


Nimrod Routing

Feedback Based Routing

Distributed Computing Seminar

Martino Trosi

trosim@student.ethz.ch



A subgraph of a BGP graph

Summary

- Limitations of actual routing system
- Feedback based routing (protocol)
 - Security, Fault tolerance
 - Scalability, Rapid convergence
- Nimrod routing (architecture)
 - Security, Scalability, New Features
 - Deployment

Actual Routing protocol

- Border Gateway Protocol - BGP
 - Distance vector
 - Hop-by-hop
 - All routers collaborate

What's wrong with that?

- BGP
 - Distance vector
 - Grows exponentially
 - Slow convergence time
 - Hop-by-hop
 - No control on the route by users
 - All routers need to collaborate
 - Tables must be consistent

Feedback Based Routing

- New routing protocol
- Link state
 - Distinguish structural and dynamic information
- Edge: routing decisions and measurements
- Core: propagate information and forward packets


Routing information

- Structural
 - Existence of links
 - Propagated to the edge of the network
- Dynamic
 - Quality of paths
 - Not propagated, but measured

Access and Transit

- Access routers
 - at the border of edge networks
 - select routes
 - insert Internet Relay Tokens (IRT)
- Transit routers
 - at the border of autonomous systems (AS)
 - don't compute routing tables
 - forward packets with respect to IRT


Normal route



Feedback algorithm

- Compute two routes to every network prefix
 - disjoint routes
 - exclusion of bad-links
- Monitor the quality of the routes
 - assign expiration time
 - sample round trip time (RTT)
- Define primary and backup routes

New Routes



Security

- Access routers always monitor the route
 - avoid black holes and bad links
 - public key to authenticate the routers
- Problems in monitoring
 - bogus TCP answers: SYN / SYN-ACK

Scalability

- Route computation is moved to the edge
 - no need to compute shortest path in core routers
- Messages reduction: only structural changes
- Transit routers
 - independent from number of address prefixes
 - can limit resources

What can be better?

- Policy routing
 - provider restrictions (cost, services, accounting)
 - user requirements (quality, provider)
- Mobility of endpoints
- Different types of communication media
 - throughput, delay
 - privacy

Nimrod

- Perfection is attained not when there is no longer anything to add, but when there is no longer anything to take away.

Antoine de Saint Exupery

- One ring to rule them all, one ring to find them, one ring to bring them all, and in the darkness bind them.

J.R.R. Tolkien

Nimrod Routing

- Nimrod: it might run one day
- RFC 1992, many other documents
- Year 1996
- Routing Architecture

Architecture vs Protocol

- Architecture provides
 - a model to understand the internet network
 - description of functionalities and interactions
- Architecture doesn't provide
 - definition of the protocols
 - definition of the algorithms

Goals

- Policies
 - access control
 - trust model
 - information hiding
 - accountability
- Service specific routing
 - best-effort, pay service
 - user requirements, provider restrictions

Goals

- Support dynamic internetwork
 - arbitrary size
 - control routing information
 - user-controlled routes
- Incremental deployment, interoperable changes
- Vendor independant

