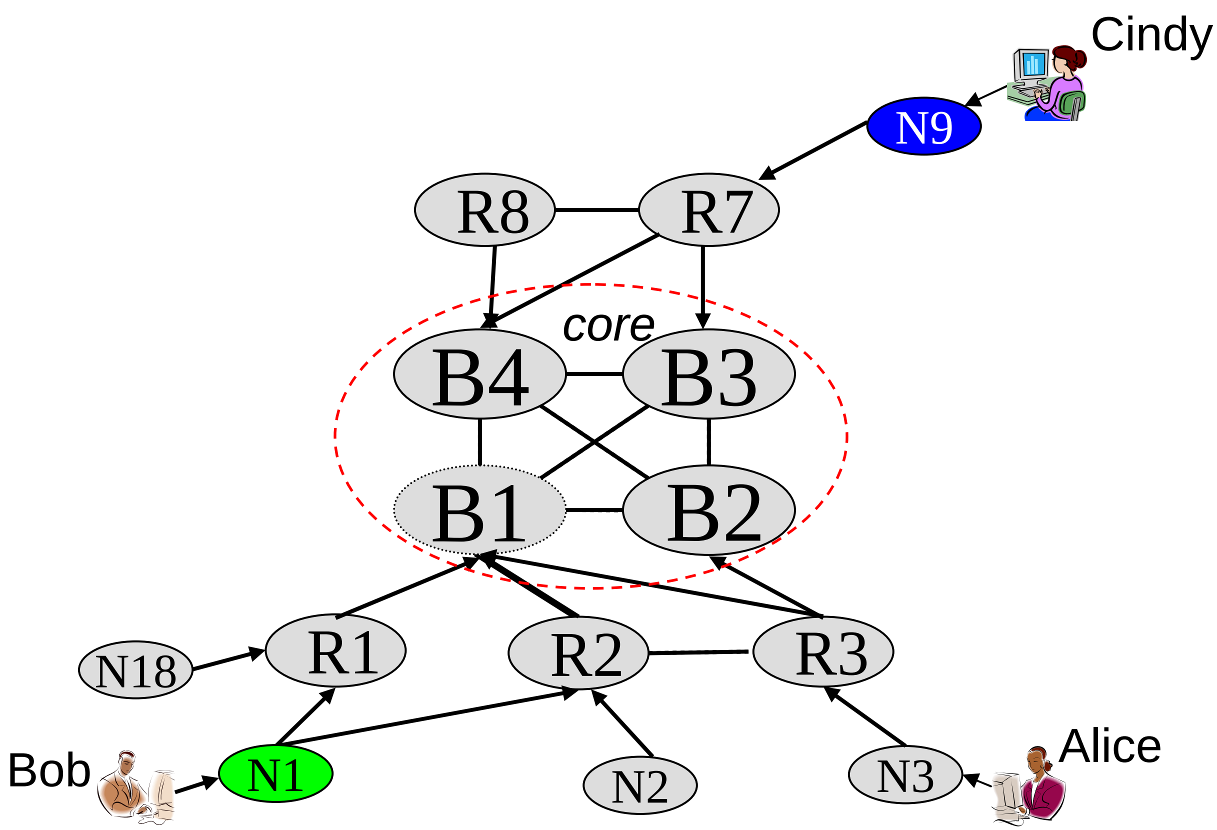

NIRA in a nutshell !

- Every node gets a path from its up-graph to the core
- All these paths get stored in a DNS-like database (NRLS)
- Path Selection:
  - Choose your up-graph as part of the route
  - Query name-to-route look-up service (NRLS) for destination’s up-graph
  - Combine the two to get a path to the destination
- User’s route not selected by the user, but by both user and destination!

Example: NIRA in a nutshell !

