HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Receiver Driven Content Routing for the Internet

20.10.2010

Prof. Sasu Tarkoma

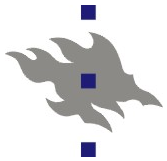

Introduction

Network endpoints are more and more interested in content

Publish/subscribe is a candidate paradigm for content distribution

Many-to-many, control distributed between subscribers (receivers) and publishers (senders)

Inherently receiver driven

A content distribution system needs to know about packet replication points and policy-compliant paths

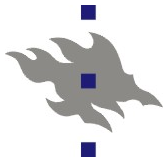

Background

Related systems include NIRA, DONA, ROFL, CCN, PSIRP, and PURSUIT

PSIRP includes a 2-tier system where a hierarchical DHT based rendezvous interconnect network joins multiple rendezvous networks together for global reachability

- Publications are identified by scopes and flat identifiers
- Typically only scopes are advertised in the interconnect
- Hierarchical structure guarantees locality for the communication

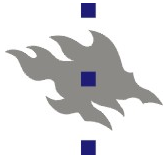

	Layer	Underlying transport	API	Namespace	Self-certification	Receiver driven	Key application
DONA	L4-L7	TCP/IP	Anycast	Flat	Yes	Yes, fetch operation	Content discovery and delivery
CCN	L3	Unreliable data transport, flow control	Name-based (with leases)	Hierarchical	Yes	Yes	Content delivery, voice
PSIRP	L2-L7, layerless	Ethernet, TCP/IP and PLA	Pub/sub and metadata-based	Various, recursive	Yes	Yes	Content delivery
Internet Indirection Infrastructu-re (i3)	L7	TCP/IP	Trigger-based	Flat	Yes	Yes (also sender driven)	Various
Haggle	Layerless	Packet-based	Metadata-based	ADU, user level names	Possible	Receiver/ sender	Mobile environment
Siena	L7	TCP/IP	Pub/sub	Content-based	No	Yes	Content delivery

