UCL DEPARTMENT OF ELECTRONIC AND ELECTRICAL ENGINEERING
COMMUNICATIONS AND INFORMATION SYSTEMS GROUP

## IFIP/IEEE CNSM 2013 Keynote Speech

# Information-Centric Networking and In-Network Cache Management: Overview, Trends and Challenges

Prof. George Pavlou
http://www.ee.ucl.ac.uk/~gpavlou/
Communications and Information Systems Group
Dept of Electronic & Electrical Engineering
University College London, UK

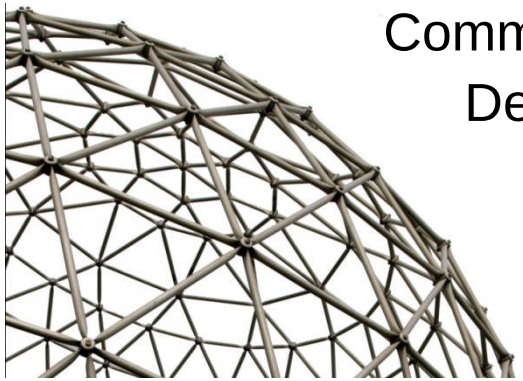

# Internet-based Content

- The vast majority of Internet interactions over the last years relate to content access
  - P2P overlays (e.g. BitTorrent, eMule, live streaming)
  - Media aggregators (e.g. YouTube, GoogleVideo)
  - Over-the-top video (e.g. Hulu, iPlayer)
  - Content Delivery Networks (e.g. Akamai, Limelight)
  - Social Networks (e.g. Facebook, MySpace)
  - Photo sharing sites (e.g. Picasa, Flickr)
  - …
- New approaches are required to cater for the explosion of video-based content
- **Continue throwing more capacity cannot work anymore!**

# The Emerging Content-Oriented Internet

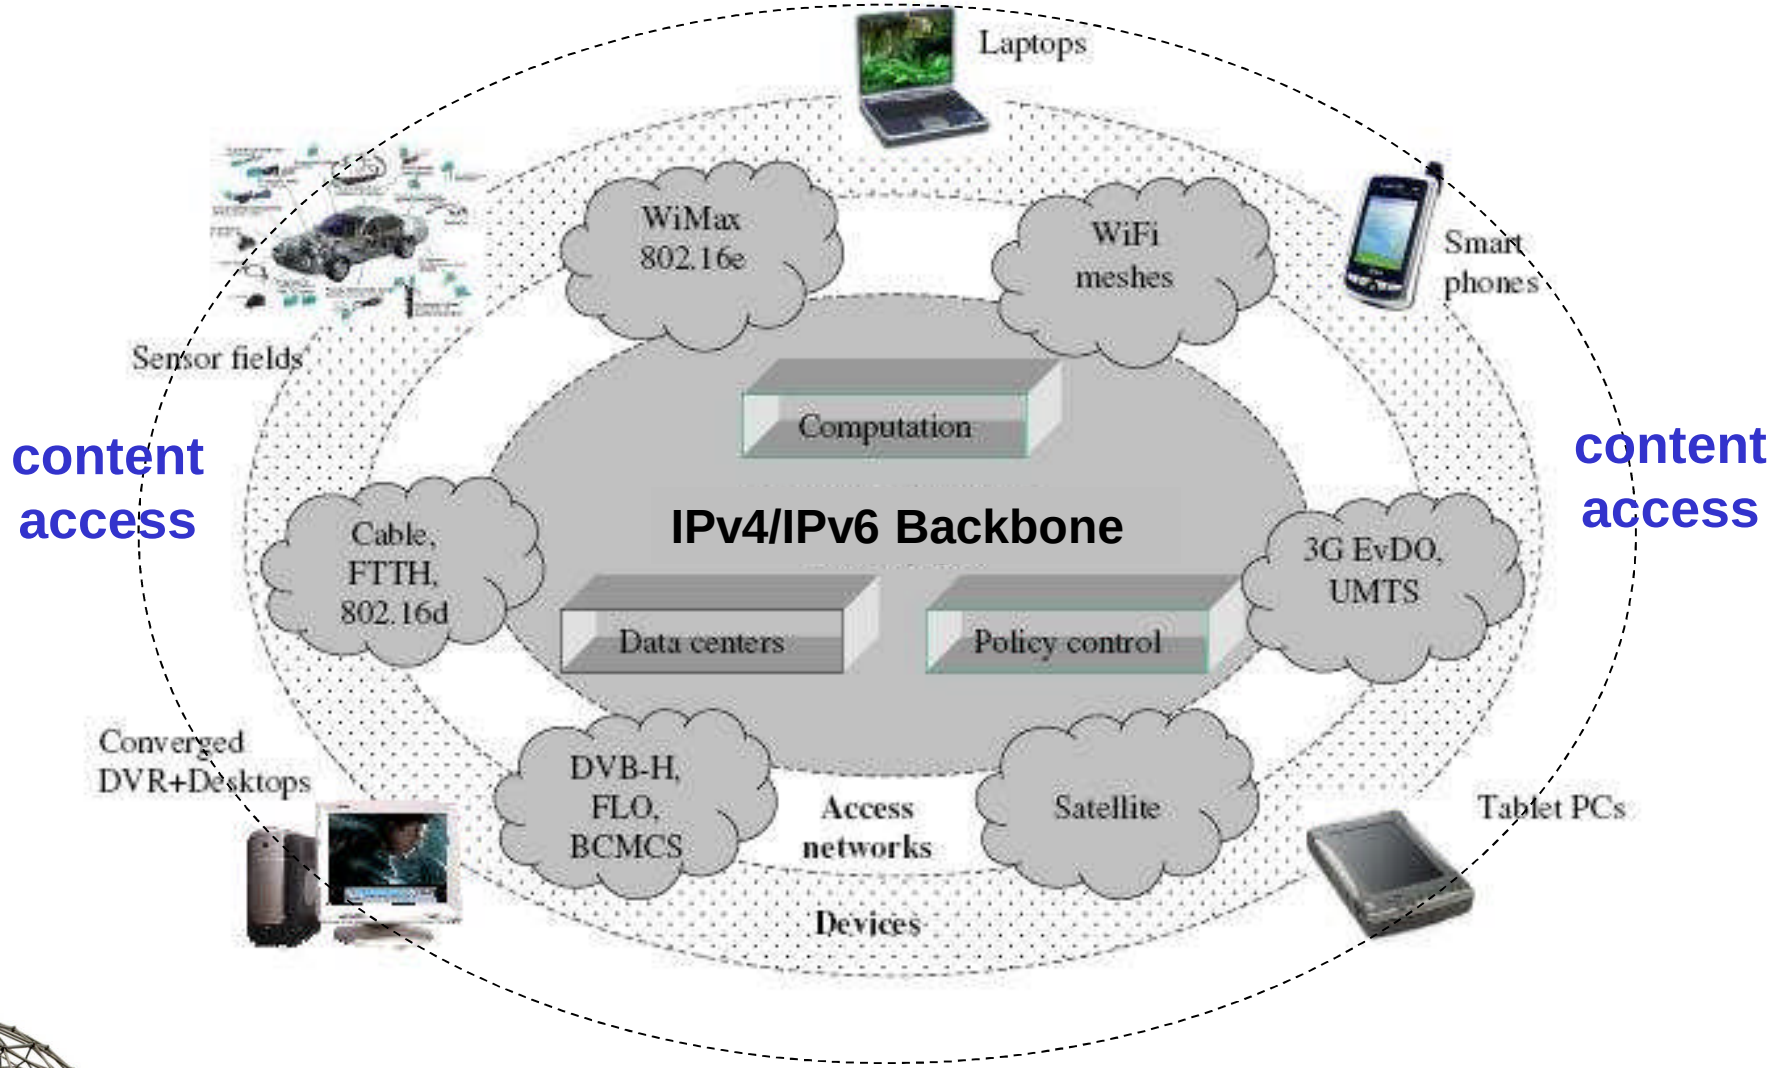

# Expected IP Traffic Growth Until 2014

- According to the Cisco Visual Networking Index (VNI) 2010:
  - Global IP traffic will quadruple until 2014
  - Global Internet video traffic will surpass P2P traffic in 2010
  - Approx. 55% of the overall Internet traffic will be video by 2014
  - Global mobile data traffic will double every year until 2014
  - Approx. 65% of the overall mobile traffic will be video by 2014