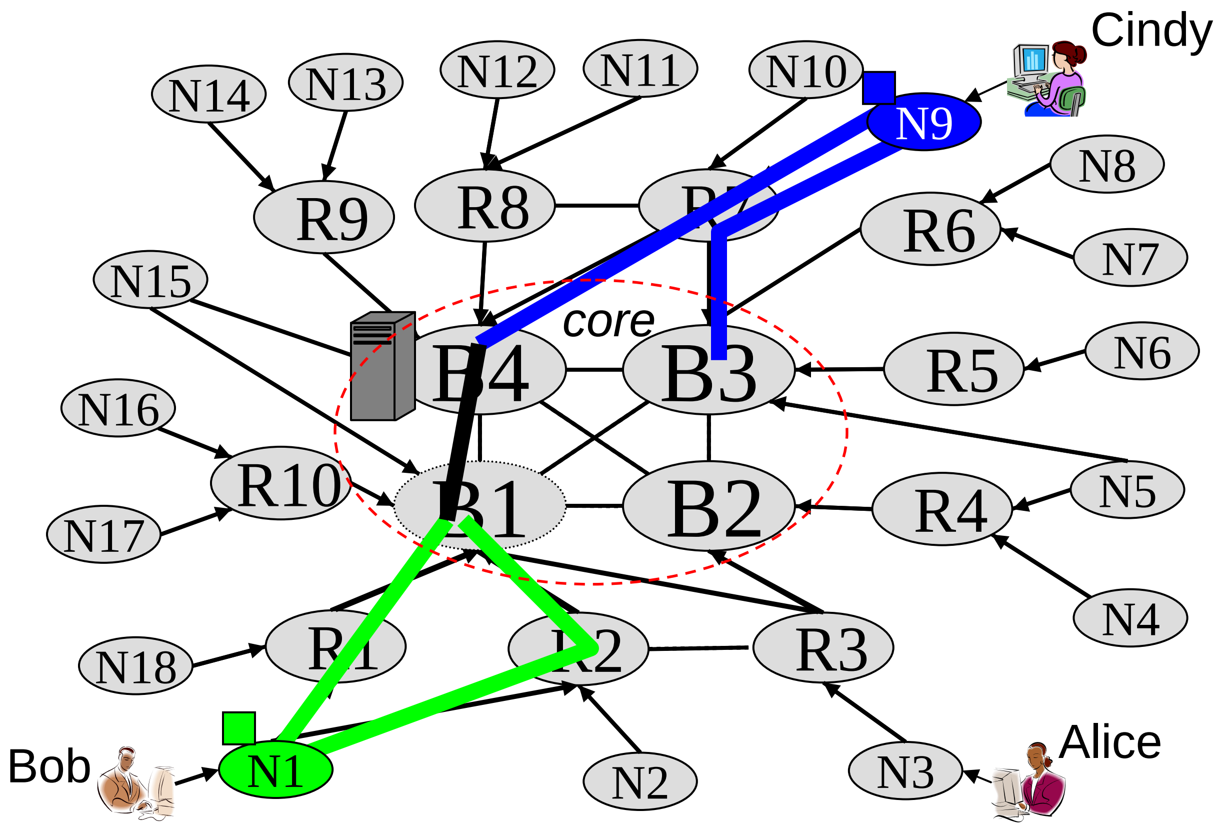

Some Interesting Details

Addressing

Addressing

- Hierarchical address assignment
- Providers in the Core obtain a globally unique address prefix
- Provider then allocates non-overlapping subdivisions of the address prefix to each of its customers

Discussion: Practical addressing scheme? One can infer ISP relationships!

