NIRA: A New Internet Routing Architecture

Xiaowei Yang

MIT Computer Science and Artificial Intelligence Laboratory

yxw@lcs.mit.edu

Why a New Internet Routing Architecture?

- Users have little control over routes.
  - BGP chooses a default route.
- User choice fosters competition.
  - Telephone system
  - A small number of local ISPs
  - Stagnation in introducing new services, e.g., lack of end to end QoS
- BGP has a number of problems.
  - Slow convergence [Labovitz00], slow reaction to failures [Feamster03], path selection based on simple metrics [Spring03]
- User selected overlay routes have better quality.
  - Detour [Savage99], RON [Andersen01]

What's NIRA about?

- ''Dumb'' routers, intelligent users
  - Routers selectively propagate topology information
  - Forward or drop packets based on policy configurations.
  - Route selection is done by end users.
    - *''User'' is an abstract entity, e.g., a software agent*
- ''Domain-level'' choices
  - Domain-level choices encourage ISP competition.
  - Individual domain's decision to offer ''router-level'' choices or not.
- NIRA is different from IP source routing.
  - IP source routing does not address route discovery.
- Enable new services: QoS route selection, multipath routing (I-ATC)

Outline

- Network Model
- Overview of design components and requirements
- Design details
- Related Work
- Conclusion

NIRA's Network Model

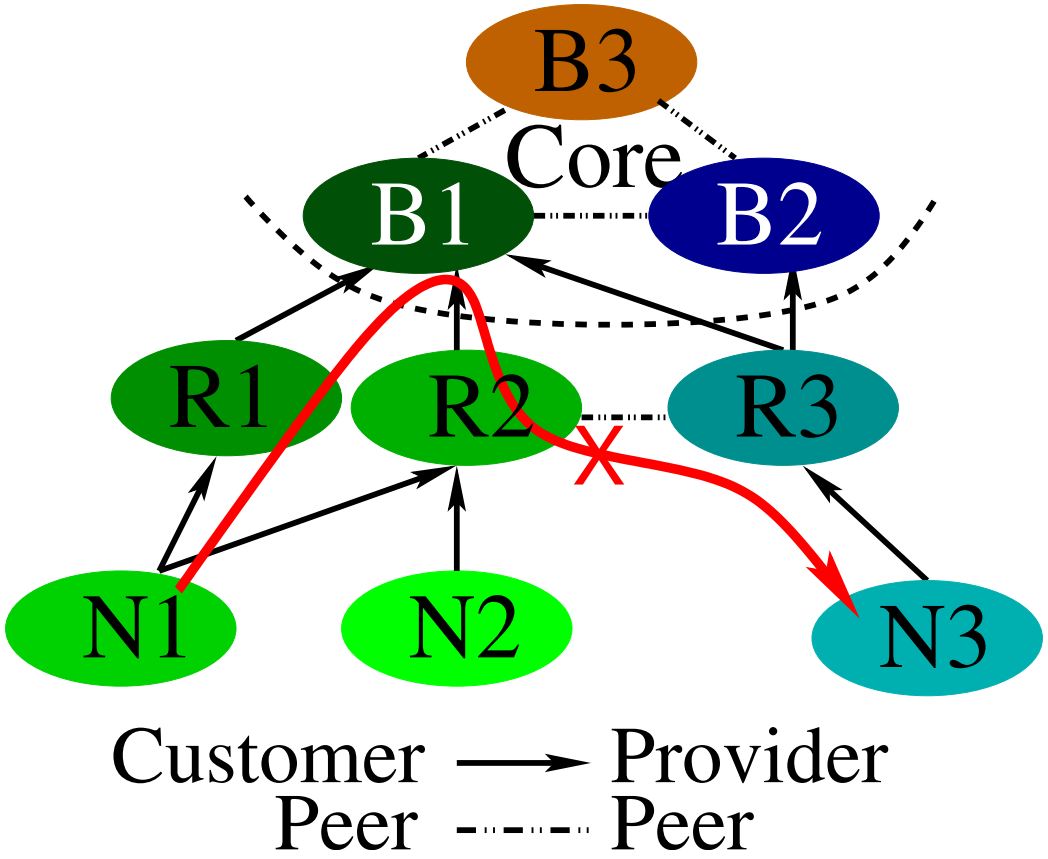

- NIRA is optimized for the special structure of the Internet.