- **Infrastructure evolution needs to be partnered with novel approaches and associated business models**

# Expected IP Traffic Growth Until 2014 (cont'd)

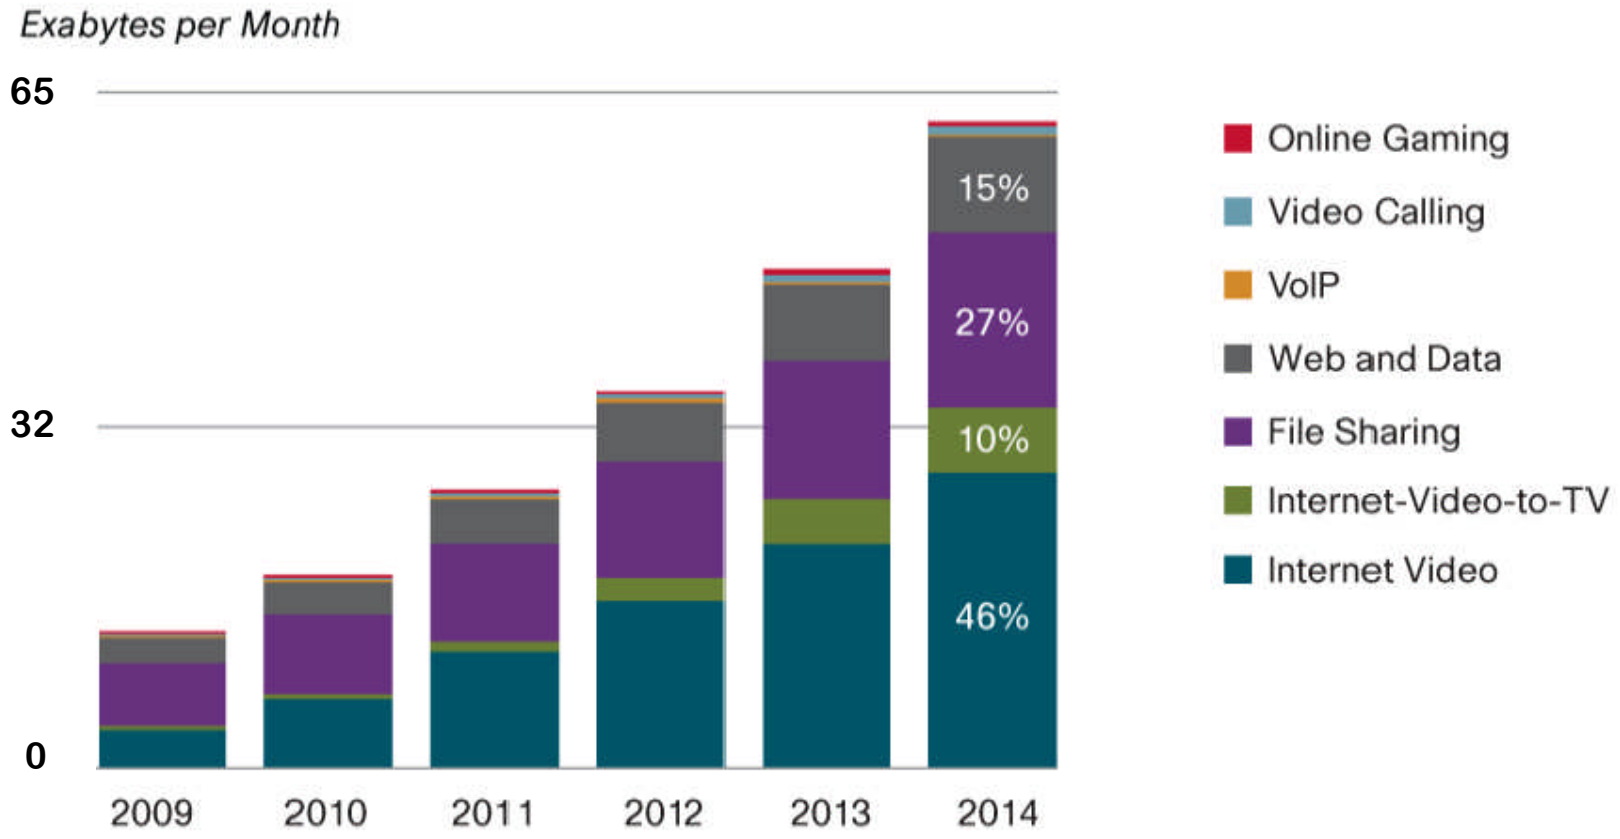

# P2P Overlays and CDNs

- Peer-to-Peer (P2P) Overlays: started from file sharing and evolved to multicast-streaming real-time video through overlay nodes
  - Content object names are resolved to candidate peers
  - Self-organized, adaptive, fault-tolerant content distribution
  - Dynamic content replica placement according to demand
- Content Distribution Networks (CDNs): they support anycast by choosing the most appropriate (i.e. topologically close) content replica to maximise user QoE
  - Use DNS-based redirection
  - Mostly offline content replica placement according to expected demand

- Both P2P overlays and CDNs make the content server transparent by accessing "named content", allowing access to cached copies
  - A first step towards an information-oriented communication model

# Current Content Naming and Security Problems

- Content URIs are effectively object locators, resolving to the IP address of the hosting server i.e. location-dependent
  - Binding breaks when object moves or when site changes domain
  - Replicas all have different URIs, appearing as different objects
  - Unique, persistent, location-transparent naming is required

- The current Internet security model provides connection endpoint as opposed to content object authentication
  - Once an object copy has left the origin server, its authenticity cannot be verified anymore, which is a problem for caching
  - In an information-centric approach it is important to be able to authenticate content objects as opposed to connection endpoints

