Future Network & MobileSummit 2011 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2011
ISBN: 978-1-905824-25-0

COMET: Content mediator architecture for content-aware networks*

Gerardo GARCÍA[1], Andrzej BĘBEN[2], Francisco J. RAMÓN[1], Adrián MAESO[1], Ioannis PSARAS[3], George PAVLOU[3], Ning WANG[4], Jarosław ŚLIWIŃSKI[2], Spiros SPIROU[5], Sergios SOURSOS[5], Eleftheria HADJIOANNOU[6]

[1]*Telefónica Investigación y Desarrollo, Ronda de la Comunicación, s/n, Madrid, 28050, Spain, Tel: +34 91318700, Fax: + 34 91 4023137, Email: ggdb@tid.es*

[2]*Warsaw University of Technology, Nowowiejska 15/19, Warsaw, 00-665, Poland Tel: +48 22 2347642, Fax: + 48 22 234 7231, Email: abeben@tele.pw.edu.pl*

[3]*Dept. of EE&E, University College London, Torrington Place, WC1E 7JE, London, United Kingdom, Tel: +44 (0)20 7679 7626, Email: i.psaras@ee.ucl.ac.uk*

[4]*University of Surrey, Guildford, Surrey, GU2 7XH, United Kingdom Tel: +44 1483 686005, Fax: +44 1483 686011, Email: N.Wang@surrey.ac.uk*

[5]*Intracom Telecom, 19.7 Markopoulou Ave., Peania, 190 02, Greece Tel: +30 210 667 7960, Fax: +30 210 667 7101, Email: spis@intracom.com*

[6]*PrimeTel PLC, The Maritime Center, Omonoias 141, Limassol, 3045, Cyprus, Tel: +357 25 867118, Email: eleftheriah@prime-tel.com*

Abstract: This paper presents the vision of the EU FP7 project COMET (COntent Mediator architecture for content-aware nETworks), which aims at defining a novel content-oriented Internet architecture for simplifying content access and supporting content distribution in the network in a content-aware fashion. The COMET architecture will deal with content as a primitive, facilitating unified access to any type of content regardless of location and way of distribution. It will provide global content naming as well as infrastructure for resolving content names to identifiers required for content access. Moreover, this content-centric approach will take advantage of network, and server awareness, as well as user and content characteristics, in order to improve efficiency of content delivery.

Keywords: mediation, content awareness, content publication, content consumption, content naming, content routing, content distribution, QoS, streaming, multimedia, Future Internet.

1. Introduction

In recent years there has been a growing proliferation of user-generated Internet content, including blogs, photos, video, etc. The increasing trend of users generating their own content has led to an abundance of intermediaries: content is published through user websites, social networks such as Facebook or MySpace, photo sharing sites such as Picasa and Flickr, pre-recorded media aggregators such as YouTube or GoogleVideo, content delivery networks such as Akamai or Limelight, or through P2P overlay networks such as BitTorrent or eMule. In the near future, massive content generation will not only come from end users. Companies like Blockbuster, Netflix and Apple are already providing movies

* *The research leading to these results has received funding from the EU FP7 COMET project - Content Mediator Architecture for Content-Aware Networks , under Grant Agreement 248784*

through the Internet. Live content is also expected to explode once the network capabilities allow it.

The main problem with the current approach to content access [1] is that, due to the increasing number of intermediaries, a lot of content tends to be accessible only by particular user communities, with **global content search and direct access being fragmented**. Specifically, the key issue lies in the lack of global content naming scheme and infrastructure to access the content, which forces end-users to search the content through the relevant intermediary, maintaining a multitude of accounts, front-ends, tools and applications in order to discover, access and consume content. Moreover, with the current approach, the **access to content needs to be machine and application-dependent**. Content access requires knowing the server hosting the content and the application/session protocol used to distribute that content. This prevents content mobility, as well as the adaptation of the type of distribution in a transparent way for the end-user. A global naming scheme providing content naming persistence would bring naturally such content mobility and adaptation, making end-users unaware of any changes in content location or application protocols. On the other hand, **today's networks are unaware of the content they are transporting** [2]. Due to this unawareness, networks cannot apply the most appropriate end-to-end transport strategy for the content in order to achieve the best quality of experience for the end users. Moreover, crowded live events that can potentially lead to network congestion peak times cannot be efficiently transported. Nevertheless, even if the networks were well prepared, intermediaries usually are not aware of the network capabilities, traffic conditions, as well as the content transfer requirements. Therefore, the content is delivered far from the most efficient way.