Overview of Design Components and Requirements

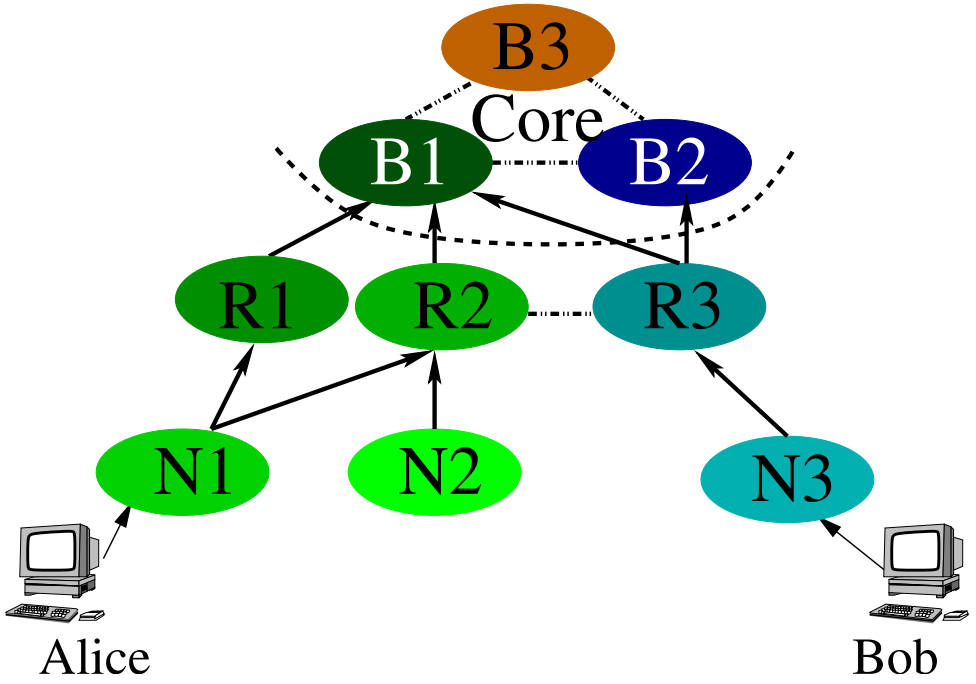

- Design components
  - Route discovery and selection
  - Route representation
  - Route failure handling
  - Provider compensation
- Design requirements
  - Scalability
  - Efficiency
  - Robustness
  - Heterogeneous user choices
  - Practical provider compensation

Provider-rooted Hierarchical Addressing

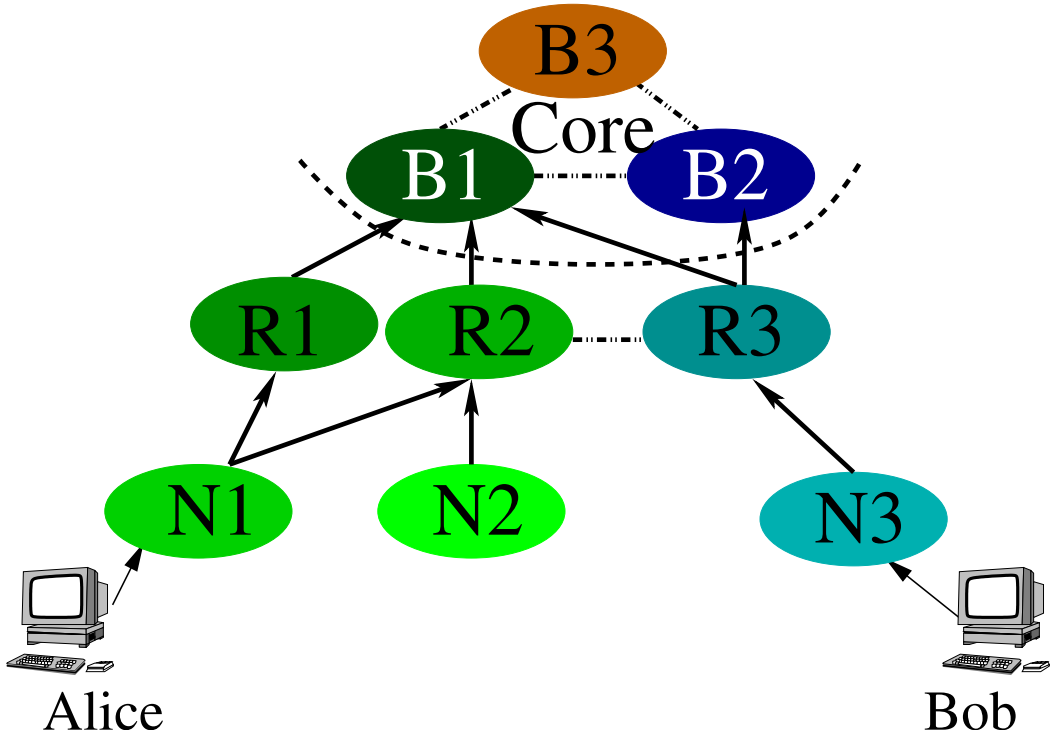

- Fixed length addresses
- Distinct address prefixes for top-level providers
- Recursive subdivision for each customer

Provider-rooted Hierarchical Addressing

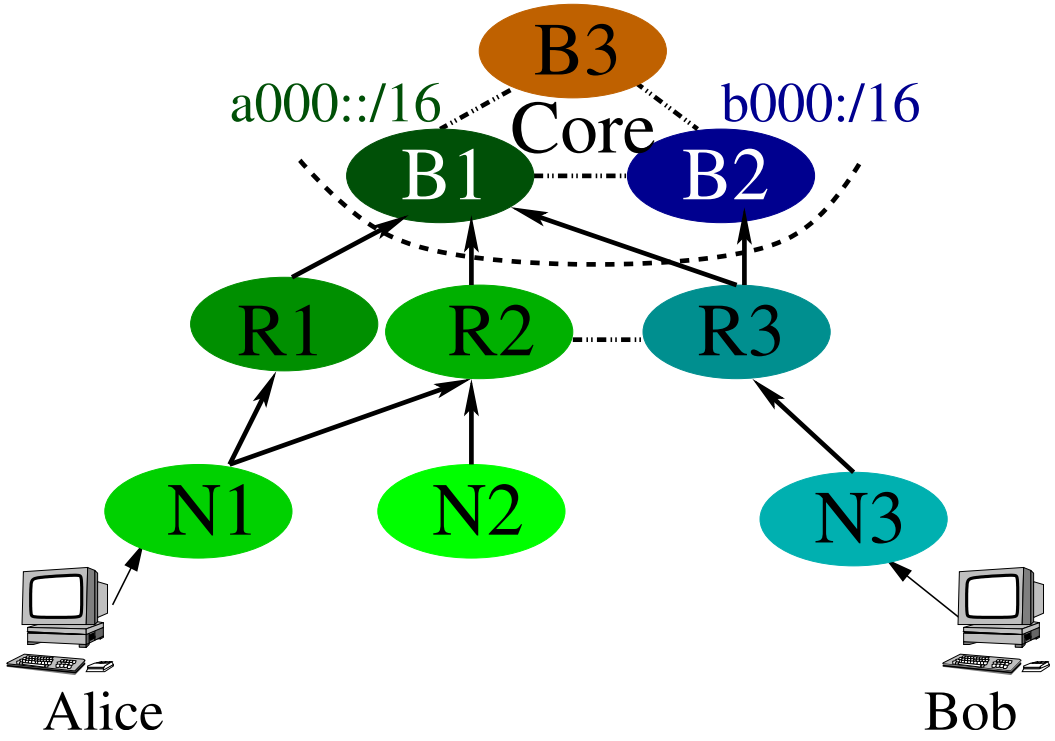

- Fixed length addresses
- Distinct address prefixes for top-level providers
- Recursive subdivision for each customer

