

Improving End-to-End QoE via Close Cooperation between Applications and ISPs

Bertrand Mathieu and Selim Ellouze, Orange Labs

Nico Schwan, Bell Labs

David Griffin and Eleni Mykoniati, University College London

Toufik Ahmed, University of Bordeaux

Oriol Ribera Prats, Telefonica I&D

ABSTRACT

In recent years there has been a trend for more user participation in Internet-based services leading to an explosion of user-generated, tailored, and reviewed content and social-networking-based applications. The next generation of applications will continue this trend and be more interactive and distributed, putting the prosumers at the center of a massively multiparticipant communications environment. Furthermore, future networked media environments will be high-quality, multisensory, multi-viewpoint and multistreamed, relying on HD and 3D video. These applications will place unprecedented demands on networks between unpredictable and arbitrarily large meshes of network endpoints. We advocate the development of intelligent cross-layer techniques that, on one hand, will mobilize network and user resources to provide network capacity where it is needed, and, on the other hand, will ensure that the applications adapt themselves and the content they are conveying to available network resources. This article presents an architecture to enable this level of cooperation between the application providers, the users, and the communications networks so that the quality of experience of the users of the application is improved and network traffic optimized.

INTRODUCTION

In recent years there has been a trend for more user participation in Internet-based applications. There has been an explosion of user-generated, tailored, and reviewed content, while social networking is beginning to replace traditional communications technologies such as email and websites. Typical examples of popular applications that only exist for and because of significant user participation are Facebook, YouTube, Flickr, Digg, eBay, Second Life, and Wikipedia. However, even though content is being created,

modified, and consumed by a large number of participants, almost all of these applications still rely on servers adequately dimensioned and carefully positioned by service providers in large data centers at strategic locations across the Internet to ensure an adequate quality of experience (QoE) for their users. These deployments require significant investment to maintain; they cannot expand beyond the selected locations and have limited flexibility to adapt to demand variations over time.

The next generation of applications will continue the trend of user-centricity where users are not just seen as consumers of a product or service, but are active participants in providing it. They will be more interactive and distributed, putting prosumers at the center of a massively multiparticipant communications environment where they can interact in real time with other users and provider resources, to provide and access a seamless mixture of live, archived, and background material. Furthermore, future networked media environments will be high-quality, multisensory, multi-viewpoint, and multistreamed, relying on HD and 3D video. These applications will place unprecedented demands on networks for high-capacity, low-latency, and low-loss communication paths between unpredictable and arbitrarily large meshes of network endpoints, distributed around the entire globe, putting additional pressure for upload capacity in access networks.

If the entire burden of supporting high volumes of HD/3D multimedia streams is pushed to the Internet service providers (ISPs) with highly concurrent unicast flows, this would require operators to upgrade the capacity of their infrastructure by several orders of magnitude to ensure end-to-end quality of service between arbitrary endpoints.

Rather than simply throwing bandwidth at the problem, we advocate the development of intelligent cross-layer techniques that, on one hand, will mobilize network and user resources

to provide network capacity where it is needed, and, on the other hand, will ensure that the applications adapt themselves and the content they are conveying to available network resources, considering core network capacity as well as the heterogeneity of access network and end device capabilities. Meeting these challenges requires a previously unseen amount of cooperation between application providers, users, and the communications networks that will transport the application data. While previous work, mainly in the telco domain, such as IP Multimedia Subsystem (IMS) [1] or Parlay/X [2] and more recently ALTO, has made progress in this direction, their use is confined to walled-garden environments or limited in terms of services the applications can request. Furthermore, they do not address the specific exchange of information in both directions so as to enable overlay applications and networks to be optimized. This article presents an architecture that enables this level of cooperation between the actors and elaborates on the related interactions.

The next section presents some example applications to illustrate the benefits of collaboration between the applications and the network. We then present an overall picture of our approach. We then elaborate on the cross-layer interactions. The functional architecture which captures the high-level building blocks of our system is detailed in the following section. We then describe research and standardization initiatives related to our approach. Finally, a summary of the article and a discussion of future work is presented.