Discussed above limitations motivate transition from the current Internet to a new "content–centric" network, which treats delivery of content as a primary paradigm. Recently, several new concepts for improving content delivery have been investigated in [2], [3], [4], [5] and [6]. They cover different approaches for routing by content name, content caching in network nodes, location independent and self-certifiable content addressing schemes, scalable content resolution algorithms, and publish/subscribe systems.

The EU FP7 COMET project (COntent Mediator architecture for content-aware nETworks) takes a unified approach to content location, access and distribution, irrespective of the intermediary used. COMET is designing a global naming scheme along with mechanisms for optimizing both content source selection and distribution, by mapping content to appropriate network resources based on transmission requirements, user preferences, and network state.

This paper describes a high-level view of the architecture of the COMET system, focusing on the concepts of content mediation and awareness. The paper is organized as follows: section 2 presents the COMET concept, section 3 highlights the main operations of the COMET system (publication and consumption), section 4 describes the basic functional blocks of the overall architecture, section 5 shows two different approaches to implement the COMET architecture, section 6 deals with the QoS and network aspects of the COMET architecture needed to provide content delivery with QoS. Finally, section 7 presents the conclusions and future work related to the COMET project.

2. The COMET concept. Rationale of the Content Mediation

We propose an approach for content access in the Future Internet architecture based on the concept of **mediation**. As presented in Fig. 1, this mediation lies in the provision by Internet Service Providers (ISP) of an intermediate plane between the world of content and the world of data transmission. On the one hand, this mediation plane will act as mediator for content publication, offering an interface for content publication and thus becoming aware of content characteristics, QoS requirements, etc. (**content awareness**), as well as the available content sources and their performance (**server awareness**). Moreover, since the mediation plane is provided by ISPs, it can be aware of the network topology and the available routing paths between content servers and clients (**routing awareness**), as well as the network conditions (**network awareness**).

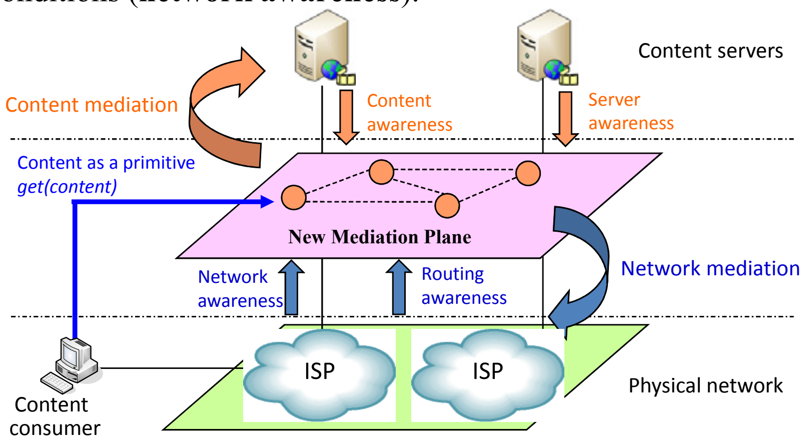

Figure 1: The concept of COMET mediation plane.

On the other hand, the mediation plane allows ISPs to act as mediators for content distribution, offering a common interface for content consumption, so that content is treated as a first-class citizen in the Internet. Thanks to the awareness achieved, ISPs will be able to decide the best sources and paths to deliver the content. Besides, since ISPs are naturally able to mediate with the network, it would be possible to instruct the network (network mediation) in order to improve content delivery in terms of quality and effective bandwidth utilization.

The mediation plane can identify QoS requirements for the content delivery, select the appropriate content server based on server and network conditions, select an appropriate distribution mode (unicast, multicast, etc.) based on server load and available bandwidth, select the most appropriate path to transmit the content from the server to the end user, and enforce QoS and path constraints onto network elements, once the server, the path and the distribution mode have been selected.

Key technical advantages that can be achieved thanks to this mediation are:

- Unified access to the content whatever its nature and location.
- Content delivery with guaranteed QoS.
- Point-to-multipoint content delivery, reducing bandwidth needs for live contents.
- Graceful handover of the content delivery path, providing more resilience and flexibility for multi-homed users.
- Advanced publication mechanisms, allowing Content Providers to update content servers on-the-fly, while switching among different ways of distribution.