Provider-rooted Hierarchical Addressing

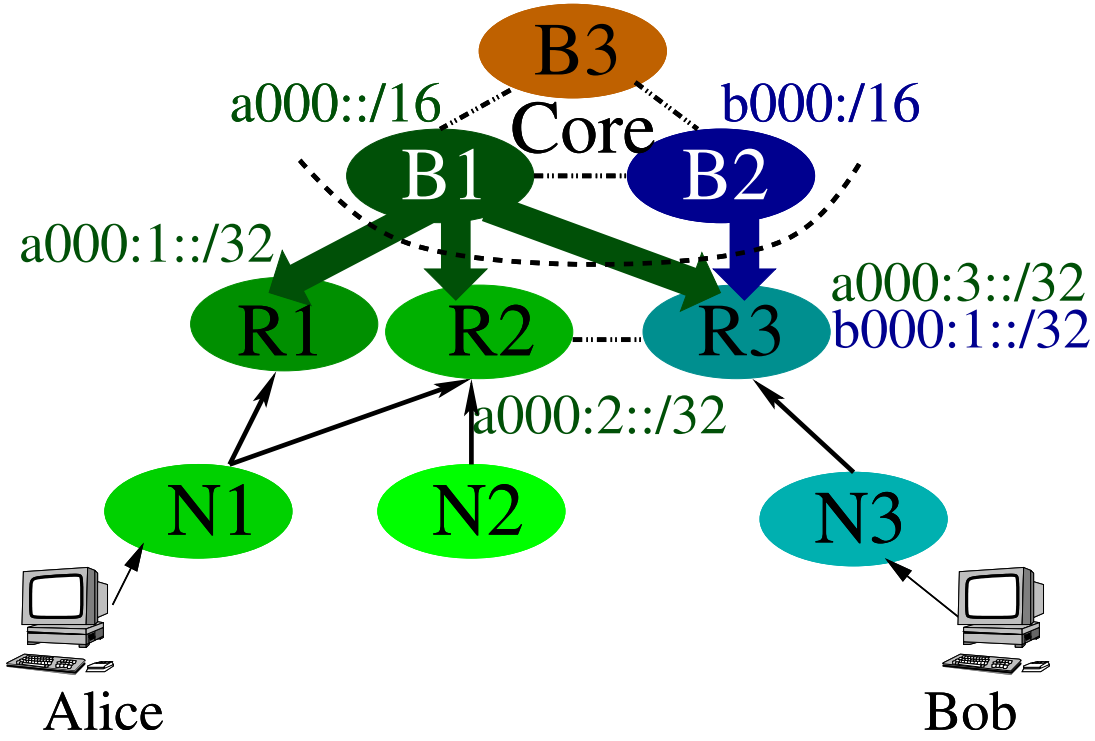

- Fixed length addresses
- Distinct address prefixes for top-level providers
- Recursive subdivision for each customer