EXAMPLE APPLICATIONS

The environment assumed in this article is one where applications are decoupled from underlying communications networks and organized as overlay networks. The application logic will be distributed over numerous overlay nodes provided by end users, application providers, and even the ISPs themselves. The complexity of the nodes and the application logic being executed at each one depends on the sophistication of the application. The architecture and cooperative approach presented in this article has been designed to be of benefit to a wide spectrum of applications. In the simplest cases the overlay algorithms may be streaming live video through swarms, while more complex applications may require on-the-fly multistream 3D video processing or real-time discovery and interactive tracking of user-generated content. Novel applications delivering rich user experiences will present new business opportunities for the consumer electronics and user applications industries beyond incremental changes to today's server-centric and web-based content retrieval services. The interaction with underlying networks will also present new opportunities for ISPs to cooperate in the delivery of media applications.

We present the following examples to illustrate the type of advanced services that can benefit from increased cooperation with underlying ISPs. Their features include low-latency and high-capacity content dissemination; manipulation of distributed and streaming content such as



Figure 1. *Bicycle use case.*

the interpolation of multiple audio-visual streams from different viewpoints; exchange of live, multisensory, and contextual information between participants; and the discovery and navigation of distributed content, information and users. The common denominator is the collaborative production, processing, and consumption of a mixture of live, archived, and background high-quality media from multiple sources with demands that can outstrip the capabilities of the underlying networks unless the applications adapt themselves, and content is tailored to meet network capacity and performance constraints, and/or specific network services such as multicast or in-network caching facilities are provisioned to support their efficient distribution.

One example of such an application is the multi-viewpoint coverage of sporting events, such as a bicycle race like the Tour de France (Fig. 1). In this scenario numerous fixed and mobile sources, such as professional media organizations, trackside spectators, as well as the cyclists themselves can generate live audio-visual streams. As each stream shows a specific view of a potentially different subject, a potentially large number of streams with overlapping content are generated. Consumers of this content, who may be distributed around the globe using various fixed and mobile end devices, can tailor their viewing experience by selecting from many streams according to their preference, or navigate between streams in real time to zoom or pan around, or follow particular cyclists. The popularity of individual streams is difficult to predict and may change rapidly; therefore, the creation and adaptation of efficient distribution trees or meshes to transmit the content to interested sets of users at the required quality levels present problems that cannot easily be solved by individual ISPs or the application overlay in isolation. ISPs are unaware of the popularity of dynamically changing content sources, the locations of the consumers of that content, or the heterogeneous end terminal capabilities in remote networks. From the ISPs' perspective the applications are simply generating large quanti-




Figure 2. 3D virtual conference.

ties of traffic between unpredictable locations. On the other hand, application logic can track and match content sources and consumers, but efficient distribution overlays can only be built with knowledge of underlying network capabilities so that caching and adaptation functions, for example, can be placed where they are required and are most effective, or advanced network services such as regional multicast distribution can be invoked where most needed to relieve congestion and improve the QoE for users.

A second example is a virtual meeting such as a 3D virtual conference, where a large number of participants, represented by virtual avatars, can meet and communicate via voice and avatar gestures, as well as share additional multimedia data such as live video, 3D models, text, and presentation slides (Fig. 2). Users will be both consumers and generators of content — custom avatars, user-generated environments, and sources of interactive video streams. Participants can move around the virtual meeting space, attend presentations, establish special interest discussion groups, socialize in coffee breaks, and so on. Tracking the participants and managing their interests and participation in various activities is the responsibility of the application overlay only; however, the distribution of content to various groups of users is something that benefits from cooperation with the underlying networks. Users in the same virtual meeting room with a similar point of view need access to similar data, such as static background material as well as dynamically changing objects that need to be synchronized between many consumers. The efficiency of the system can therefore be greatly improved by organizing the overlay with regard to the position of content within the virtual space and making use of ISP-provided network services such as localized in-network content caching or multicast for distributing state changes of common objects to reduce latency in live updates, network load, and therefore costs.

It can be seen that both of the above examples could generate huge amounts of data to be transmitted to sets of receivers that range from small groups to many hundreds or even millions of consumers, with some specific constraints of

quality levels such as maximum latency. Our proposed solution of increased cooperation between ISPs and the overlays will assist in exchanging rich information between the application and the network so that the overlays can be organized efficiently and can adapt to network constraints by avoiding high cost or highly congested areas, or adapting the quality of the streams to match available network capacity, or enabling specific network services such as multicast distribution or caching to be invoked in areas of densely populated receivers.

