European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND
Co- financed by Greece and the European Union

Reducing Forwarding State in Content-Centric Networks with Semi-Stateless Forwarding

Christos Tsilopoulos, George Xylomenos, Yannis Thomas

Mobile Multimedia Laboratory,
Athens University of Economics and Business

Outline

- The problem: Does the PIT scale in CCN?
  - As is, very unlikely!
- A semi-stateless forwarding scheme
  - Interest tracking
  - Bloom filter-based Data forwarding
- Evaluation
  - Unicast & Multicast
- Conclusions

A CCN Primer

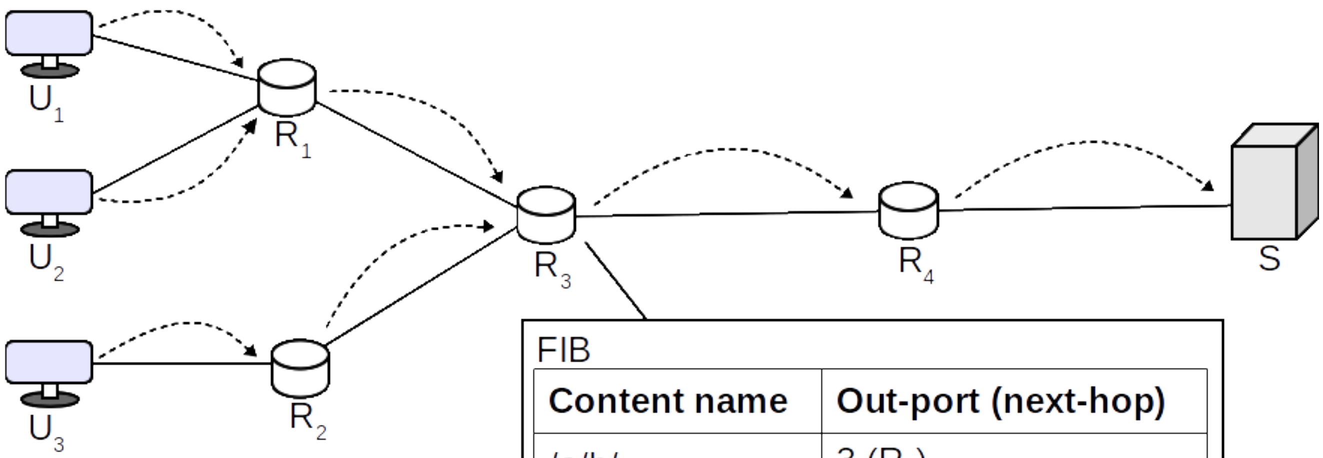

FIB

Content name	Out-port (next-hop)
/a/b/	3 (R_4)

PIT

Content name	Requested from port
/a/b/c.jpg	1 (R_1), 2 (R_2)

- Hop-by-hop Interest forwarding (FIB)
- Pending Interests Table (PIT)
- Stateful forwarding advantages

PIT Scalability

- # entries: bandwidth x RTT / packet
  - Example: 40 Gbps, 80 ms, 1 Kbyte
  - 400K PIT entries, longer than IP addresses
  - Even more in live streaming and pub/sub
  - Realistic estimates indicate millions of entries
- Large memory requirements for PIT
  - Does not fit in on-chip memory
  - Bad performance with main memory

Reducing PIT size

- DiPIT: combine entries with Bloom filters
  - Drops per Interest information
- ENPT: uses a trie to encode Interests
  - Lookup time depends on name size
- CONET: Interests accumulate source routes
  - Drops all benefits of stateful forwarding
  - No multicast or adaptive forwarding
  - No dropping of unwanted Data

Semi-Stateless Forwarding

- Track Interests at *some* on-path routers
  - On average, every *d* hops
  - For N-hop paths, Interest tracked at N/d routers
- Stateless forwarding between stateful routers
  - Using in-packet Bloom filters

