



# Accelerating File Downloads in Pub/Sub Internetworking with Multisource/Multipath Transfers

I. Thomas, C. Tsilopoulos,  
G. Xylomenos, G.C. Polyzos

Mobile Multimedia Laboratory,  
Athens University of Economics and Business

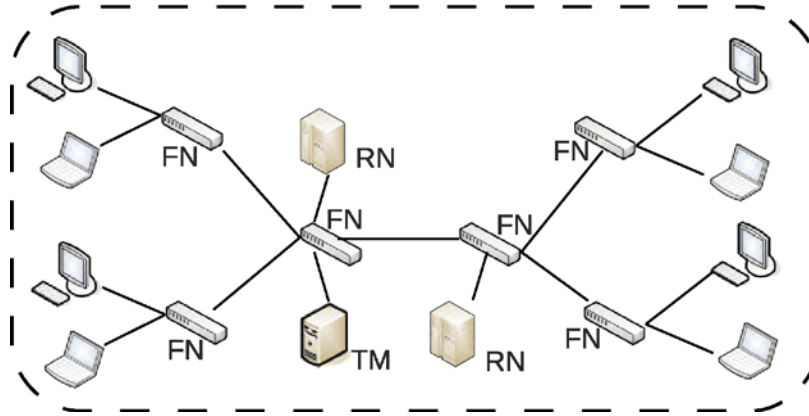
# Outline

- Motivation
- PSI Basics
- Design Choices
- Route Setup and Protocol Operation
- PlanetLab Results
  - Multisource/multipath gains, resilience
- Conclusions and Future Work

# Motivation

- Information-Centric Networking (ICN)
  - Emphasis on named content, not endpoints
  - Designed for content distribution
- MMTP: multisource/multipath transport
  - Designed for PSI (PSIRP/PURSUIT)
    - Exploits centralized routing and explicit forwarding
  - Implemented in the PSI prototype
  - Tested over PlanetLab

# PSI Basics




- Rendezvous function (RN nodes)
  - Matches publishers and subscribers
- Topology Manager (TM node)
  - Calculates routes between endpoints
- Forwarding (FN node)
  - In-packet Bloom filter with link labels

# Design Choices

- Receiver-driven operation
  - Subscriber requests packets from publisher
- Multisource downloads
  - File retrieved from multiple publishers
- Multipath delivery
  - Uses multiple paths to each source
- Centralized path selection
  - PSI TM selects paths to sources

# Route Setup

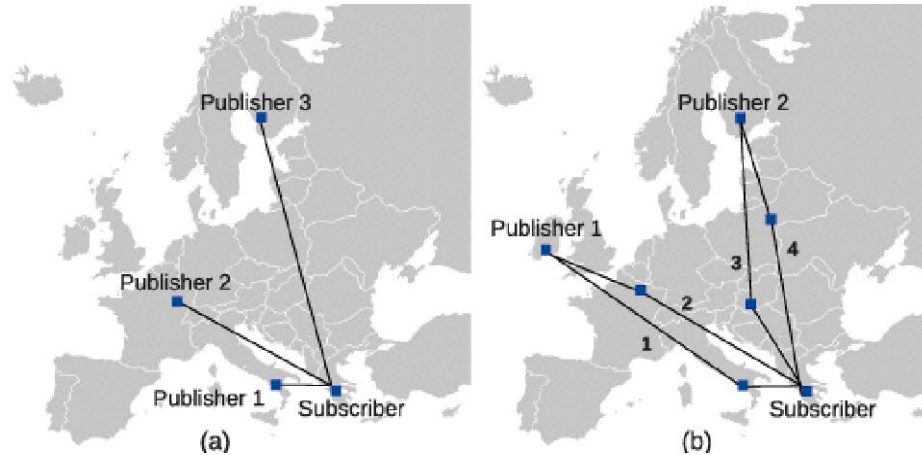


- RN matches subscription to publications
- TM computes paths to publisher(s)
  - Many pairs of forward/reverse paths
  - May use multipath and/or multisource
  - Path pairs are sent to subscriber

# Protocol Operation

- Subscriber requests individual packets
  - Each request carries reverse path
  - May be served by on-path caches
  - Retransmission after a timeout
- Congestion control
  - Different sub-flow over each path pair
  - TCP-like congestion control over disjoint paths
  - Less aggressive behavior over shared paths

