

**ΟΙΚΟΝΟΜΙΚΟ
ΠΑΝΕΠΙΣΤΗΜΙΟ
ΑΘΗΝΩΝ**



**ATHENS UNIVERSITY
OF ECONOMICS
AND BUSINESS**

Information-Centric Networks

Section # 10.3: Publish/Subscribe

Instructor: George Xylomenos

Department: Informatics




Funding

- These educational materials have been developed as part of the instructors educational tasks.
- The **“Athens University of Economics and Business Open Courses”** project only funded the reformatting of these educational materials.
- The project is being implemented as part of the Operational Program “Instruction and Lifelong Learning” and is co-financed by the European Union (European Social Fund) and national funds.



Licensing

- These educational materials are subject to a Creative Commons License.



Week 10 / Paper 3

- Illustrating a publish-subscribe Internet architecture
 - Nikos Fotiou, Dirk Trossen, George C. Polyzos:
 - Telecommunication Systems (to appear)
- Main point
 - PSIRP = Publish-Subscribe Internet Routing Paradigm
 - Totally clean slate architecture
 - Supports availability, security and mobility
 - Opportunities for innovative applications

Introduction

- The Internet has not changed a lot since its inception
 - Still reminiscent of the telephone network
 - Connect two endpoints via the network infrastructure
- User needs have changed a lot in the meantime
 - Endpoints do not trust each other
 - Receivers cannot avoid sender traffic
 - Content is delivered via CDNs and P2P applications
 - Applications do not care which endpoint provides the content
 - Middleboxes try to fix each problem separately
- The publish-subscribe paradigm
 - Centered on information, not endpoints
 - Receivers only get what they asked for (via subscriptions)

Introduction

- Components of a publish-subscribe system
 - Publishers: feed information to the system
 - Subscribers: consume published information
 - Brokers: forward data from publishers to subscribers
 - Rendezvous Nodes: match publications with subscriptions
 - Initiation of delivery from publishers to subscribers
- Advantages of publish-subscribe
 - Publishers and subscribers are decoupled in time and space
 - Publishers are not aware of subscribers and vice versa
 - Publications and subscriptions are not synchronized
 - Multicast can be taken advantage of
- PSIRP: EU funded project that builds a pub-sub system
 - Continued by PURSUIT

The PSIRP architecture

- Publications
 - Information becomes available via publications
 - Any size, from data chunks to entire movies
 - Identified by a Rendezvous Identifier (RId)
 - Flat and unique within a scope
- Scopes
 - Containers for publications
 - Any type of grouping, from social networks to network areas
 - Identified by a Scope Identifier (SId)
 - Flat and unique within a scope
 - Hierarchical organization of scopes
 - Access control is based on scopes
 - Who can issue or access publications within a scope?

The PSIRP architecture

- Publishing data
 - A RId is derived by an application specific function
 - The proper SId is selected to organize and control information
 - The responsible Rendezvous Node for the scope is contacted
 - This node is called the Rendezvous Point (RP) for the scope
 - The publication metadata is sent to the RP
 - Metadata can indicate that a publication is divided into pieces
 - The RId's for the pieces may be algorithmically related
- Subscribing to data
 - Need to know the SId/RId of the publication
 - The RP for the scope is contacted
 - The RP initializes forwarding from publisher to subscriber

The PSIRP architecture

- Forwarding in PSIRP
 - A topology formation function creates delivery paths
 - A path is denoted by a Forwarding Identifier (FId)
 - The RP makes sure the FId is sent to the subscriber
 - The Publisher sends the publication with the FId
 - Brokers forward the publication using the FId
- Caching in PSIRP
 - Any node can cache publications it is forwarding
 - It then becomes a source for these publications
- Multicasting
 - When multiple subscribers exist, publications are multicast

Current solutions

- Note: this reflects the current PURSUIT status
- Rendezvous
 - Use of a hierarchical DHT to assign scopes to RPs
 - Ideally the RP should be close to the publisher
 - DHTs have a tendency to violate routing policies though!
- Forwarding
 - Use of the LIPSIN source routing mechanism
 - Unicast is a special case of multicast
 - How are interdomain paths handled?
- Topology management
 - Link-state protocol (like OSPF) for topology discovery
 - Dijkstra's algorithm for intra-domain shortest paths
 - How do multiple domains co-ordinate routing?

Application development issues

- Many open issues at the fundamental level
- Publication identifiers
 - Not obvious how the RId can be determined
 - It has to be unique within a scope (or more!)
 - How will algorithmic identifiers work?
 - Does the network know how they are generated?
- Publication scoping
 - How visible are scopes to applications?
 - They may be used only at a low level
- Trust in transactions
 - Which nodes can be trusted to provide services?
 - Which RP is the correct one?

Development status

- Two prototypes available
 - Blackhawk: node prototype for FreeBSD
 - Blackadder: network code prototype for Linux
 - Both migrating to use the Click modular router framework
 - Click is portable and can run in user or kernel space
 - Both available on an open source basis
- Many things are being tested
 - Rendezvous: different DHTs for inter-domain rendezvous
 - Forwarding: inter-domain and large group solutions
 - Topology: inter-domain topology formation
 - Transport: pull-based protocols with multiple sources
 - Caching: on-path and off-path solutions
 - Mobility: exploit caching to hide disconnections

**ΟΙΚΟΝΟΜΙΚΟ
ΠΑΝΕΠΙΣΤΗΜΙΟ
ΑΘΗΝΩΝ**



**ATHENS UNIVERSITY
OF ECONOMICS
AND BUSINESS**

End of Section # 10.3

Course: Information-Centric Networks, **Section # 10.3: Publish/Subscribe**

Instructor: George Xylomenos, **Department:** Informatics

