ICN Support for Multimedia Services: Overview, State 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

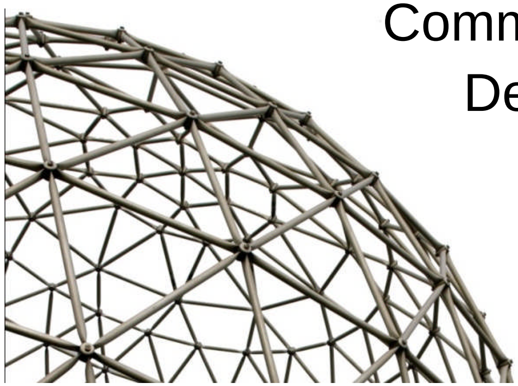

Paradigm Shift Towards ICN

Node-centric design: sharing network resources, addressing nodes

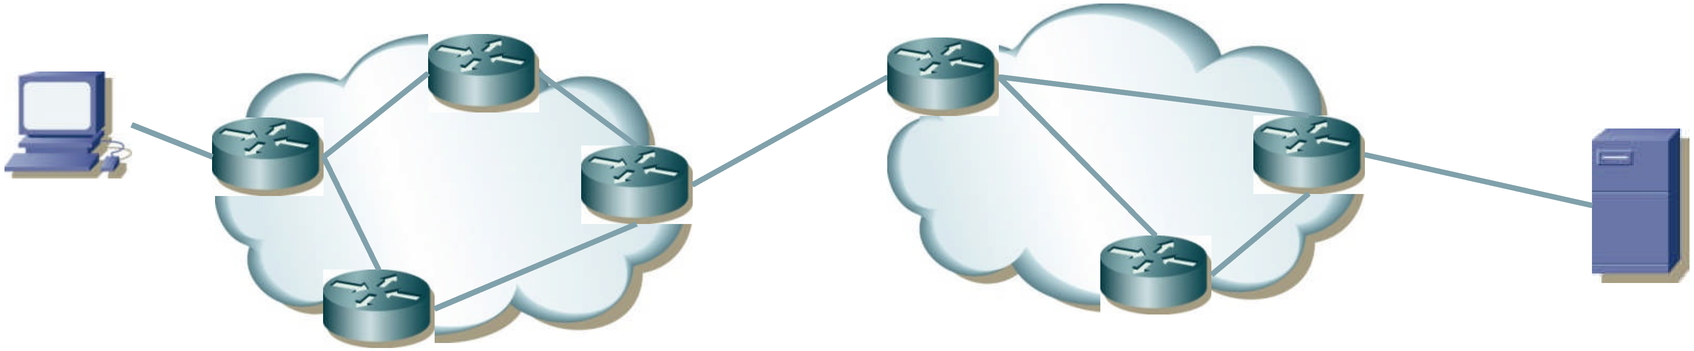

Information-centric design: content access and distribution, naming content

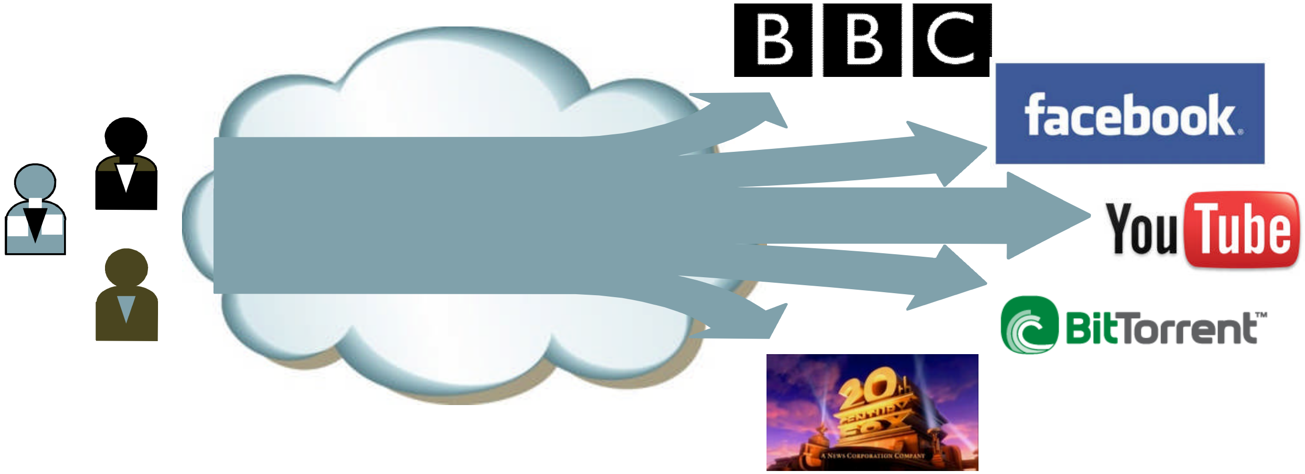

Business Case and Initiatives

- P2P traffic of the early 2000s was replaced by video traffic from 2005 onwards and gave rise to CDNs
  - Both these paradigms promote “named content”
  - Massive revenue streams for content and CDN providers
- ISPs/INPs not happy to just continue “pushing bits”
  - No sustainable network growth, try to get into content market
- A number of research projects and initiatives towards ICN
  - US: UCB DONA and PARC CCN/NDN
  - EU 4WARD/SAIL NetInf
  - EU PSIRP/PURSUIT PubSub and COMET CMP

Information-Centric Networking

- Given that users are interested in named content, consider a clean architectural approach to redesign the current Internet in order to better support content access & distribution
  - Provide a P2P/CDN-like paradigm within the network

- **Information-Centric Networking (ICN)** targets 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 that supports natively multicast
  - Also location transparency, mobility & intermittent connectivity