OVERALL APPROACH

Because applications will be more participatory and interactive, today's model of centralized or replicated servers in large data centers is likely to be replaced by a highly distributed model where processes run in user equipment and interwork with one another in an overlay layer and can be enhanced via the invocation of network services. In our approach, we advocate close and strong cooperation between ISPs and overlay applications for optimized delivery of content to end users.

This cooperation is achieved via the comprehensive, media-aware and open Collaboration Interface between Network and Applications (CINA), which bridges the gap between ISPs and application overlays and aims at:

- Increasing the degree of cooperation between the network layer and the applications through mutual exchange of information
- Optimizing application overlay networks with respect to the capabilities of the underlying networks and the participant end users
- Providing the means by which service providers can request the activation of specialized network services or resources to achieve efficient distribution of highly demanding content streams
- Enabling dynamic adaptation of the content to meet the abilities of the underlying networks and user requirements

The overall picture of the system is illustrated in Fig. 3, which highlights the CINA interface. The overlay application network consists of nodes provided by one or more service providers (SPs), the users themselves, and, optionally, ISP nodes. There will be separate overlay networks for each application. The different applications may be more or less dependent on SP nodes, with peer-to-peer (P2P) applications running entirely on user nodes being the extreme case.

Given that the applications are global in coverage and require end-to-end traffic optimization involving multiple hops in different networks, it is necessary to collect information from many underlying networks via CINA. Since data from one network may conflict with that provided by another, or the quantity and quality of the information may differ from ISP to ISP, the harmonization of the information gleaned from the ISPs is required in the overlay. The overlay could also aggregate the information collected from different ISPs, with additional data collected by measurements of the overlay itself for the global

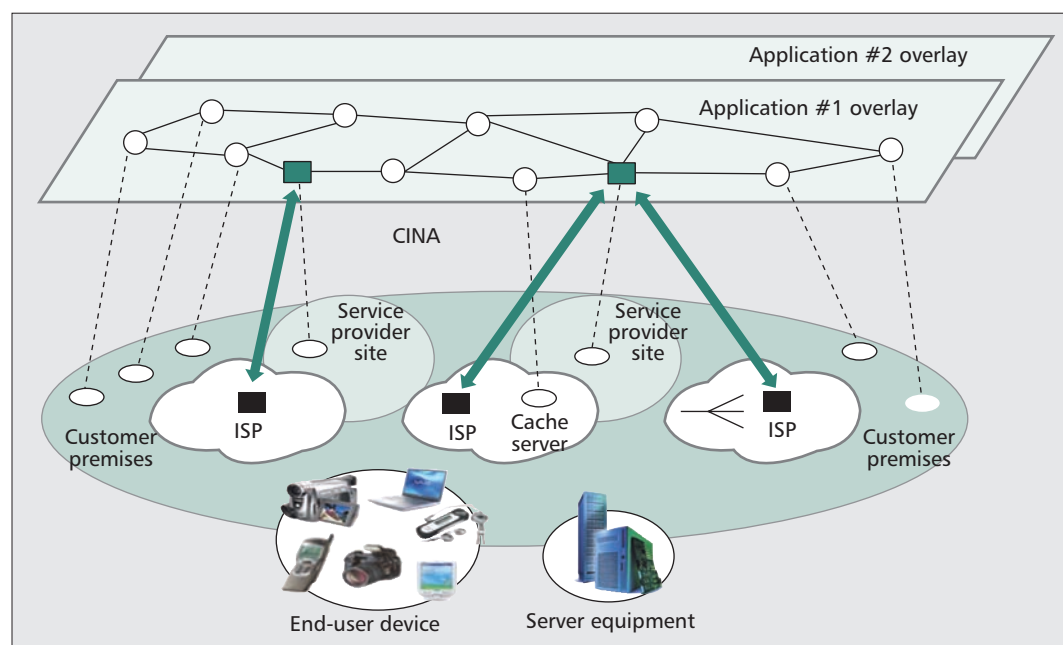


Figure 3. Overview of the system and its relationship with users, ISPs, and overlay applications.

Content adaptation has two dimensions: personalizing and tailoring the content for the subjective viewpoint of the user(s); and encoding content in a flexible way to match the capabilities of the network.

optimization of the application. The use of this information will benefit algorithms that are needed for optimizing the distribution of the content. These algorithms determine which application resources need to be involved, and how to best interconnect the participants and distribute the load and the content to achieve the best QoE given the available resources.