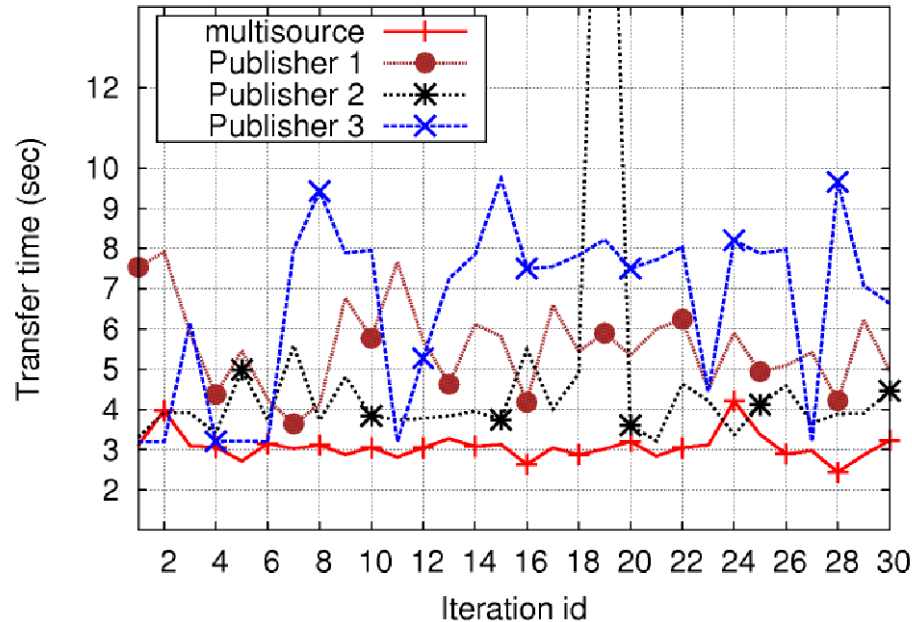
# PlanetLab Topology



- PSI prototype over UDP tunnels
  - Multipath topology
    - Paths to three different publishers
  - Multisource topology
    - Two paths to each of two publishers


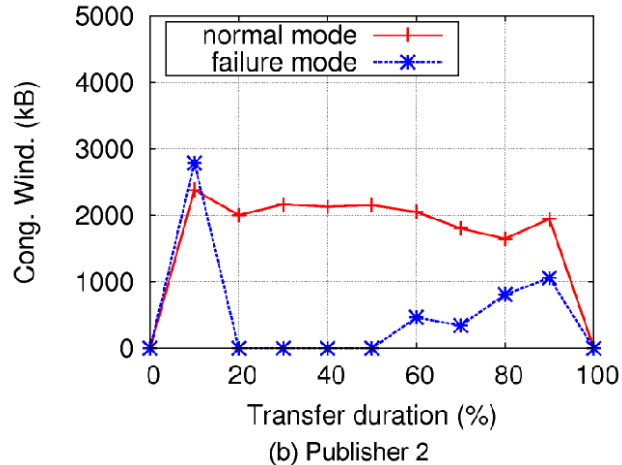



# Multisource Gains



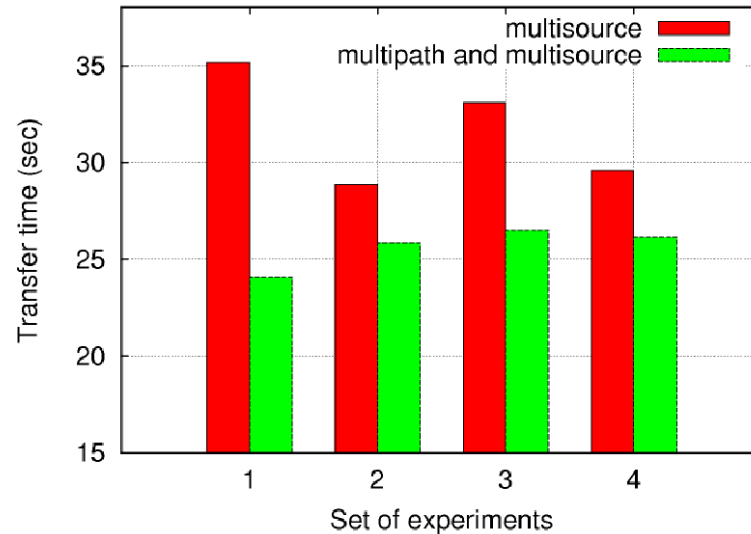
- Topology 1: Three publishers, one path each
  - Single: 3.9 MB/s (average), Multisource: 4.8 MB/s
  - Multisource has much lower variance

# Resilience to Failures



- Topology 1: Three publishers, one path each
  - Publisher 2 stops responding for 7 s
  - No failures: 4.09 MB/s
  - With failures: 3.49 MB/s
    - Automatic switch from Publisher 2 to Publisher 1

# Multipath Gains



- Topology 2: Two publishers, two paths each
  - Multisource: one path per publisher
  - Multipath: two paths per publisher
  - 17% average gain over four experiments

# Conclusions and Future Work

- mmTP: multisource and multipath transport
  - Implemented for the PSI prototype
  - Relies on centralized path computation
  - Exploits explicit packet forwarding
- What next?
  - Congestion algorithm depends on paths
    - Are they disjoint or shared?
  - Experiments with competing TCP flows
    - Fine-tune TCP friendliness



# Thank you

xgeorge@aueb.gr