HAIR: Hierarchical Architecture for Internet Routing

Re-Architecting the Internet – ReArch ’09

Wolfgang Mühlbauer
ETH Zürich / TU Berlin
wolfgang.muehlbauer@tik.ee.ethz.ch

Anja Feldmann
Deutsche Telekom Labs / TU Berlin

Luca Cittadini
Università Roma Tre

Randy Bush
Internet Initiative Japan

Olaf Maennel
Loughborough University

Re-Architecting Internet Routing

- **Routing problems**:
  - routing table growth
  - high update rates
  - address shortage
  - mobility
  - multi-homing
  - traffic engineering
  - lack of security
  - ...

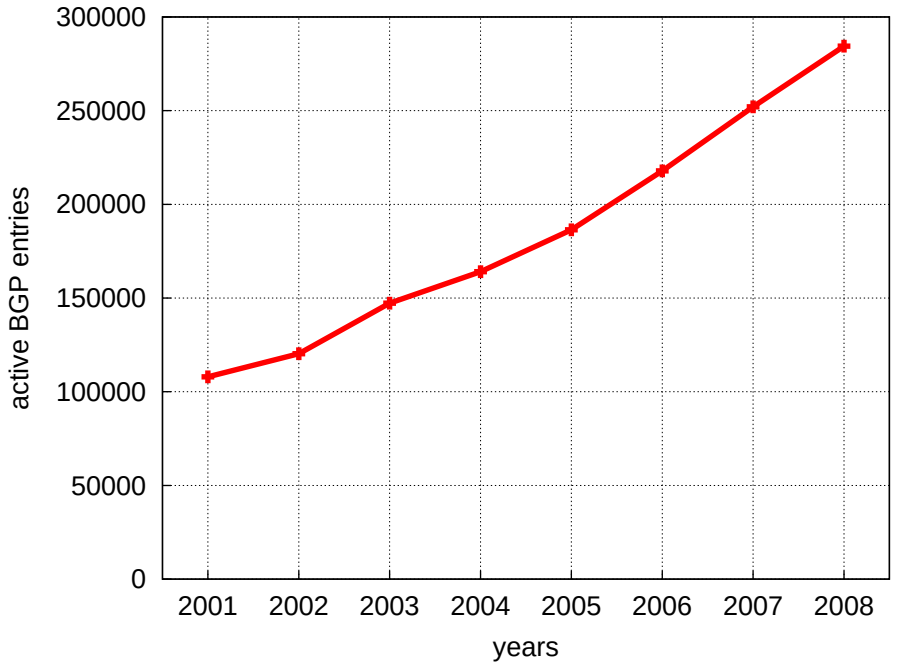

- **Clean-slate approach**: assume we could start from scratch
  - ideas may be incrementally applicable to current Internet
- **Our work:**
  - Hierarchical Architecture for Internet Routing (HAIR)

Outline

Outline

Related Work

- **LISP**
  - addressing: separate
    - *locators* (RLOC)
    - *identifiers* (EID)
  - packet forwarding:
    - map EID to RLOC
    - tunnel packet through core based on RLOC
  - multihoming easier
  - routing table size ↓
- **shim6**
  - multihoming for IPv6-enabled sites
  - hosts control which locator is used
- HLP, HIT, and many others