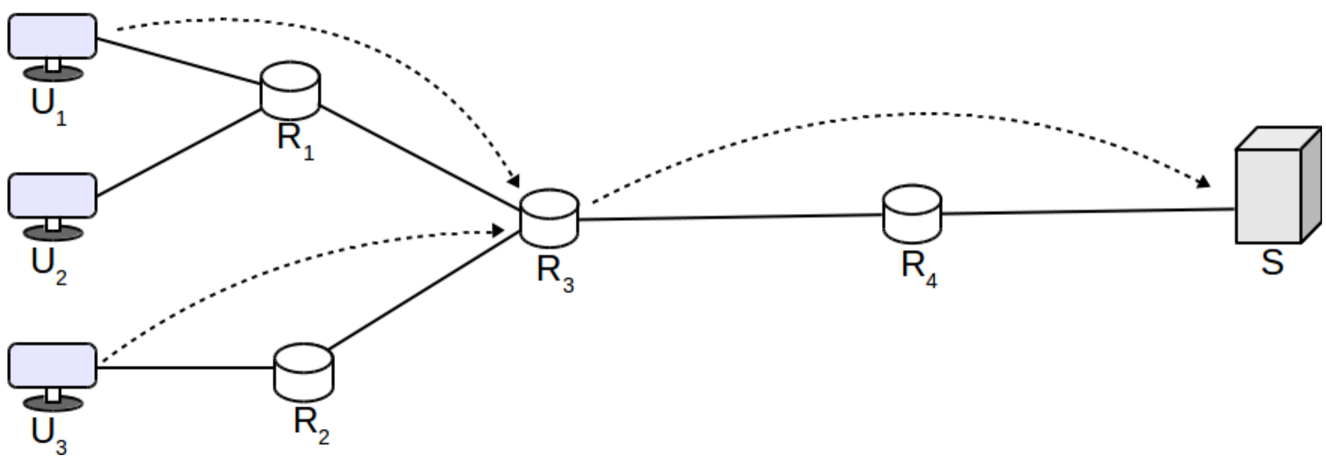

Interest Tracking: Requirements

- Spread state across routers
  - Avoid bottleneck points
  - Routers should track 1/d of forwarded Interests
- Efficient multicast rendezvous
  - Aggregate Interests for the Same Data

