

Global Rendezvous Scale Issues for an Information-Centric Future Internet



George C. Polyzos

polyzos@acm.org

Department of Computer Science and Engineering
University of California, San Diego

&

Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business, Greece



with

Konstantinos Katsaros

and the AUEB MMLab PURSUIT team

Motivation



Clean-Slate Internet Design

- At the **beginning**...
 - ◆ Cooperation/no competition...
 - ◆ NO commercial traffic!
 - ◆ Endpoint-centric services
- **Now**...
 - ◆ Content distribution...
 - >50% of traffic today is video↑
 - ◆ P2P, Overlays... DPI by ISPs...
 - ◆ Trust? Endpoint trust?
 - viruses, phishing, DoS attacks...
 - ◆ E2E?
 - NAT, firewalls, middleboxes, CDNs
 - ◆ Current net economics favor sender
 - ◆ Tussles...
 - e.g.: privacy vs. accountability

Information-Centric Networking


• **Connecting Wires**

- ◆ the past...



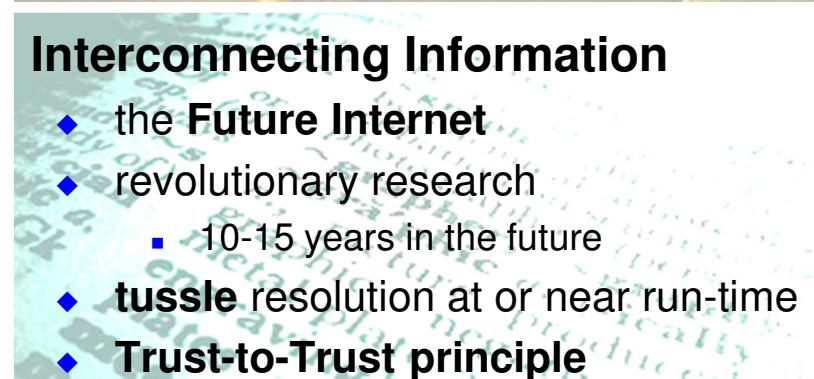
• **Interconnecting Computers**

- ◆ the current **Internet**
- ◆ evolutionary development
- ◆ ... started decades earlier



• **Interconnecting Information**

- ◆ the **Future Internet**
- ◆ revolutionary research
 - 10-15 years in the future
- ◆ **tussle** resolution at or near run-time
- ◆ **Trust-to-Trust principle**



Information-Centric Networking (ICN)



- The network becomes aware of **what** it delivers
 - ◆ Focus on information objects, not end-hosts
e.g., name-based routing, DHT-like routing
- Users simply denote the desired piece of information
 - ◆ network responsible for locating the information & delivering it
- Benefits:
 - ◆ Decoupling location from identity
 - Inherent mobility support
 - ◆ Enabling caching (at multiple levels: object, chunk, packet...)
 - ◆ Efficient delivery structures, e.g., multicast
 - ◆ Restoring imbalance of powers between sender and receiver
 - Resistance to Spam, DoS attacks



ICN research efforts



- FP7 **PSIPR/PURSUIT** ([http:// www.fp7-pursuit.eu/](http://www.fp7-pursuit.eu/))



- NSF **Named Data Networking** (<http://www.named-data.org/>)



- NSF FIND, **Cache-and-Forward Architecture**
(<http://www.winlab.rutgers.edu/docs/focus/CNF.html>)

- FP7 **4WARD** (<http://www.4ward-project.eu>)

- FP7 **SAIL** (<http://www.sail-project.eu/>)



- FP7 **COMET** ([http:// www.comet-project.org](http://www.comet-project.org))



- FP7 **COAST** (<http://www.coast-fp7.eu>)