Provider-rooted Hierarchical Addressing

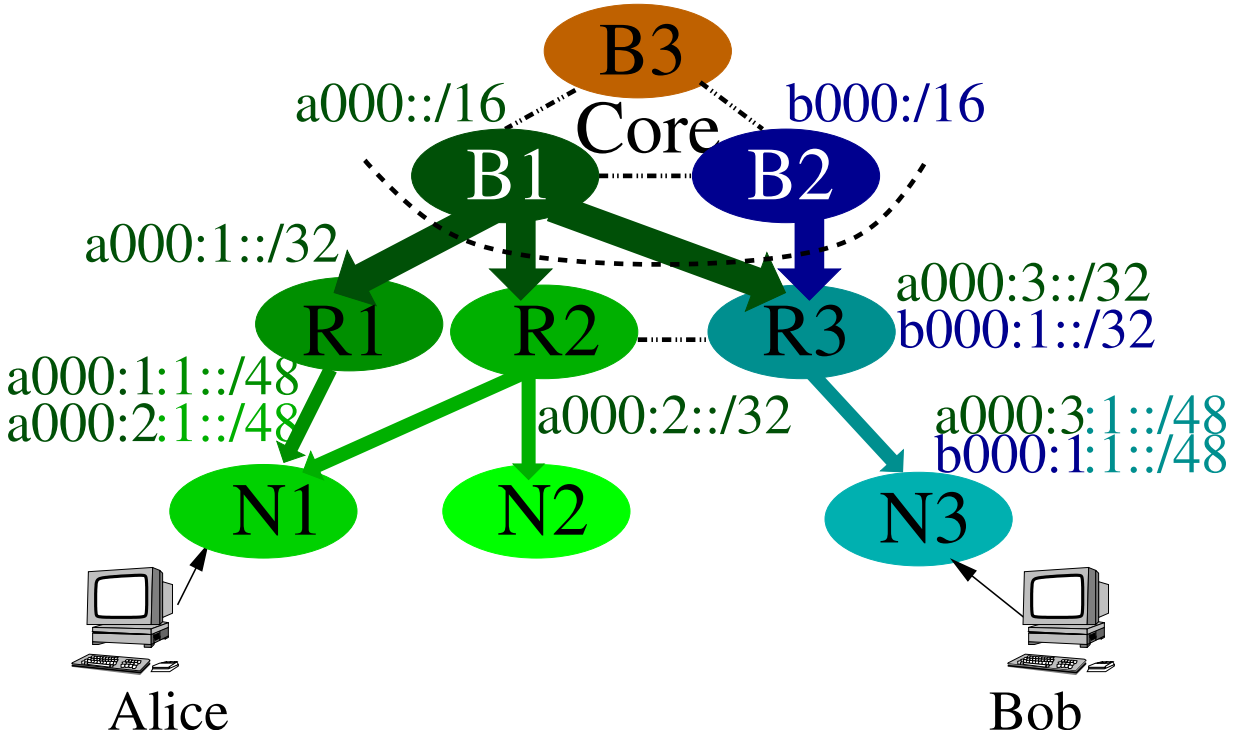

- Fixed length addresses
- Distinct address prefixes for top-level providers
- Recursive subdivision for each customer

Provider-rooted Hierarchical Addressing

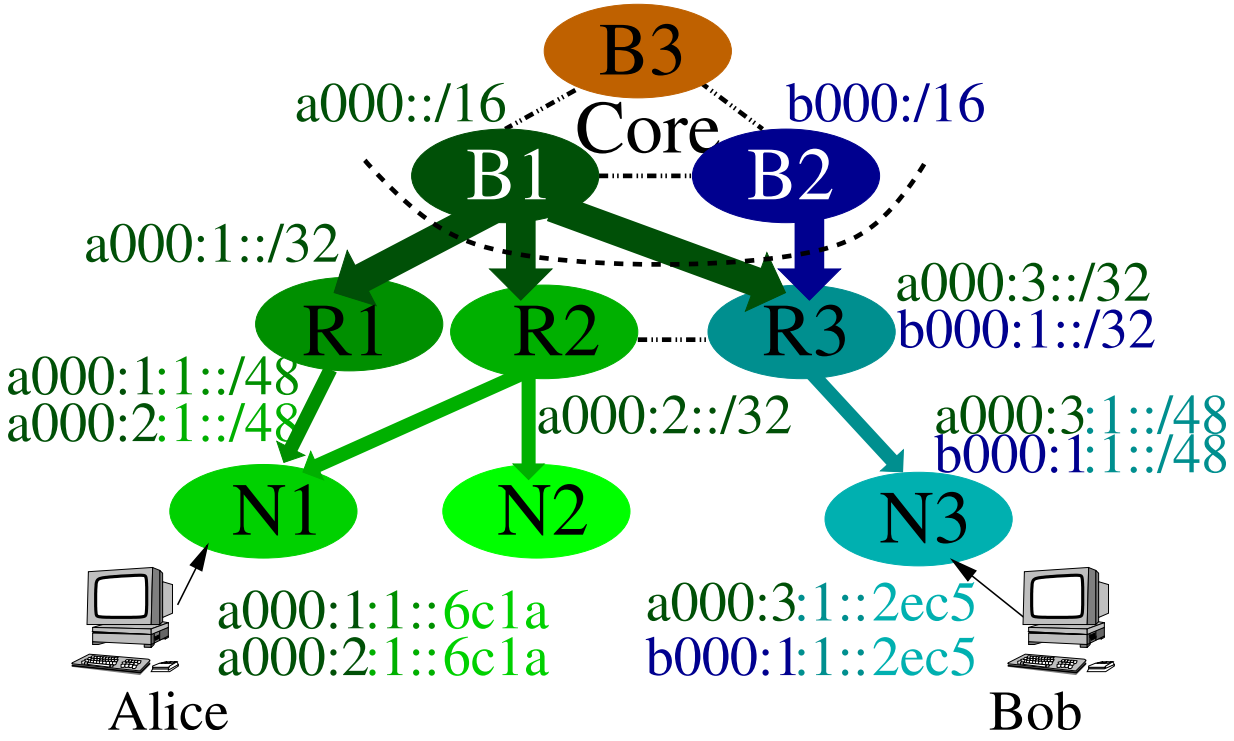

- Fixed length addresses
- Distinct address prefixes for top-level providers
- Recursive subdivision for each customer