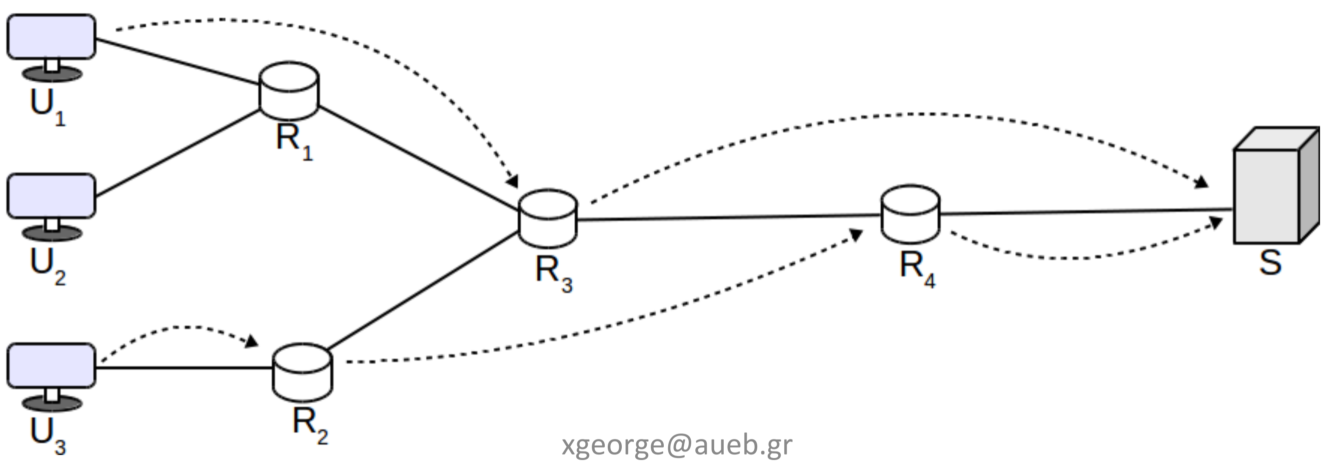

1. Probabilistic Tracking

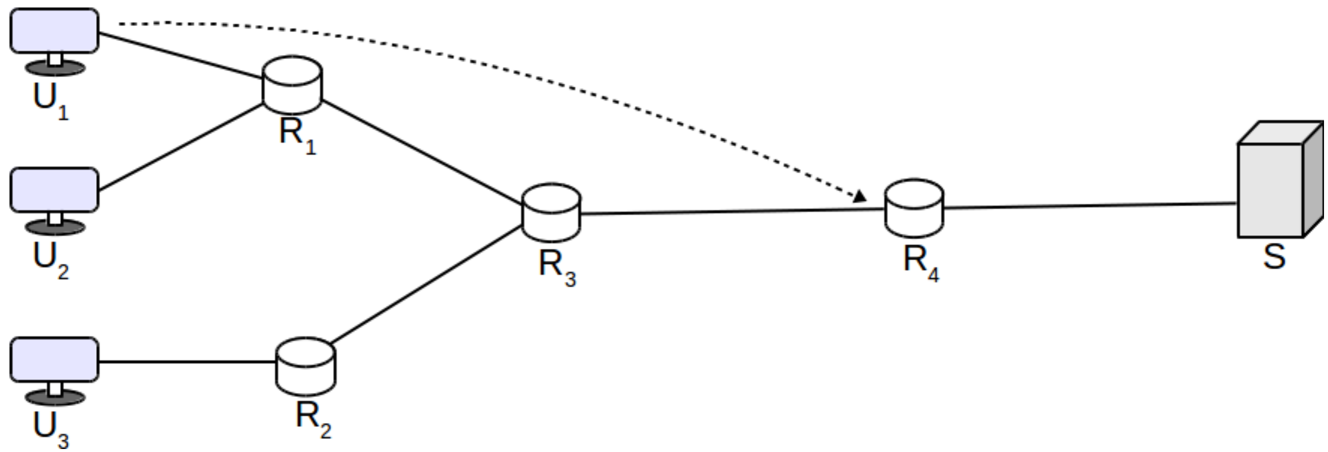

- Router tracks Interest with probability 1/d
  - Interest aggregation: check if entry exists
  - Upper bound on stateless hops
    - Avoid large stateless parts
    - Uses hop counter in Interests

1. Probabilistic Tracking

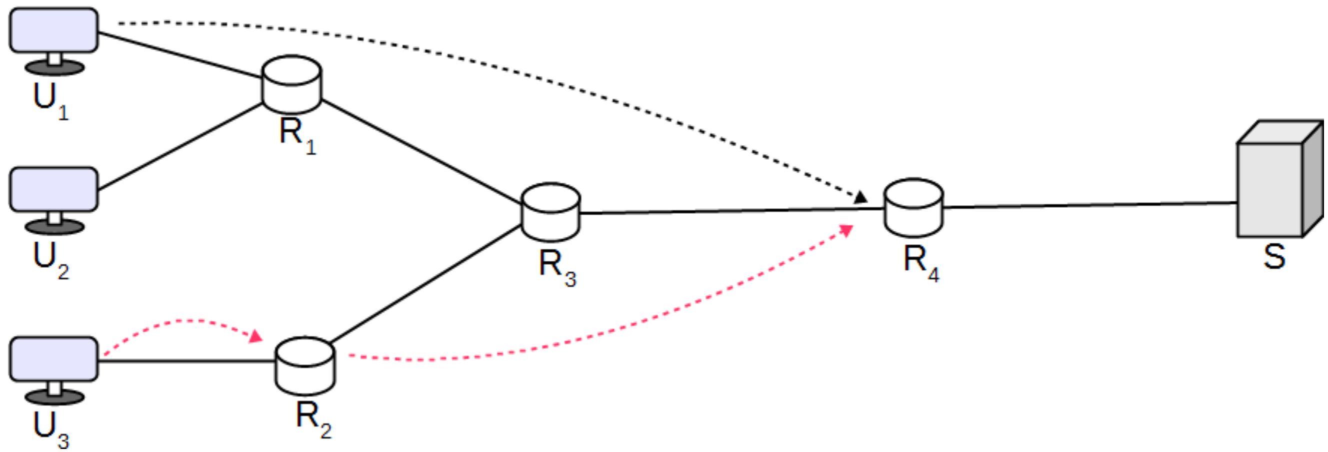

- Router tracks Interest with probability 1/d
  - Interest aggregation: check if entry exists
  - Upper bound on stateless hops
    - Avoid large stateless parts
    - Uses hop counter in Interests