Example: Addressing

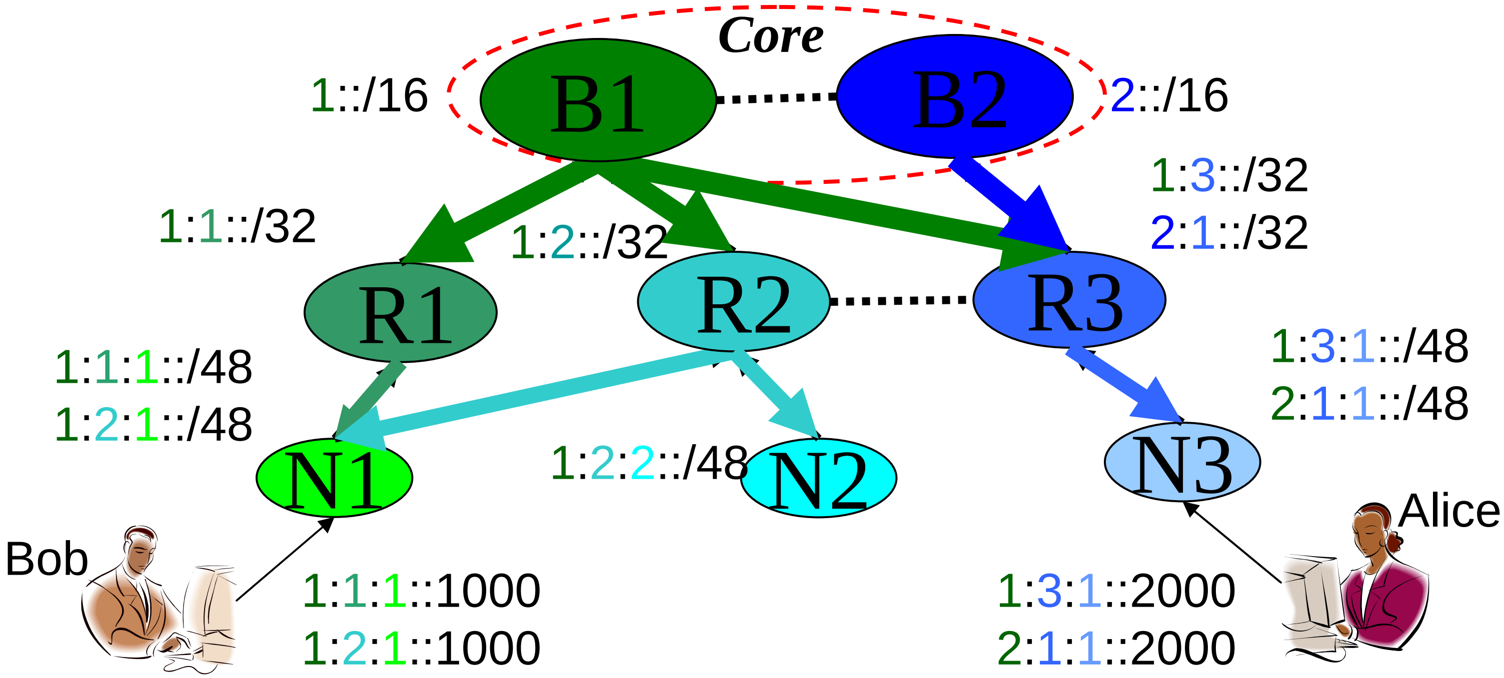

- Note: An address represents a valid route to the core.