Flash-Crowd Effect Due to Content Popularity

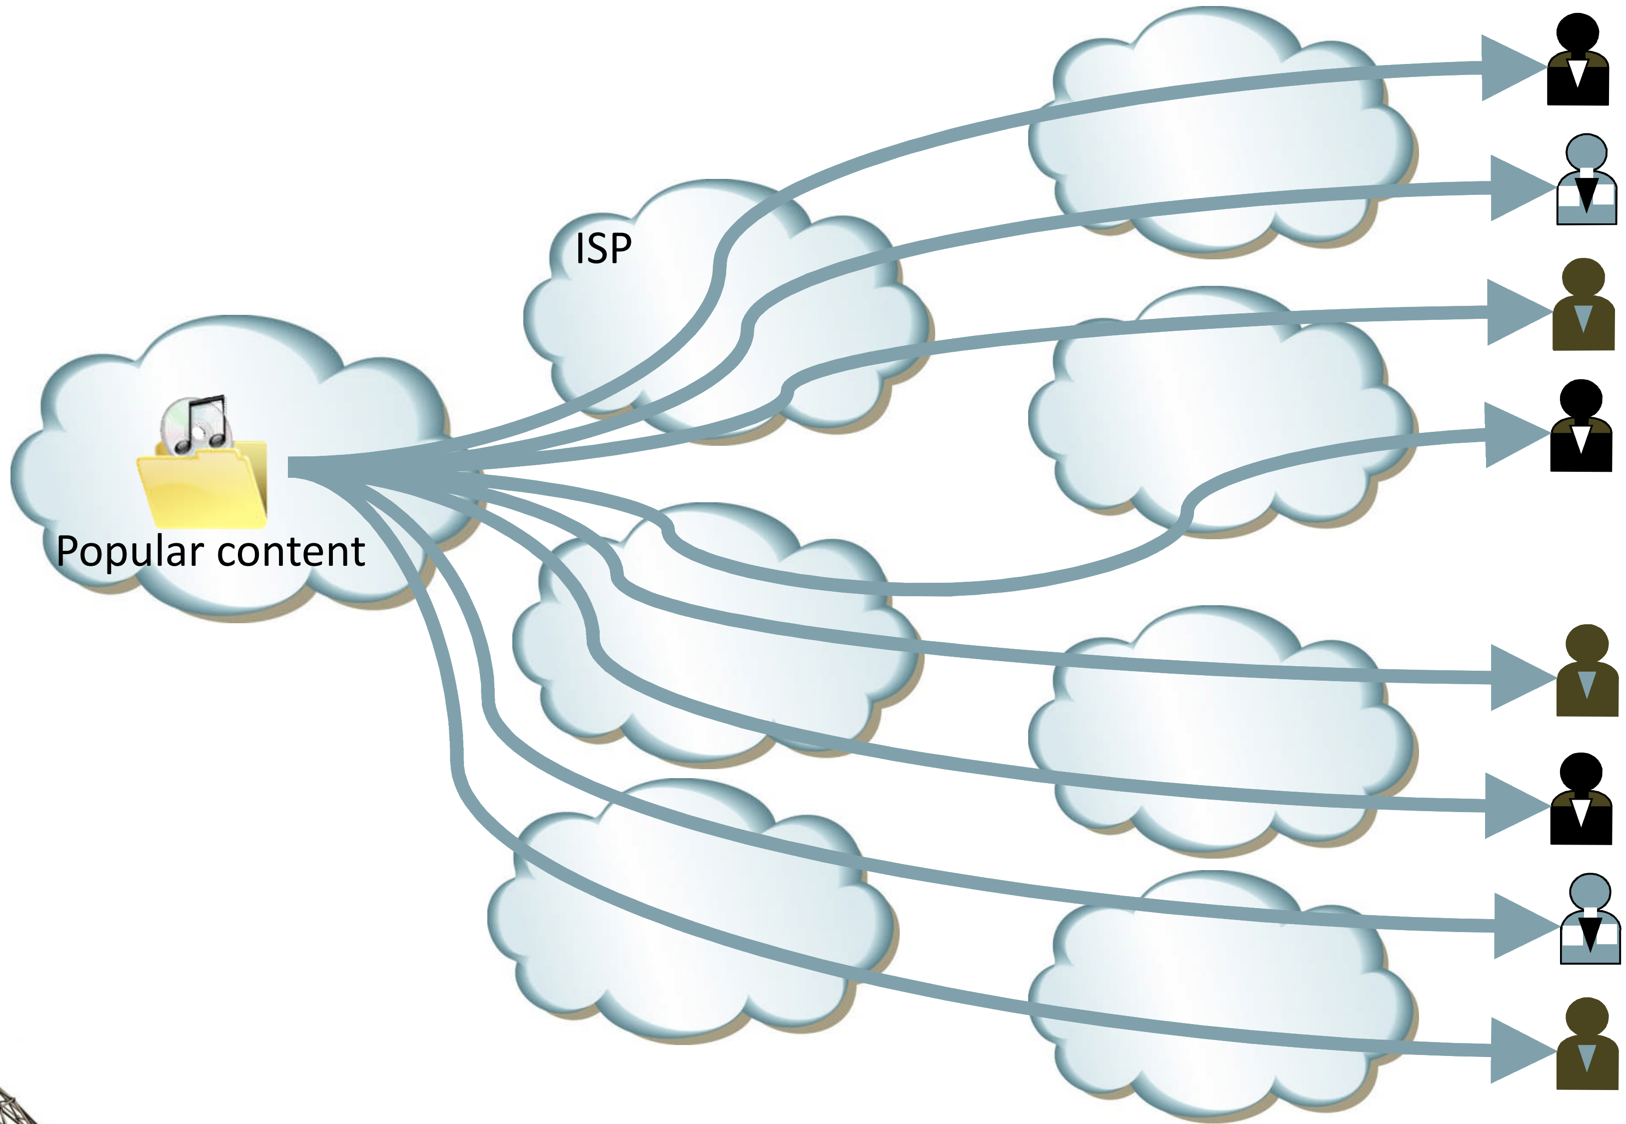

Scalable Cache-based Content Distribution

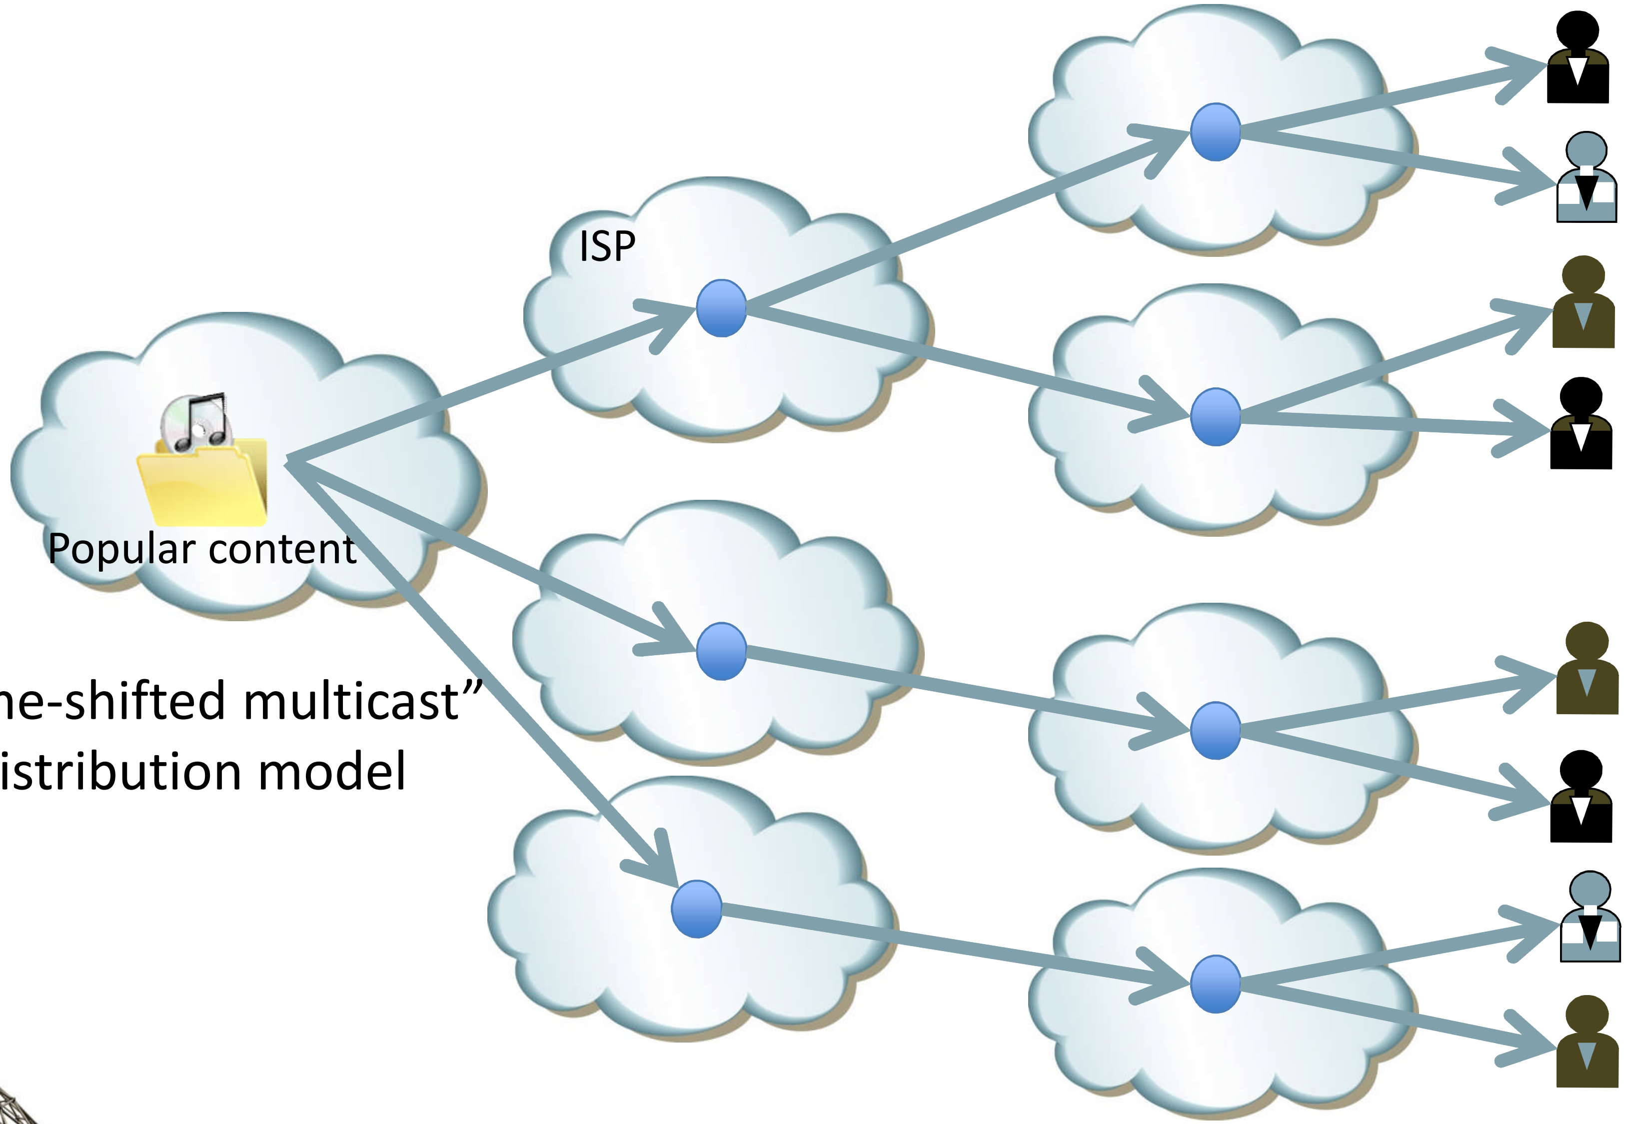

"Time-shifted multicast" distribution model

ICN Caching Benefits for Video Streaming

- What we can achieve by ICN-based in-network caching for video streaming:
  - Maximum caching benefits: ↓ 22% traffic, ↓ 34% server load
  - Significant caching benefits for large object caches i.e. > 100 GB

Grandl, Su, Westphal, "On the Interaction of Adaptive Video Streaming with Content-Centric Networking", *Proc. Packet video*, Oct 2013.