# Current Paradigm Shift

Node-centric design: sharing network resources

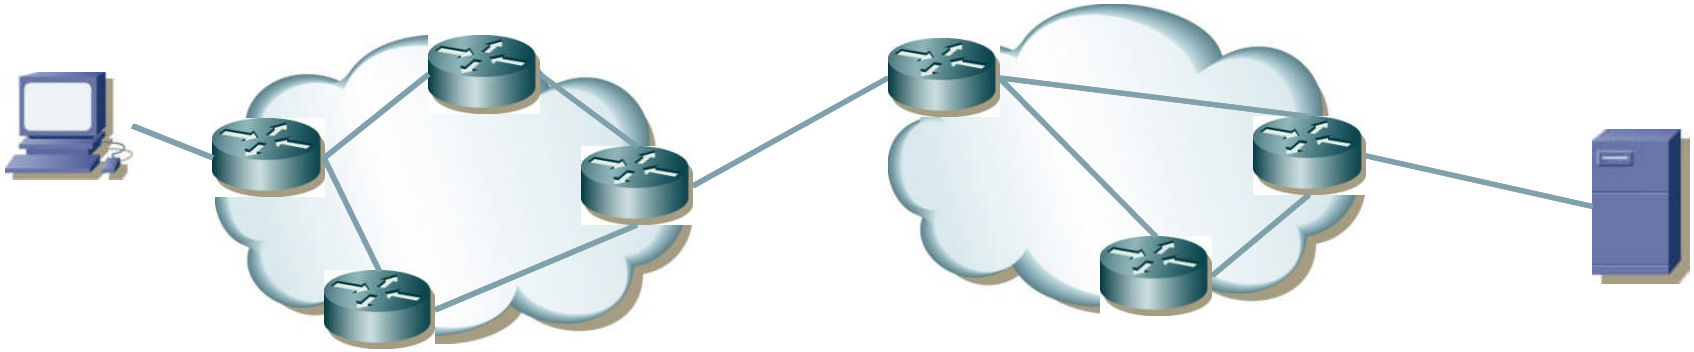

Information-centric design: content access and distribution

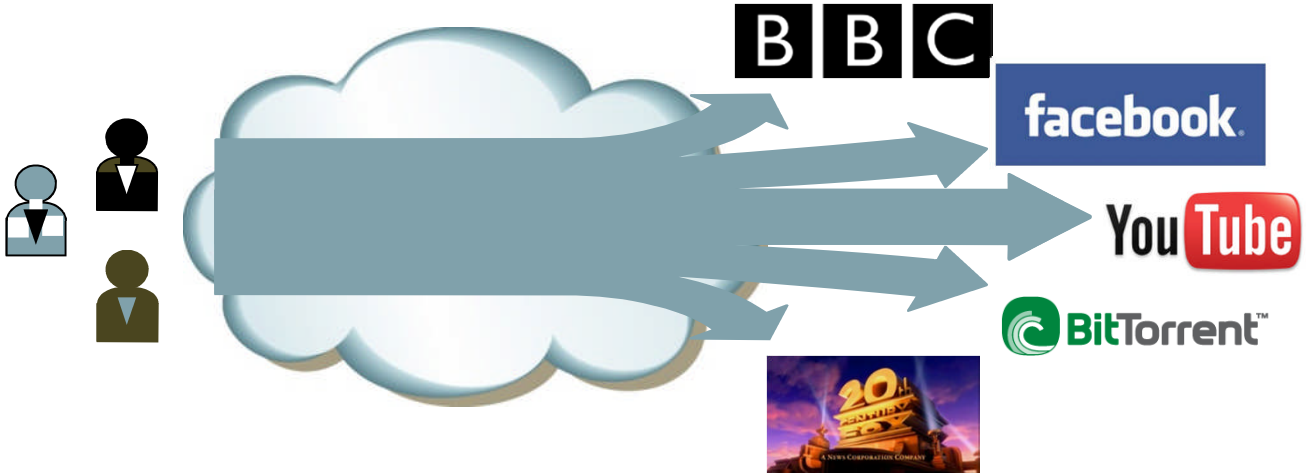

# Information-Centric Networking

- Given that users are interested in named content and not in node endpoints, is there a clean architectural approach to address the relevant requirements?
  - All encompassing instead of domain-specific application-level add-ons
  - Provide an enhanced P2P/CDN-like paradigm within the network

- **Information-Centric Networking (ICN)** targets general infrastructure that provides in-network caching so that content is distributed in a scalable, cost-efficient & secure manner
  - Exploits the large amounts of DRAM in current generation routers
  - Receiver-driven model – subscribe/get objects of interest
  - Support for location transparency, mobility & intermittent connectivity

