# Publish/Subscribe for the Internet

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

**AMICT'2010, Petrozavodsk, 2010-05-25**
**Dr. Arto Karila, Helsinki Institute for Information Technology**

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# New paradigm

**PSIRP**

- 2.5-year FP7 project
- Design, implement, and validate an entirely information-centric pub/sub-based Internet architecture through all the layers

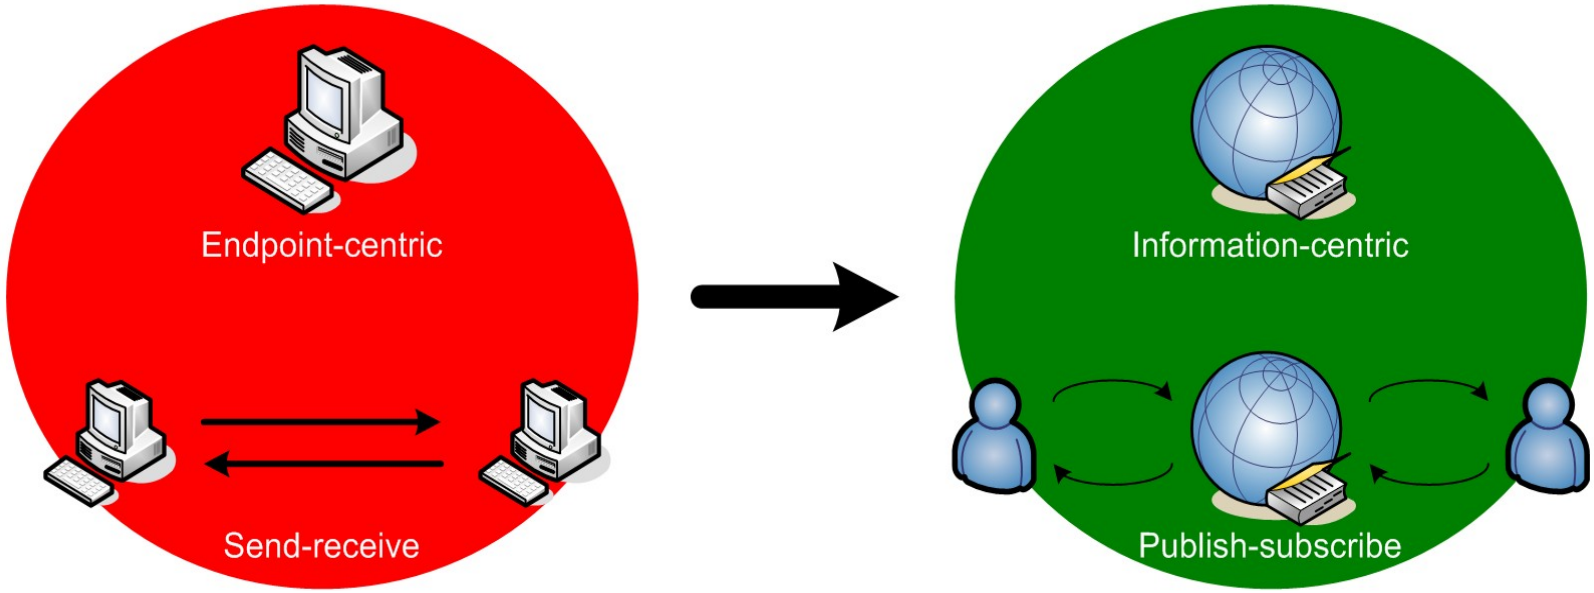


**ARPANET, current Internet, telephony etc.**

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Clean slate approach

Data published once, received multiple times

- Asynchronous multicast, *timely* separated requests
- Data delivery from caches instead of the actual source
  - Caching becomes an essential function