Key ICN Features

- **Request-response per content chunk**
  - Similar to HTTP-GET but <u>per chunk at the network layer</u>
- **Explicitly named content objects/chunks**
  - Named content enables in-network caching
- **Name-based routing**
  - Hierarchical, flat or mixed names
  - Still unsolved problem because of content & Internet scale

Content-Centric Networking (CCN)

- *Interest*/*Data* packets flowing in a “reverse ack/data” style
  - Data packets are cached everywhere along the delivery path (Cache Everything Everywhere – CE^2) for other consumers
  - Least Recently Used (LRU) packet discard policy
- Pending Interest Table (PIT) keeps reverse path state, aggregates interests and supports multicast
- Forwarding Information Base (FIB) implements the routing table based on hierarchical content names

- CCNx implementation publicly available, the NSF NDN project looks at CCN-related research issues

CCN-like Universal In-Network Caching

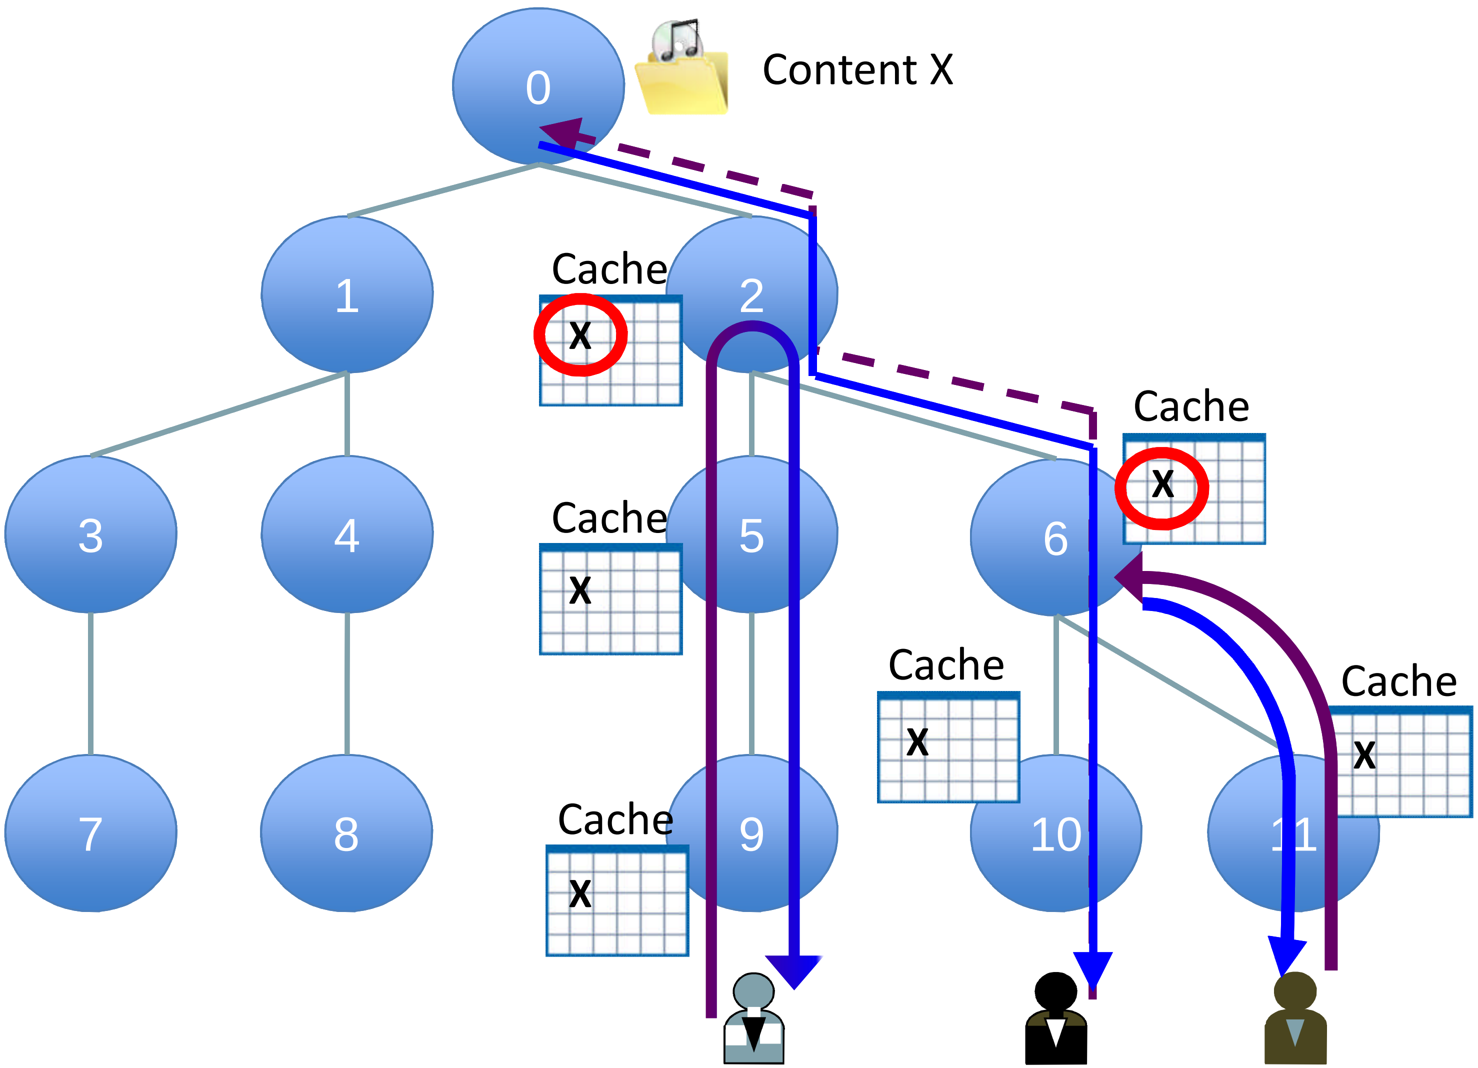

Universal In-Network Caching Issues

- 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	$\sim$ 0.15 GBs	$\sim$ 64 secs
OC-48	2,4 Gbps	$\sim$ 0.31 GBs	$\sim$ 32 secs
OC-192	9,9 Gbps	$\sim$ 1.25 GBs	$\sim$ 4 secs
OC-768	39,8 Gbps	$\sim$ 5 GBs	$\sim$ 2 secs
OC-1536	79,6 Gbps	$\sim$ 10 GBs	$\sim$ 1 sec
OC-3072	159,2 Gbps	$\sim$ 20 GBs	$\sim$ 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**

In-Network Caching Challenges

- **Router Cache Placement**
  – Where is it best to put caches?
  – Not all nodes can possibly be cache-capable
- **What to cache and where?**
  – Cache everything or most popular content?
  – Which content should be cached where?