Efficient Route Representation

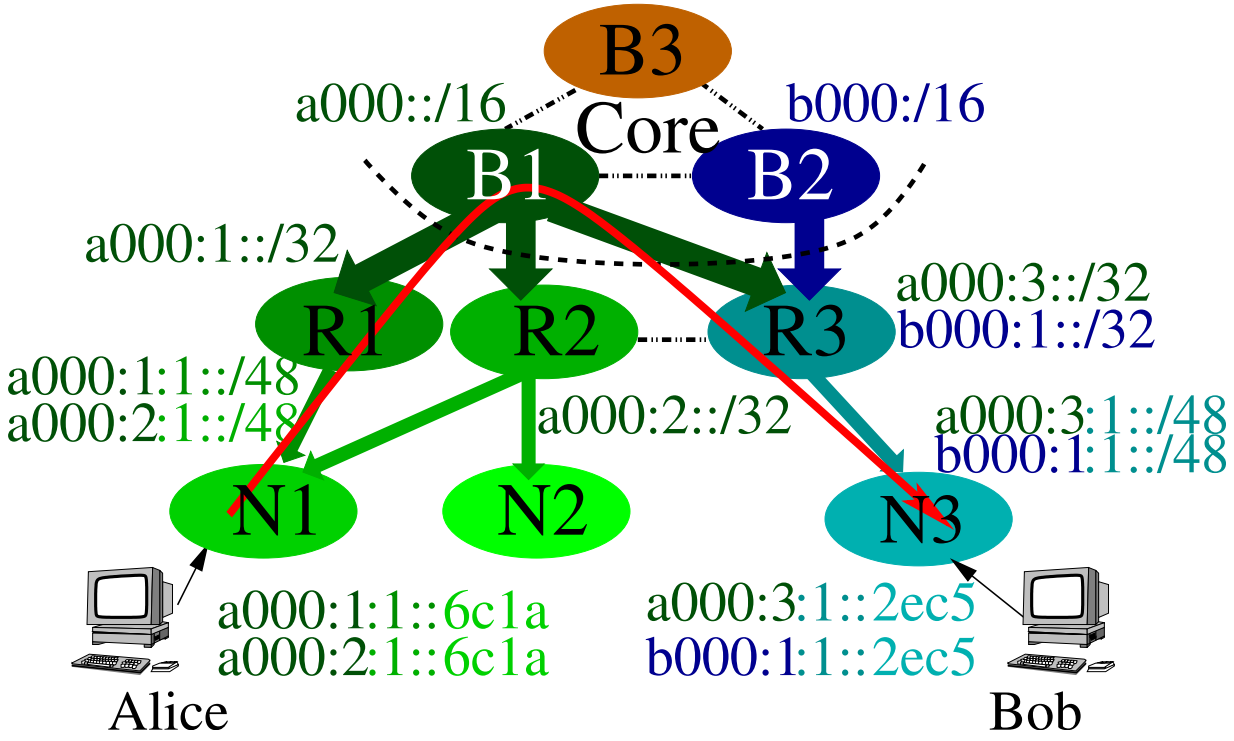

- A ''valley-free'' route can be represented by two addresses.
- Route ''N1 → R1 → B1 → R3 → N3'' can be represented as a000:1:1::6c1a and a000:3:1::2ec5
- Source routing header for more complicated cases

Scalable Route Discovery

- Route discovery is divided into two halves.

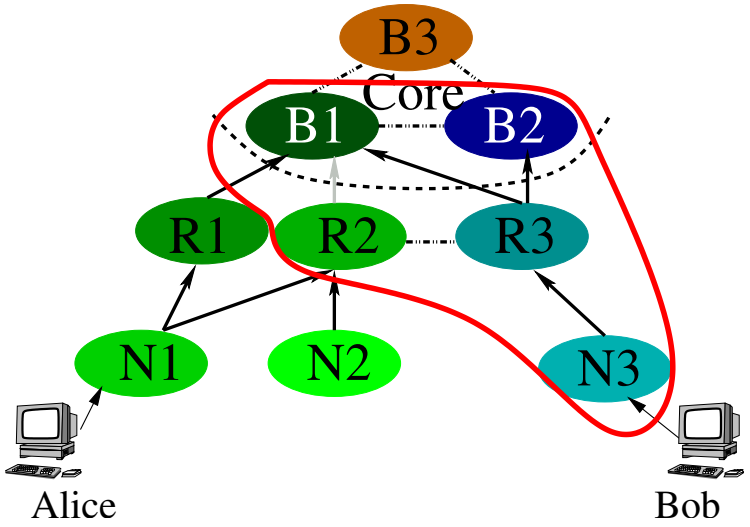

- *Topology Information Propagation Protocol (TIPP) tells Bob his ''up-hill'' routes.*

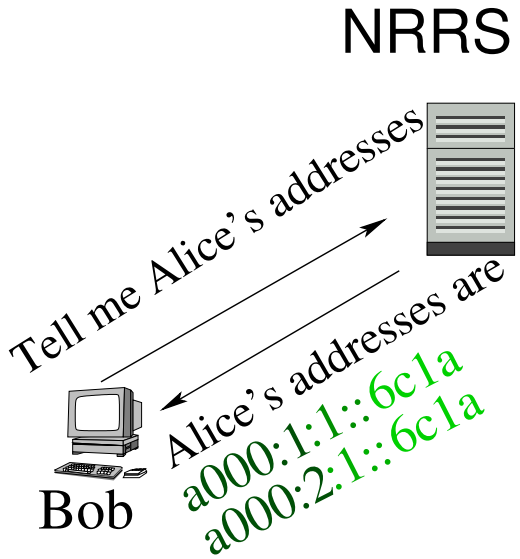

- *Name-to-Route-Resolution service (NRRS) tells Bob the addresses of Alice.*

Topology Information Propagation Protocol (TIPP)