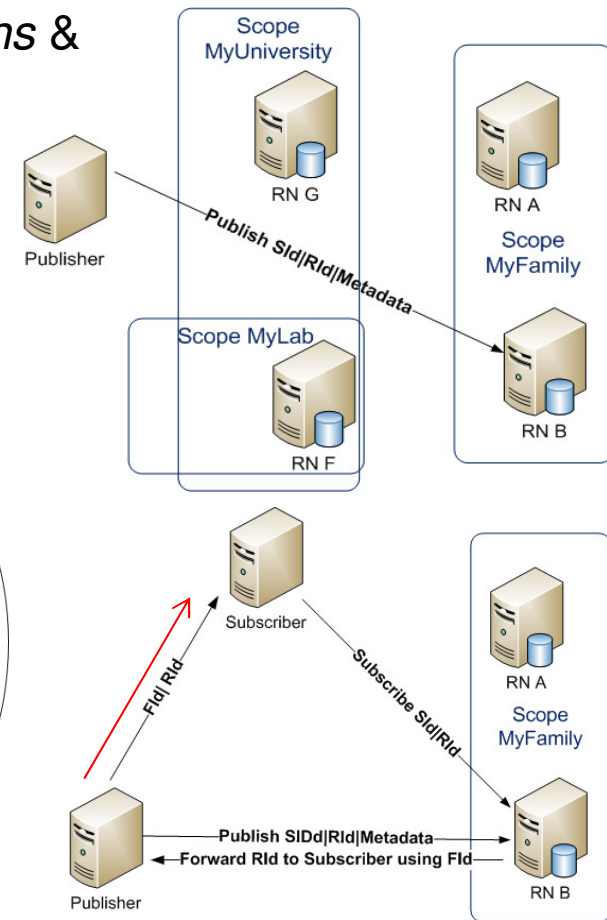
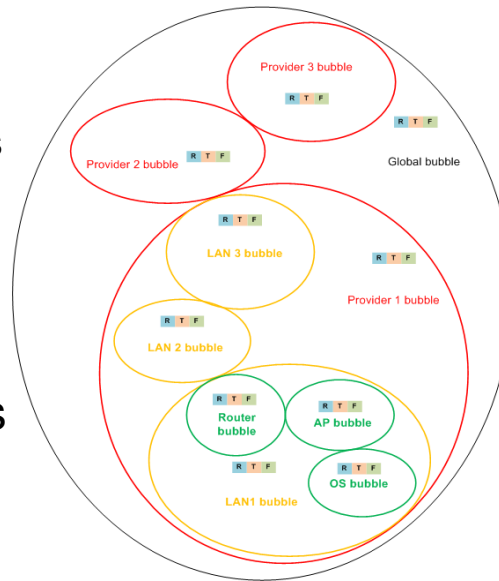
- FP7 **Convergence** (<http://www.ict-convergence.eu>)



OceanStore (SIGARCH'00), **I3** (SIGCOMM'02), **FARA** (FDNA'03), **VRR** (SIGCOMM'06),
ROFL (SIGCOMM'06), **DONA** (SIGCOMM'07), **CCN** (CoNEXT'09), **LIPSIN** (SIGCOMM'09)

The Publish-Subscribe Internet Architecture

- **Rendezvous**: Matches *publications* with *subscriptions* & initializes the forwarding process
- **Topology**: Monitors the network & creates information delivery paths
- **Forwarding**: Implements information delivery
- Applied recursively...
 - ◆ local, global rendezvous
 - ◆ slow path/fast path rendezvous
- IDs: Rendezvous ID, Scope ID, Forwarding ID...
- **Separation** of functions
- 2 prototype implementations
 - ◆ Blackhawk (**PSIRP**)
 - ◆ Blackadder (**PURSUIT**)



➤ N. Fotiou, G.C. Polyzos, D. Trossen, “**Illustrating a Publish-Subscribe Internet Architecture**,” *Telecommunication Systems*, Springer, vol. 52, no. 3, Special Issue on ‘**Future Internet Services and Architectures: Trends and Visions**,’ Online: 23/2/2011.