2. Hash-based Tracking

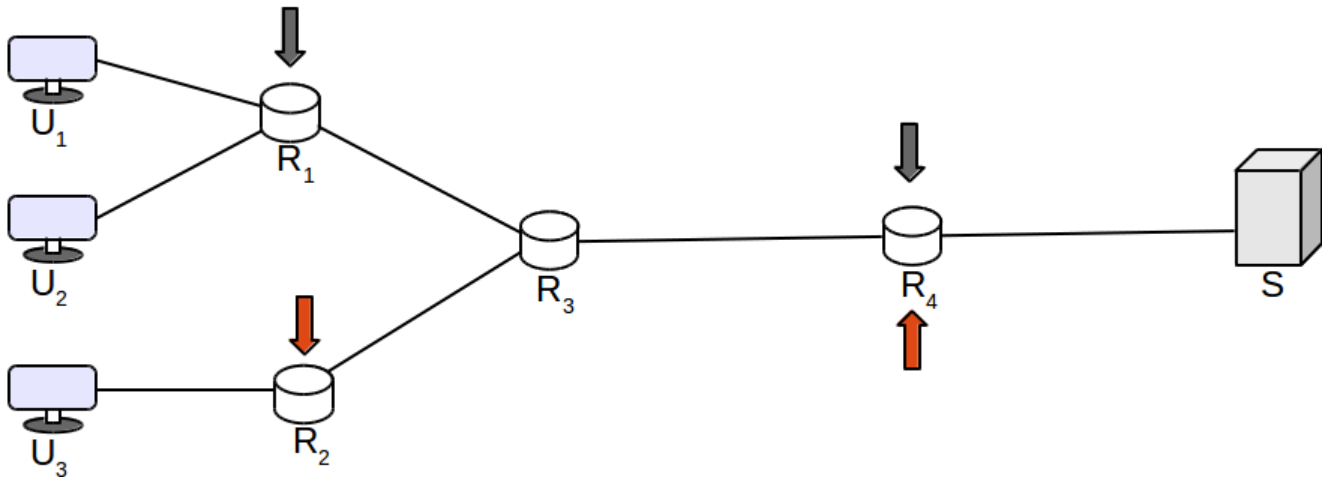

- Router tracks Interest of `h mod d == 0`
  - `h=hash(content_name+router_suffix)`
  - Deterministic selection of storage points
  - Also needs upper bound on stateless hops

3. Hop Counter-based Tracking

- Both previous policies require hop counters
- Why not use *only* the hop counters?
  - Much simpler to implement
  - Stateless paths never get long
- Select initial value randomly in [0, d-1]
  - Spread state across routers
- Check PIT for pre-existing entry
  - Aggregate entries as soon as possible

Data Forwarding: Bloom filters

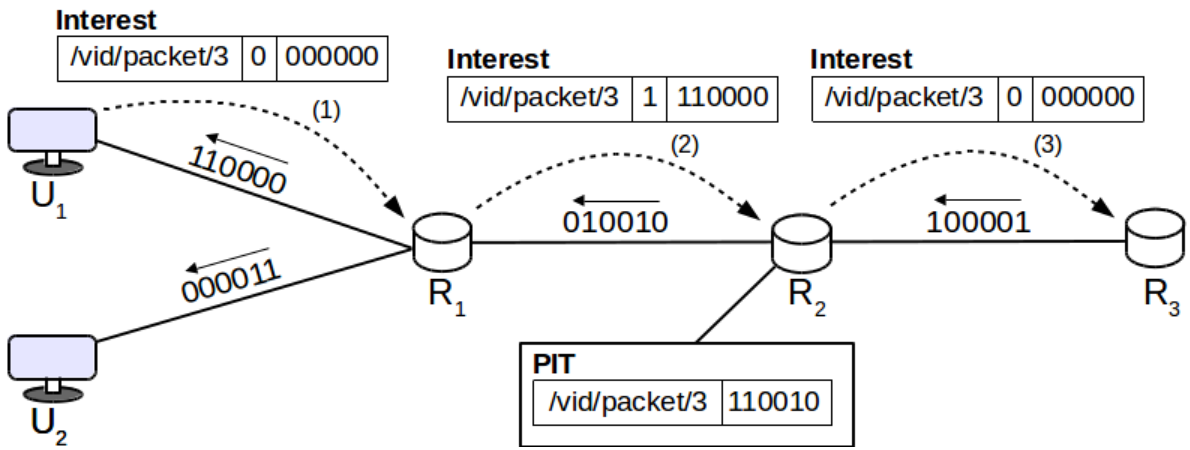

- Track reverse path information in routers
  - Each link has a random ID (LID)
  - Add (OR) reverse LID at each hop
  - Store the Bloom filter (BF) in the PIT
  - Add (OR) the BFs to aggregate paths