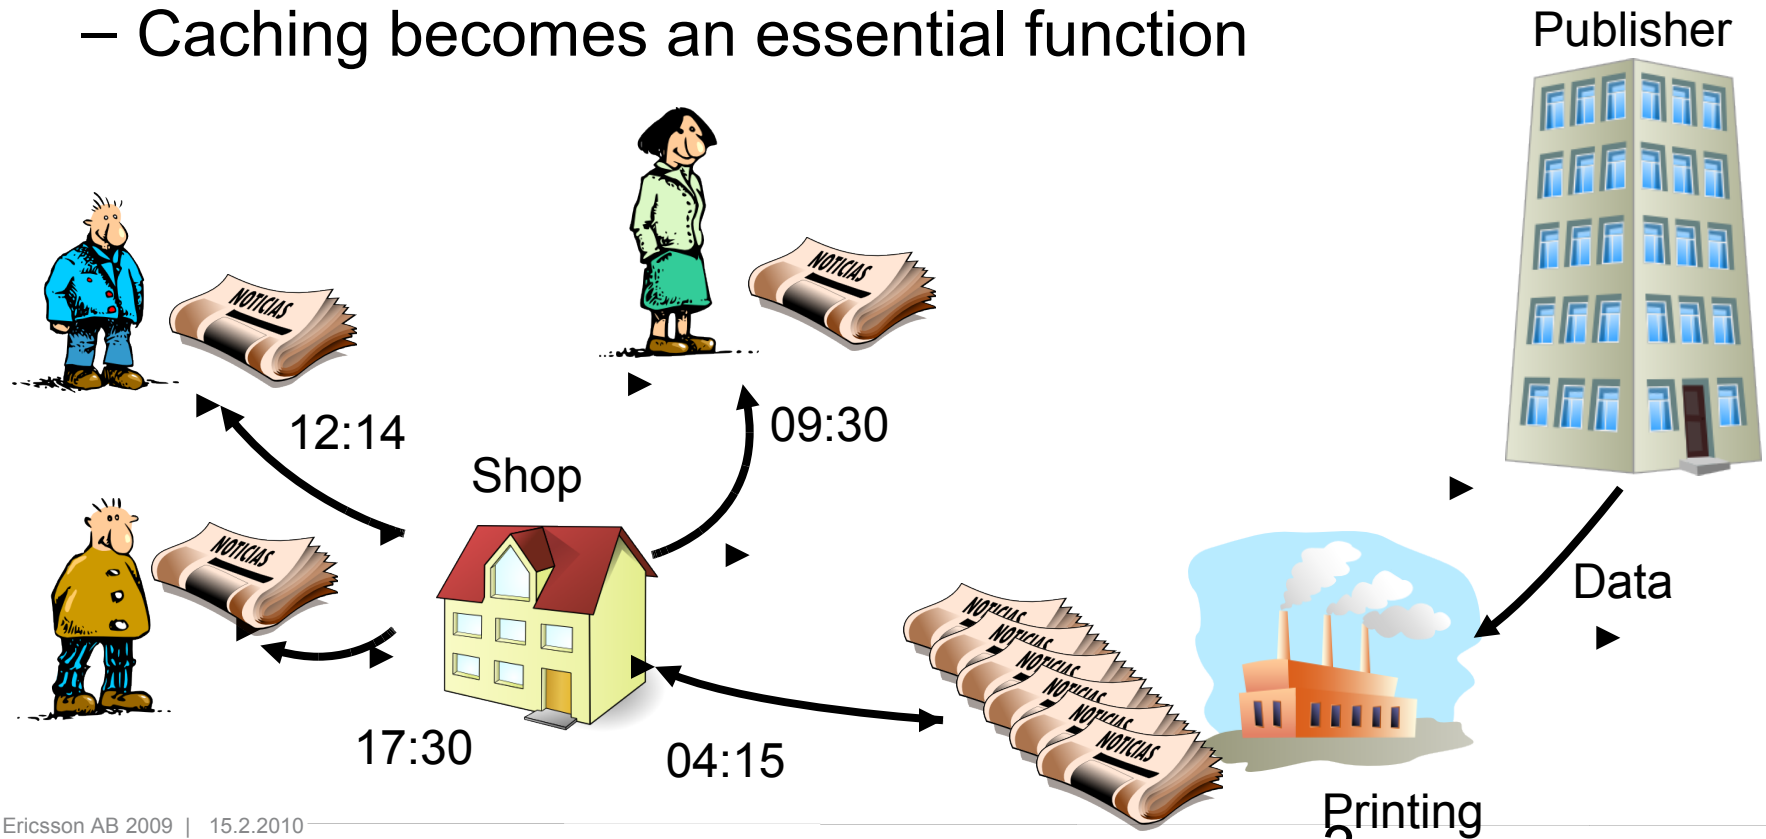

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Data-centric publish/subscribe

- The publisher and the subscriber are **decoupled** in time and space, with the publication in the middle
- Publication = persistent, immutable association between an ID and the data value created by the publisher
- Immutability of publications enables caching in the network
- Structured identifier-space
  - Streaming media with **Algorithmic Identifiers**