# Flash-Crowd Effect Due to Content Popularity

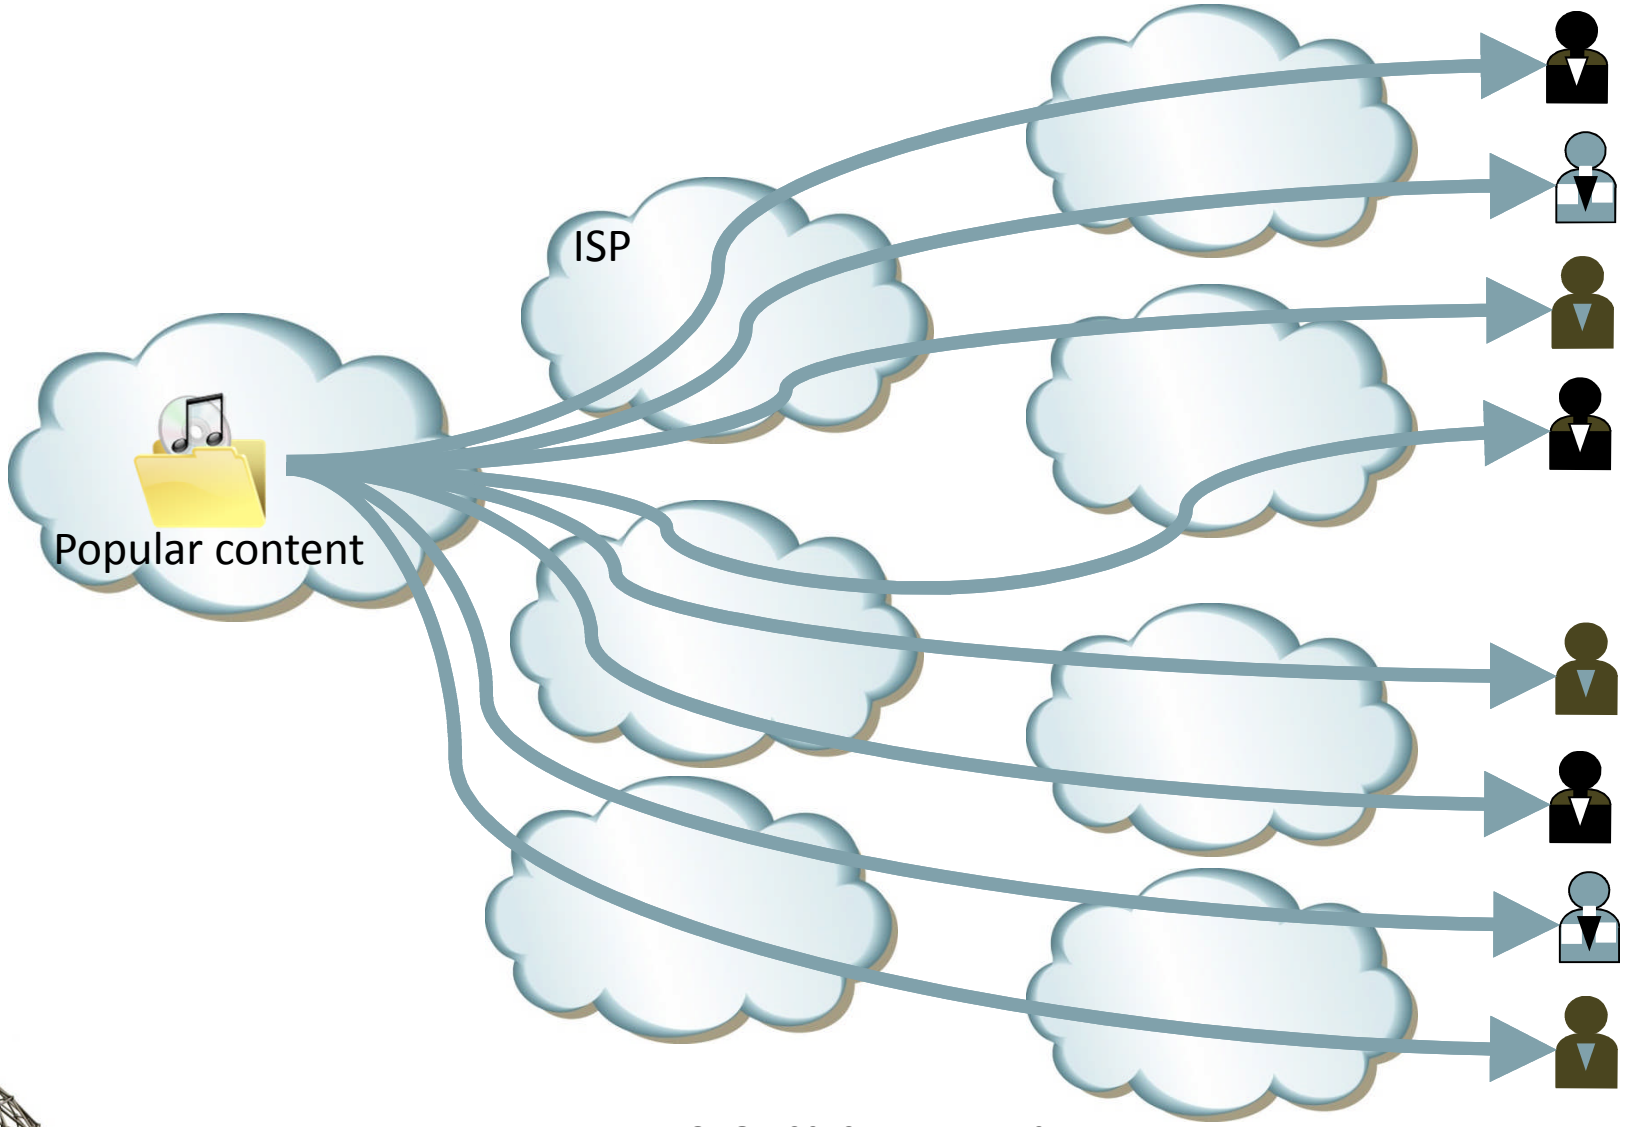

# Scalable Cache-based Content Distribution

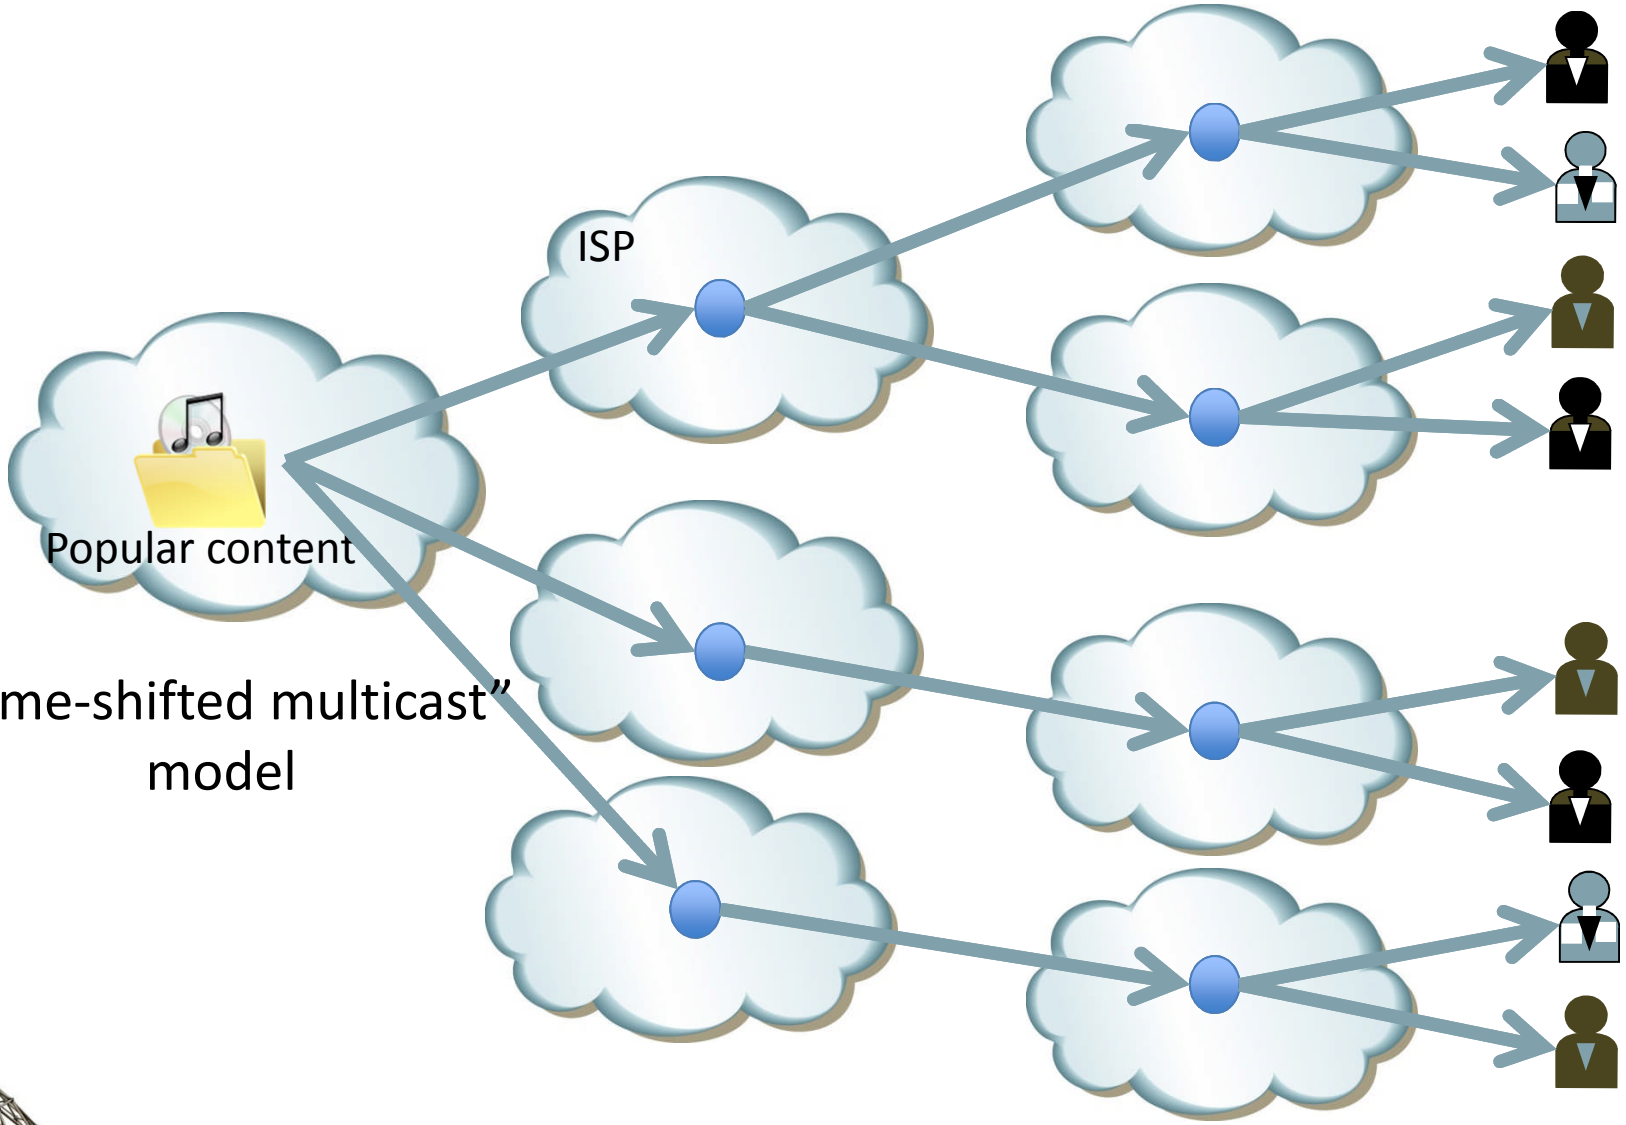


"Time-shifted multicast" model

# In-Network Caching Approaches

- In-network caching is dynamic reactive (as in P2P overlays) and not offline proactive (as in CDNs) because of required line speed operation
- Different options for the granularity of caching:
  - Chunk-level: caching information chunks – an information object is split into a number of fixed-size information chunks
  - Packet-level: caching individual packets as transmitted through the network – typical size 1.5Kb each – as proposed by CCN
  - But **not at the file/information object level** as in overlay approaches e.g. CDNs

- Distributed line speed decision making is required w.r.t. what/where to cache and what to drop to maximize gain