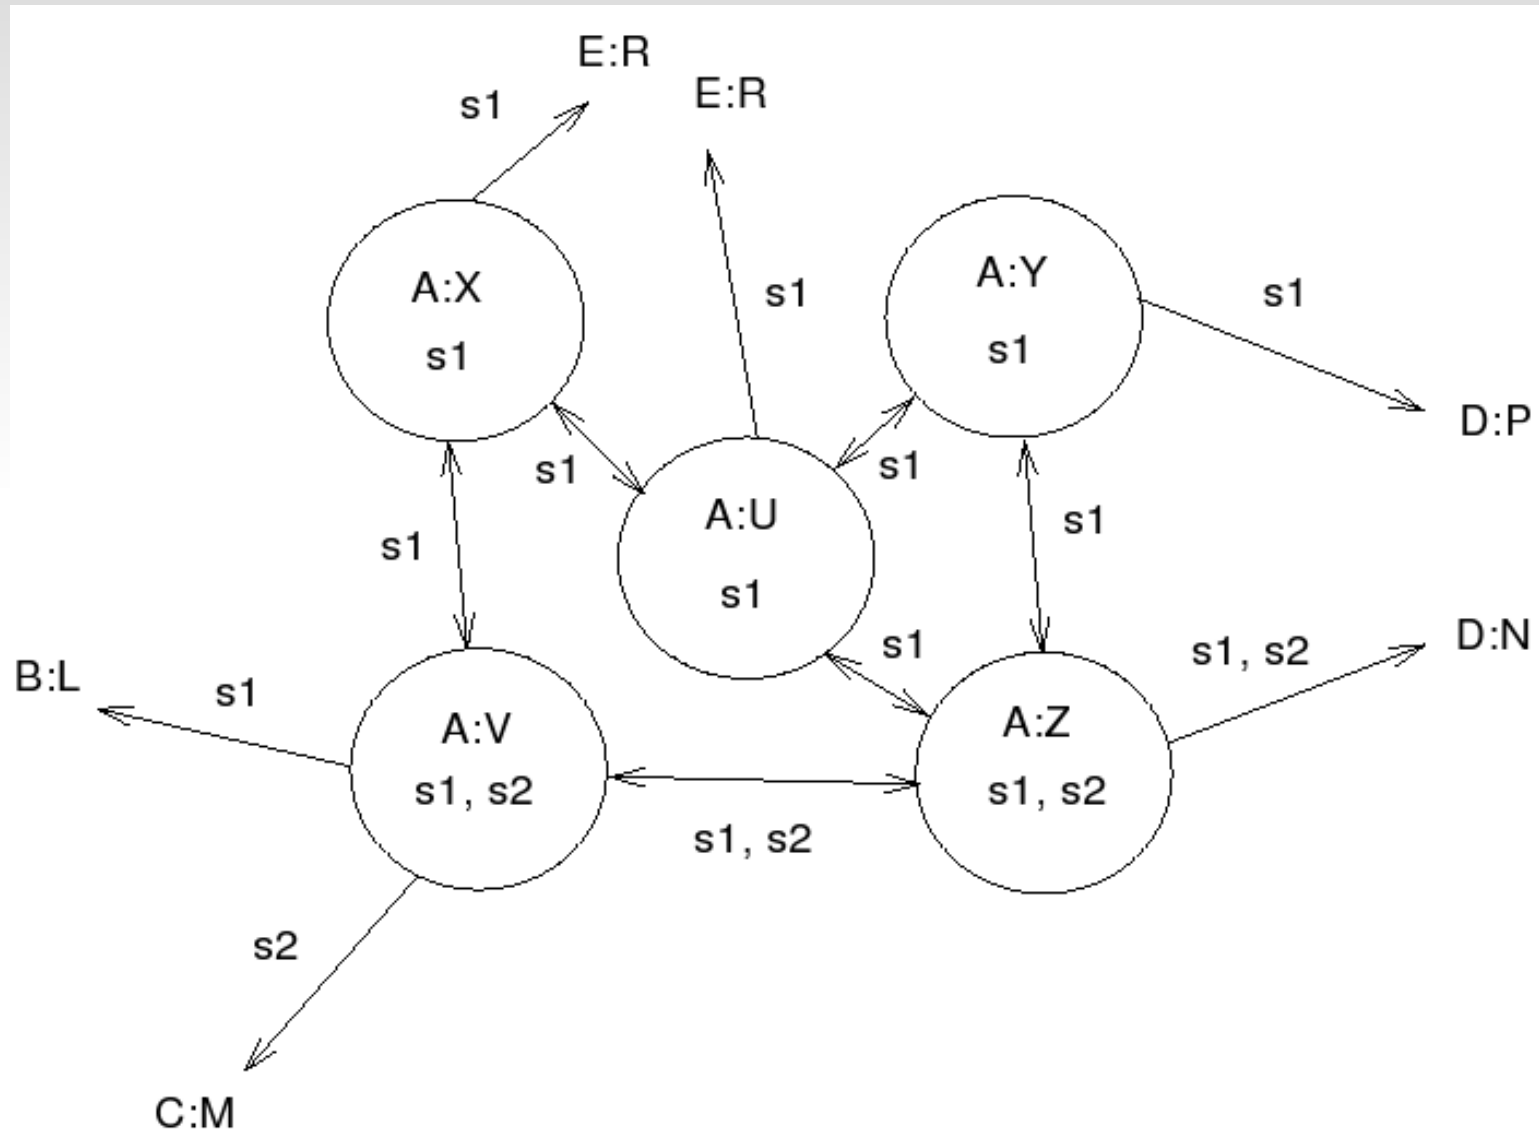
Architecture

- Node
 - region of the internetwork
 - host, continent, process
- Adjacency
- Map
 - graph composed by nodes and adjacencies
 - used for routing
 - different levels of abstraction

Architecture

- Endpoint Identifiers (EID)
 - identify host or user
 - globally unique
 - no topological significance
- Locators
 - identifies a location
 - hierarchical
 - used for routing decisions


Raw Map



Scalability

- Clustering
 - reduce the number of visible entities
 - hierarchy with unique universal root
 - must satisfy connectivity
- Abstraction
 - reduce the amount of information
- Either algorithms are not imposed by Nimrod

Abstract Map



Scalability

- Information hiding
 - each node decides what information to disclose
 - advertising of specific portions
 - advertising to specific users
- Limit forwarding information
 - multiplex traffic flows
 - install information only for active flows

Scalability

- Local selection of routes
 - compute special routes only if needed
 - protocol not imposed by Nimrod
- Caching
 - temporary routes used to generate the final ones
 - expected duration of useful information

Forwarding

- Connectivity Specification Sequence
 - sequence of nodes to visit
 - doesn't specify physical path
- Connectivity Specification Chain
 - continuous list of nodes to visit
 - doesn't specify physical path