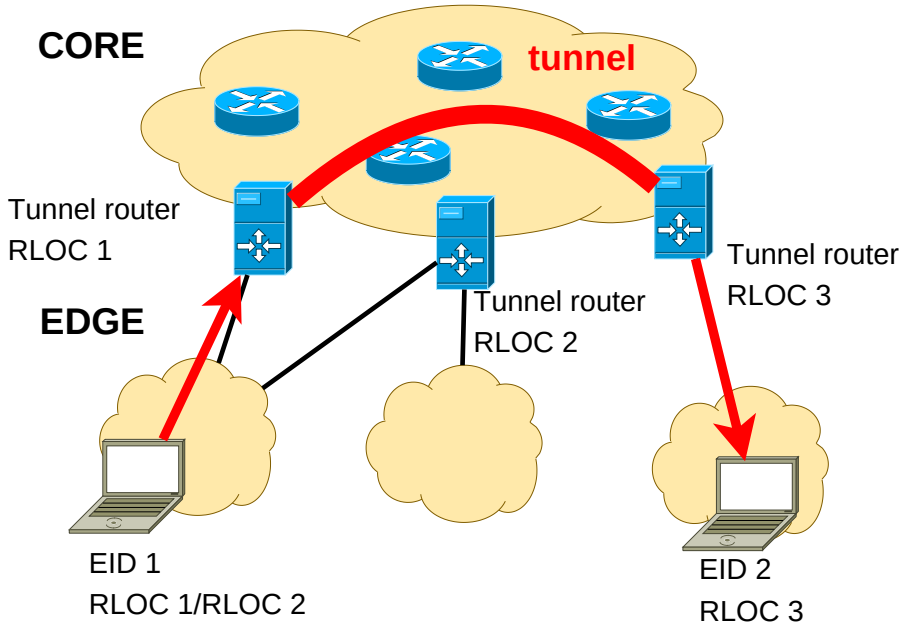

Clean-Slate: Design Choices

- **Separation of locators/identifiers (LOC/ID split)?**
  - *no*: current Internet
  - *yes*: LISP
- **Flat/structured namespaces for LOCs and IDs?**
  - *flat*: Routing On Flat Labels (ROFL), SIGCOMM'06
  - *structured*: current Internet
- **Host- or network-based solution?**
  - *host-based*: Shim6, no state in the network
  - *network-based*: LISP, mapping done at tunnling router
- ...

Outline

HAIR: Key Ideas

1. **Separation of locator/identifier** function of IP address
2. Use of **hierarchical** routing *and* mapping system
3. **Edge-based**: if possible transfer tasks to edge hosts, keep network simple

HAIR: Hierarchical Routing

- **Why hierarchical?**
  - to provide scalability
    $\rightarrow$ graph theory
- **Leverage Internet hierarchy**
  e.g., 3 levels:
  - `core`:
    large transit provider
  - `intermediate`:
    small providers
  - `edge`:
    access networks, LAN

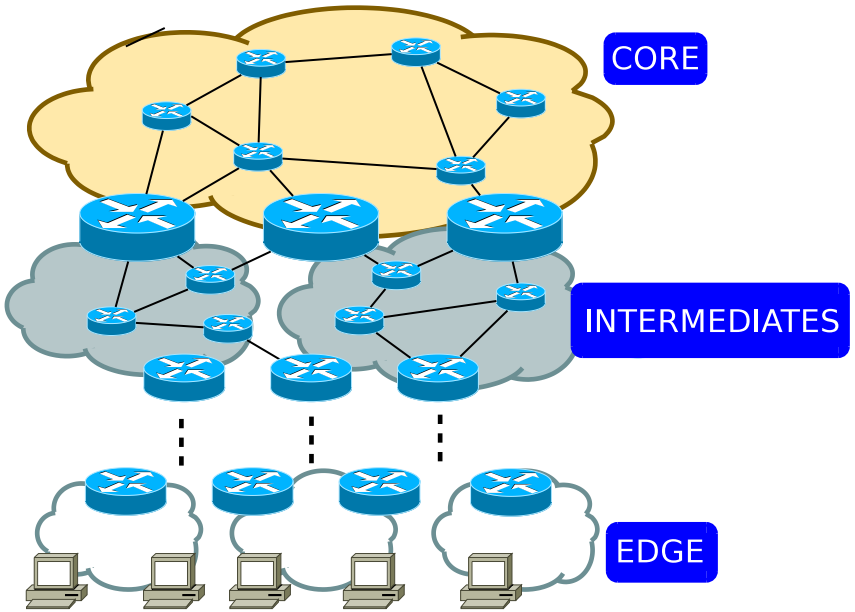

HAIR: Packet Forwarding

- **Locator**: 3 parts
  - core exit point
  - intermediate exit point
  - identifier (ID)
- **Forwarding**:
  1. send packet to `core` (direct peerings supported)
  2. forward along “exit points”
- **Local routing scope** within hierarchy levels