Forwarding Tables

Uphill table

1::/16	B1

Downhill table

1:1:1::/48	N1
1:1::/96	self

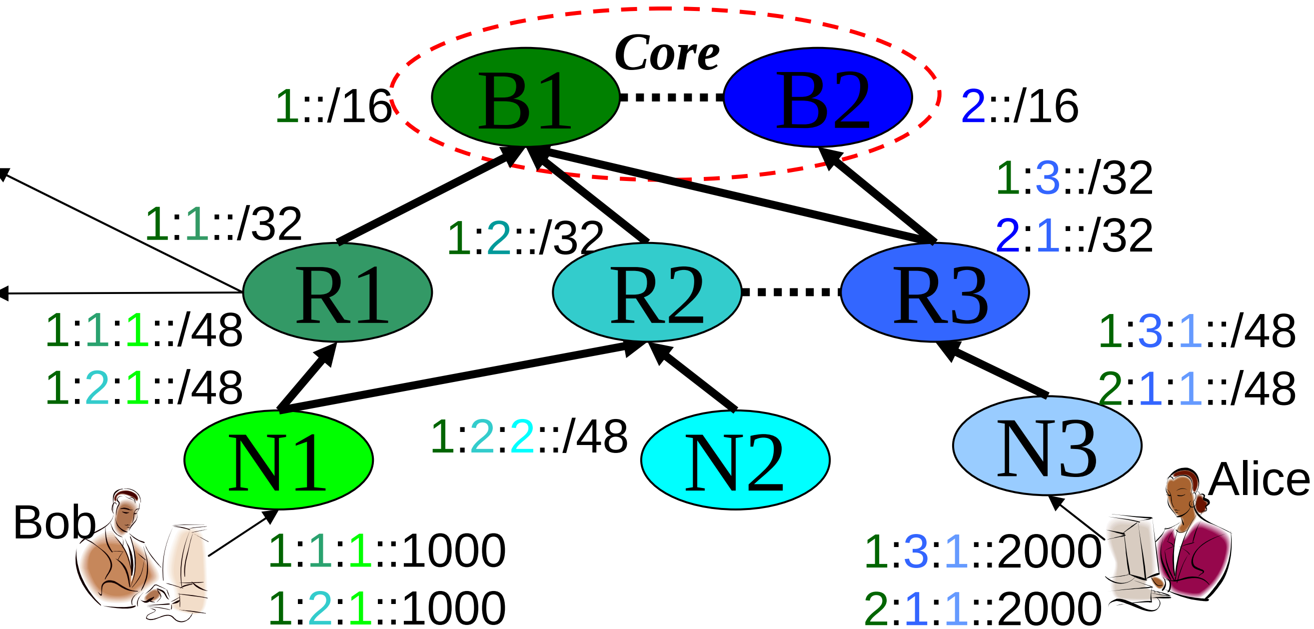

- Uphill table: providers
- Downhill table: customers, self
- Bridge table: all others

- Scalability: Size of core limited (financial factors), Provider hierarchy is shallow (domains have limited number of providers)

Hierarchical Addresses

- Provider-rooted hierarchical address
  - User can use a source and a destination address to compactly represent a “valley-free” route
  - Switch routes by switching addresses
  - Both source and destination addresses used for forwarding
    - Limits source address spoofing
      - Router may not find an address with an arbitrary source address