Data Forwarding: Bloom filters

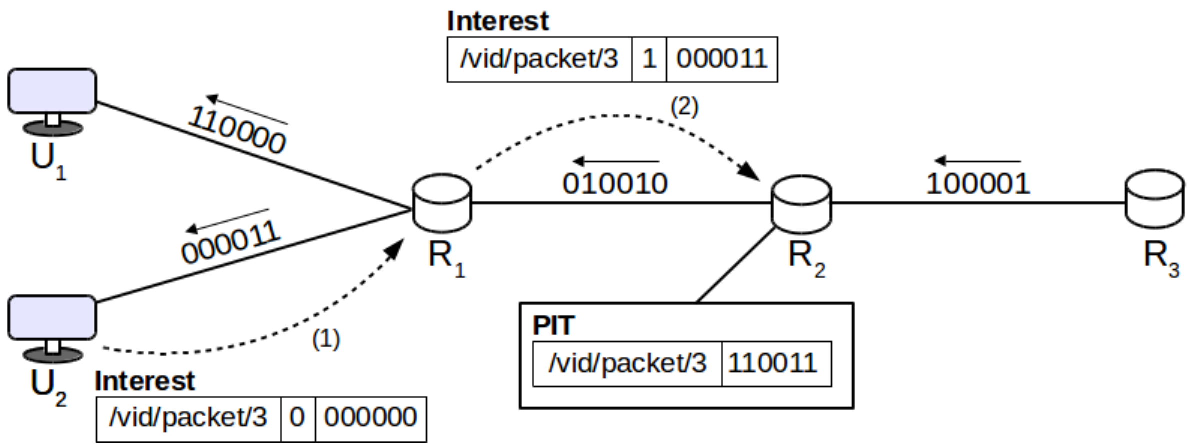

- Track reverse path information in routers
  - Each link has a random ID (LID)
  - Add (OR) reverse LID at each hop
  - Store the Bloom filter (BF) in the PIT
  - Add (OR) the BFs to aggregate paths

Tradeoffs

- Larger Interest & Data packets
  - Need to carry BFs in headers
- Additional Interests with multicast
  - Aggregation beyond first common router
  - Does not influence Data transmissions
- Additional Data
  - False positives in BFs
- Fewer but larger PIT entries

Evaluation

- Simulations using AS 224
  - Similar results with other ASes
  - 500 hosts distributed to access routers
  - 128 bit Bloom filters for forwarding
- Metrics
  - Reduction in PIT entries
  - Reduction in PIT size (bytes)
  - Additional Interests and Data