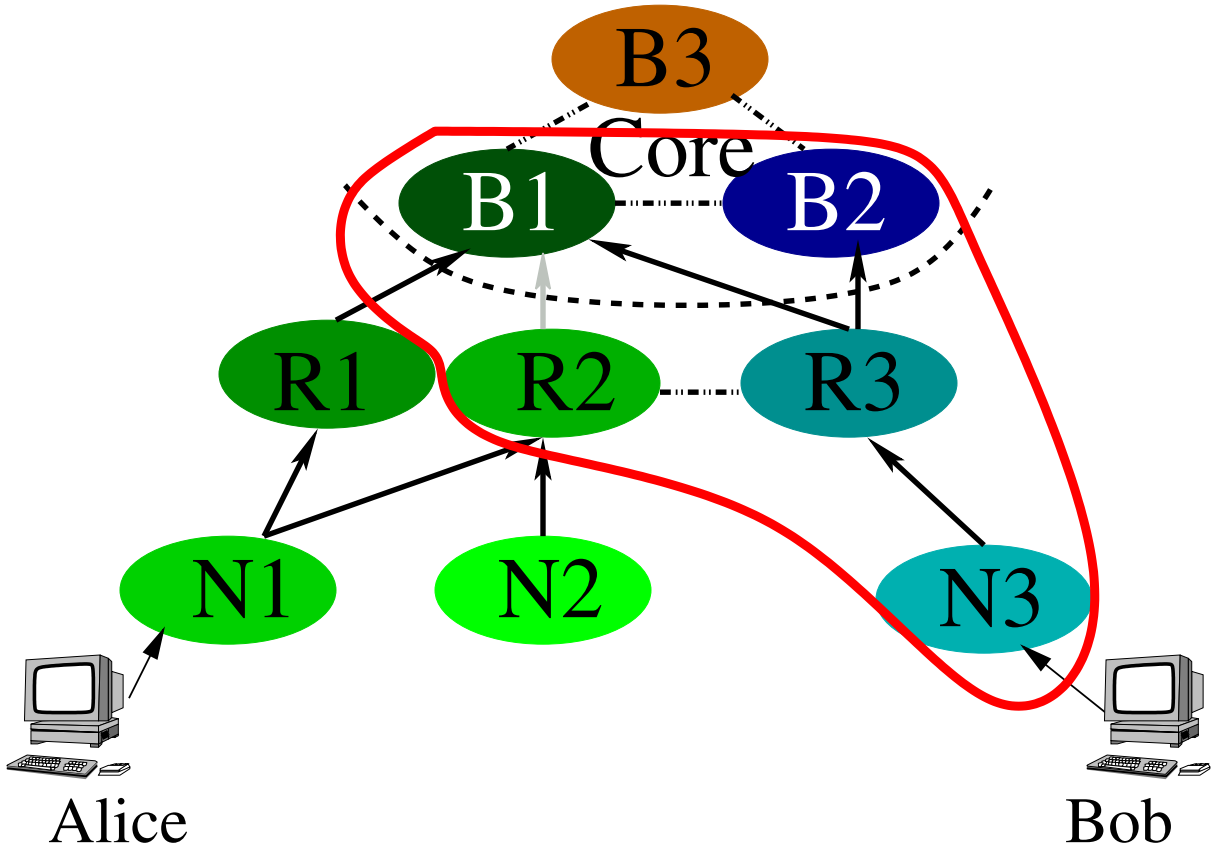

- Policy-based, link-state like protocol
- Routers selectively propagate topology information.
- Messages do not propagate globally.
- An average user has less than 20 address prefixes.

Topology Information Propagation Protocol (TIPP)

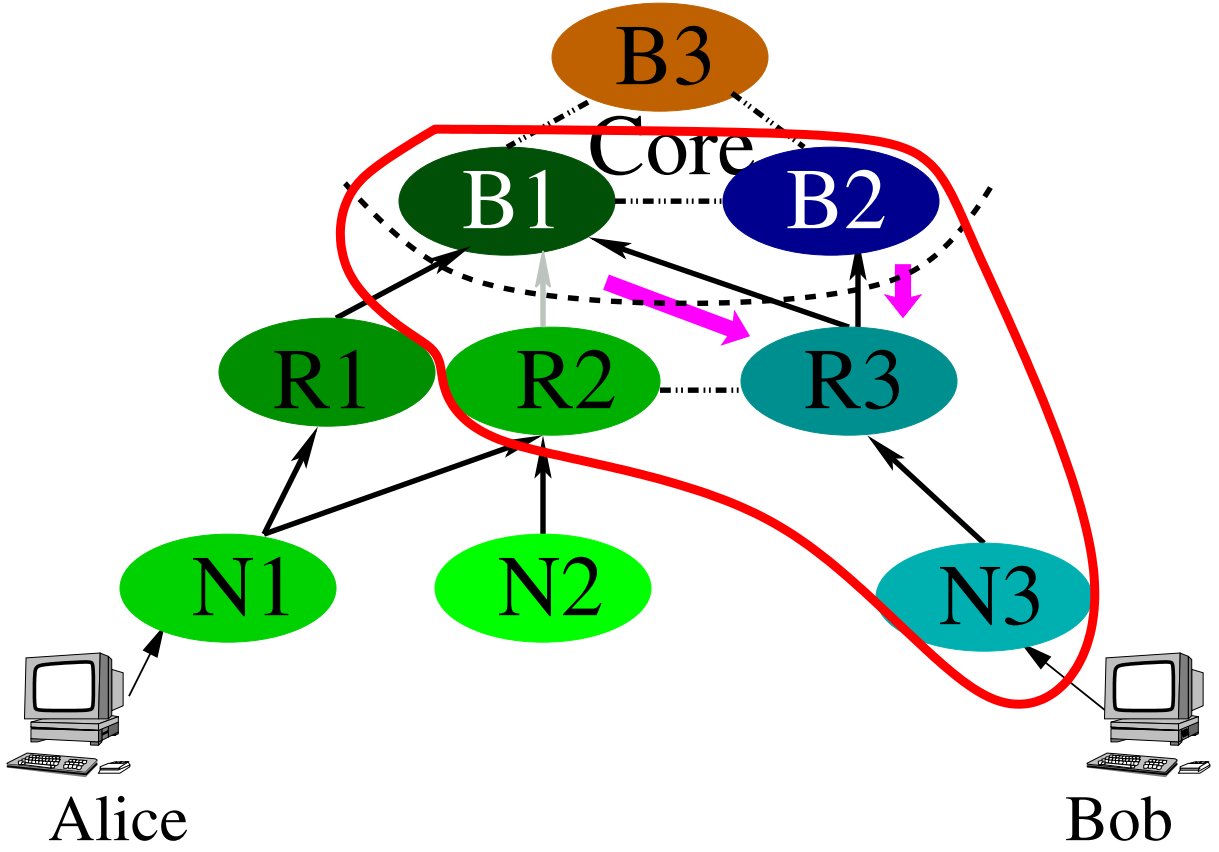

- Policy-based, link-state like protocol
- Routers selectively propagate topology information.
- Messages do not propagate globally.
- An average user has less than 20 address prefixes.