3. Content publication and consumption in COMET

This section details how content publication and consumption work in COMET. On the one hand, Content Publication in COMET starts with the **Content identifier allocation**. The COMET System identifies the content by using Content-IDs or Content Names. The Content-IDs are machine readable, while Content Names plays the role of human readable aliases. Note that the allocation of the Content-IDs and Content Names may be performed in a different way, as discussed in [3] and [4], but it has to result in a globally unique Content-ID and/or Content Name allocated for particular content. Then, one or several content copies are placed in the Content Servers (**Content Storage).** Although, this process is out of the project scope, it is worth to mention that content usage statistics collected in COMET system may be helpful for optimisation of content distribution. Note that, nevertheless, the network location of the content copies has to be made known to the COMET system, something that is done by the next sub-operation, the Content Registration. **Content registration** consists in the creation by the COMET System of a relation between Content-ID and network locations of the content copies. It can be performed only after the previous two sub-operations. The outcome is a situation where particular content may be used in Content Consumption.

Content Consumption consists of one sub-operation performed all the time, independently from consumption requests (**Awareness**), and two other sub-operations invoked per Content Consumption request (**Content Resolution** and **Content Delivery**).

The **Awareness** sub-operation gathers the current information about routing (topology), server conditions and network conditions. The outcome of this sub-operation is the knowledge, which is essential in making correct decisions during Content Consumption. Awareness is covered by three COMET processes. First, the process of providing **routing awareness** observes the network topology by new and existing mechanisms, e.g., business relations, BGP routing tables and other protocols for gathering of network reachability information. Second, the process of providing **server awareness** gathers the information about servers' conditions. This may include CPU load, number of active streams, traffic load over network interface cards, etc. Note that this process may interact with existing solutions for server monitoring. Third, the process of providing **network awareness** gathers the information about conditions of links and/or paths in the network, e.g., link load and packet transfer delay metrics. This process may rely on and interact with specific solutions for network monitoring, external to the COMET system.

The **Content Resolution** sub-operation consists in the location of the content copy and the preparation of the network for Content Delivery. It starts with the request from the Content Client of a specific content from a Content-ID and/or a Content Name. The outcome is the situation where the COMET System decides which server and which path in the network should be used for particular Content Consumption. The whole sub-operation is accomplished by four processes. First, the process of **name resolution** locates the content by using Content-ID and/or Content Name. Second, the process of **path discovery** obtains, for each specific requested content, the properties of the paths from Content Servers to the Content Client, using the output of the process of providing routing awareness. Third, the **decision process** combines the results of above two processes to select the best server and path for particular Content Consumption. Fourth, the process of **path configuration** is responsible for enforcing the decision at the network level.

Finally, Content Delivery deals with the delivery of the content to the Content Client according to the decisions made during Content Resolution. This sub-operation relates mainly to the content-aware forwarding process, which involves low-level functions in the network nodes forwarding the content.

4. Overall COMET architecture

To achieve the above operations, COMET follows a two-plane approach. The upper plane is the **Content Mediation Plane** (CMP), which is mainly responsible for name and content resolution as well as the preparation of path used for content delivery. The lower plane of the COMET system is the **Content Forwarding Plane** (CFP) and is mainly responsible for delivering the content back to the content consumer. This is done based on mediation performed by the CMP taking into account the information about server and path conditions.

Figure 2 depicts the COMET Architecture, which consists of the following elements:

- **Two planes**, the upper Content Mediation Plane (CMP) and the lower Content Forwarding Plane (CFP), as already discussed before. A large part of the COMET control functionalities are in the CMP, providing the intelligence needed for the network