Also, content adaptation services will profit from the cooperation. Until now digital coding and encoding systems have been designed following the client/server paradigm, but now applications will have to deal with the fact that content may come from several sources and terminal devices with different capabilities, residing in networks that offer different service levels. Applications need to adapt and select *quality layers* with a brand new set of constraints and circumstances. Pre-adaptation of content using offline pregeneration strategies on the subset of the most popular content is inefficient and cannot meet the continuously increasing number of users' changing interests and heterogeneous terminal capabilities.

Content adaptation has two dimensions: personalizing and tailoring the content for the subjective viewpoint of the user(s); and encoding content in a flexible way to match the capabilities of the network. The latter could also allow the upload capacity of participants, especially those that act as content sources, to be boosted by making use of parallel connections across several available access networks.

Adaptation is achieved through mechanisms that dynamically adjust the content via the use of scalable video coding, video layer coding, or adaptive streaming being performed at the source or end-user side, or by intermediate nodes as appropriate to match the requirements of the set of users receiving the content while adhering to network capabilities and restrictions such as congested access links.

CROSS-LAYER INTERACTIONS

The Internet Engineering Task Force (IETF) ALTO working group defined an architecture and a protocol for communication between ALTO clients and ALTO servers. Possible ALTO information to be exchanged are an abstracted network map and an associated cost map or list of peers ranked according to ISP preferences. In our study we aim to go further, allowing network operators to provide any kind of network information it wishes or that it accepts to provide to overlay applications, after agreement between the two actors. For instance, the ISP could inform the application about the capabilities of its network: access networks (type, link capabilities, coverage, etc.), current network services status (availability of multicast groups, caches, etc.), as well as other possible metrics (load of routers, bandwidth, delay, etc.). Since some information is critical for network operators and they do not want to reveal it (e.g., internal detailed topology or BGP policies), the CINA interface is designed having in mind the agreements between the applications and the ISP, and is adaptable to allow any kind of agreed information to be exchanged.

Furthermore, via the CINA interface, the network operators can also get information from the overlay so that they can optimize the traffic in their networks, mobilize resources, and adapt to the overlay applications, eventually transparently; this is not covered by ALTO. Typically, the application could inform the ISP about its traffic demand: information related to users (e.g., user location and estimated traffic matrix) or content (quantity of sources, their bit rates, adaptive coding, etc.).

This information exchanged between the application and the ISP goes further than information reflecting the preferences and policies of the involved business entities as currently defined in the ALTO working group.

No one can force both the overlay applications and the network operators to collaborate, but with the possible win-win approach of both (optimized QoE for application, reduction of traffic load as well as monetization of network services), both actors will benefit from cooperation.

CINA also goes further than ALTO in the way that it enables future networked media applications to make use of advanced network services in a dynamic and flexible way to achieve cost-efficient delivery of high QoE for their users. It is known that an ISP could offer information to applications such as the location of users or some user profile information, but in our approach we go further via the offering of advanced network services. For example, such possible network services can be:

- Multicasting: Possibly with hybrid application layer and native IP multicast since the applications will usually be spread over several ISPs, or the use of high fan-out nodes located in the network.
- Caching: Via the use of specialized nodes, provided either by ISPs or third-party entities, to optimize delivery and save bandwidth in the network.
- Bandwidth on demand: To enable delivery toward end users over multiple access networks simultaneously, and provide bandwidth on demand over aggregated access networks.
- Dynamic quality of service (QoS) mapping: Invocation and mapping of application QoS requirements, network capabilities, end-user devices, and access networks.
- Ad/text insertion: In order to offer added-value services that might be monetized by network operators.
- Content adaptation: The presence of heterogeneous end-user devices and network infrastructures will require multiple versions of the same resource that can be efficiently generated using content adaptation.

Our work mainly relates to the open interface that lets applications request the dynamic activation of such services and does not focus on redefining them. It is the ISP's responsibility to list the network services it can offer and make agreement with an application to decide if it is for free or provided with appropriate billing. The final decision of activating the service remains under the control and management of the network operator. No one can force the overlay applications and network operators to collaborate, but with the possible win-win approach of both (optimized QoE for application, reduction of traffic load, as well as monetization of network services), both actors will benefit from cooperation.