Efficient Route Representation

Example: Route Representation

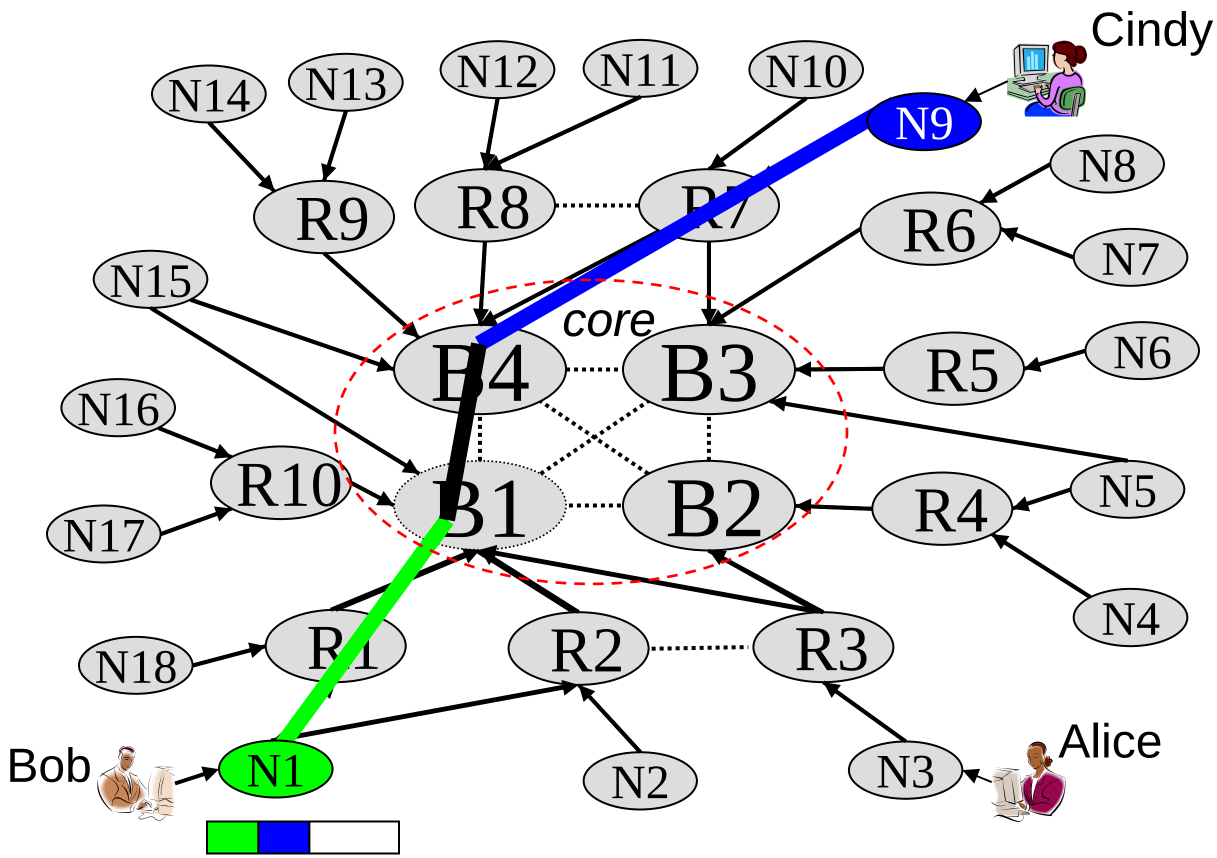

Efficient Route Representation

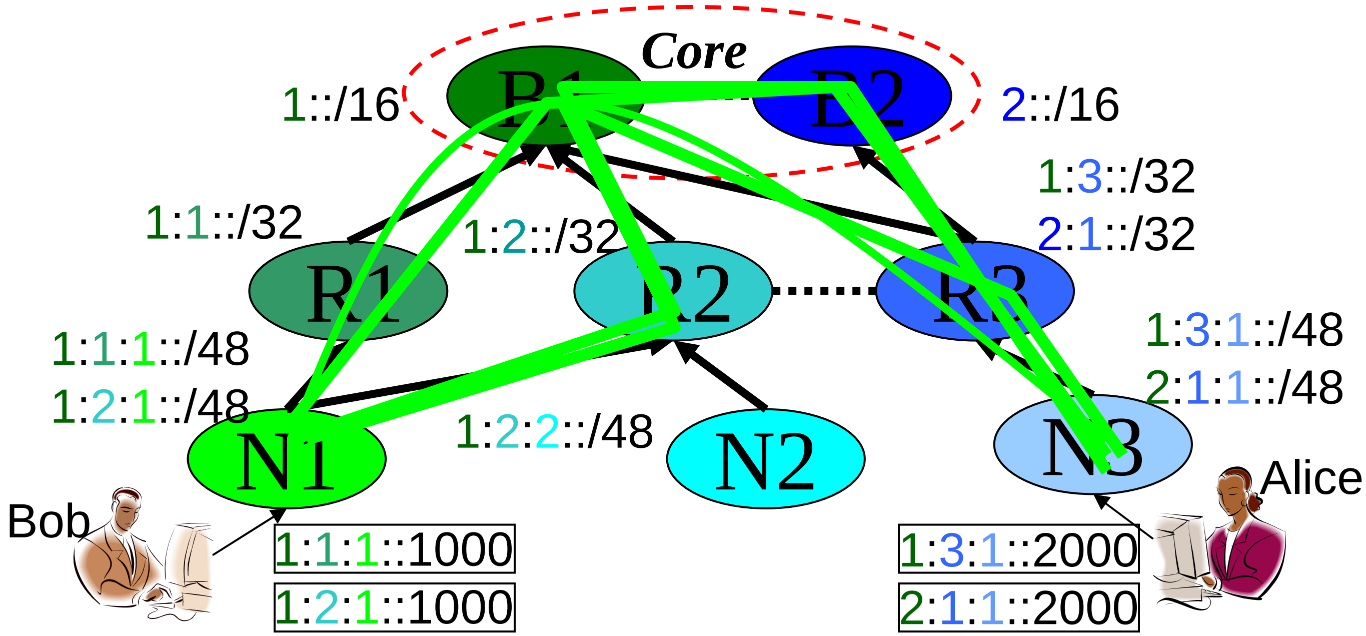

- A source and a destination address unambiguously represent a route.

Forwarding

Overview

- Packet first forwarded along the sequence of domains that allocate the source address
- Within the core (from source’s provider to destination’s provider)
- Finally, along the sequence of domains that allocate the destination address