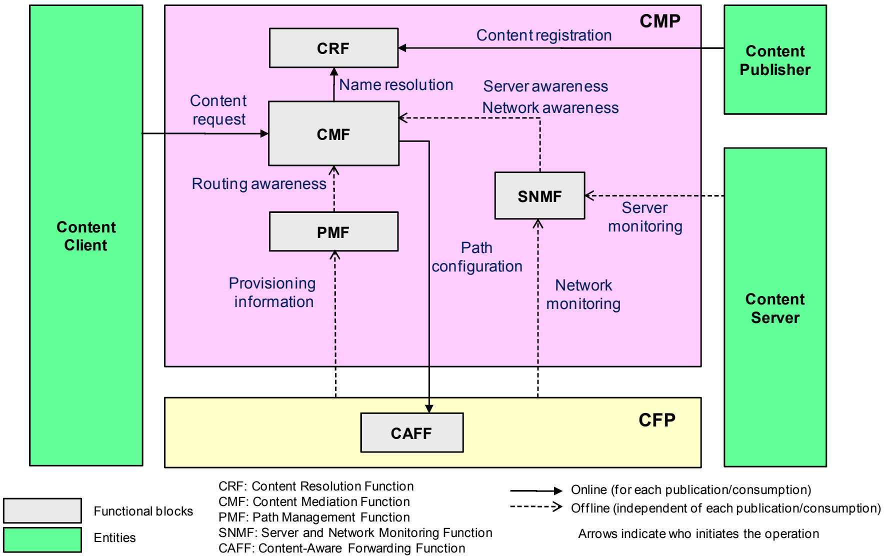

Figure 2: COMET Architecture.

to become content-aware, while the CFP implements the data plane functionalities to deliver content according to instructions provided by the CMP.

- **A number of functional blocks**, which accomplish the operations of the COMET system from content publication to content consumption. These functional blocks are grouped into one or more entities, as we explain later on in next section.
- **A number of entities** (i.e., Content Client, Content Server and Content Publisher), which interact with the functional blocks in the COMET system in order to i) publish content (i.e., register content to the COMET system), ii) request for content and iii) provide content-server monitoring information, as depicted in Fig. 2.

Next, we describe the functional blocks that build up the COMET architecture. The functional blocks included in the **Content Mediation Plane** are the following:

- **Content Resolution Functional block (CRF)**: This function is triggered during the content publication process as well as content consumption operations. During the content publication operation, the CRF component is contacted by the Content Publishers every time a new content is to be published to the COMET system. Therefore, the main task of the CRF is to keep track of the content records or indexes. It is clear that for scalability purposes not all CRFs have knowledge of all the contents of the Internet/COMET system, but each CRF is keeping track of the records of its area of responsibility. This includes allocation of content identifiers, as well as the registration of new unpublished content. The second main responsibility of the Content Resolution Function is to resolve Content Names to Content Properties. These Content Properties may include Content Records, metadata, or even translation of the Content Name to Content-ID.
- **Path Management Functional block (PMF)**: The PMF is one of the functional blocks of the CMP that interacts with the underlying network and the functional blocks belonging to the CFP. In particular, the PMF is responsible for gathering Network Reachability Information, which relates mainly to the underlying topology (i.e., inter-domain connections/links, link availability etc.); that is, routing information regarding the delay, loss and bandwidth (in terms of network load) of underlying network paths. This information is provided upwards to the Content Mediation Functional block (CMF).
- **Server and Network Monitoring Functional block (SNMF)**: The responsibility of SNMF is pretty straightforward. It is responsible for collecting data for the status of: i) content servers, namely their availability and load and ii) the underlying network conditions, namely ingress and egress load on peering links, and load on access links

(e.g., for admission control). This information is then fed to the CMF for it to make the appropriate decisions during the Content Consumption operation.

- **Content Mediation Functional block (CMF)**: Content Mediation Function is considered to be the central function as "decision maker" in the CMP. It gets input from CRF, SNMF and PMF blocks, and it also interacts with Content Client during Content Consumption operation. It also interacts with the CAFF block in the CFP for the purpose of configuring delivery paths during the Content Consumption. Its main functionality is to make decisions regarding the selection of the best possible copy of the content, based on information about server and network conditions received from SNMF and the information about the available paths from the PMF, and then to setup and configure the content delivery paths.

The only functional block of the **Content Forwarding Plane** is the **Content-Aware Forwarding Functional block (CAFF)**. CAFF in the COMET system is a sophisticated forwarding function, which allows delivering content through paths selected in CMP (specifically by CMF). The CAFF will have enhanced capabilities in order to provide the required QoS for end-to-end content delivery. These capabilities include traffic classification, point-to-multipoint forwarding, Network Address Translation (NAT) functionality in order to hide the Content Server's IP address to the Content Client, etc.

5. Two implementations of the COMET architecture: decoupled and coupled approach

The functional blocks described in the previous section are the main components of our architecture. These components can be mapped into entities, which interface with other entities to communicate each other. The mapping of functional blocks into entities has led to two different implementations namely decoupled approach and coupled approach. These approaches differ mainly in specific mechanism and algorithms used to perform mentioned above functions as well as the strategy to deploy entities in the network.