Locating Content

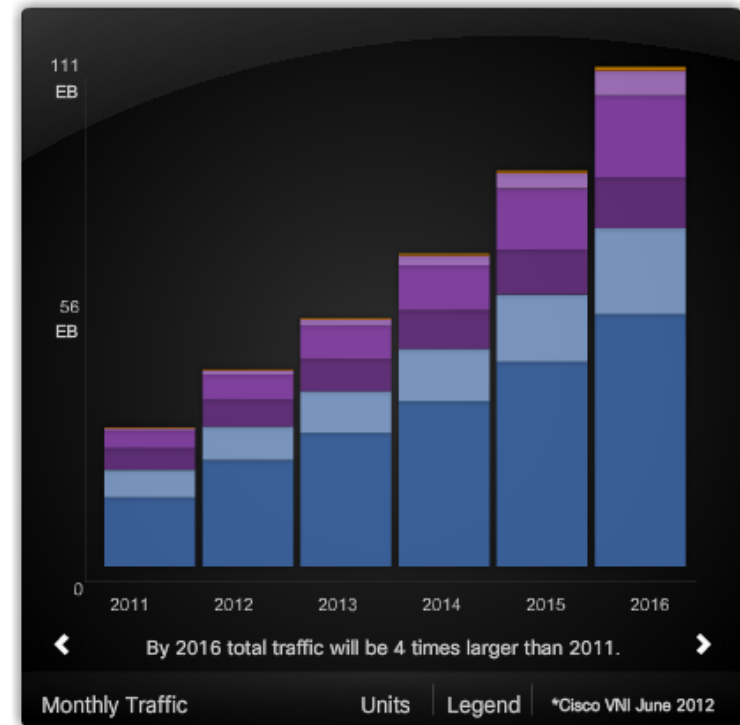


- Searching for information
 - ◆ Crawling sources of information
 - ◆ Indexing for lookup purposes
- *Locating* information
 - ◆ Name resolution:
 - From information identifiers to the content itself
 - Various approaches
 - Directory lookup service e.g., DNS
 - Integrated with routing e.g., CCN architecture

Scale



- Huge volumes of content
 - ◆ Google indexed web pages: ~1 trillion [1]
- Huge Volumes of traffic, [2]:
 - ◆ *“Global IP traffic has increased eightfold over the past 5 years, and will increase threefold over the next 5 years”*
 - ◆ *“Global IP networks will deliver 12.5 petabytes every 5 minutes in 2016.”*
 - ◆ *“Globally, Internet video traffic will be 55 percent of all consumer Internet traffic in 2016, up from 51 percent in 2011.”*




[1] Google Official Blog, *“We knew the Web was big”*, July 2008. URL: <http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>


[2] Cisco, *“Visual Networking Index: Forecast and Methodology, 2011-2016”*, May 2012

Complexity comes with Size

- Diversity of access devices
 - ◆ Desktops, Mobile phones, sensors, home appliances, TVs, etc.
 - ◆ *“The number of devices connected to IP networks will be nearly three times as high as the global population in 2016 “ [1]*
- Large number of Networks
 - ◆ More than 40K ASes [1]



Internet Users in the World
by Geographic Regions - 2011



Source: Internet World Stats - www.internetworldstats.com/stats.htm
 Estimated Internet users are 2,267,233,742 on December 31, 2011
 Copyright © 2012, Miniwatts Marketing Group

[1] Cisco, *“Visual Networking Index: Forecast and Methodology, 2011-2016”*, May 2012
 [2] CAIDA, *“Internet Topology”* <http://www.caida.org/research/topology>.



ICN Scaling Issues

- Information-Centric Networking
 - ◆ scale wrt information objects (content)
- Global Rendezvous
 - ◆ rendezvous at multiple levels (recursive architecture)
 - e.g.: OS, local, AS/network, ... global
- Future Internet
 - ◆ scale also wrt devices in a Future Internet
 - IoT/sensors, info appliances, etc.
- Multi-homing and Multicast need also be considered
- most current FI research looks at local/small scale
- most critical issues relate to scaling at the global level

Open Questions & Current approaches



- Open Questions
 - ◆ How does (inter-)network heterogeneity impact mechanisms employed to locate the content & performance?
 - ◆ What is the impact on network/ICN performance of the name resolution system (existing or new)?
 - ◆ How would/does the exchange of information between different actors affect the operation of such a system in terms of reachability of content?
- Current approaches
 - ◆ Confined in specific technological contexts
 - Neglecting diversity across multiple axes
 - e.g., mobile vs. fixed, CDNs, etc.
 - ◆ Investigations restricted in size
 - Number of networks, end hosts...

The need for richer modeling



- Need to *simultaneously* consider:
 - ◆ Content generation processes
 - ◆ Popularity evolution: both in time and space
 - ◆ Locality characteristics of end users
 - ◆ CDN practices
 - ◆ Inter-domain AND intra-domain topologies
 - ◆ Routing policies
 - ◆ Access technology heterogeneity
 - ◆ ...
- All this in a non trivial scale
 - ◆ So as to also reveal interactions