Topology Information Propagation Protocol (TIPP)

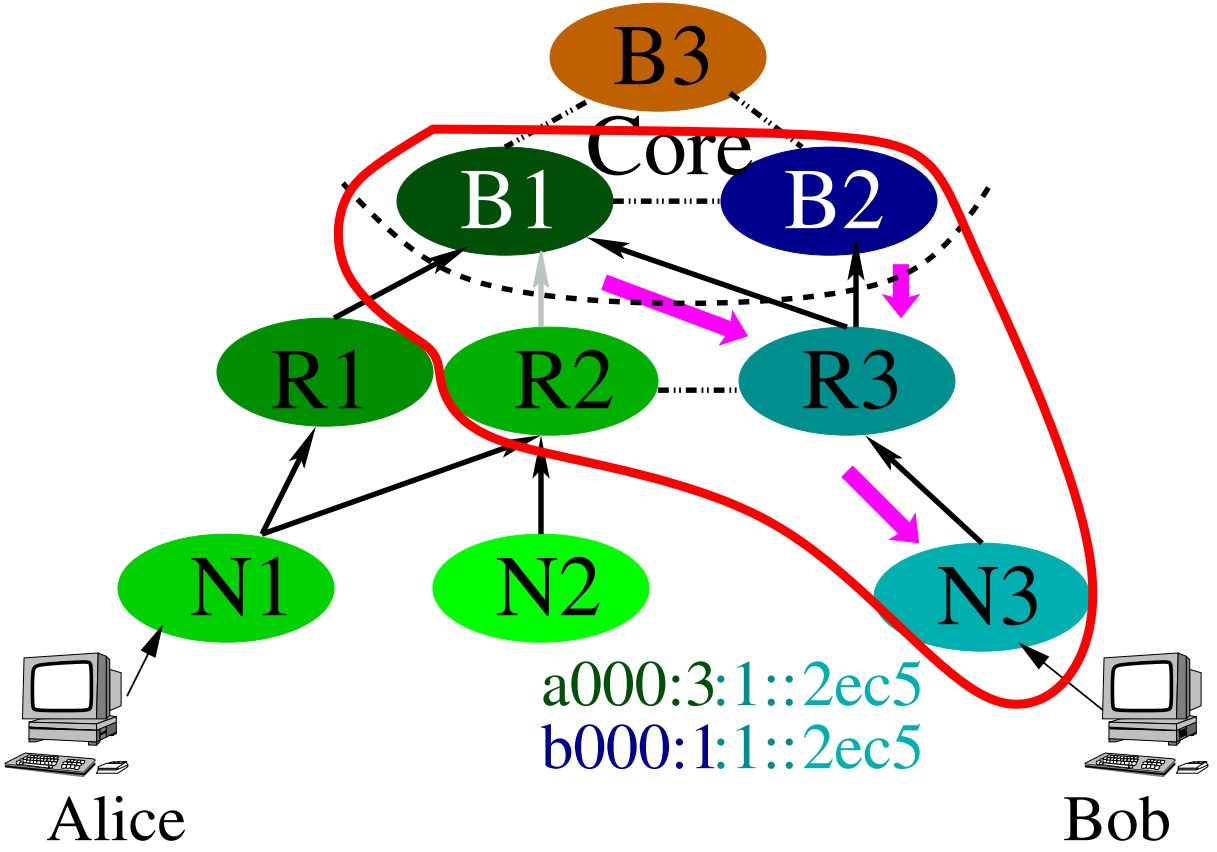

- Policy-based, link-state like protocol
- Routers selectively propagate topology information.
- Messages do not propagate globally.
- An average user has less than 20 address prefixes.

Name-to-Route Resolution Service (NRRS)

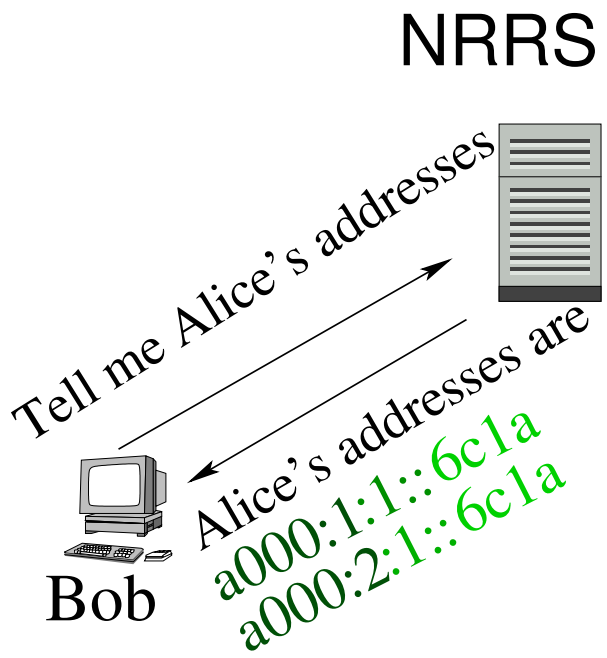

- Hierarchical namespace
- Hard-coded addresses for bootstrapping
- Route information associated with addresses is optional.