# Information Objects

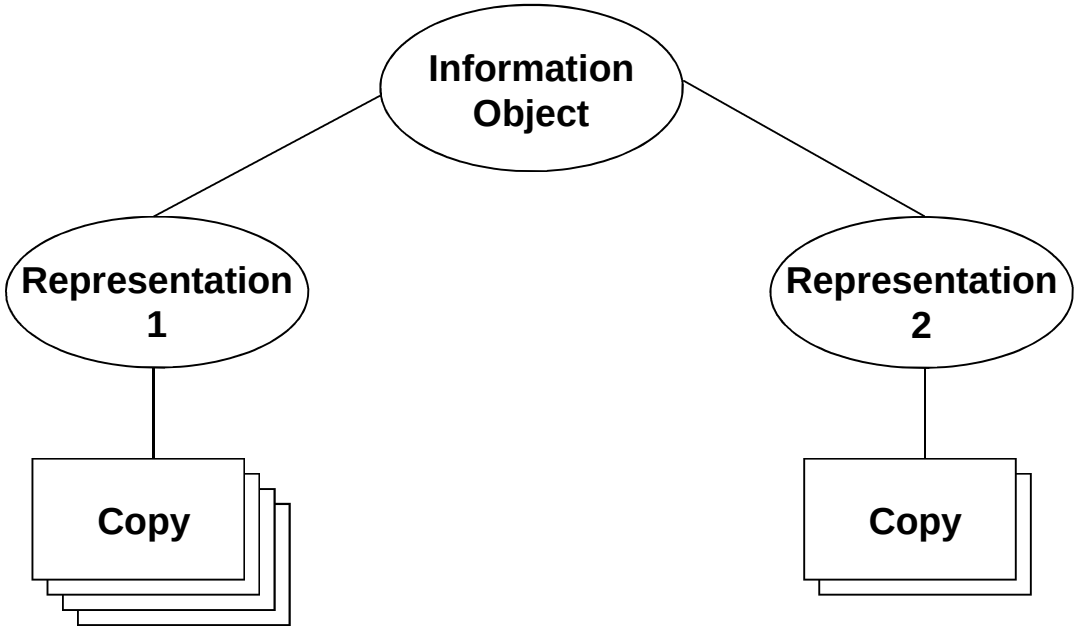


The representations of an information object (e.g. different encodings) and their cached copies all share the same ID

# Content Naming

- Information objects are named by location-independent IDs
- Given that in ICN security applies to information, object IDs in many ICN architectures incorporate security tags
  - Non human-friendly IDs
  - But human-friendly names can also be associated with IDs
- There can be hierarchical names, flat self-certifying names or combined schemes

# Name Resolution and Routing

- Two key approaches: two-phase and one-phase
- In two-phase approaches (evolutionary), name resolution takes place first by mapping the content ID to locators, with the most suitable one selected (anycast)
  - Content name resolution servers are required e.g. DNS++
  - Routing to the content source and subsequent content delivery simply use locators i.e. IP addresses
- In one-phase approaches (radical/revolutionary), content ID-based routing is used natively within the network
  - Content-ID based routing uses content state in the network ("breadcrumps") and supports anycast to the closest copy
  - The content delivery path is typically the reverse path of the request, i.e. symmetric as opposed to the current asymmetric routing

# Name Resolution Scalability

- Let's look at the two phase evolutionary approach
- A vast amount of information objects
  - Currently more than 1 trillion unique URLs (Google 2008)
  - 26 billion web pages (www.worldwidewebsize.com)
  - 119 million 2nd level domain names in the DNS (end of 2010)
- Possible to operate DHTs with >2 million nodes
  - For 1000 trillion objects ($2^{15}$) with 100 bytes per record and no replication/caching, 50Gb of DRAM is necessary
  - With 10 replicas/cached copies and 1Kb per record, 5Tb of RAM is necessary and can be supported with SSD, albeit expensively

# Key Projects and Approaches

Europe:

- 4WARD/SAIL **NetInf** - Network of Information
- PSIRP/PURSUIT **PubSub** - Publish Subscribe Routing
- UCL **COMET CMP** - Content Mediation Plane

US:

- Berkeley **DONA** - Data-Oriented Network Architecture
- Xerox PARC **CCN/NDN** - Content-Centric Networking / Named Data Networking

Also various other research activities worldwide

# Network of Information (NetInf)

- Started in the EU project 4WARD and was continued in its follow-on project SAIL
- Both one-phase and two-phase approaches
  - Two-phase evolutionary approach uses a hierarchy of DHTs – well developed
  - One-phase radical approach less-well developed

# Content-Centric Networking (CCN)

- Originated by Van Jacobson
- One-phase approach through *Interest*/*Data* packets flowing in a “reverse ack/data TCP-style”
  - Data packets are cached everywhere along the delivery path (Cache Everything Everywhere – $CE^2$) as they may be useful to other consumers
  - Least Recently Used (LRU) packet discard policy
  - Hierarchical naming scheme

- CCNx implementation is publicly available while the NSF NDN project looks at more general CCN-related research issues