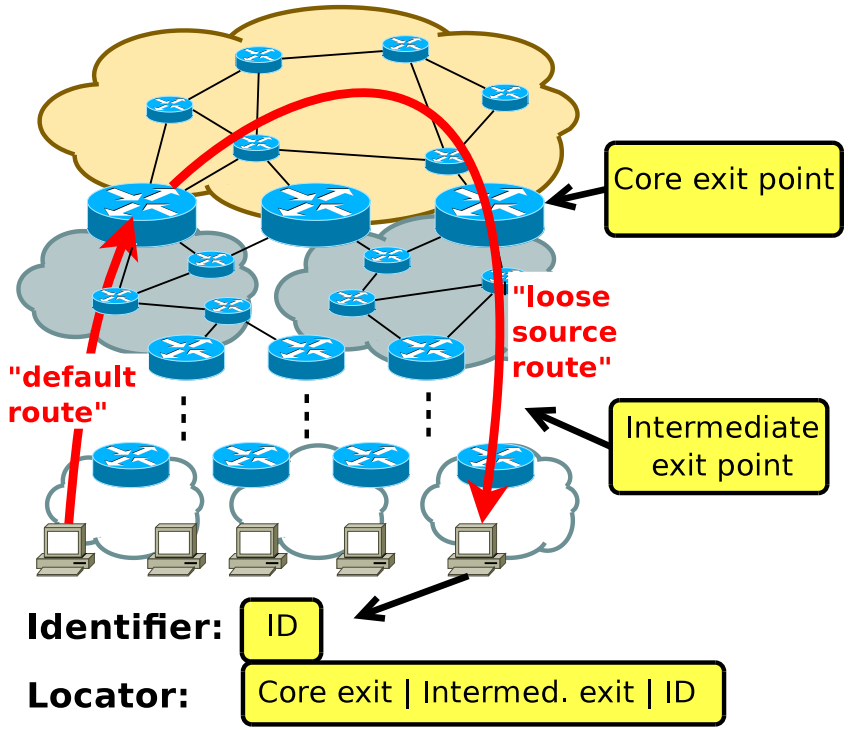

HAIR: Hierarchical Mapping System

- **Design requirements**
  - scale with number of hosts
  - fast response times
- **Hierarchical directories:**
  - *local*: `intermediates`
  - *global*: `core`
- **Resolve mapping**
  1. get pointer to local directory
  2. get actual mappings
- **Edge-based**:
  - request sent by end host
  - no action needed from e.g., `core` exit points

