Pursuing a Pub/Sub Internet

Scalability Issues in ICN

Christos Tsilopoulos, Xenofon Vasilakos and
<u>George Xylomenos</u>

Mobile Multimedia Laboratory
Department of Informatics
AUEB, Greece

Motivation

- ICN involves named data
  - Names may be hierarchical or flat
  - But they are definitely too many!
  - Orders of magnitude more than hosts
  - How can we handle such numbers?
- At least two issues need work
  - Name resolution scalability
  - Forwarding scalability
- Discussion starter presentation
  - Problems & some ideas

Name resolution

- Obvious problem: the namespace is huge
  - Trillions of named items to begin with
  - Plus tons of data from the Internet of Things
  - Can we maintain huge PITs?
  - Can we maintain huge DHTs?
- Less obvious problem: the names are also huge
  - Simply because there are too many
  - Can we use huge names in packets?
  - Can we survive with shorter versions?

Pursuing a Pub/Sub Internet

Name resolution

- DNS relies on locality and caching
  - Similar names lead to similar name servers
  - Name server addresses can be cached
- CCN/NDN names are hierarchical, but different
  - Prefixes do not lead to unique locations
  - Caching pointers based on prefixes does not work
- Are name prefixes expected to be controlled?
  - For example, /aueb is assigned to, well, AUEB
  - If they are, maybe the problem is tractable
  - But this is a back door for location/identity binding!

Name resolution

- DHTs scale well and distribute load, but…
  - They assign names to potentially unwanted servers
  - They violate routing policies
  - They stretch resolution paths
  - Hierarchical DHTs only partially help
  - Caching not nearly as effective as in DNS
- DONA-like solutions work better than we thought
  - Huge amounts of storage are needed at the top
  - But, we do have huge server farms
  - We may be able to throw hardware at this problem

Forwarding

- The actual packet transfer is also problematic
  - Assume that we aim for native transfers
  - Otherwise we are passing the problem to IP!
- We use single-source multicast as an example
  - Multicast is supposed to be an ICN strength
  - Single-source is commercially more interesting
- Two general options
  - Hop-by-hop (as in CCN/NDN)
  - Source routing (as in PSIRP/PURSUIT)

Forwarding

- Hop-by-hop forwarding (CCN)
  - \+ Forwarding state distributed to routers
  - \+ Fast join/leave for multicast
  - \- Large forwarding table (PIT) size
    - Especially with very fast routers
    - Trickier with variable-length names
    - State compression ideas need to be evaluated
  - \- Strictly symmetrical paths only
    - The resolution part fixes the forwarding part
    - Can be either a feature or a bug

Pursuing a Pub/Sub Internet PURSUIT

Forwarding

- Source-routing (PSIRP)
  + Flexible path selection
    - Content/cache awareness
    - Steiner tree multicast
  + Stateless router operation
  - Large multicast trees hard to handle
    - Processing delay at the path computation engine
    - Slow join/leave for multicast
  - Forwarding inefficiencies due to false positives
    - Bigger problem with larger groups
    - Solutions for large trees need to be evaluated