PSIRP Architecture

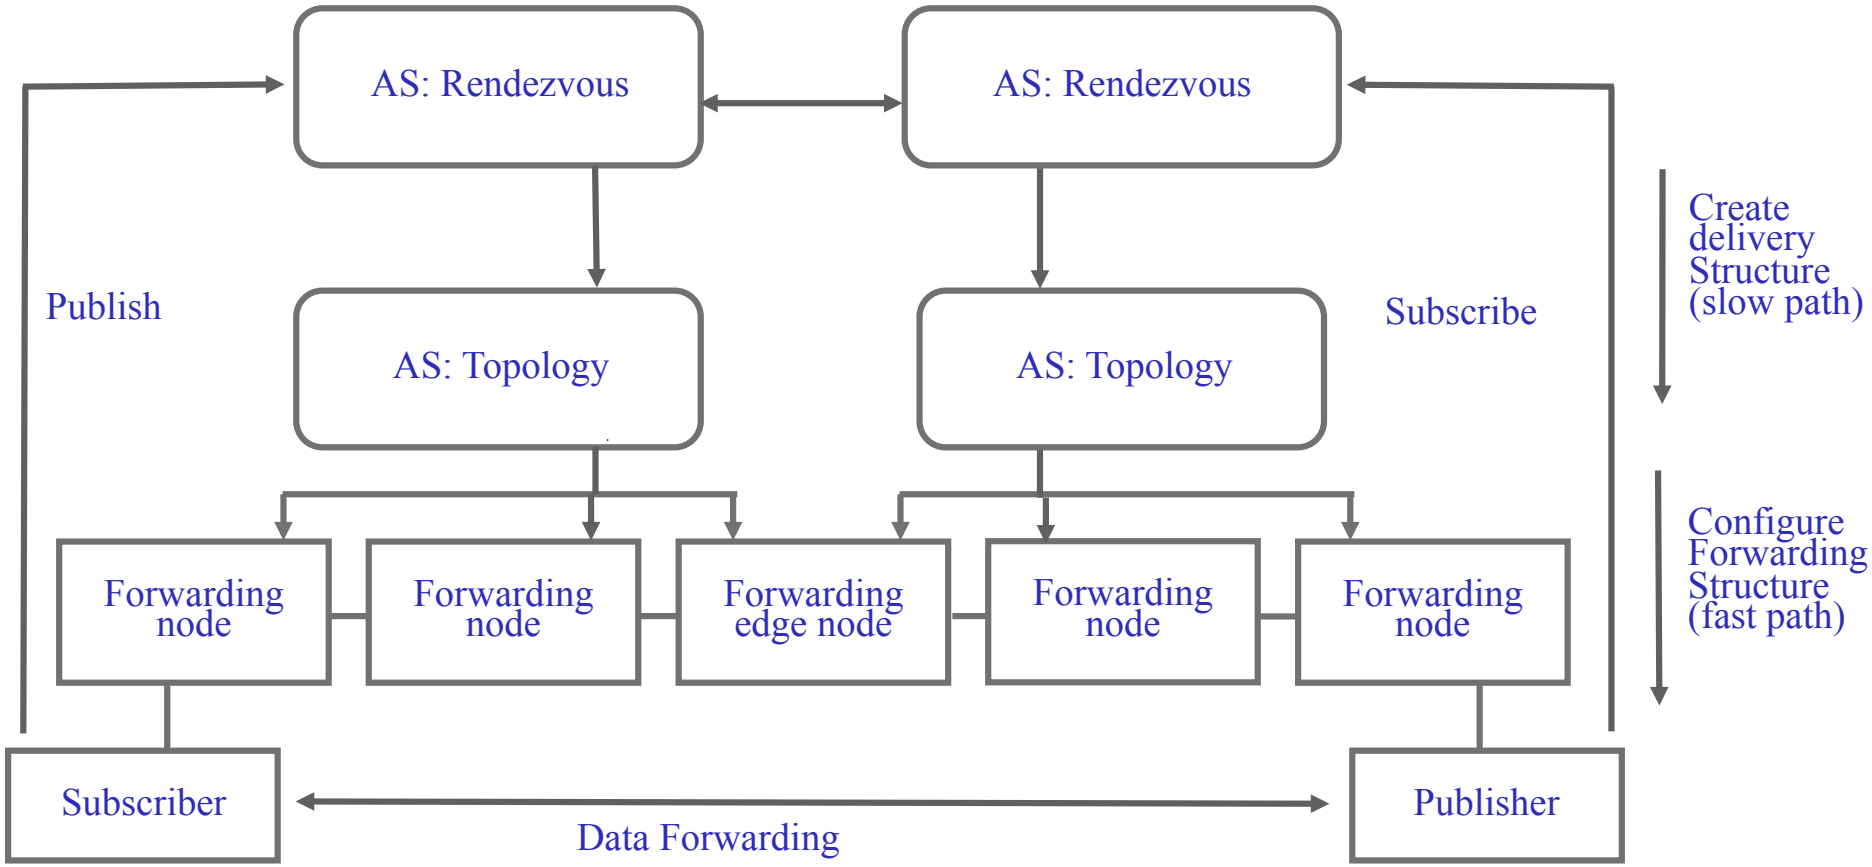

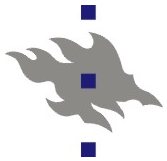

Canopy: Using Upgraphs for Pub/Sub

The NIRA system used upgraphs for
- Allowing more control for receiver
- Finding best paths for unicast

Canopy uses upgraphs for pub/sub
- Upgraphs combined at publisher-side rendezvous point
- Can take both subscriber & publisher policies into account
- Supports multi-path routing
- Result is a policy-compliant multicast structure
- Can be used for both overlays and on the network layer
- Works with in-packet Bloom filter-based forwarding

Reference: S.Tarkoma, M. Antikainen. Canopy: Publish/Subscribe with Upgraph Combination. Global Internet Symposium 2010.