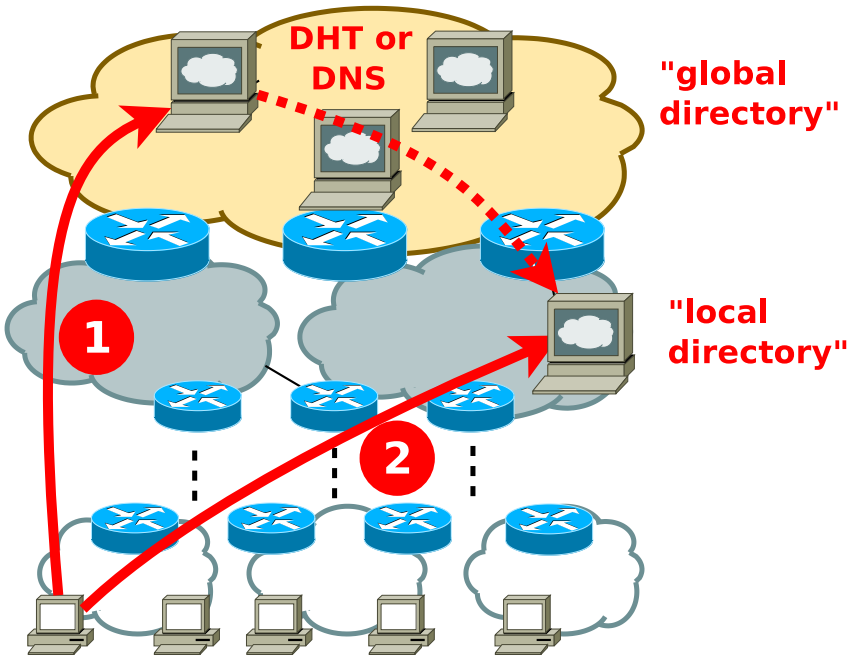

Dynamics

- **Link/router failure** inside `core` or `intermediate`:
  - find alternative route between all pairs of exit points
  - updates are localized in scope to `core` or `intermediate`
- **Failing or unreachable exit point**:
  - e.g., monitor reachability of exit points
  - update all affected locators in the mapping system
- **Change of locator**:
  - “intra-domain”: update local directory
  - “inter-domain”: update global directory, move locators

Outline

Estimating the Benefits

- **Goal**: benefits if HAIR was deployed in today's Internet?
  - how much can we scale the DFZ routing table?
  - `core` isolated from update churn originated by "edge"?
- **Data sources**
  - BGP updates and table dumps
  - classification of ASs according to business type e.g., transit provider, enterprise networks
    → Dhamdhere et al., IMC 2008

Estimating the Benefits – Results

- DFZ table size: reduction by more than a half
- Updates: majority of current updates from “edge”, see plot

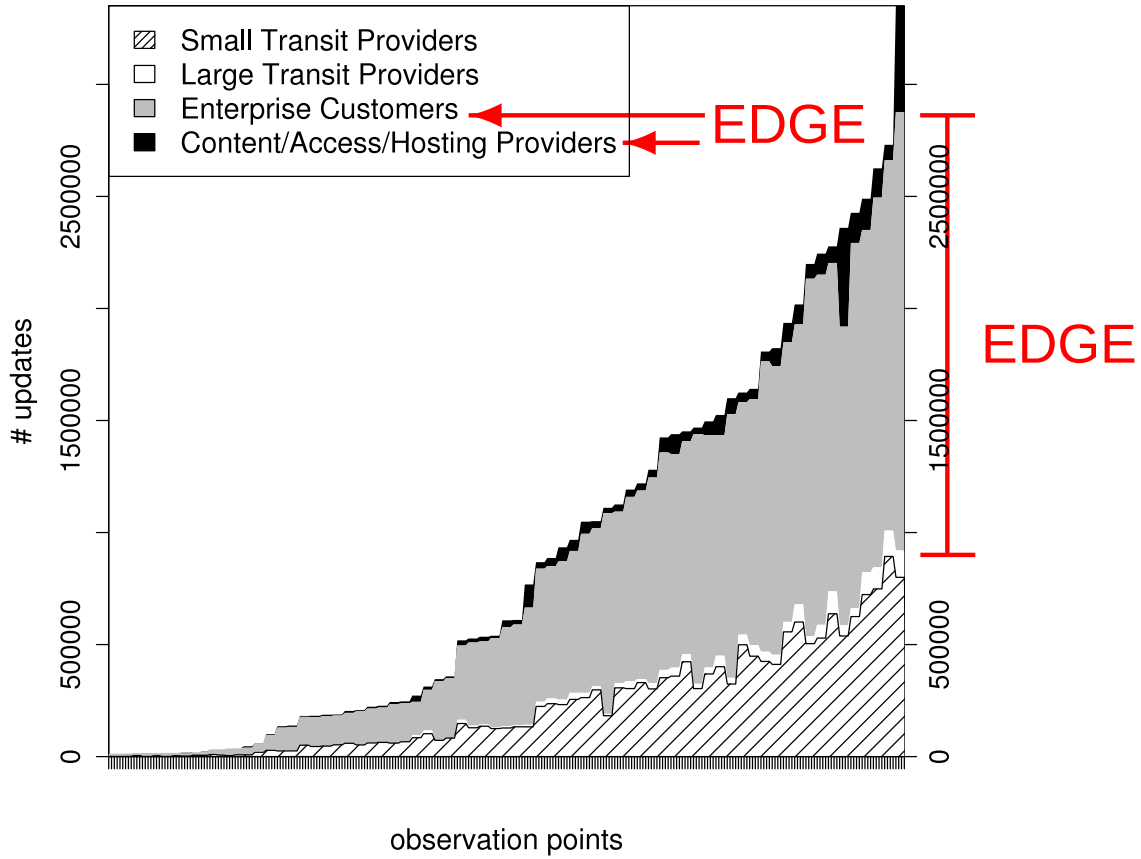

Scalability since Internet mainly grows at the “edge”