# CCN-like Universal In-Network Caching

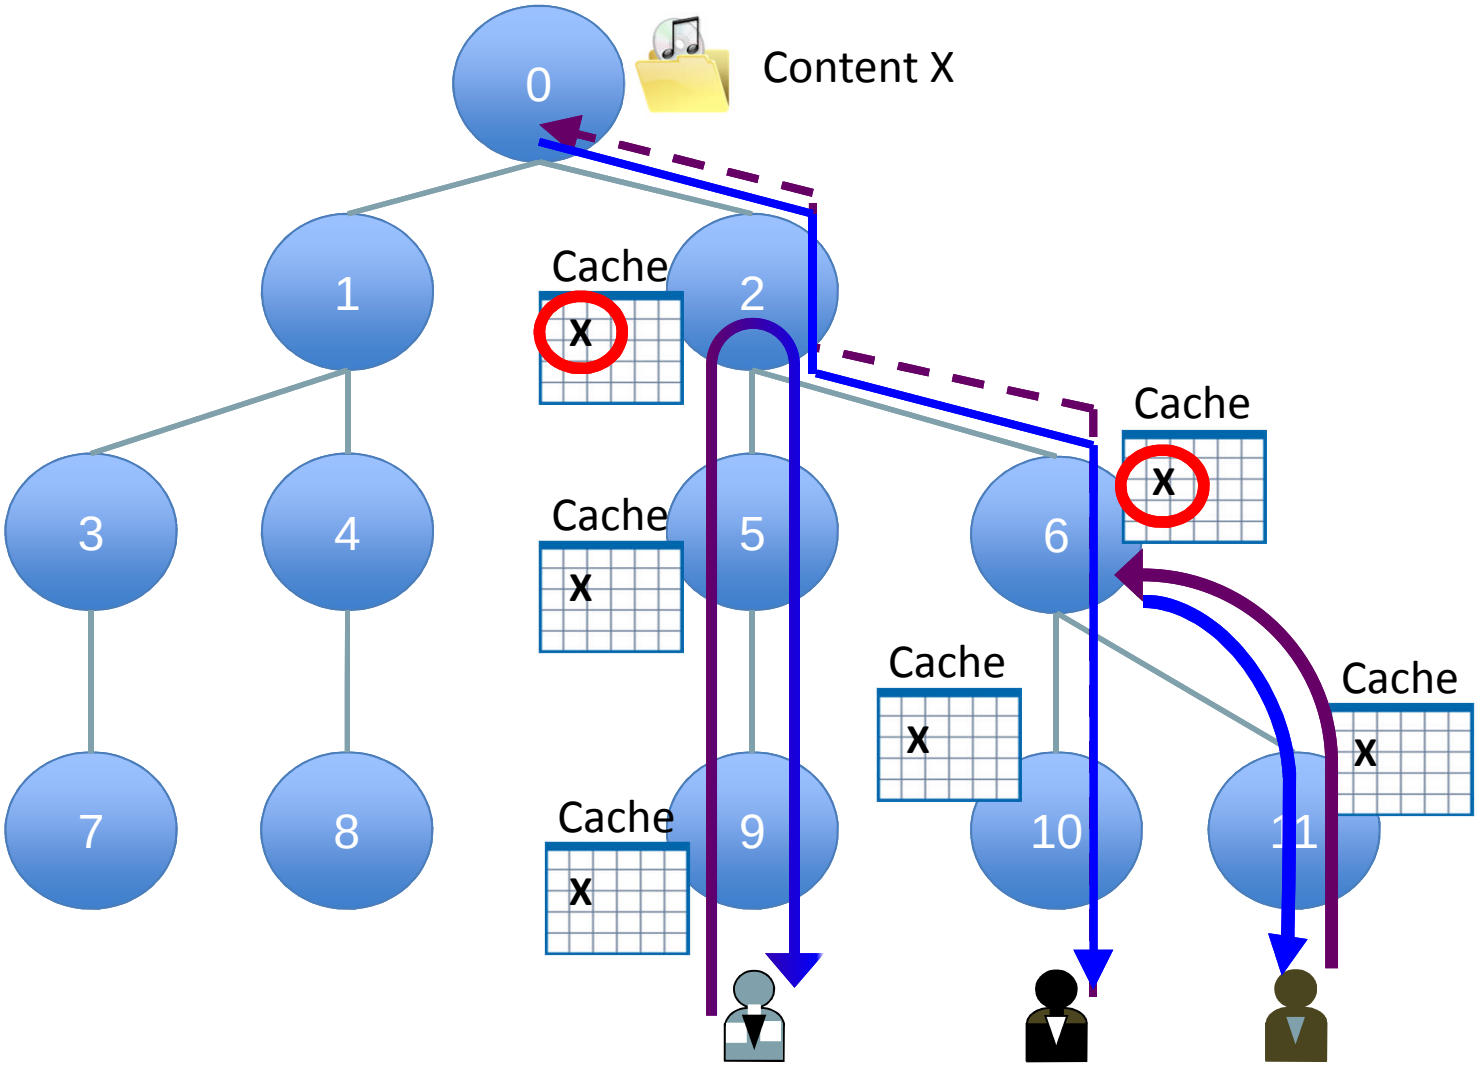

# Universal In-Network Caching Issues

- Assuming an average cache size of 10Gb per router, the time in cache before dropping is as follows:

| LINK NAME | LINK SPEED | 1-SEC OF TRAFFIC | SECS OF TRAFFIC IN A 10GB CACHE |
|---|---|---|---|
| OC-24 | 1,2 Gbps | ~ 0.15 GBs | ~ 64 secs |
| OC-48 | 2,4 Gbps | ~ 0.31 GBs | ~ 32 secs |
| OC-192 | 9,9 Gbps | ~ 1.25 GBs | ~ 4 secs |
| OC-768 | 39,8 Gbps | ~ 5 GBs | ~ 2 secs |
| OC-1536 | 79,6 Gbps | ~ 10 GBs | ~ 1 sec |
| OC-3072 | 159,2 Gbps | ~ 20 GBs | ~ 0.5 secs |

- Indiscriminate universal caching as in CCN ($CE^2$) can be unnecessarily costly and suboptimal
  - **Cached content may be replaced before getting a hit**

# Basic ICN Conventions

- **Request-response at the chunk/packet level**
  - Similar to HTTP-GET but at the network layer
- **Explicitly named content chunks/packets**
  - Named content enables in-network caching
- **Name-based routing** (in native approaches)
  - Hierarchical, flat or mixed names

# In-Network Caching Challenges

- **Cache Placement**
  - Where is it best to put caches?
- **Content Placement**
  - Which content should go where?
- **Request-to-cache routing**
  - How to find cached contents?

# Placement: Cache Less for More (CL4M)

- Considering the simple topology above, only caching at node $v_3$ is meaningful in terms of cache hits
  - ❖ Content cached elsewhere will be simply eventually evicted
- By strategically caching the content at "better" node(s), we can decrease the cache eviction rate and increase cache hit
  - ❖ Note that node $v_3$ is on all the shortest paths from all clients to the server, same as $v_1$ and $v_2$ further upstream

# Placement: Cache Less for More (cont'd)

- ***Betweenness centrality***: measures the number of times a specific node lies on the content delivery path between all pairs of nodes in a network topology.

$$betweenness\ centrality, C_B(v) = \sum_{i \neq v \neq j \in V} \frac{\sigma_{i,j}(v)}{\sigma_{i,j}}$$

  - The network management system calculates the betweenness centrality $C_B$ of every node and "tells" the node about it
  - The highest $C_B$ is recorded in the header of the request packet
  - Response packets/chunks are cached only in nodes whose $C_B$ matches the highest value recorded in the header

- Substantial performance improvement w.r.t. "cache everything everywhere", best paper award in IFIP Networking 2012

# Placement: Probabilistic Caching

- The aim is to achieve fair content multiplexing along a path and make sure all flows get "equal" cache treatment
  - By CE$^2$ we can accommodate content of only few flows
  - Ideally, we would only like to cache contents of a flow only once along the path so that we achieve "fairness"
- A mathematical formula based on the cache capability (path resources) decides probabilistically where to cache
  - Excellent improvement, highly cited paper in SIGCOMM ICN 2012

***ProbCache*: Probabilistic In-Network Caching**

$$ProbCache(x) = \underbrace{\frac{\sum_{i=1}^{c-(x-1)} N_i}{T_{tw} N_x}}_{TimesIn} \times \underbrace{\frac{x}{c}}_{CacheWeight}$$

**Caching Capability of a Path** **Weight-based Caching**

# Placement: Probabilistic Caching (cont'd)

- **TSI: Time Since Inception** *(c in equation)*
  - Increased per hop in the *Request* packet
  - Denotes *hop-distance* server-client
- **TSB: Time Since Birth** *(x in equation)*
  - Increased per hop in the *Data* packet
  - Denotes current hop-distance from server

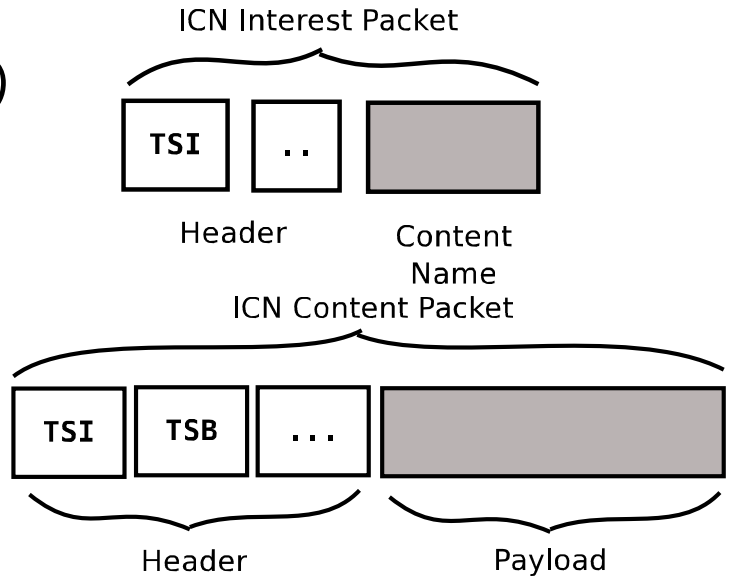


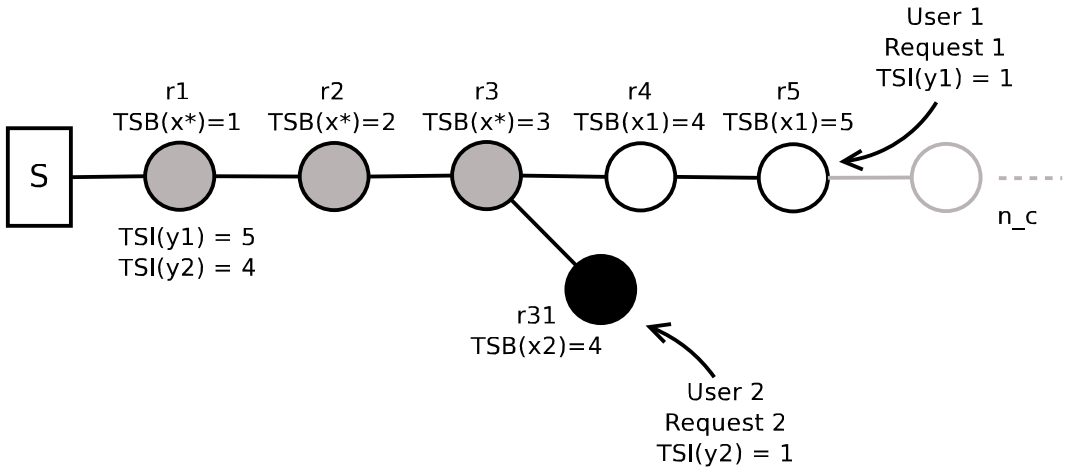

# Routing: Hybrid Hash Routing

- Hash calculated at edge and request redirected accordingly
  - If content not in cache, request forwarded further towards source
- **HR Symm:** Request and content are following the same (non shortest) path – content gets cached in responsible node
- **HR Asymm:** Content follows shortest path back to client – content is not cached in this domain
- **HR Multicast ("Twocast"):** content is replicated at the edge router and sent both to client through shortest path and to responsible cache