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# PSIRP identifiers

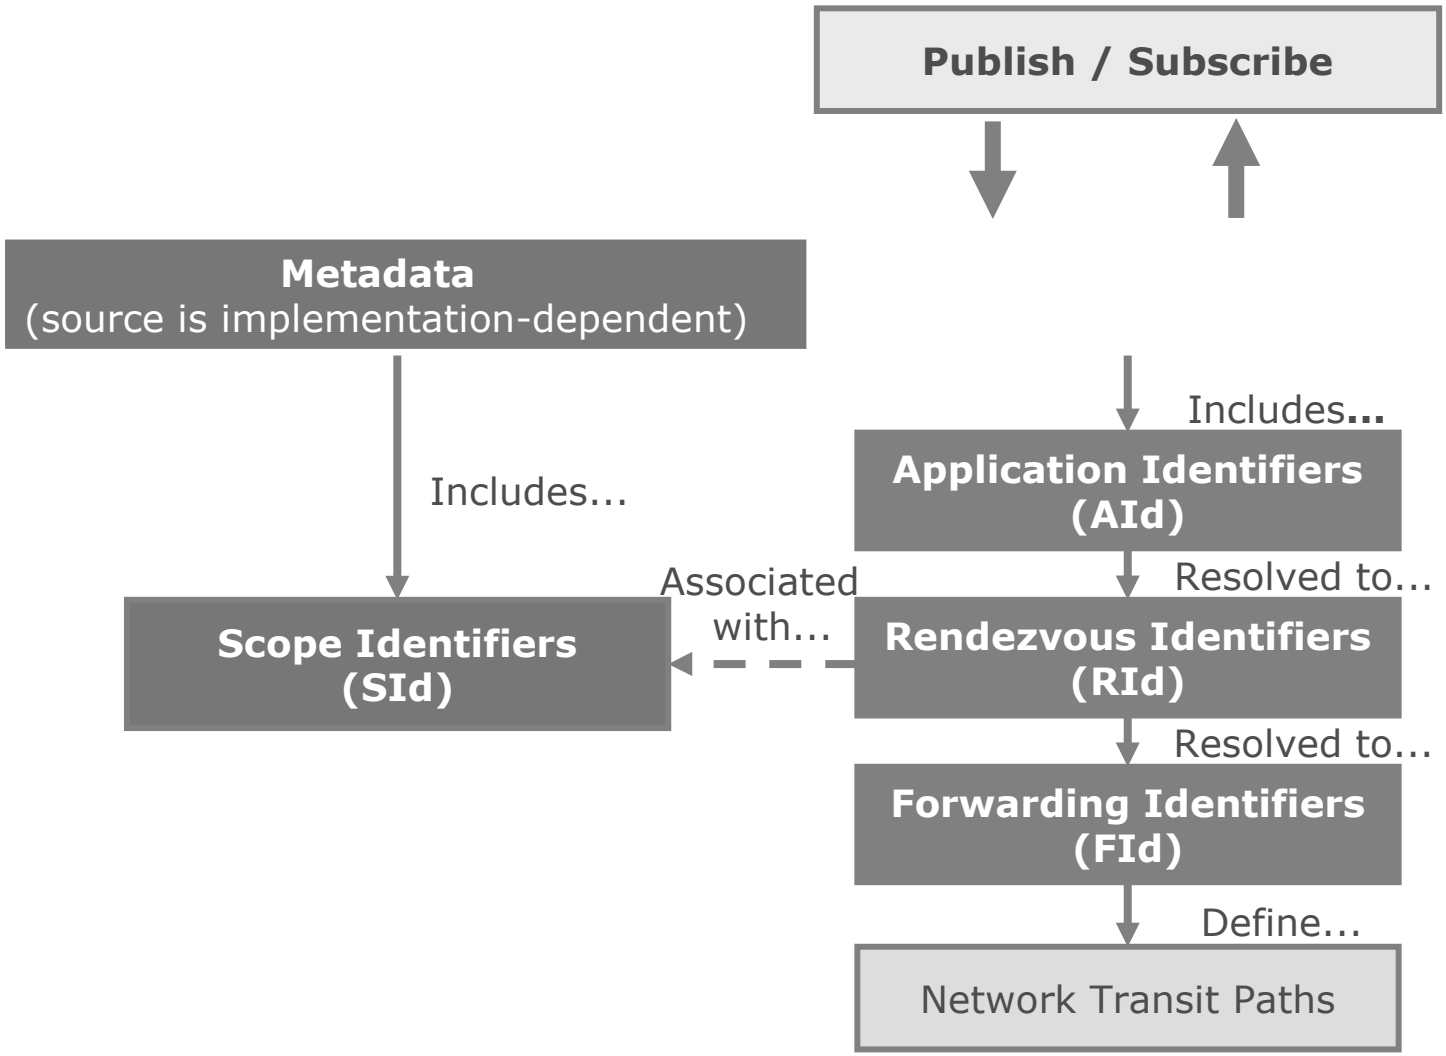

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# PSIRP component wheel

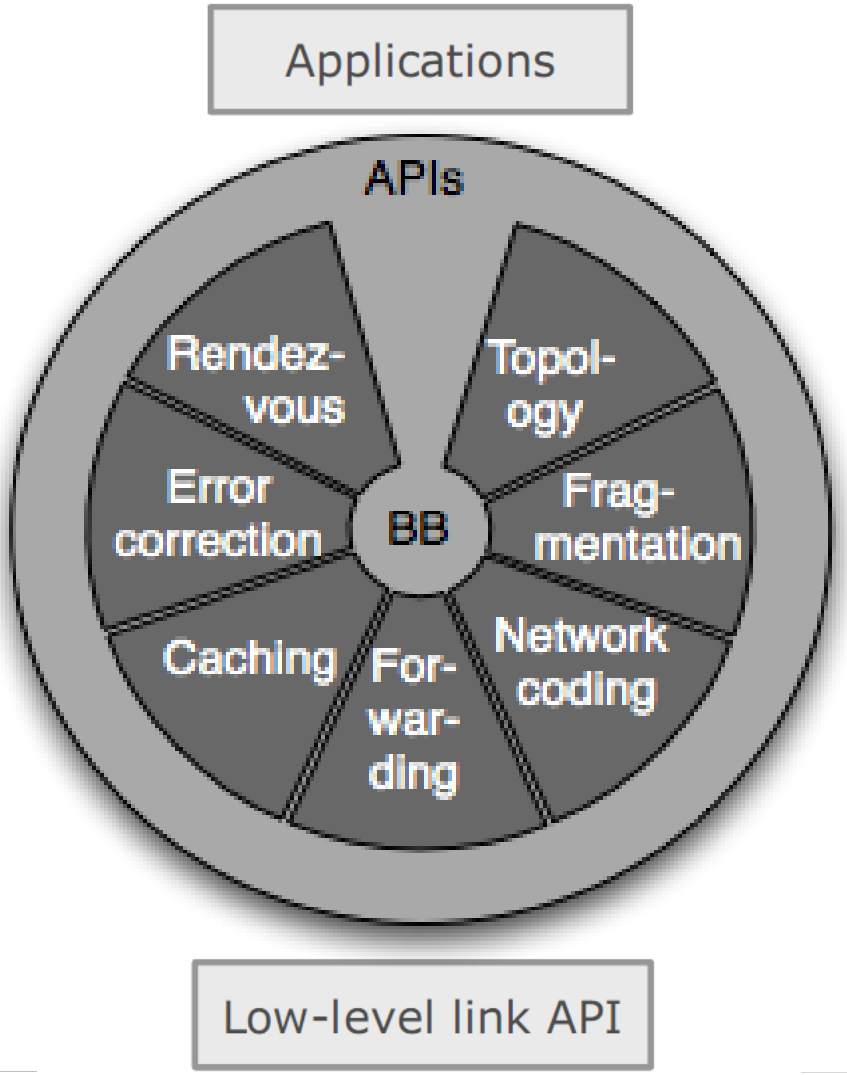

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# RTFM architecture

- Rendezvous - matching publish and subscribe events
- Topology - network topology knowledge, path creation
- Forwarding - fast delivery