Placement: Cache Less for More (CL4M)

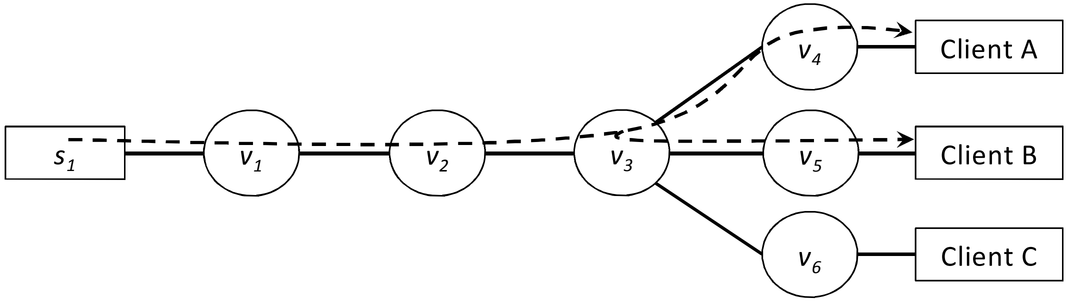

- Considering the simple topology above, only caching at node v_3 is meaningful in terms of cache hits
  - ❖ Content cached elsewhere will be simply eventually evicted

- By strategically caching the content at “better” node(s), we can <u>decrease the cache eviction rate and increase cache hit</u>
  - ❖ 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**

Why Evolutionary ICN Architectures

- IP represents too much investment and cannot be easily changed, especially with a radical content-oriented protocol!
  - Even IPv6 penetration has been quite small until now
- No significant changes since BGP policy routing in 1991 and CIDR in 1993
  - MPLS has been a compatible addition rather than replacement
- Evolutionary = not just overlay approaches, but retain IP as thin “hourglass waste” and BGP inter-domain routing
- Cannot see radical approaches been deployed but can see gradual changes towards this direction

CURLING

- Can direct requests to replicated content objects in servers “a la CDN” and also to the cached content chunks in Content-Aware Routers (CARs) at the network edge
- Supports multicast for real-time content streaming
- Identifies popular content and, as such, can scale to the expected massive size of the future content echosystem
- It is IP-compatible and allows for incremental deployment
    1. First as pure overlay for “global CDN” operation
    2. Then through CARs at network edges to support multicast and exploit in-network caching
    3. Eventually, in the long term, through native ICN operation

Content Mediation Plane (CMP)

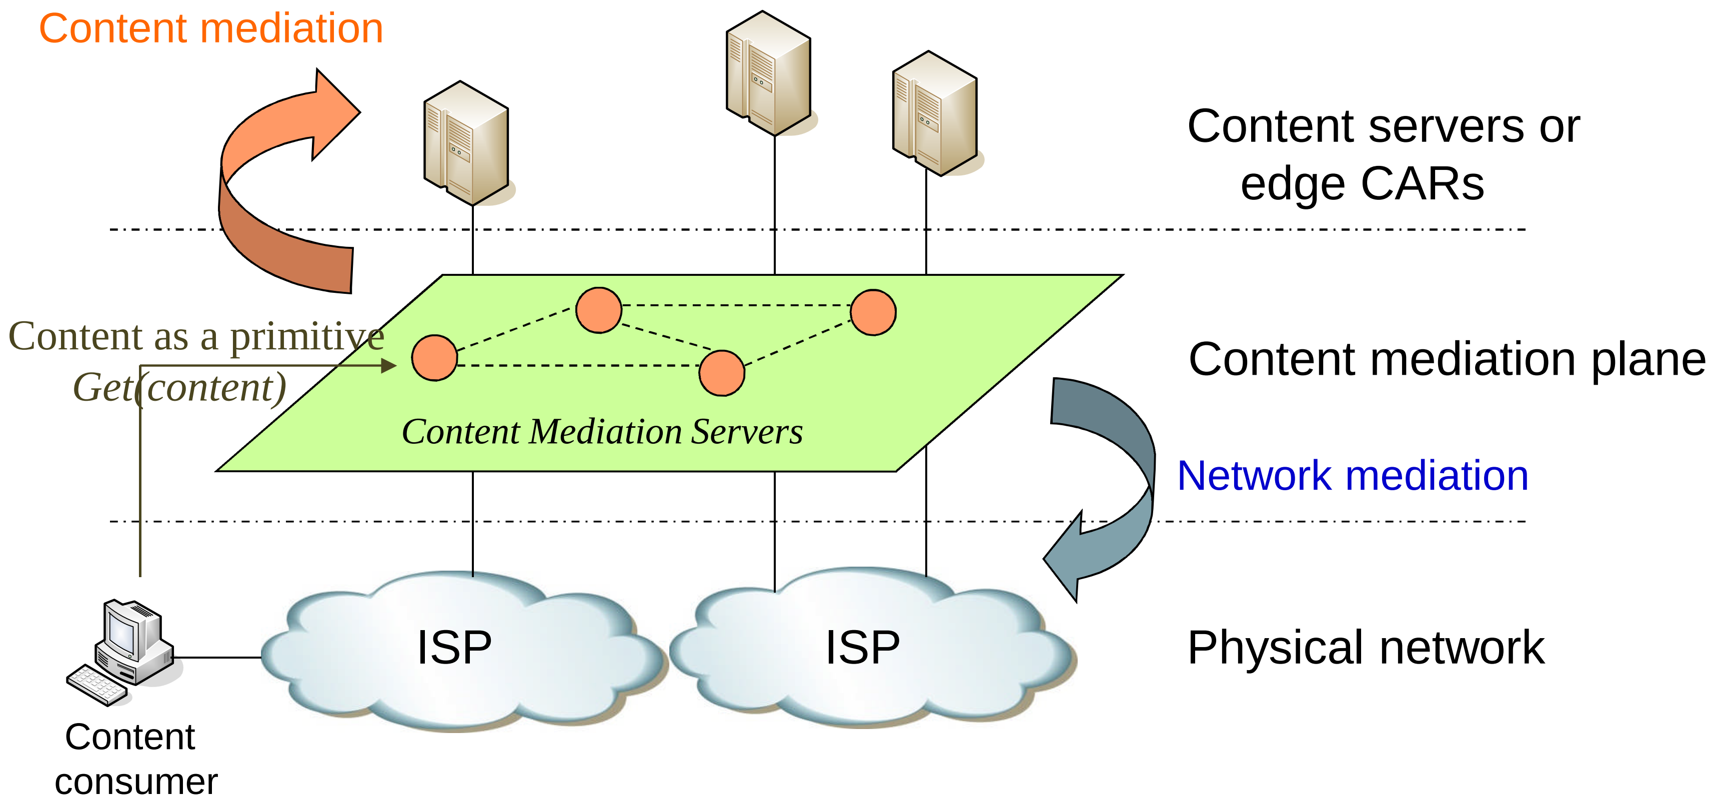

The CMP finds the right content copy (content mediation) and possibly prepares the network for contentID-based transmission (network mediation for popular content caching and live content)