To illustrate the potential benefits of cooperation between the application overlay and the underlying ISPs, consider the following example. The bicycle race use case introduced earlier depends on the streaming of live audio-visual streams from multiple sources to many consumers. The efficiency of the delivery scheme depends on the popularity of the streams and the distribution of the consumers across the Internet. While low-popularity feeds can be delivered efficiently by unicast streaming techniques, the source node needs significant upload network capacity for this to scale to a large number of consumers. Overlay P2P swarming techniques reduce the need for high-bandwidth links at the source node, but they do not reduce the load on the underlying network, and can even

increase traffic.

However, if the application overlay is able to invoke network services such as multicast or caching services through CINA, the load on the underlying network can be reduced significantly. In addition, there is less overhead in terms of swarming control operations required at the overlay.

As an example, the Abilene¹ core network topology of 11 nodes and 14 links was assumed for the ISP. Assuming equal link metrics, a stream bandwidth of S , a single source located at one of the Abilene nodes and stream consumers at each of the other nodes, for unicast streaming the maximum link load ranges from $4S$ to $8S$ and the total traffic (sum of link load) ranges from $21S$ to $30S$ depending on the node to which the source is attached. In all cases the source requires an upload capacity of $10S$. If the overlay uses a P2P swarming distribution scheme such as BitTorrent, the source node requires an upload capacity of no more than S resulting in a total load on the network (sum of link load) of approximately $26.8S$, each peer retrieves an equal fraction of the stream from each of the others. If, however, the overlay is able to cooperate with the ISP through CINA and invoke network-layer multicast to distribute the stream to all consumers, the total load on the network is reduced to $11S$, a reduction of 58.9 percent of the total traffic. If we assume there are even more consumers at each of Abilene's core routers, the total traffic on the ISP's network increases linearly with the average number of peers per core router to which they are attached, while the load in the multicast case is not increased.

OVERALL FUNCTIONAL ARCHITECTURE

In this section we describe the functional architecture defined to model the new interfaces and interactions required to enable the users, third parties, application providers, and ISPs to contribute and allocate resources in a dynamic and coordinated way, allowing the cross-layer optimization between the independent application and the network processes. The proposed architecture is built in order to define the boundaries of responsibility between the involved actors and the required interactions that are required across these boundaries.

The ISPs (lower layer blocks in Fig. 4) operate the network and provide Internet connectivity services at particular locations defining their network domains. Within their domains, the ISPs can provide enhanced network services, including multicast and prioritized traffic treatment, and they can operate application layer services such as content caching on behalf of the application, or for improving their particular traffic optimization objectives. On top of the network infrastructure, application providers (middle layer blocks in Fig. 4) operate their service infrastructure, which may span several locations and the network domains of many ISPs. Application providers may be assisted by third-party providers who would operate additional infra-