References



- M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, “ROFL: routing on flat labels,” Proc. ACM SIGCOMM 2006. DOI=10.1145/1159913.1159955 **[not efficient to impose a DHT without some underlying shortest path substrate]**
- T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, I. Stoica. “A data-oriented (and beyond) network architecture,” *ACM SIGCOMM Computer Communication Review* , vol. 37, no. 4, August 2007. **[key paper on name resolution in ICN]**
- P. Jokela, A. Zahemszky, C. E. Rothenberg, S. Arianfar, and P. Nikander, “LIPSIN: line speed publish/subscribe inter-networking,” Proc. ACM SIGCOMM, Barcelona, Spain ,2009. DOI=10.1145/1592568.1592592 **[PSIRP forwarding]**
- V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard (PARC), “Networking Named Content,” Proc. ACM CoNEXT '09, 2009. DOI=10.1145/1658939.1658941.
- N. Fotiou, G.C. Polyzos, D. Trossen, “Illustrating a Publish-Subscribe Internet Architecture,” *Telecommunication Systems*, Springer, vol. 52, no. 3, Special Issue on ‘Future Internet Services and Architectures: Trends and Visions,’ (published online: February 2, 2011; DOI: 10.1007/s11235-011-9432-5). **[introduction to the overall PSIRP/PURSUIT model and architecture]**
- J. Rajahalme, M. Särelä, K. Visala, and J. Riihijärvi, “On name-based inter-domain routing,” *Computer Networks*, vol. 55, no. 4 March 2011.
- K.V. Katsaros, N. Fotiou, X. Vasilakos, C.N. Ververidis, C. Tsilopoulos, G. Xylomenos, G.C. Polyzos, “On Inter-domain Name Resolution for Information-Centric Networks,” Proc. 11th International Conference on Networking (IFIP/TC6 Networking), Prague, Czech Republic, May 2012.
- G.F. Marias, N. Fotiou, G.C. Polyzos, “Efficient Information Lookup for the Internet of Things,” Proc. First IEEE WoWMoM Workshop on the Internet of Things: Smart Objects and Services (IoT-SoS), San Francisco, CA, June 2012 .
- N. Fotiou, K. Katsaros, G.C. Polyzos, M. Sarela, D. Trossen, G. Xylomenos, “Handling Mobility in Future Publish-Subscribe Information-Centric Networks,” *Telecommunication Systems*, Springer, Special Issue on ‘Mobility Management in the Future Internet’ (to appear).
- D. Perino and M. Varvello, “A reality check for content centric networking,” Proc. ACM SIGCOMM workshop on Information-Centric Networking (ICN 2011), Toronto, Canada, August 2011. **[on feasibility]**
- A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox, “Information-centric networking: seeing the forest for the trees,” Proc. 10th ACM Workshop on Hot Topics in Networks (HotNets '11), Cambridge, Massachusetts, USA. DOI=10.1145/2070562.2070563 **[challenges ICN]**



Thank you!



Global Rendezvous Scale Issues for an Information-Centric Future Internet

George C. Polyzos

polyzos@acm.org

<http://mm.aueb.gr/>

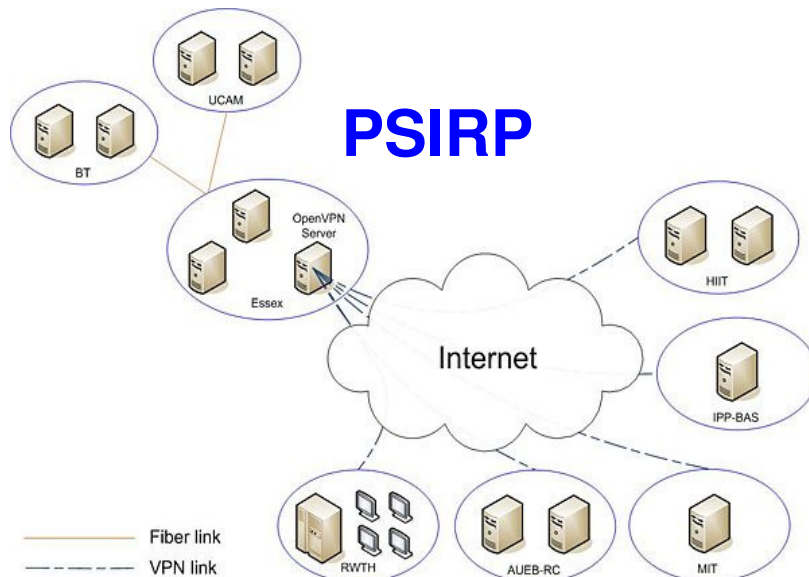


Department of Computer Science and Engineering
University of California, San Diego

&

Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business, Greece

Prototype Implementations & Testbeds



PSIRP Testbed (w/ Blackhawk)

- 6 countries: UK, FI, GR, D, BU, US
 - In addition: Belgium during ICT demos
- Tunneled over the public Internet
 - **+dedicated fiber** where available

Current & future work

- Socket emulator, multipath transport, voice & multimedia conferencing over PSI
- Wireless/mobile testbed extensions

PURSUIT Testbed (w/ Blackadder)

- 25 nodes
- 5 countries: UK, FI, GR, D, US
- Tunneled (VPN)
 - over the public Internet