Key Features

- Content is publicised following the inter-domain “provider route” forwarding rule
  - Like DONA but only uphill and not “p2p” – better scalability
- Content requests are resolved uphill to tier-1 domains and then downhill to domain where content lies
  - If a content chunk is cached, it is served from that CAR
- For live content, CMSs install states in edge CARs so that transmission is based on contentID and IP tunnels
- Provider route rule has a “tromboning” effect, so route optimisation follows with shortcut paths learned from BGP
  - Prefer customer over peering and peering over provider routes
- Network mediation is only for popular content – only such content is cached in CARs for performance / scalability

Content Publication

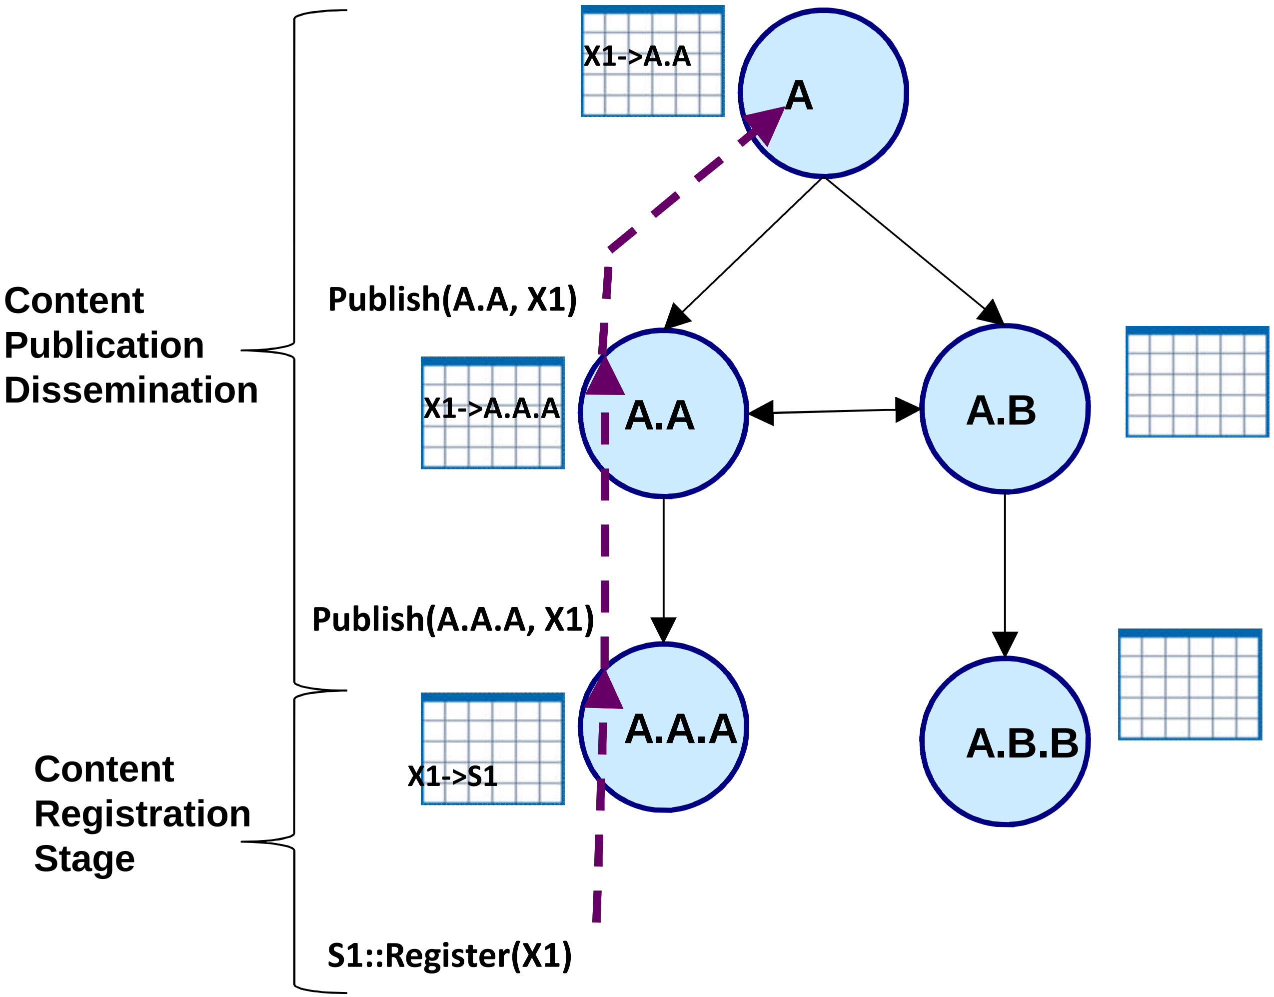

Content Resolution

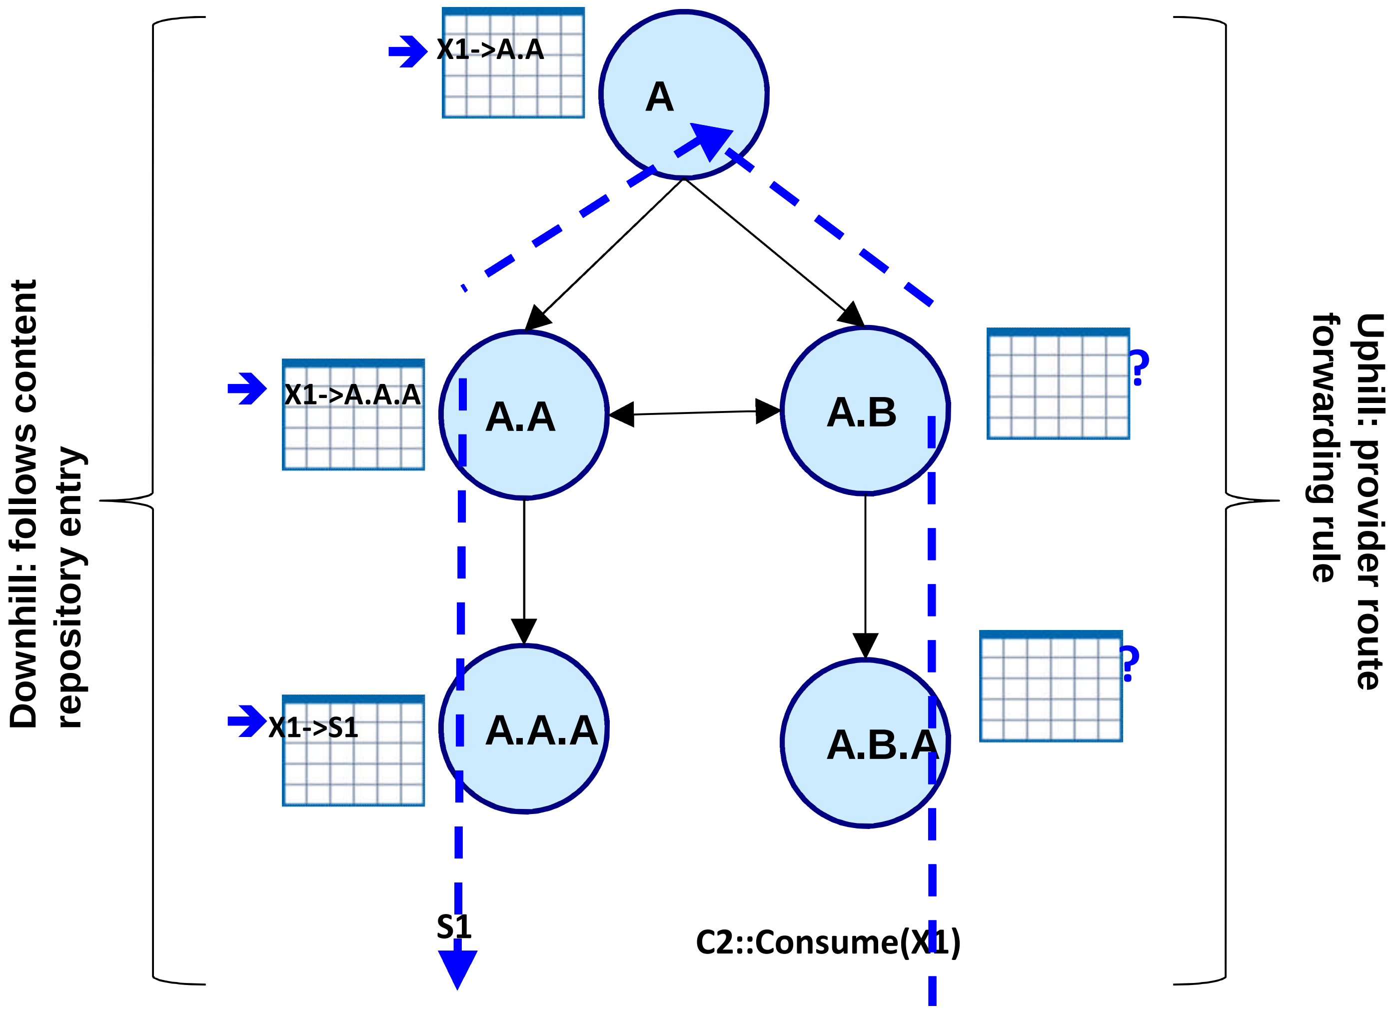

Incremental Deployment Stages

1. The CMP is deployed as a DNS++, in a pure “content-resolution-only” fashion for “advanced CDN operation”
   – No network changes at all, only changes to end systems