Unicast

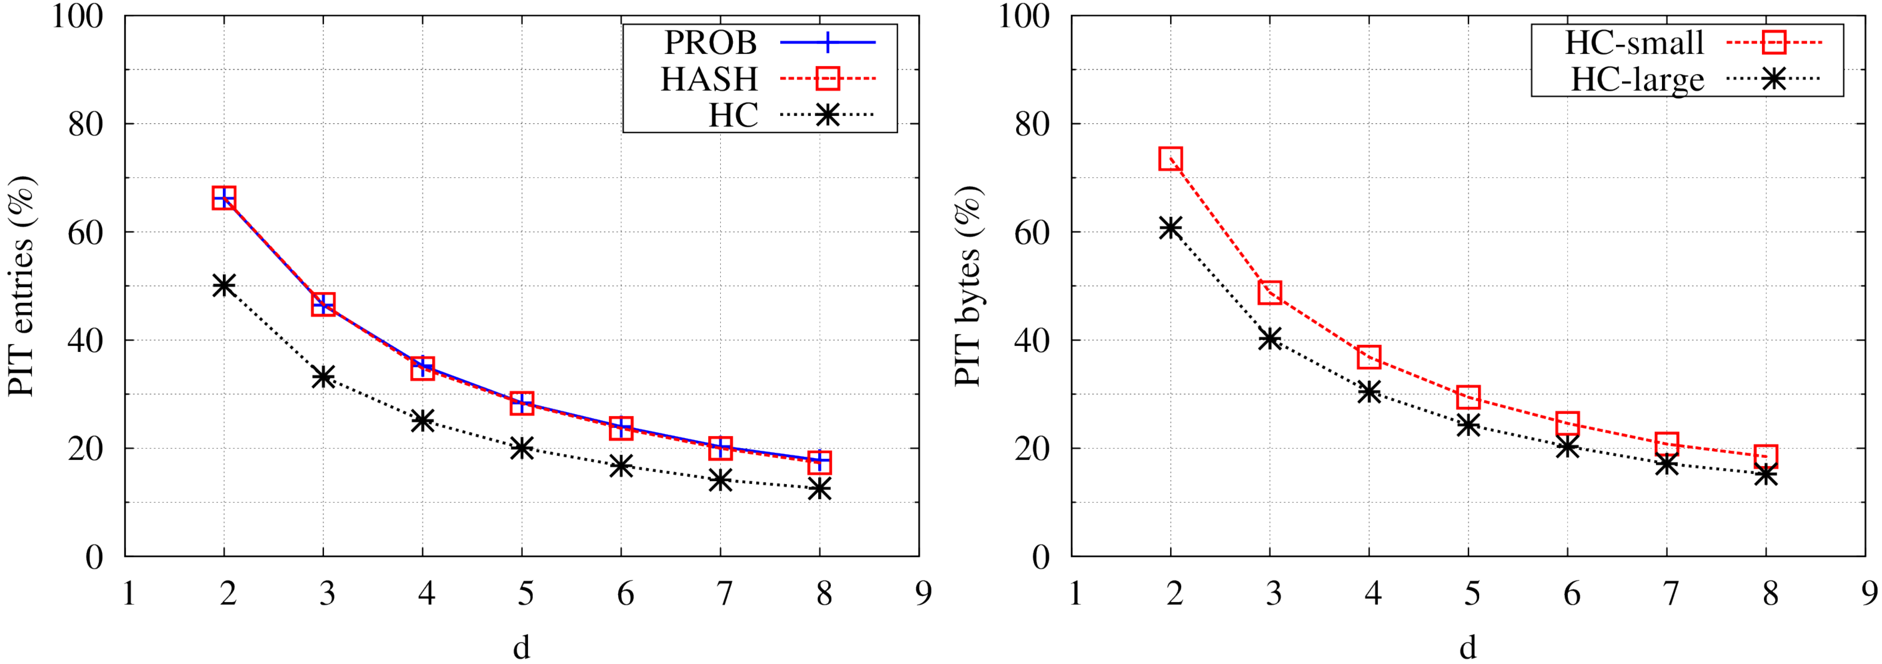

- Simple file transfer
  - 66% reduction in PIT entries (HC, d=3)
  - 54-61% reduction in PIT size (HC, d=3, small/large)

Multicast: PIT entries

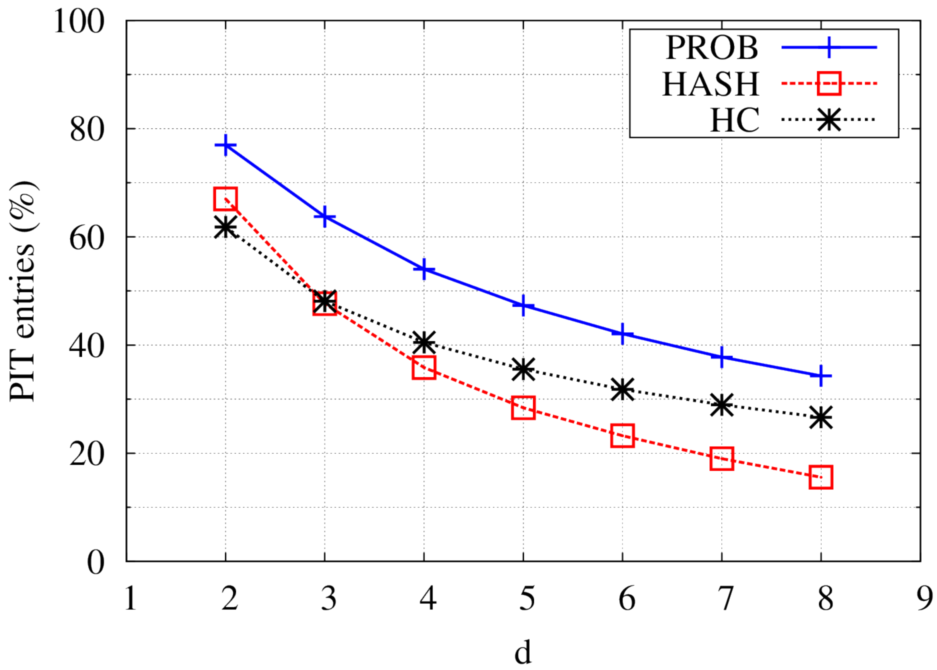

- Live media streaming
  - 1000 groups, Zipf distribution of sizes
  - 52% reduction in PIT entries (HC, d=3)