Proof-of-Concept Implementation

- **Requirements:**
  - support existing applications
  - standard IP forwarding
  - user space
  - bootstrapping
- **Use existing software**
  - *IPv6*
  - *Click*
  - *Scapy*
- **Setup in testbed**
  - latency: `ping`
  - throughput `iperf`
  - mobility scenarios

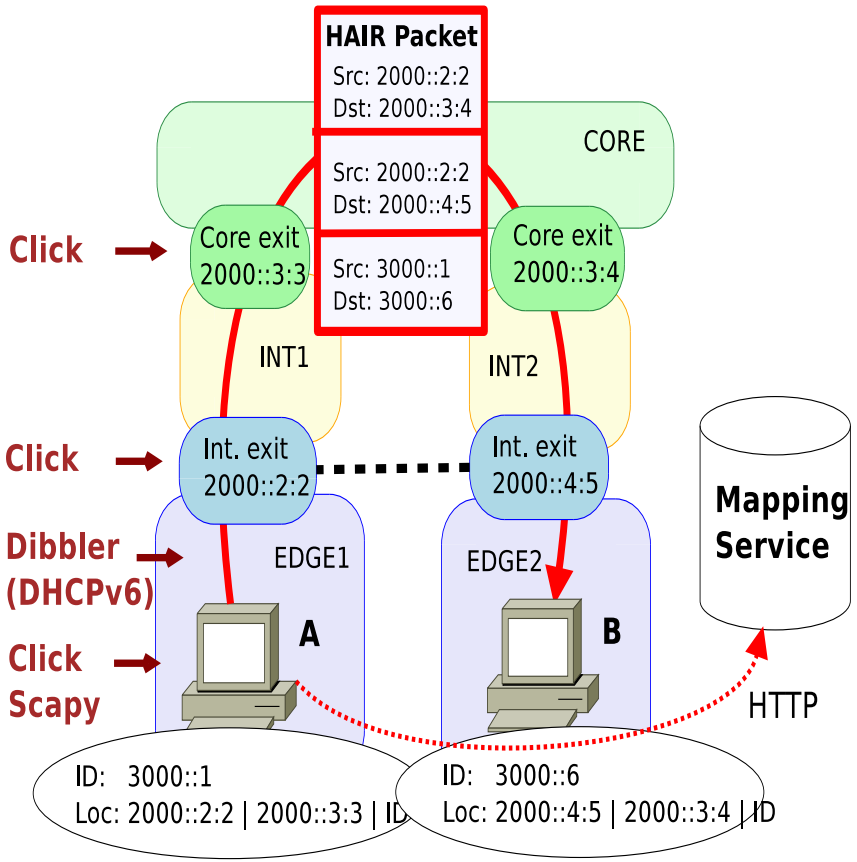

Discussion/Summary

- **Scalability**
  - routing AND mapping largely on a local scope
  - HAIR captures growth of Internet at the edge
- **Multihoming, multipath, inbound traffic engineering:**
  - can be supported by mapping system
- **Migration path**
  - support legacy hosts via NAT-like boxes

Outline

Conclusion

- **Key ideas:**
  - *separation of locator/identifier* function of IP address
  - use of *hierarchical* routing **and** mapping system
  - *edge-based*: if possible transfer tasks to edge hosts, keep network simple
- **Current status:**
  - architecture specified
  - proof-of-concept implementation demonstrates feasibility
- **Future work:**
  - mapping system
  - security model and analysis

Thank you!

Conclusion

- **Key ideas:**
  - *separation of locator/identifier* function of IP address
  - use of *hierarchical* routing **and** mapping system
  - *edge-based*: if possible transfer control to edge hosts, keep network simple
- **Current status:**
  - architecture specified
  - proof-of-concept implementation demonstrates feasibility
- **Future work:**
  - mapping system
  - security model and analysis