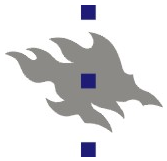

Canopy Overview

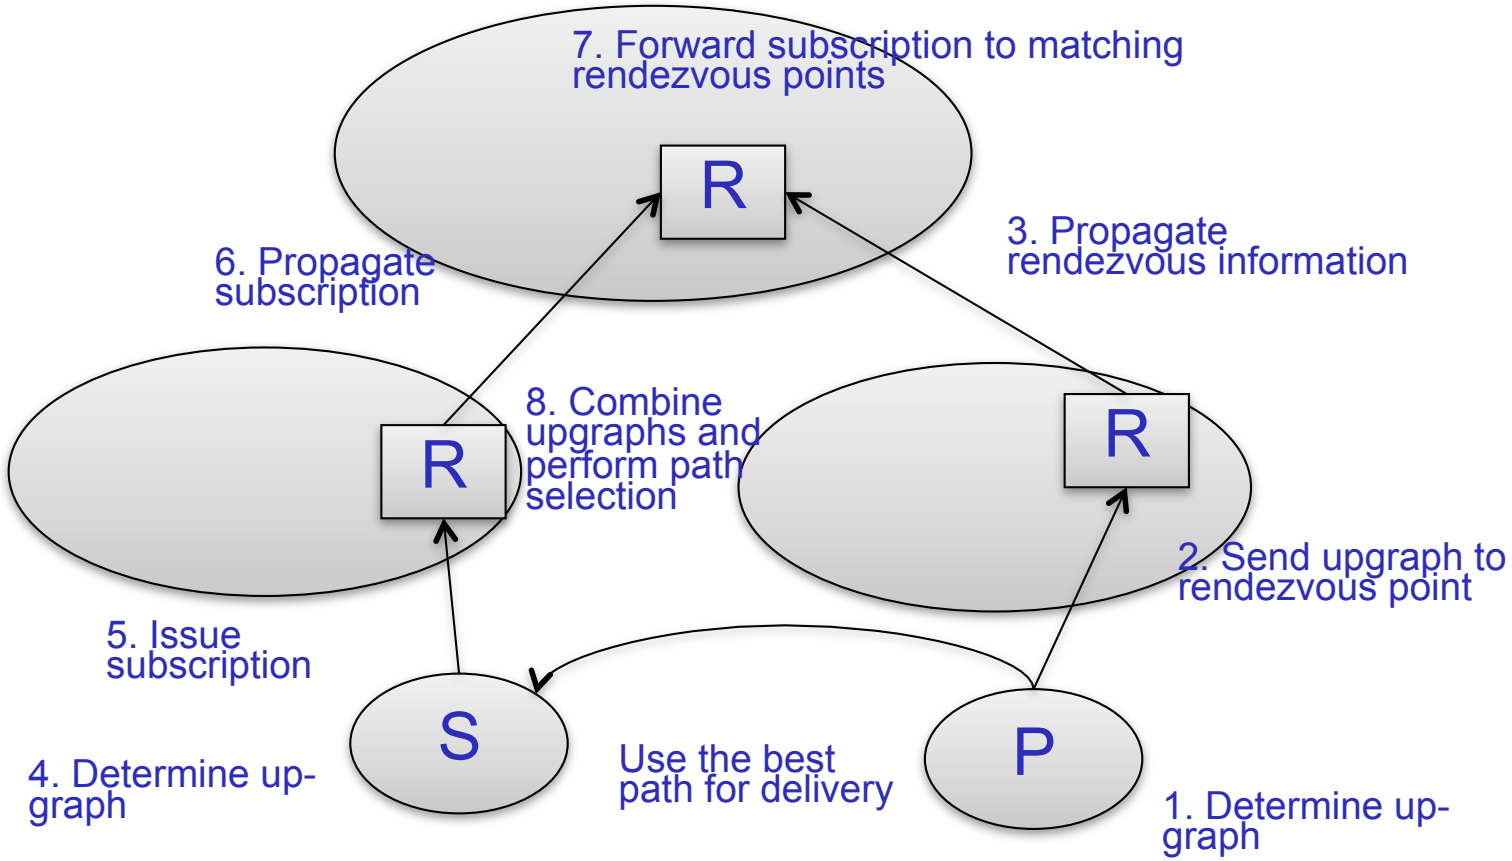

An **upgraph** is a dag that contains all the possible paths from a given node to the tier-1 networks

Publisher upgraph roots from the publisher and contains all possible uphill and peer-to-peer paths

Subscriber upgraph does not contain peer-to-peer links, but only all possible uphill paths

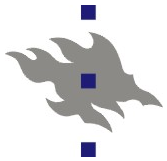

Conclusions

Receiver driven communications has favourable properties

- Pub/sub is inherently receiver driven

Research question: How to make content supply and demand meet in the network and how to make it scale?

Canopy uses upgraphs for determining policy-compliant multicast structures for publish/subscribe

- Representing packet replication points and network paths in upgraphs
- Combining subscriber and publisher upgraphs at designated rendezvous points in order to determine policy-compliant data distribution strategies