2. CARs are deployed *at the domain edges* only to support edge-network caching and multicast for popular content only and contentID-based forwarding
   – Routers use IP/BGP and intra-domain edge-to-edge tunneling
3. CARs are also deployed *within the network* to enable native in-network content caching, resolution & delivery
   – Support for native information-centric networking functionality, possibly with a protocol like CCN
   – The mediation plane effectively “collapses” onto the network

ICN-based Streaming Challenges I

- Adaptive video coding cannot be easily supported in multisource transfers
  – Retrieving content from multiple in-network caches (or replication points) cannot support multiple video rates
  – The situation gets more complicated in case of multiple users requesting the same content but with different access bandwidth
  – Example: (Client-to-Cache)←→(Client-to-Server) oscillations

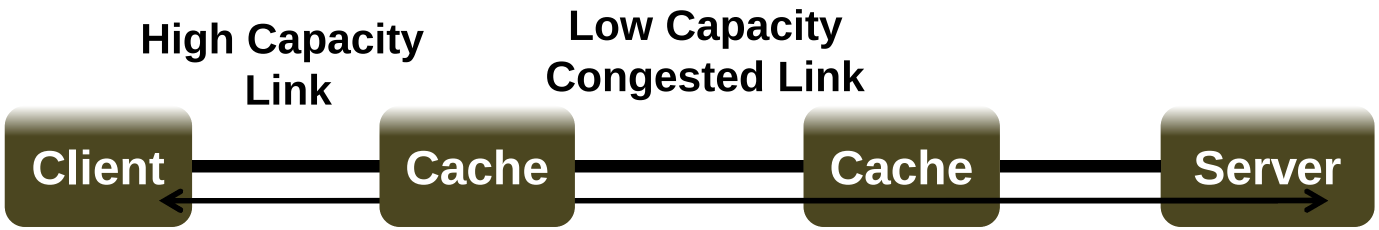

Yu, Bronzino, Fan, Westphal, Gerla, "Congestion-Aware Edge Caching for Adaptive Video Streaming in Information-Centric Networks", Proc. IEEE CCNC, Las Vegas, USA, Jan 2015