Multicast: PIT Size

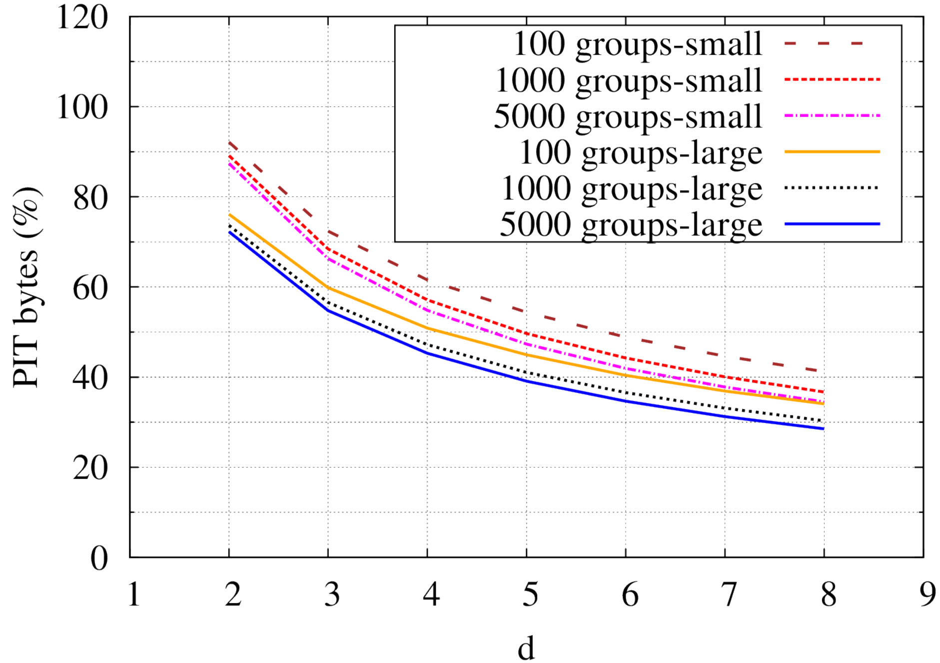

- Live media streaming
  - 34% smaller PIT size (5k groups, d=3, small names)
  - 45% smaller PIT size (5k groups, d=3, large names)