- A fundamental tradeoff: topology change will cause address change
  - Root servers reside in top-level providers.
- TIPP propagates address allocation information
- Route record updates are likely to be manageable
  - AS birth and death rate $<$ 15 per day
  - AS-level link birth and death rate $<$ 50 per day [Chen02]

Failure Handling

- A combination of proactive notification and reactive feedback

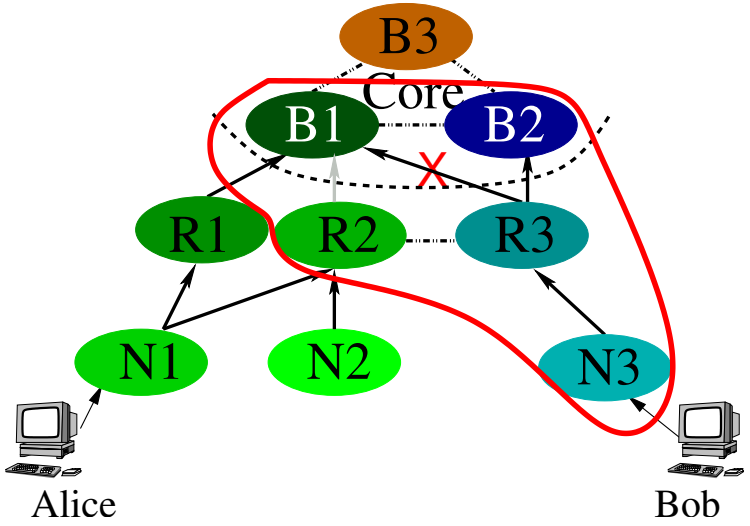

- ''Uphill routes''
- Proactive notification via TIPP

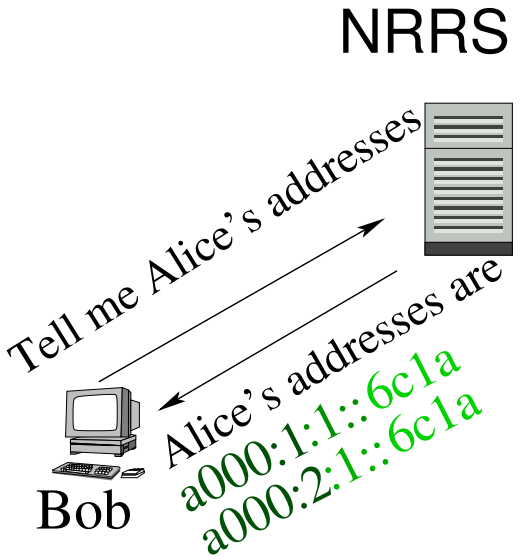

- ''Downhill routes''
- Reactive discovery via router feedback or timeout

Failure Handling

- A combination of proactive notification and reactive feedback

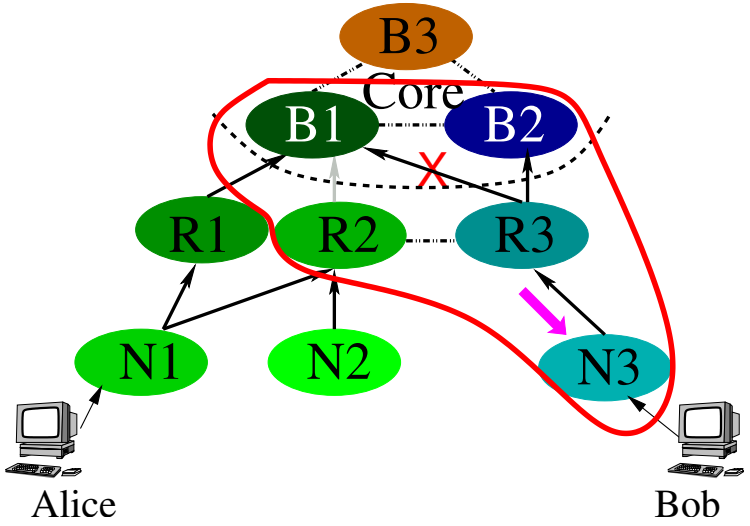

- ''Uphill routes''
- Proactive notification via TIPP

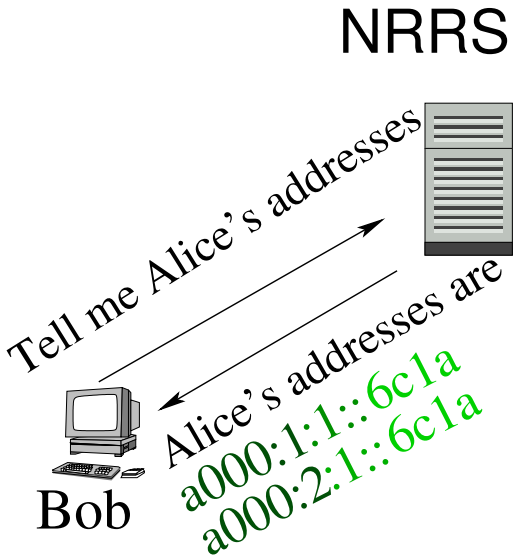

- ''Downhill routes''
- Reactive discovery via router feedback or timeout

Practical Provider Compensation

- Contractual agreements
- Providers use policy checking to prevent illegitimate route usage.
  - Direct business relationships
    - *Common cases: verify source address*
    - *General cases: match source routing header against policy filters*
  - Indirect business relationships: end users pay remote providers
    - *Policy is made upon the originator or consumer of a packet.*
    - *An open and general question: how to tell who sends a packet?*
    - *e.g, overlay providers, second hop provider sells QoS service*
- Various billing schemes are possible.
  - Flat fee, usage-based billing
  - Some routes are more expensive.

Related Work

- NIMROD
- IDPR, SIDRA
- PIP, SIPP
- TRIAD
- Feedback Based Routing [Zhu02]
- Our work is different:
  - Scalable route discovery
  - Efficient route representation
  - Heterogeneous user choices

Conclusion

- The design of NIRA
  - Scalability
    - *Divide the task of route discovery into two halves*
    - *Topology Information Propagation Protocol*
    - *Name-to-Route Resolution Protocol*
  - Efficiency
    - *Provider-rooted hierarchical addressing for efficient route representation*
  - Robustness
    - *A combination of proactive notification and reactive feedback for failure discovery*
  - Heterogeneous user choices
    - *Scalable route discovery*
  - Practical provider compensation
    - *Contractual agreements*