Lederer, Mueller, Rainer, Timmerer, Hellwagner, "Adaptive streaming over content centric networks in mobile networks using multiple links," IEEE ICC workshop, Jun 2013

ICN-based Streaming Challenges II

- Video traffic volume inflates caching resource requirements
- Caching redundancy (a general problem in in-network caching) affects more severely real-time video transfers
- Caching redundancy in video transfer along the same path effectively cancels any in-network caching benefit

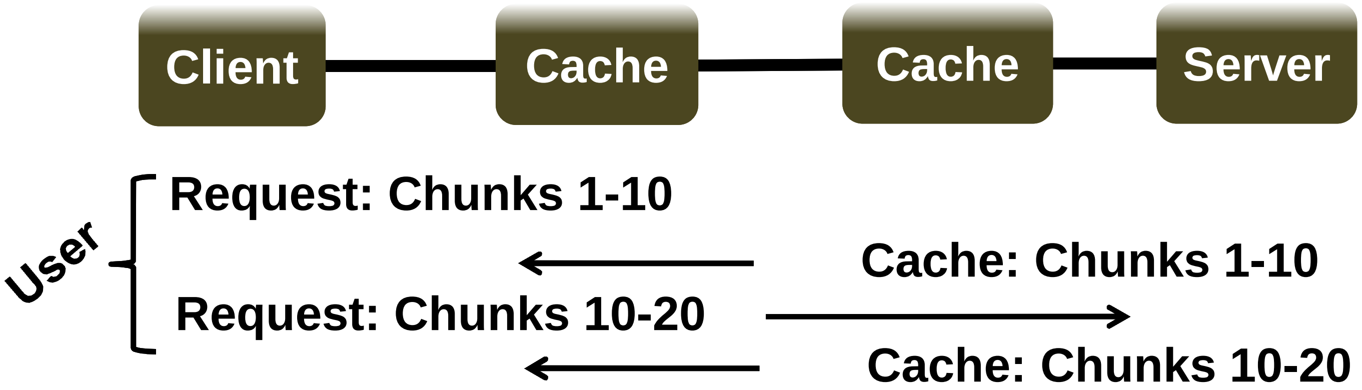

Zhe Li et al., "Time-Shifted TV in Content Centric Networks: The Case for Cooperative In-Network Caching," IEEE International Conference on Communications (ICC), vol., no., pp.1,6, June 2011

ICN-based Streaming Challenges III

- Access logging by content providers not supported when accessing video content from in-network caches
  - A more general problem of ICN, but more pressing for video delivery given the huge revenue streams of IPTV, VoD platforms, etc.
- Header overheads i.e., impact of content names, signatures
- Missing appropriate pipelining mechanism
  - Interplay with yet-to-be-specified ICN transport layer

Benefits / Opportunities I

- Multicast natively supported
- Receiver driven communication enables estimation of traffic conditions ➔ network support for media adaptation
- Devices can switch between different interfaces (cellular, WiFi etc.) and request data according to the fastest incoming link – this is more difficult with IP

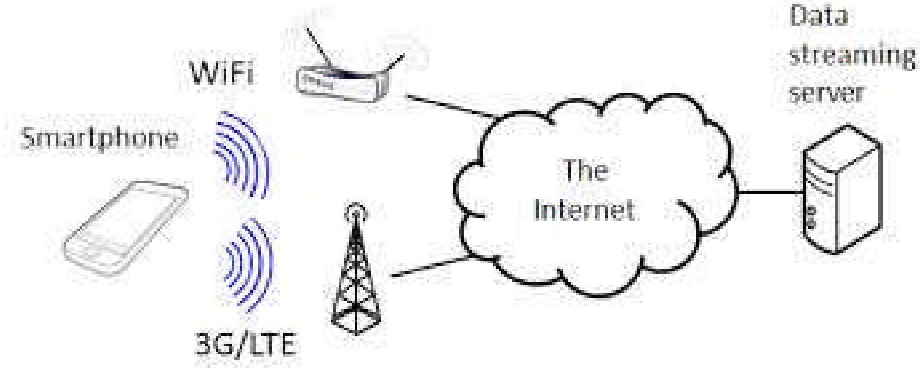

Detti A. et al. "Offloading cellular networks with Information-Centric Networking: The case of video streaming“, Proc. of IEEE WoWMoM 2014

Detti A., Ricci B., Blefari-Melazzi N., "Peer-to-peer live adaptive video streaming for information centric cellular networks“, Proc. Of IEEE PIMRC 2013

Benefits / Opportunities II

- Application-level semantics can be gracefully integrated within the network (i.e. at the network layer)
  - Automatic embedding of subtitles or adverts
- In-network service instantiation
  - Named-function networking
  - Service function chaining

Arumaithurai M., Chen J., Monticelli E., Fu J., Ramakrishnan K.K., “Exploiting ICN for flexible management of software-defined networks”, Proc. of ICN 2014

Sifalakis M., Kohler B., Scherb C., Tschudin C., “An information centric network for computing the distribution of computations”, Proc. of ICN 2014

Key Publications

Architectural papers:

- *CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next Generation Services*
  - IEEE Communications, Vol. 49, No. 3, March 2011
- *Internet-Scale Content Mediation in Information-Centric Networks*
  - Annals of Telecommunications, Vol. 68, No. 3-4, April 2012

Caching papers:

- *Modelling and Evaluation of CCN-Caching Trees*
  - Proc. IFIP Networking, Valencia, Spain, May 2011
- *Cache "Less for More" in Information-Centric Networks*
  - Proc. IFIP Networking, Prague, Czech Rep., May 2012, ***Best Paper***
- *Probabilistic In-Network Caching for Information-Centric Networks*
  - Proc. Sigcomm ICN, Helsinki, Finland, August 2012
- *Hash Routing Schemes for Information-Centric Networks*
  - Proc. Sigcomm ICN, Hong Kong, August 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
  - https://www.ee.ucl.ac.uk/comit-project/

Summary

- ICN is an emerging new network paradigm that exploits in-network caching for scalable content distribution
  - Can also result in better user QoE
- Can really benefit multimedia streaming but there is still a number of challenges to overcome
  - Also benefits and opportunities
- Radical emerging ICN approaches may take a long time to be deployed
  - Hence we should strive for evolutionary approaches so that we can enjoy the benefits in the short-to-medium term