Multicast: Overhead

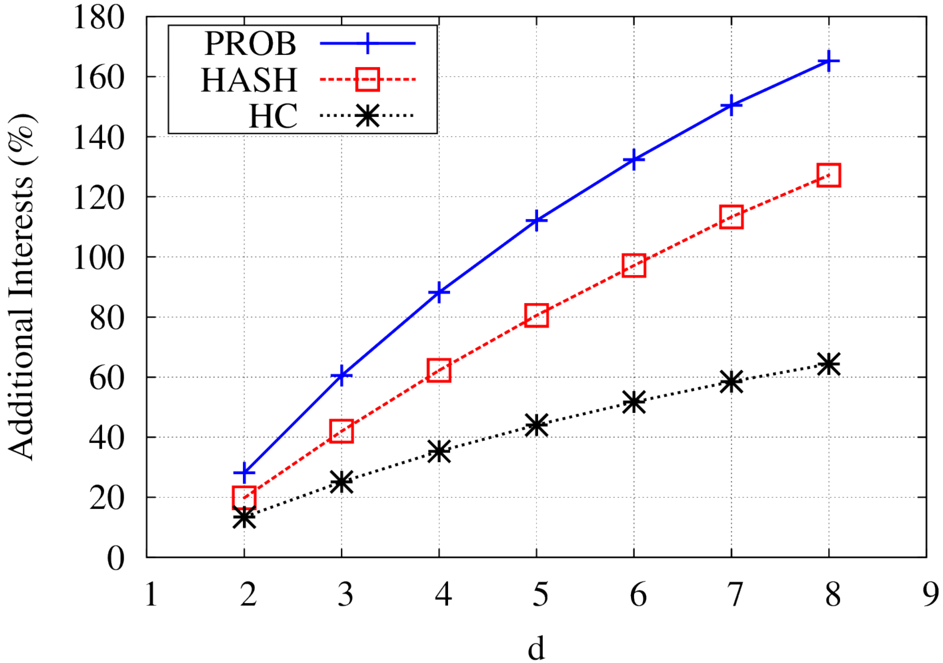

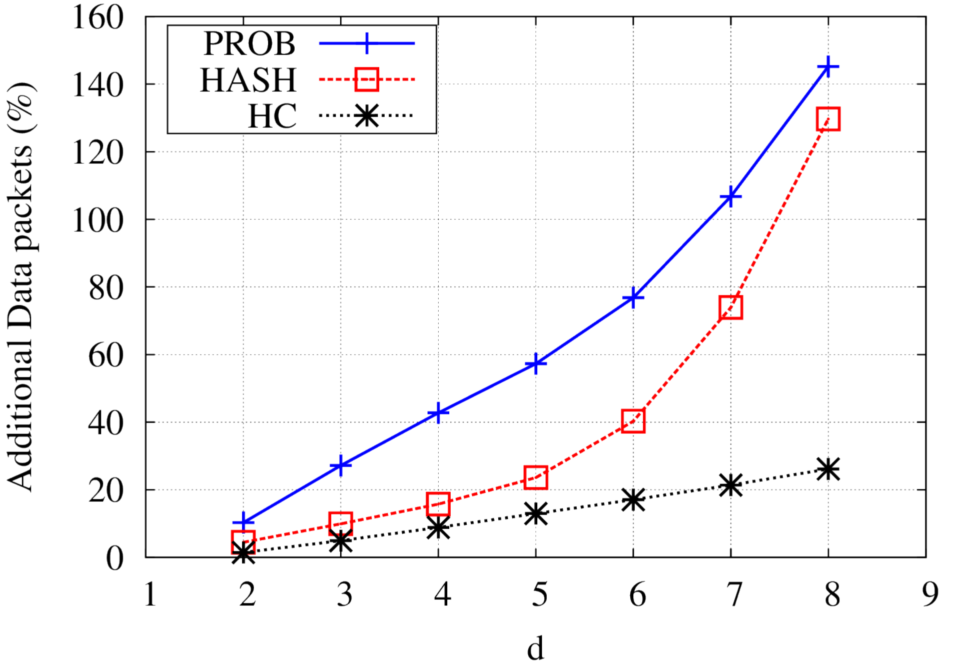

- Live media streaming
  - 25% additional Interests (HC, d=3)
  - 5% additional Data (HC, d=3)

Multicast: Total Overhead

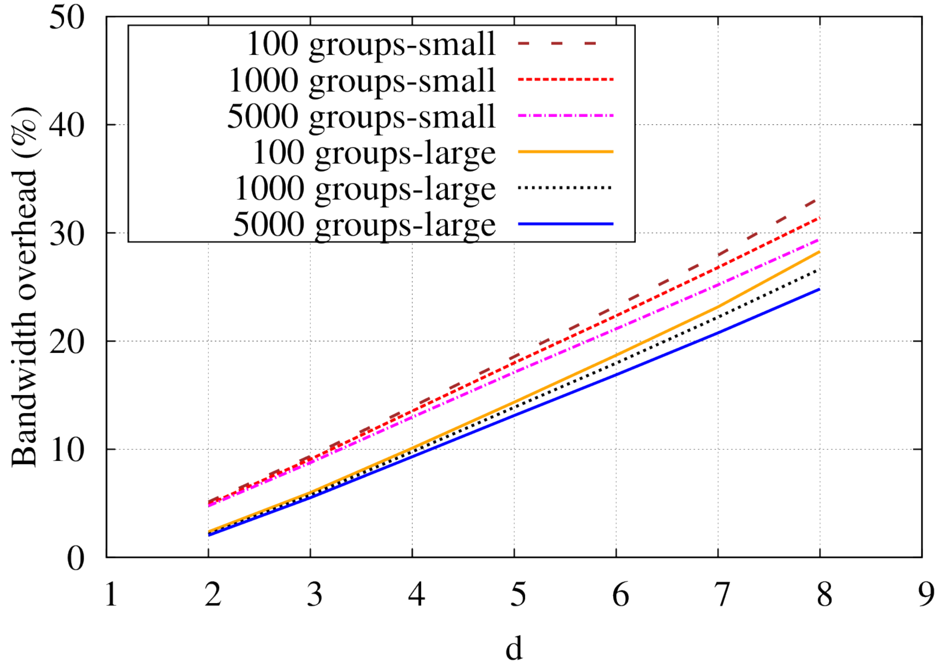

- Live media streaming
  - 9% more bandwidth (5K groups, d=3, 1.5K pkt)
  - 6% more bandwidth (5K groups, d=3, 8K pkt)

Evaluation overview

- Significant PIT size reduction achieved
  - 54-61% compared to CCN for unicast
  - 34-45% compared to CCN for unicast
- Small bandwidth penalty
  - No penalty at all for unicast
  - 6-9% extra bandwidth on multicast
- The HC policy works best with *d*=3 or 4
  - *d* trades off PIT size against bandwidth overhead

Conclusion

- Semi-stateful forwarding reduces PIT size
- Qualitative aspects of CCN are maintained
  – Multicast, dropping unwanted packets
- Only Interest & Data processing changes
  – LIDs are constructed autonomously
- Future work
  – Resort to semi-stateless only as PIT fills up
  – Vary *d* per router to minimize overhead

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Thank you

xgeorge@aueb.gr