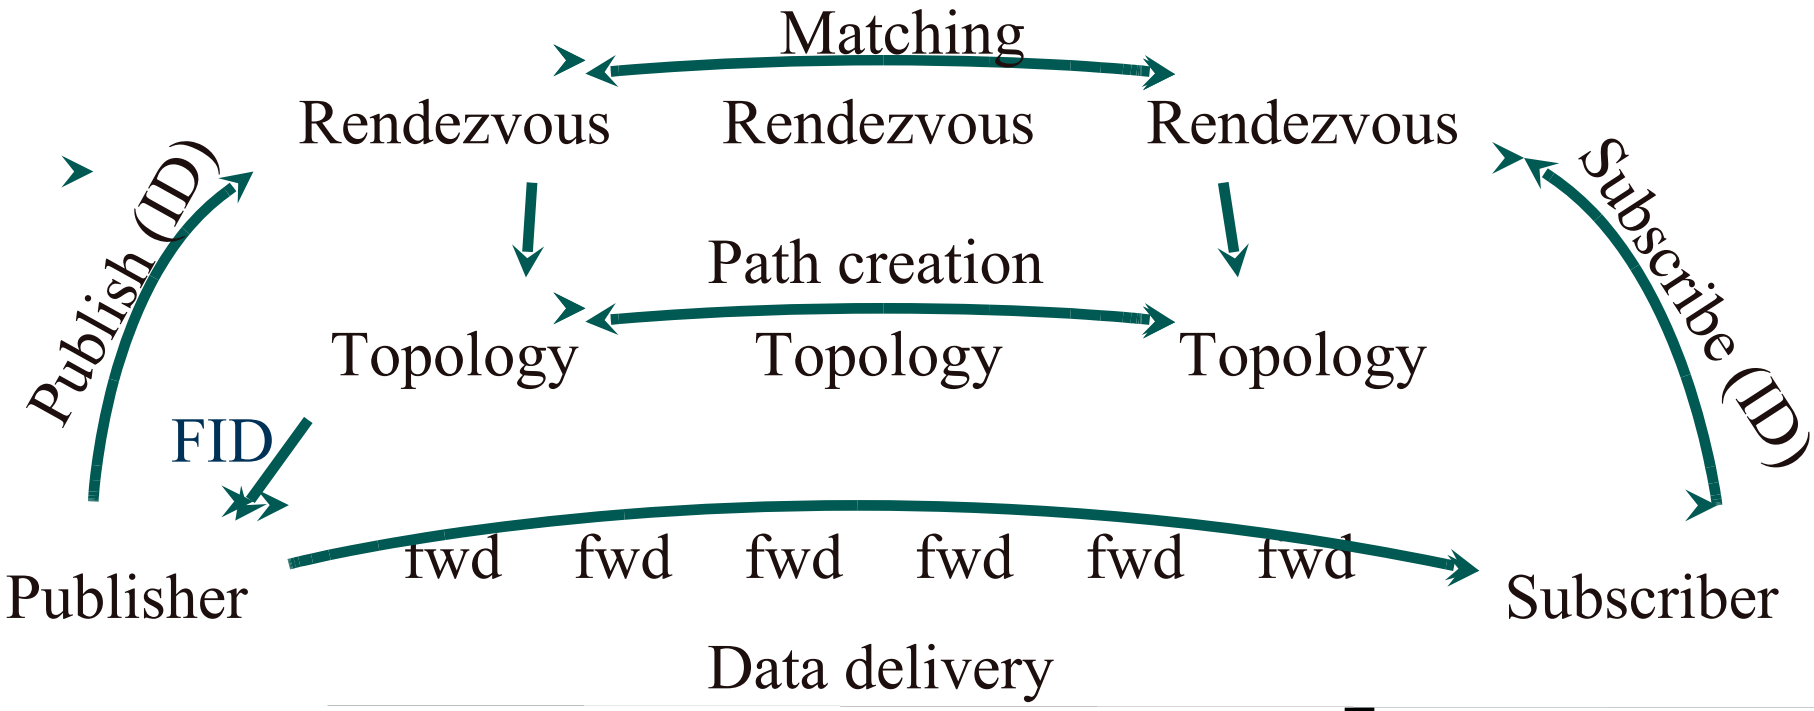

# Concrete Results

- An open-source implementation of the core architecture:
  - FreeBSD-based **BlackHawk** prototype (end node)
  - **NetFPGA** based fast forwarding and encryption
  - Working **Rendezvous Server** and **Topology Manager**
- Some test applications, including:
  - **Multicast BitTorrent**
  - **Firefox plug-in**
  - **Socket emulator**
- All of the components above are already working but the integration is still not complete
- An integrated system, including some applications, will be demonstrated at ICT 2010 in Brussels in September

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Implementation

- NS3 simulator
- FreeBSD 7.x : end-host + forwarding
- NetFPGA : Forwarding
- BSD & NetFPGA Implementations available at http://www.psirp.org

# Testbed

- The PSIRP implementation is validated on a relatively large and real-life-like testbed
- The University of Cambridge and the University of Essex are testing PSIRP over an optical connection between them
- VPN connections will be created to the other partners: Aachen, Athens, Helsinki and Sofia
- PSIRP is cooperating with OneLab2 under FIRE
- The University of Essex is integrating the Blackhawk with campus-wide wireless network (using a dedicated SSID) accessible to 2500+ students in their dorms
- Helsinki University of Technology is running a code camp where students develop applications on the BlackHawk
- Interested parties are invited to join in the testbed – the hardware is inexpensive and the software is downloadable at http://www.psirp.org