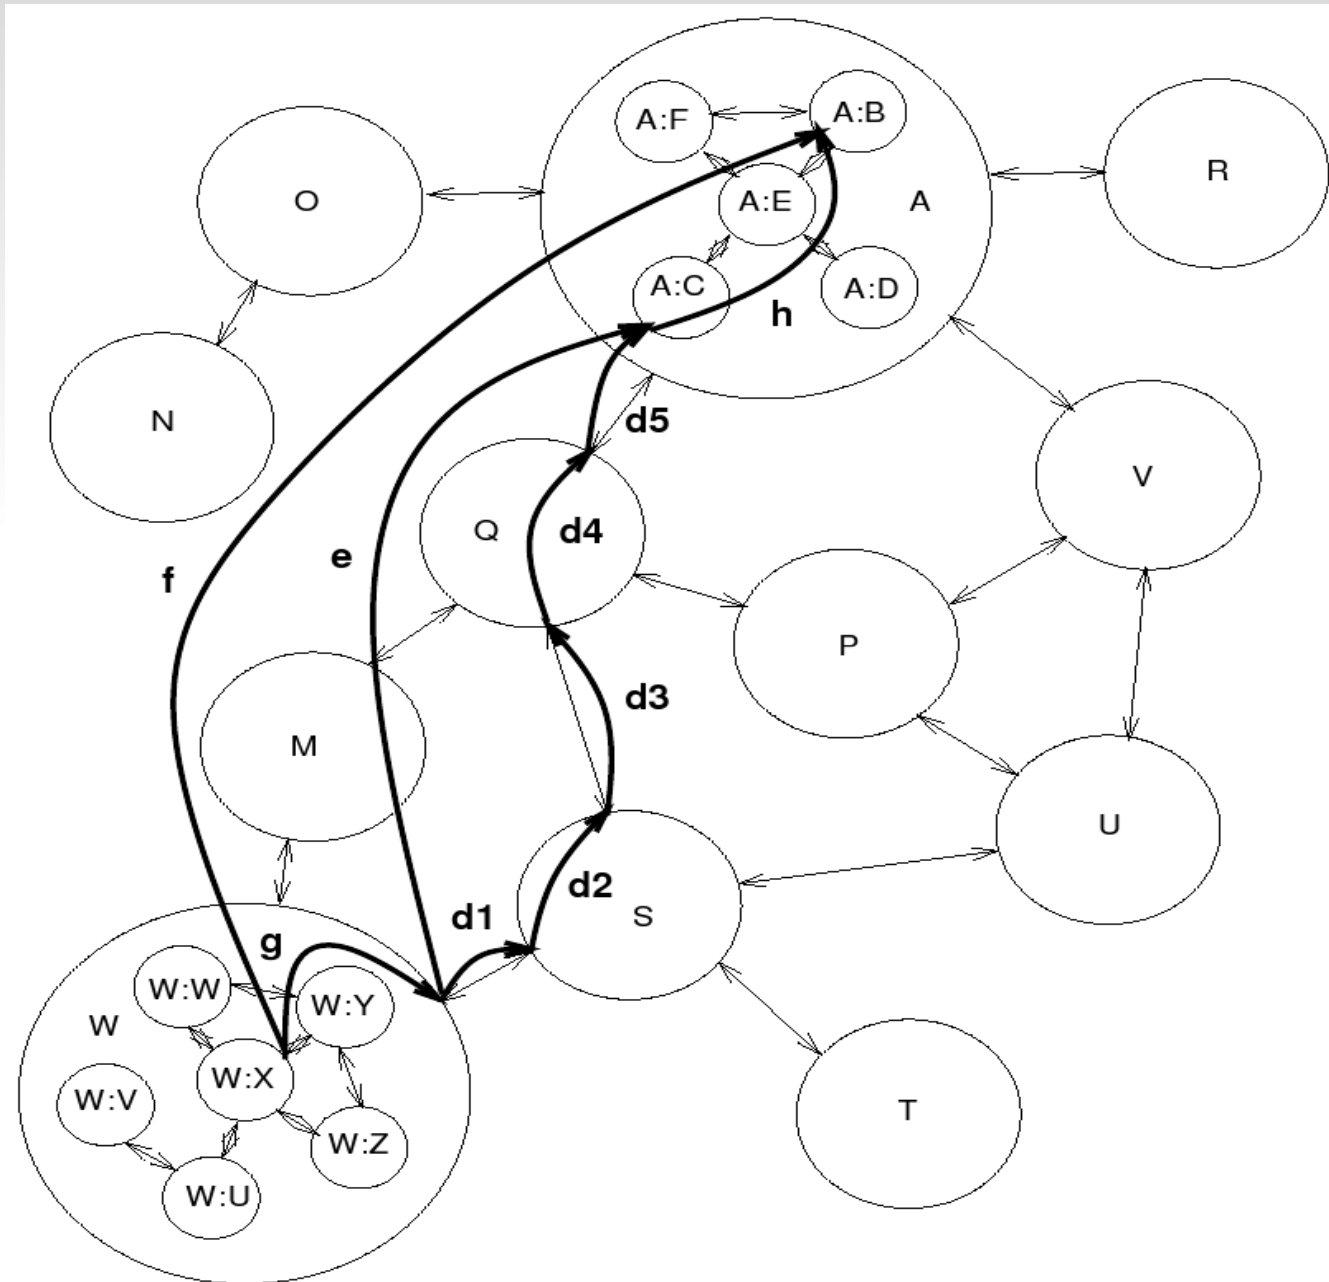
Forwarding

- Flow mode
 - initially installs a path in the routers
 - for each packet, just follow the installed instructions
- Datagram mode
 - every packet carries source and destination locators
 - no information is saved in routers