Forwarding

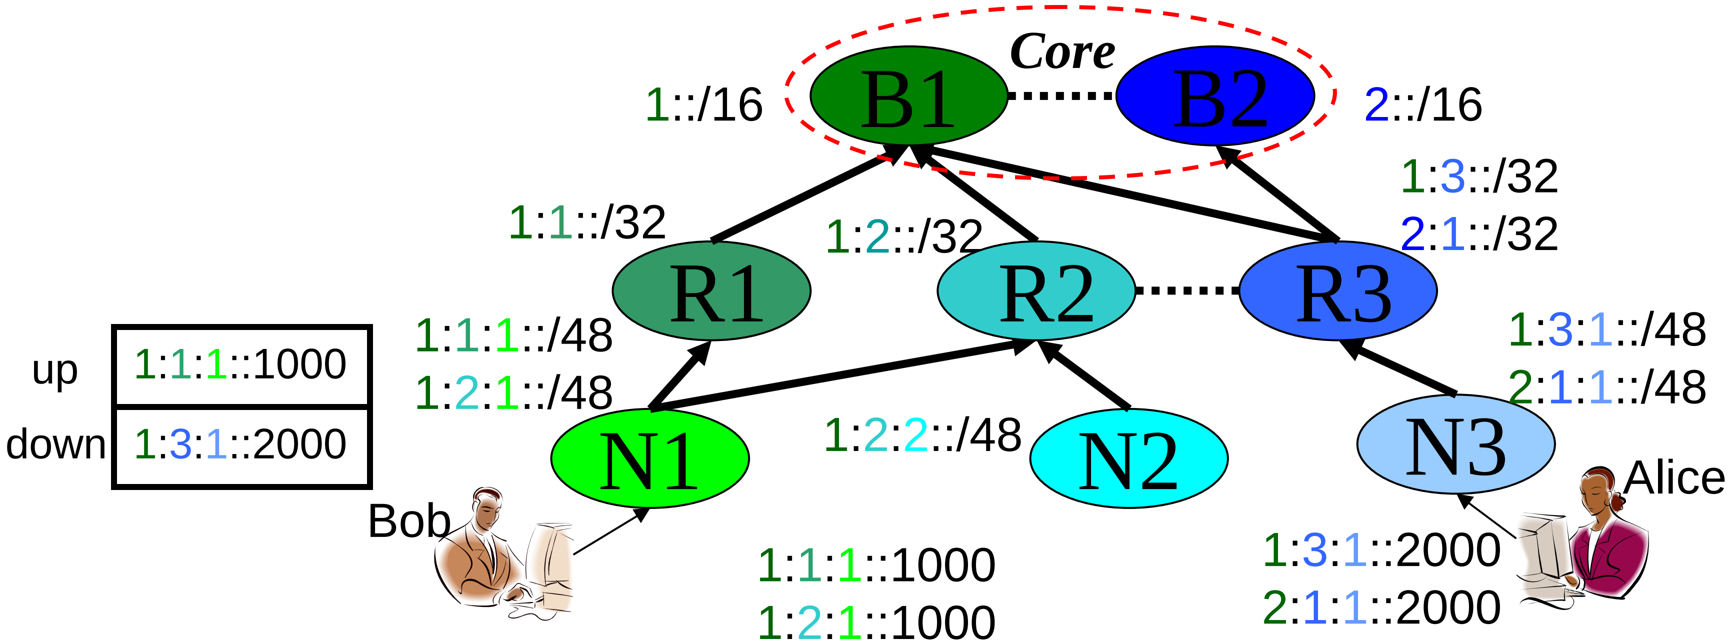

- Look up destination address in the downhill table. If no match:
- Look up the source address in the uphill table.

Discussion

- Scalability?
  - Consider each ISP having two providers. An user at level ‘k’ will have $O(2^k)$ paths.
- User control?
- How to exploit this control?
  - How to measure “goodness” of a domain-level route?
- Security:
  - Does “stronger users” necessarily mean “stronger attacks”?
- Mobility?

Back-up slides

(TIPP and Route Failures)

Topology Information Propagation Protocol (TIPP)

- Path-vector component
  - Propagating domain level routes
  - Providers propagate routes to their customers, which in turn propagate routes to their customers
  - No route selection (no policy-enforcement)

- Link-state component
  - Information about dynamic network changes
  - Link-state messages could potentially be propagated only down the hierarchy (no message from a customer to provider required)

Handling Route Failures

Route Failures

- Problem:
  - TIPP messages do not propagate globally
    - The sender might not have up-to-date information about destination’s path (when the destination does not update its routes in NRLS very frequently)
- Solution:
  - If the route in the packet header is unavailable, inform the sender!
  - If no information received, use timeout!