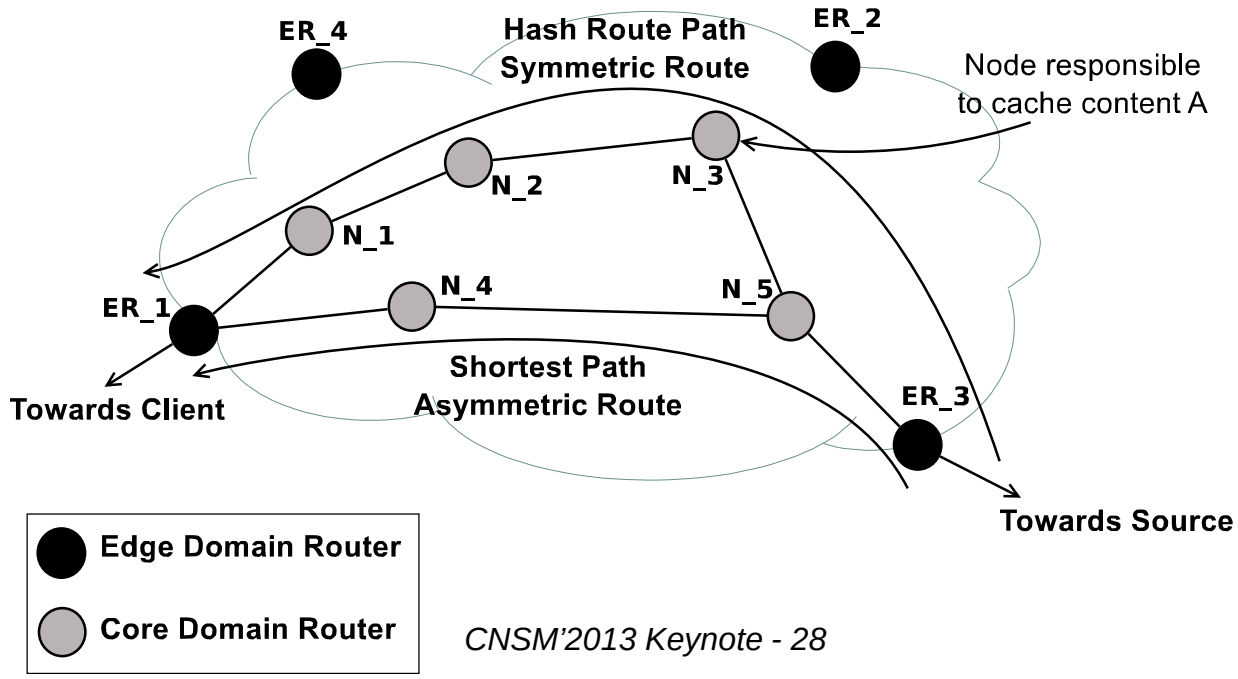

# Routing: Hybrid Hash Routing (cont'd)

- **HR Hybrid Multicast:**
  - If DSM = (d(ER1-N3) + d(N3-ER3)) – d(ER3-ER1) > 0, => HR_Multicast
  - If DSM = (d(ER1-N3) + d(N3-ER3)) – d(ER3-ER1) < 0, => HR_Symmetric
- Very good performance, recent paper in ICN SIGCOMM 2013

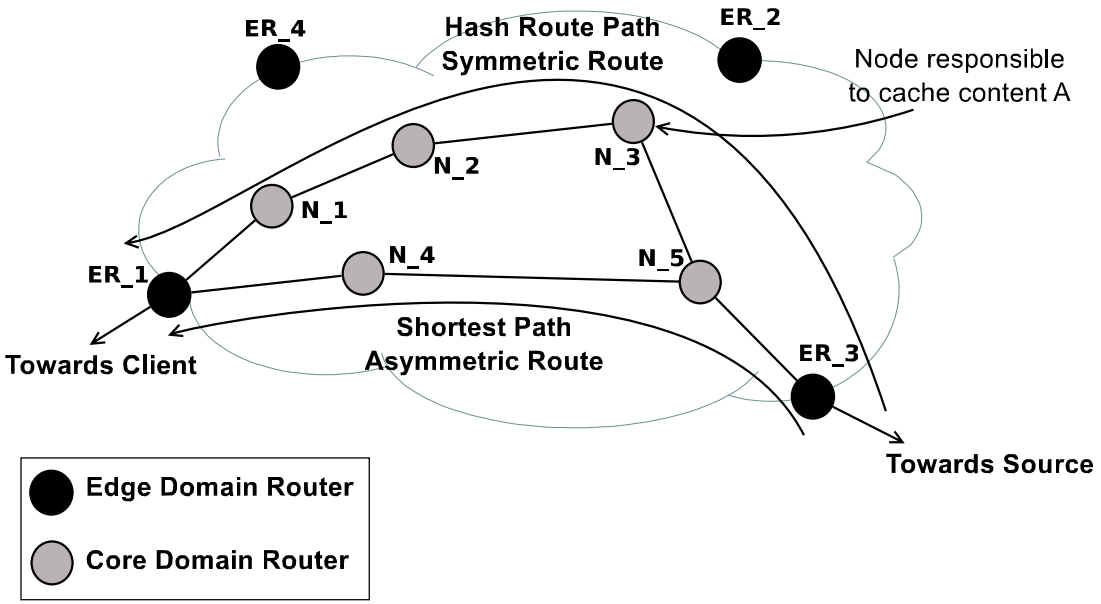

# Placement: ISP as a Limited Capacity CDN

- Limited caching at network edge can strategically cache popular content, alleviating network load and improving QoE

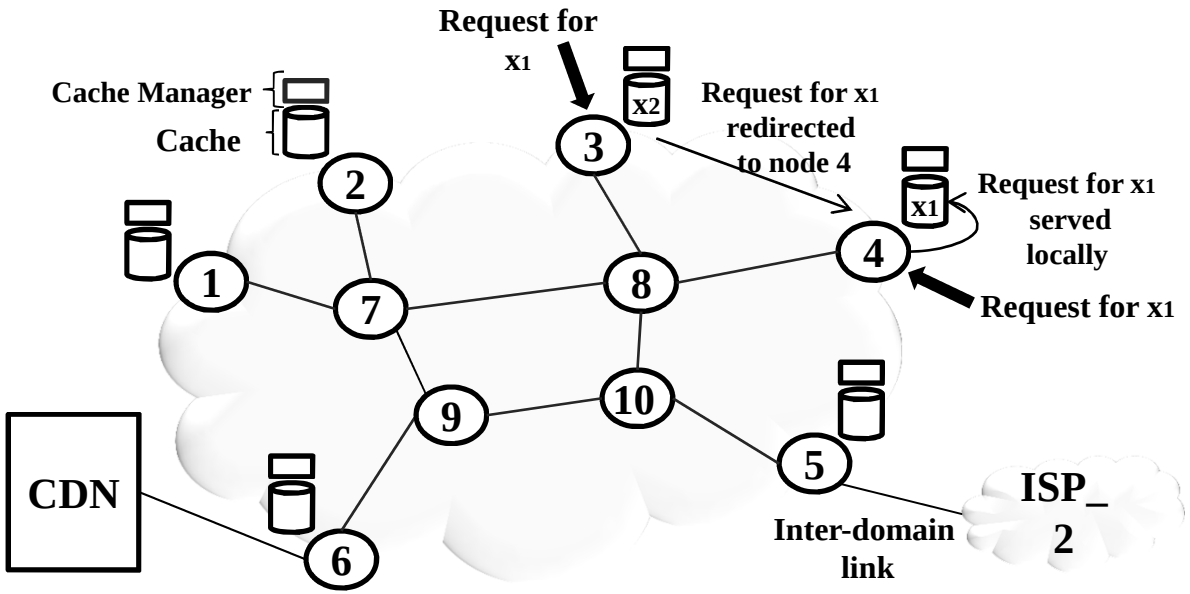


- This requires CDN-ISP cooperation, already happening
  - ❖ Paper in *this* CNSM in the Performance Management session