¹ <http://abilene.internet2.edu/>

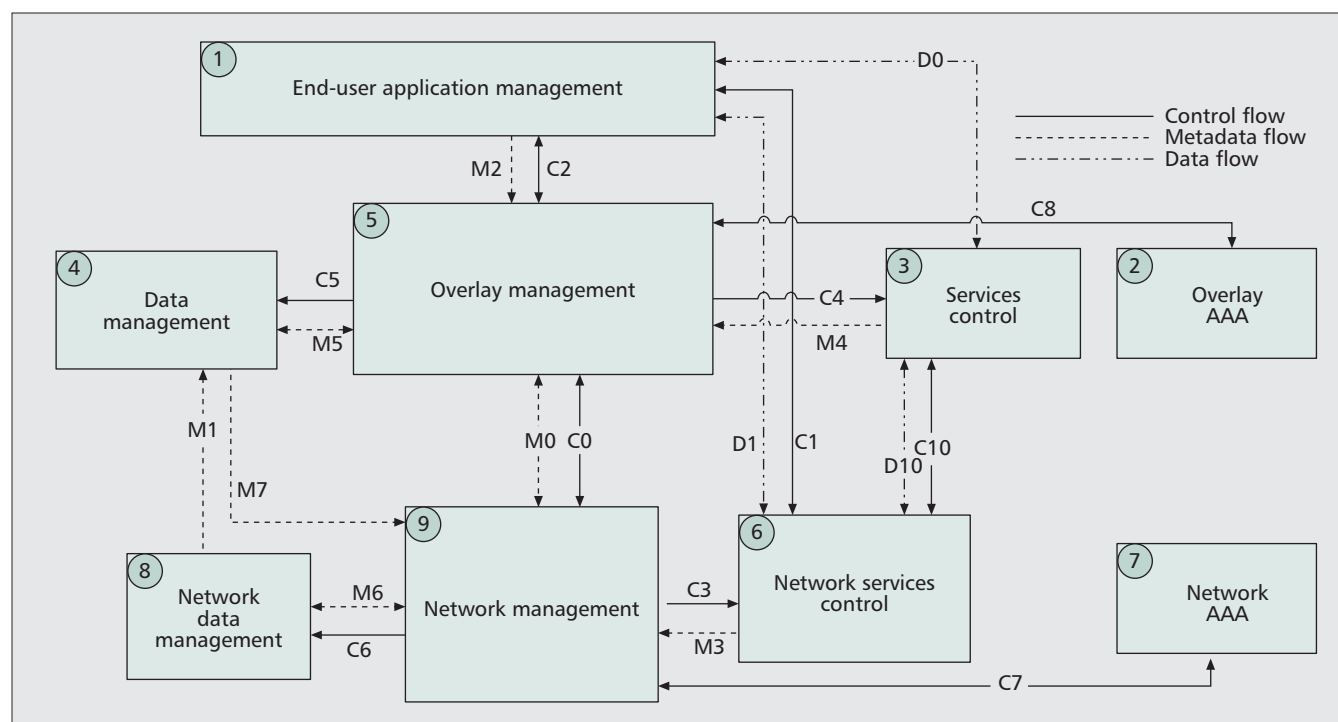


Figure 4. Functional architecture.

structure at additional locations or provide specialized added-value services, offered to application providers under flexible agreements. Finally, the end users (upper layer block in Fig. 4) access the applications and offer their resources to enhance the application infrastructure in a dynamic way, similar to third-party providers but on a smaller scale and possibly under the full control of the application providers.

The rest of this section describes the blocks and interfaces that model the functionality and the interactions between these actors. The interactions between the network and the application providers in particular (interfaces between the middle and bottom layer blocks) define CINA.

The first functional block, *end-user application management*, models the functionality at the end user. It includes functions for:

- Content generation, consumption, search, and so on
- Data flow handling (e.g., transmission, reception, synchronization)
- Interest and profile management
- Providing QoE feedback to the application

At the application level, the *overlay management* block includes the application optimization logic, and communicates with *services control* to dynamically invoke services and resources where they are required, and with *data management*, which is responsible for maintaining up-to-date information about the network and the application.

The data management block collects, consolidates, maintains, and provides network and application level information. It allows the application to access the information provided by ISPs regarding network performance and capabilities, and the network and application services they may provide, and the ISPs to access application information regarding the traffic demand

and quality requirements at particular locations.

The services control block manages the basic functions and value-added services that build the application overlay following the instructions from the overlay management block. Such basic services include data uploading to multiple receivers, content adaptation, caching, content personalization like picture-in-picture, ad placement, and so on. These services could be provided using dedicated application provider equipment, by third-party infrastructure or service providers, or by the end users themselves, contributing their own resources to the application.

The overlay management block implements the overlay optimization algorithms, taking into account information regarding the particular content characteristics and adaptation options, the end-user access means and QoE requirements, the available network resources and offered network services (e.g., multicast) and the available application services (caching, content adaptation, etc.). This block decides which QoE is sustainable for the users and how to allocate resources to their service requests. It dynamically invokes resources and services, requesting multicast transmission within a particular ISP, performing content adaptation for a set of end users, activating a set of cache servers, and so on. Finally, it interacts with the underlying ISPs to coordinate the use of the network resources in a mutually beneficial way (e.g., by reducing the data rate for a specified region under congestion conditions) or to indicate where caching resources could be allocated to reduce the load on the network.

At the network level, similar blocks model the network optimization, service control, and data management functionality. In particular, the *network data management* block maintains

Since our approach can be seen as an extension or evolution to ALTO, we investigate if solutions defined in this group can be applied to ours. Typically the use of a RESTful HTTP based interface that uses JSON encoding is under consideration.

information about the network topology, offered network services, ISP policies, and preferences to be communicated to the application, the status of the network resources, the current network performance, and so on. This block interacts with the application to provide information regarding the particular network domain to the data management block.