5.1 – Decoupled approach

The decoupled approach follows the basic paradigm of the current Internet by allowing the physical signalling routes for content resolution and the physical content delivery paths to be separated but coordinated. By decoupling the location of the content and the content delivery, it is possible to have different architectures for both operations, thus implementing the most appropriate one for each purpose.

This approach relies on the existence of a global directory system that stores Content Records, which are data structures containing content properties such as the QoS requirements for the content, the list of available content sources and the application protocols to be used to retrieve the content from each Content Server. This global directory system resolves from a Content Name to a Content Record, using for that resolution a hierarchical architecture similar to the one used in DNS.

In this approach, there are specific entities that hold the CRF functional block (namely Content Resolution Entity, CRE) and act as that global directory system for contents. These entities are different from the ones holding the CMF functional block (namely Content Mediation Entity, CME). While the entity hosting the CRF is not linked to any specific network domain, the entity hosting the CMF is associated to a specific network domain.

5.2 – Coupled approach

The coupled approach follows a disruptive paradigm with respect to the current Internet, with the physical signalling routes for content resolution and the corresponding content delivery being coupled. The main strategy of this coupled approach is to simplify the process by combining the two operations of the content resolution and content delivery into one single operation, thus avoiding two round-trip requests.

For this purpose, the content resolution follows a hop-by-hop basis across intermediate domains: the content request is forwarded from CMF to CMF, approaching with each hop to the destination domain where the content is located. At the same time, content delivery paths in CFP are maintained with content states installed during the resolution phase in CMP. In the end, the domain-level content delivery path exactly follows the reverse direction of the original resolution path for each content consumption request.

Unlike the decoupled approach, a unified entity (Content Resolution and Mediation Entity, CRME) holds both CRF and CMF, and it is associated to a specific network domain.

6. QoS engineering and content-aware forwarding

One of the main advantages that can be achieved thanks to the mediation is the possibility to deliver contents taking into account their QoS requirements and QoS capabilities offered by the network. The COMET approach for Quality of Service (QoS) engineering assumes that content is delivered using content delivery paths, which support COMET Classes of Service (CoS). The COMET CoSs are defined in the global scope, and then they are mapped into intra-domain CoSs, that are provided by particular domains, e.g. using DiffServ or NGN architectures. The idea of investigated approach is illustrated in Figure 3.

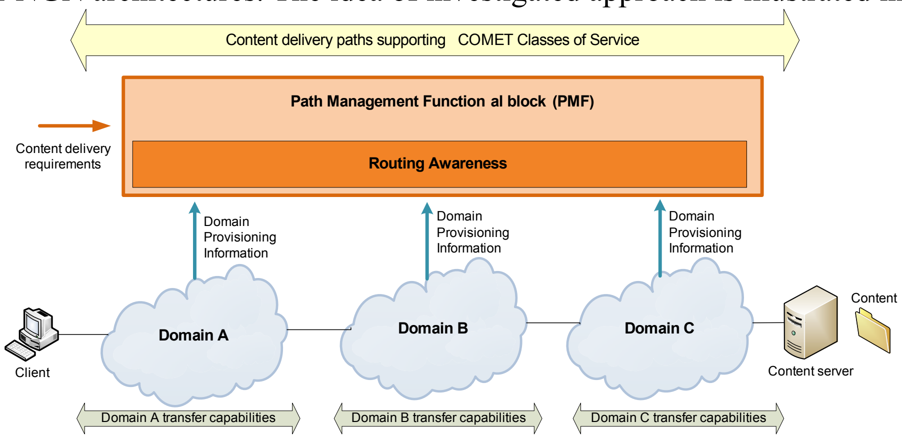

Figure 3: Routing awareness and creation of content delivery paths.

The content delivery paths are organized and maintained by the Routing Awareness Entities (RAEs), which are the entities encompassing the Path Management Functional block (PMF). The RAE entity is the same in both decoupled and coupled approaches. This entity, implemented in each domain, is in charge of gathering network reachability information across domains from other remote RAEs, thus providing routing awareness to the entity encompassing the CMF functionality (i.e. CME in the decoupled approach and CRME in the coupled approach). This routing awareness is provided in terms of the COMET CoSs; this means that the RAE is responsible for matching the content delivery requirements of each COMET CoS with the transfer capabilities offered by particular domains. This mapping should be performed as a part of configuration of COMET system, after the domain provisioning information is available, e.g., domain owner policies, available resources and transfer capabilities.