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Experimentation

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Bloom filters – basic idea

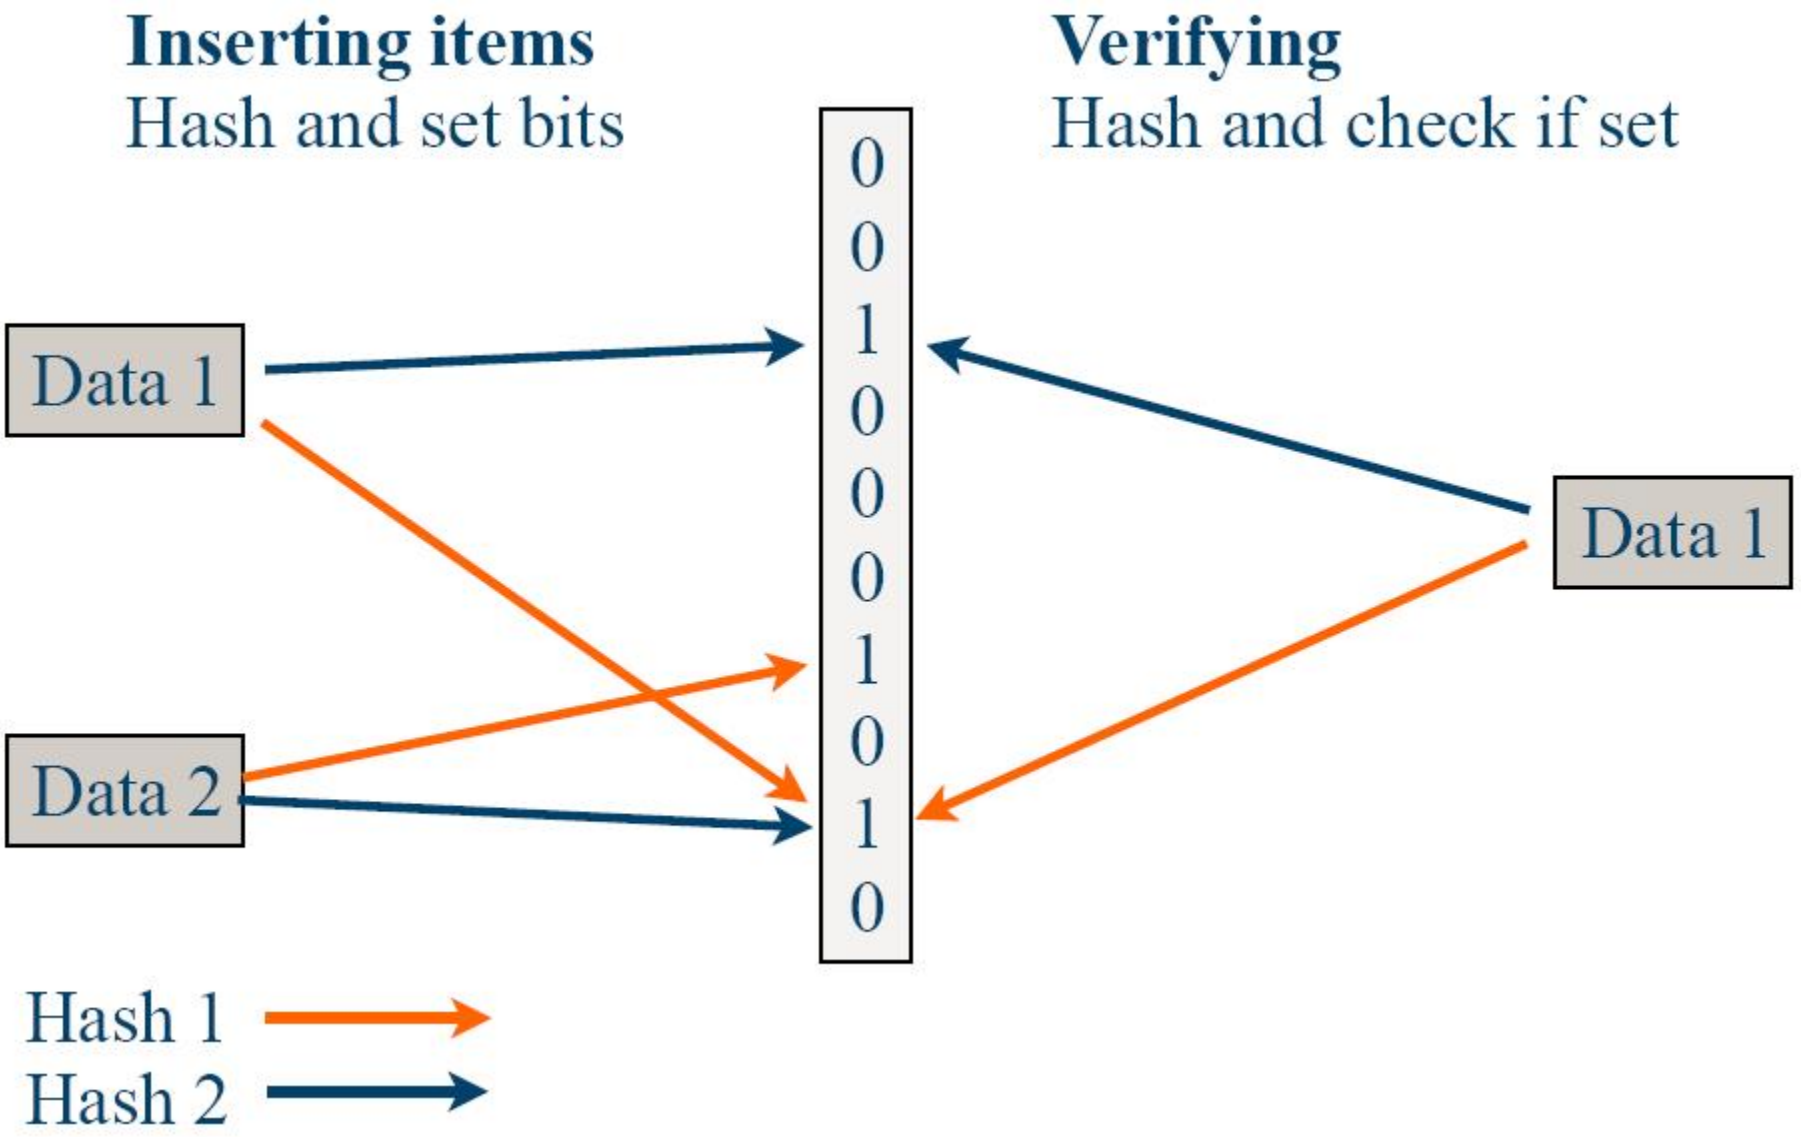

# Link IDs and forwarding Bloom filters (zFilters)

- No names for nodes
  - Each link is identified with a unidirectional Link ID
- Link IDs
  - Statistically unique
  - Periodically changing
  - Size e.g. 256 bits
  - Local or centrally controlled
- Source routing
  - Include all Link IDs into a Bloom filter
  - Multicasting supported
- “Stateless”