# Key Publications

## Architectural papers:

- ***CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next Generation Services***
  - » IEEE Communications, Vol. 49, No. 3, March 2011
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Journal-papers/Chai-11a.pdf
- ***Internet-Scale Content Mediation in Information-Centric Networks***
  - » Invited paper that opens the Annals of Telecommunications special issue on Networked Digital Media, Vol. 68, No. 3-4, April 2012
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Journal-papers/Pavlou-12a.pdf
- ***COMET: Content Mediator Architecture for Content-Aware Networks***
  - » Future Network and Mobile Summit 2011, Warsaw, Poland, June 2011
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Conference-papers/Garc-11.pdf

# Key Publications (cont'd)

## Caching papers:

- ***Modelling and Evaluation of CCN-Caching Trees***
  - » Proc. IFIP Networking, Valencia, Spain, May 2011
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Conference-papers/Psar-11.pdf
- ***Cache "Less for More" in Information-Centric Networks***
  - » Proc. IFIP Networking, Prague, Czech Rep., May 2012, ***Best Paper Award***
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Conference-papers/Chai-12a.pdf
- ***Probabilistic In-Network Caching for Information-Centric Networks***
  - » Proc. Sigcomm ICN, Helsinki, Finland, August 2012
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Conference-papers/Psar-12.pdf
- ***Hash Routing Schemes for Information-Centric Networks***
  - » Proc. Sigcomm ICN, Hong Kong, August 2013
  - » http://www.ee.ucl.ac.uk/~gpavlou/Publications/Conference-papers/Saino-13c.pdf
- ***More Control Over Network Resources: an ISP Caching Perspective***
  - » Proc. CNSM, Zurich, Switzerland, October 2013

# Previous and Ongoing UCL Projects

- EU FP7 **COMET** (Content Mediator Architecture for Content-Aware Networks): Jan 2010 – Mar 2013, masterminded by UCL who was the technical leader
  - http://www.comet-project.org/

- EU-Japan FP7 **GreenICN** (Green Information Centric Networking): Apr 2013 – Mar 2016, led by Univ. of Goettingen
  - http://www.greenicn.org/

- UK EPSRC **COMIT** (Active Content Management at Internet Scale): Jan 2014 – Dec 2016

# ICN Research Group in the IRTF

- Recently established group brings together ICN researchers, exchange research results, create a common ICN framework and feed input to existing IETF WGs
- Research topics that are addressed:
  - ICN naming schemes
  - Scalable name resolution for flat names
  - Scalable routing
  - Protocol framework
  - Security
  - API / application design
  - Business, legal and regulatory framework
- Also ITU-T recommendation Y.FNDAN “Framework of Data Aware Networking for Future Networks”

# Key ICN Challenges

- Naming – intricately linked with resolution and ID-based routing, so essential to get it right
- Scalability - cope with at least $10^{15}$ information objects
- Security per object, privacy concerns given that the network "sees" the information objects, spam control
- Manageability, real-time usage data to drive e.g. opportunistic caching based on content popularity
- Incremental deployment, the ability to gradually migrate without obliterating existing IPv4/v6 infrastructure
- Incentives and novel business models to engage involved stakeholders

# ICN Could Make This Much Better!

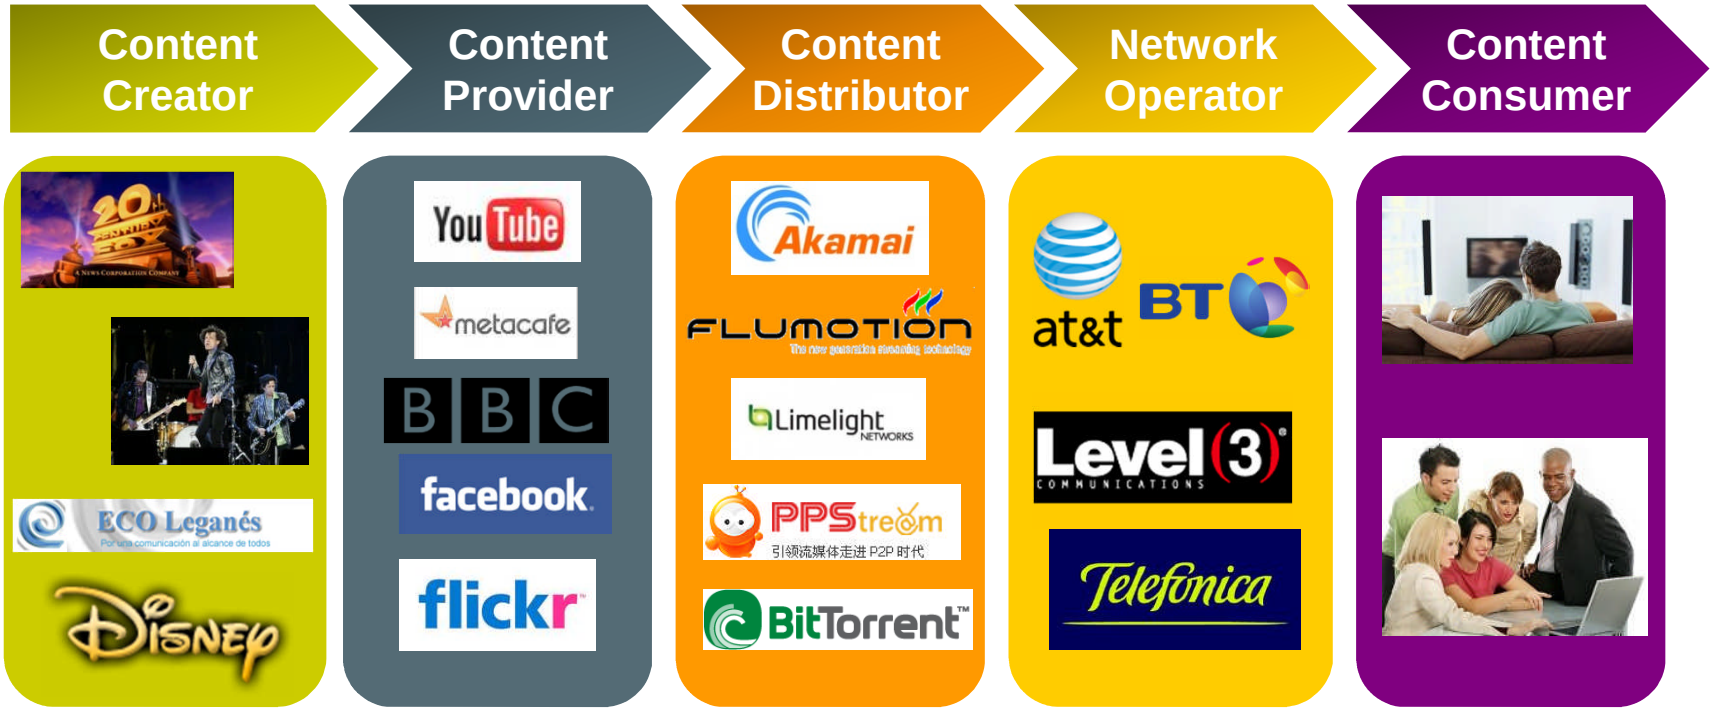


- ICN can provide tangible benefits to most stakeholders in an Internet that will be engineered according to its prevailing use