Nimrod Agents

- Entity representative
 - entity attributes
 - assign locators
- Association agent
 - answer queries
 - propagate information
- Route agent
 - collect maps
 - generate routes
- Forwarding agent
 - initiate relationships
 - request routes

Example of Path Setup



Conclusions

- Feedback based routing
 - routing resources independent from network size
 - offer better convergence times
 - solves some security issues

Conclusions

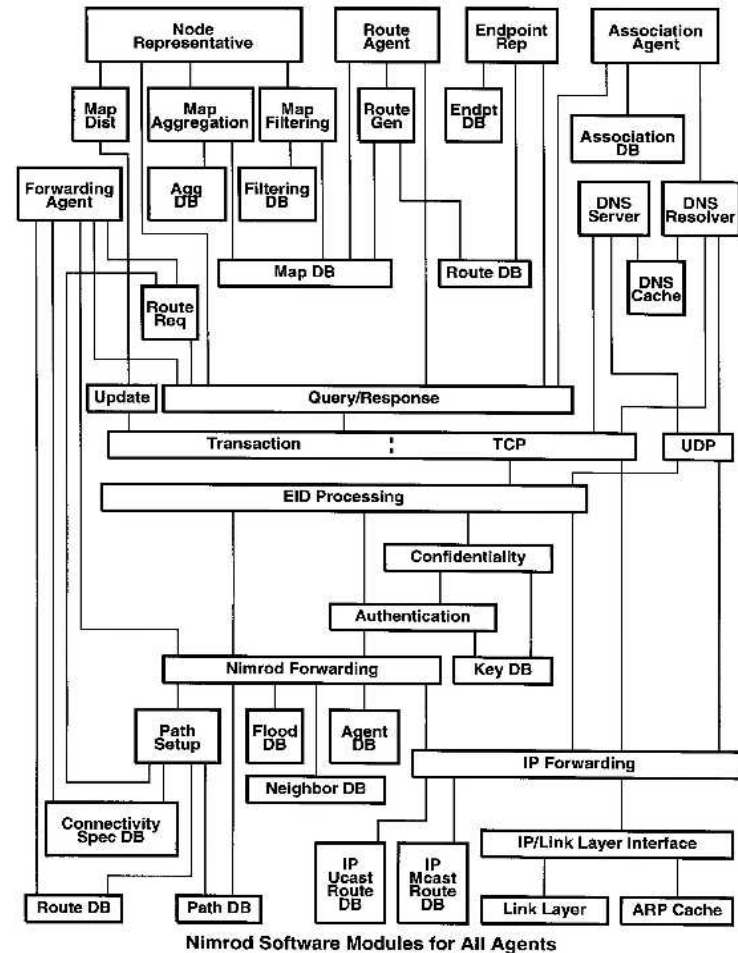
- Nimrod provides a new architecture
 - based on actual and future requirements
 - provides base for diversity (networks, services, ...)
- Transition is taken into account
 - can be integrated with actual architecture
 - changes are interoperable
 - vendor specific algorithms

Criticisms

- Feedback based routing
 - Addressing needs
 - End-to-end principle
 - Lack of support for new features

Criticisms

- Nimrod deployment
 - IP as locators and EID?
 - Global registries
 - Implementation



Questions?

References

- *Feedback Based Routing*
D. Zhu, M. Gritter, D. R. Cheriton
- *The Nimrod Routing Architecture, 1996*
I. Castineyra, N. Chiappa, M. Steenstrup
- *A Perspective on Nimrod Functionality, 1994*
M. Steenstrup; internet draft
- *A New IP Routing and Addressing Architecture*
N. Chiappa; draft

Nimrod Routing

Feedback Based Routing

Thank you