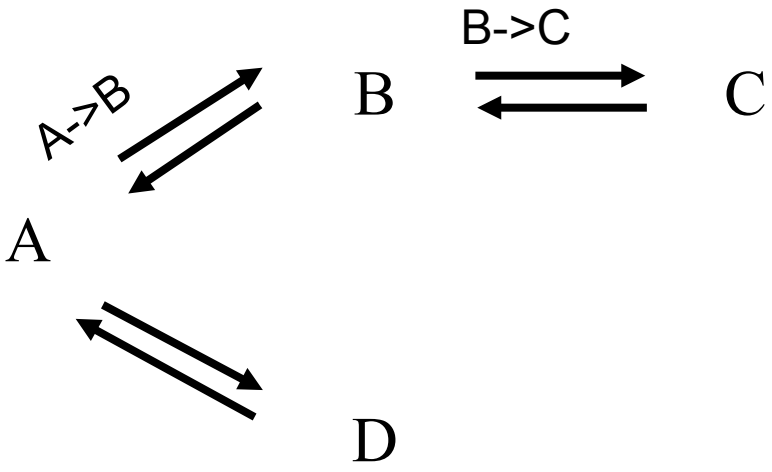


| | |
|---|---|
| A->B | 0 1 0 0 0 1 0 0 1 |
| B->C | 1 0 0 0 0 1 1 0 0 |
| zF: A->B->C | 1 1 0 0 0 1 1 0 1 |

# Forwarding Decision

Forwarding decision based on binary AND and CMP

- zFilter in the packet matched with all outgoing Link IDs
- Multicasting: zFilter contains more than one outgoing links

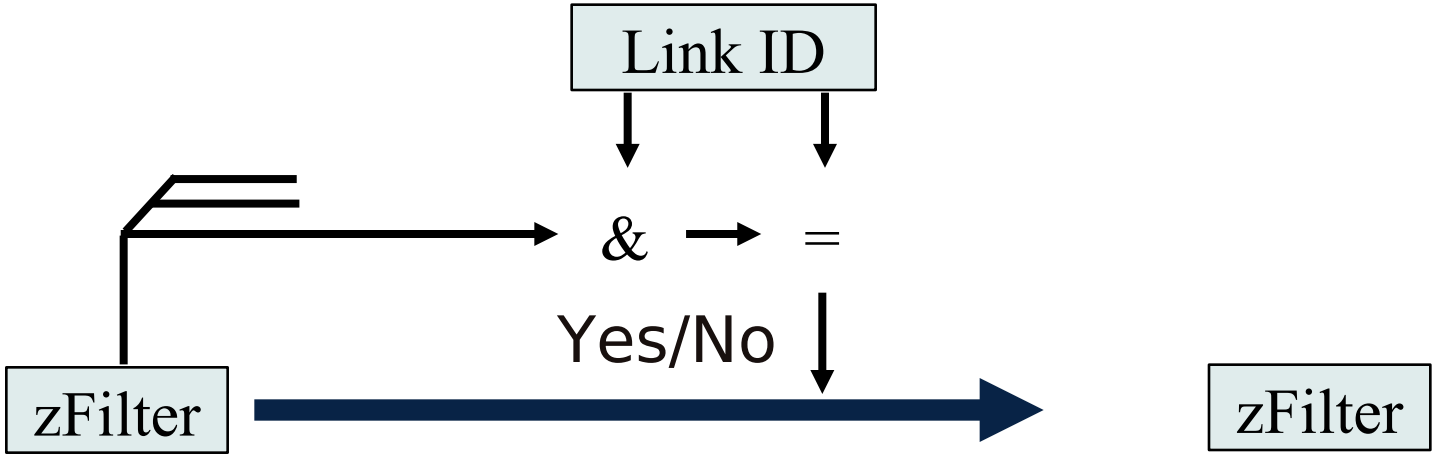

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

# Using Link Identity Tags (LIT)

- Make results better with a simple trick
  - Define d different LITs instead of a single LID
  - LIT has the same size as LID, and also k bits set to 1
  - [Power of choices]
- Route creation and packet forwarding
  - Calculate d different candidate zFilters
  - Select the best performing zFilter, based on some policy