Other important elements enabling the QoS engineering in COMET are the Content Aware Forwarding Entities (CAFEs), the entities encompassing the CAFF functionality. CAFEs handle the packets transferring the content by using a dedicated COMET header appended to the packets as shown in Fig. 4. This allows taking advantage of three mainfeatures. First, CAFEs may enforce the content delivery paths regardless of existing routing. This feature not only relax constraints of standard routing protocols where single shortest path is used (e.g., BGP), but it also allows to use content delivery path on demand for particular content consumption (using previously established paths). Second, CAFEs may perform adequate packet classification/marking for mapping from COMET CoSs into transfer capabilities used in particular domain. Although this function seems to be simple, it is essential to provide end-to-end QoS across multi-domain networks. At last, CAFEs may support multicast distribution of particular content at intra- and inter-domain levels.

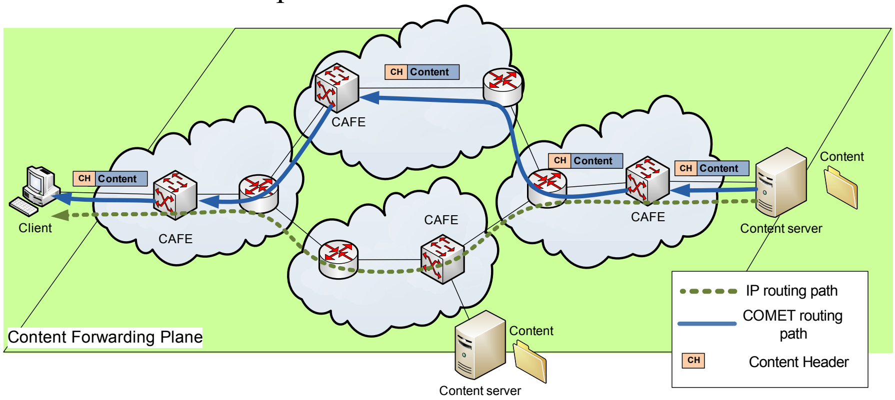

Figure 4: Enforcement of the content delivery path.

Summarising, even though the content delivery in COMET relies on transfer capabilities provided by particular domains, the RAE and CAFE allow applying specific traffic control and engineering algorithms in order to optimise content delivery paths. Thanks to them, the COMET may improve both the quality perceived by content consumers and the efficiency of resource utilisation (network and server resources).

7. Conclusions

This paper has presented ongoing work within the COMET project towards defining a novel content-oriented architecture for simplifying content access and for supporting network-aware content distribution. The proposed architecture lies in the concept of mediation, which lies in the provision by Internet Service Providers of an intermediate plane between the world of contents and the world of data transmission. This plane becomes aware of content location, server and network conditions, and available routing paths, and, based on that info, decides the best source and delivery path for the content.

The COMET architecture is flexible enough to accommodate different implementations, allowing not only the deployment of content-aware facilities over the current Internet, but also its deployment over multi-service environments that favour differentiated content delivery, and over revolutionary Future Internet approaches that put content at centre stage.

Further study and design are currently in progress so that the COMET architecture would be implemented and validated over a small-scale testbed and a real ISP network.

References

[1] EC FIArch Group, Fundamental Limitations of Current Internet and the path to Future Internet, report ver. 0.9, 2010, http://www.future-internet.eu/uploads/media/FIArch_Current_Internet_Limitations_V0_9.pdf

[2] V. Jacobson, et al., “Networking Named Content”, CoNEXT 2009, Italy, December 2009, pp. 1–12.

[3] T. Koponen, et al., “A Data-oriented (and Beyond) Network Architecture”, Proc. ACM SIGCOMM ’07, Kyoto, Japan, August 2007, pp. 181–92.

[4] M. D. Ambrosio, et al., “Second NetInf architecture description”, Deliverable 6.2 of project 4WARD, FP7-ICT-2007-1-216041-4WARD/D-6.2, January 2010

[5] P. Jokela et al., “LIPSIN: Line Speed Publish/Subscribe Inter-networking”, Proc. ACM SIGCOMM ‘09, Barcelona, Spain, August 2009

[6] I. Stoica et al., “Internet Indirection Infrastructure,” Proc. ACM SIGCOMM’02, Pittsburgh, USA, August 2002, pp. 73–86.