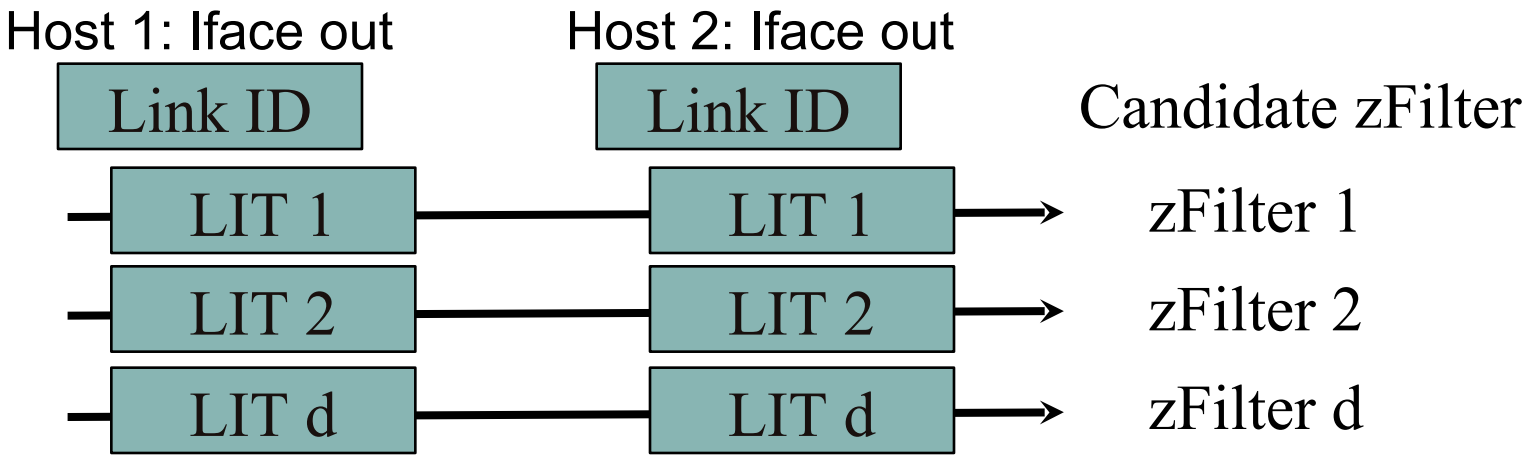

# Virtual Trees

- Popular paths can be merged into virtual trees
  - A single Link ID for the tree
  - Additional state in the forwarding nodes
  - Increase scalability

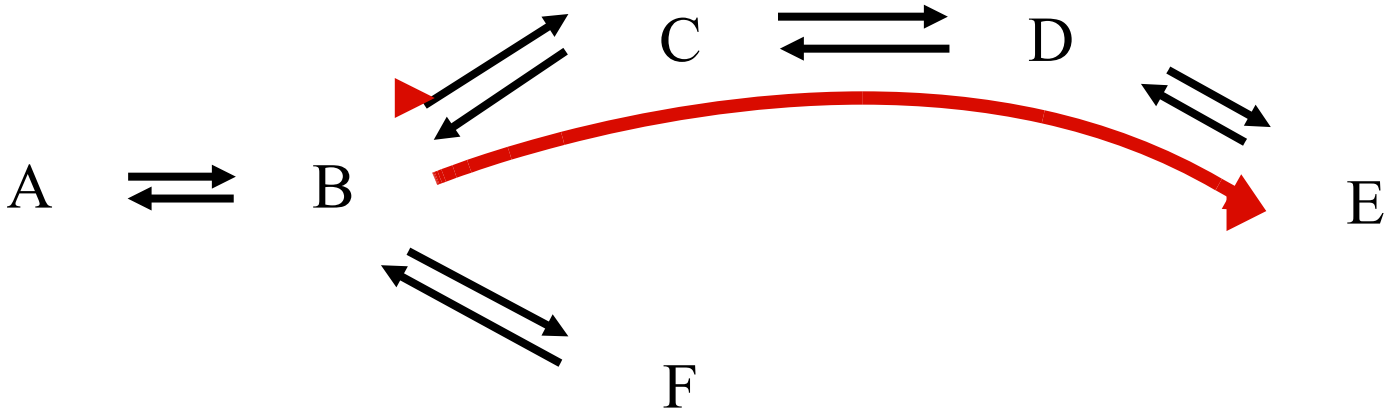


Virtual B->C->D->E  0 0 1 0 1 0 0 0 1

# Summary of LIPSIN and zFilters

- Link-identity-based source routing
- Stateless small-group multicast and unicast
- Small forwarding table
- Very simple forwarding decision
- Preventing unwanted traffic
  - No possibility to send data by guessing the destination
- Forwarding implemented in both software and hardware


# Conclusions

- During the 2.5 years we have defined, implemented and tested an entirely pub/sub based internet architecture
- The system appears to work, scale, and offer some clear benefits over conventional end-point-centric networking
- Because of the iterative ("life-cycle") mode of working, the current prototype is not the latest version of the architecture
- In the new PURSUIT project we'll expand the work towards:
  - Applications
  - Wireless
  - Optical
- We need human-understandable naming
- Academic dissemination appears to be the way to ultimately get these new ideas into deployment