The *network services control* block operates the available network services, including multicast, caches, ad insertion, and so on. These services are invoked by the *network management* block to optimize the data distribution as a response to explicit requests received by overlay management, or independently to optimize network optimization objectives.

The network management block is responsible for the management of the network resources within the ISP network domain. It receives information about the network from network data management, information about the traffic demand and quality requirements of different applications from data management, explicit requests for allocation of resources or invocation of services from overlay management, and finally, feedback regarding the performance of network services from network services control. Based on this information, the network management block determines the appropriate allocation of resources and invocation of network services, and specifies the preferences that are communicated to the overlay application.

Finally, the interactions between different actors create security considerations that necessitate the introduction of authentication, authorization, and accounting (AAA) functionality. The *overlay AAA* block handles the authentication of users joining the overlay, ISPs cooperating with the overlay, and potentially third-party service providers. The *network AAA* block enables the authentication of overlays cooperating with the ISP. Both blocks offer standard AAA functions, such as accounting facilities, access authorization, and profile management with security mechanisms, and can be controlled by the overlay management and network management blocks, respectively.

RELATED WORK

Overlay applications are currently agnostic of the underlying network infrastructure and thus perform end-to-end measurements to gain some knowledge [3], but as this is not in cooperation with the ISP, it can lead to undermining the routing policies of ISPs [4, 5]. To avoid this, several initiatives have promoted cooperation between overlay applications and underlying networks.

The P4P initiative [6] and later the IETF ALTO working group have investigated how overlay networks and ISPs can cooperate to optimize traffic being generated by P2P applications and transported over the ISP's infrastructure. In their approach the ISP is able to indicate preferences on which peers should exchange data. Our solution is related to ALTO; however, it proposes a much richer interface that will allow true cross-layer cooperation, in terms of both information exchange and the possibility for overlays to dynamically request ISPs'

network services. Since our approach can be seen as an extension or evolution of ALTO, we investigate if solutions defined in this group can be applied to ours. Typically, the use of a RESTful HTTP-based interface that uses JSON encoding is under consideration.

Some research work, such as that investigated in [7, 8], aims at building a P2P framework or a delivery platform for live TV. In contrast to our approach, they do not consider multi-viewpoint applications and highly interactive applications between participants in the network; also, [7] does not deal with content adaptation.

Research work studied P2P support in the context of massively multiparticipant virtual environments such as [9] for virtual environments and [10] for 3D streaming. Those applications are close to the ones we focus on, but in these solutions the volume of the content is rather small compared to HD video from potentially a large number of sources we address in our approach. Also dynamic activation of network services via the cooperation between overlay and ISPs for improving QoE is not taken into consideration.

The definition of interfaces between the underlying network and the application or control level has been investigated for some years now, and some solutions are deployed such as the well-known IMS [1]. IMS designs an interface between the control entity and the network entities. However, this interface is still closed and only usable by the ISPs in a walled-garden fashion. Also, the possible network services that might be dynamically activated are limited. In another initiative, the Parlay/OSA framework was created some years ago for the telecommunications networks and telecom services (call control, call redirection, etc.), and has more recently been adapted for the Internet (web services) with the Parlay/X specification [2]. The objectives of this group were similar to ours; however, the defined interfaces depend on the network service to be activated, do not permit the exchange of dynamic information between the network and the application (in both directions), and the supported services were limited to more traditional telecom applications rather than the media applications we aim to support. Furthermore, those two initiatives do not really address issues for the massive multiparticipant distributed applications we envision.

CONCLUSIONS AND FUTURE WORK

In this article we have presented a new architecture, fostering cooperation between overlay applications and ISPs for optimized delivery of services to end users. Overlay algorithms are optimized thanks to information provided by ISPs, and the delivery QoE is further improved via the activation of network services provided by ISPs.

While the approach presented in this article has been developed to support interactive, multi-party, high-capacity media applications such as those presented earlier, the architecture, the interface, and the principles of cross-layer cooperation can also benefit existing applications. Content distribution networks (CDNs) providing

video on demand, for example, could have greater awareness of network capabilities and also make use of other network services provided by CINA, such as capacity reservation for background distribution of content to CDN nodes.

Ongoing work is focused on the internal functions of each block, including the development of overlay optimization algorithms, source selection algorithms according to context information, dynamic activation logic for network services such as multicast (and multi-ISP multicast), and caching/adaptation functions. Finally, evaluation through both simulation and testbed deployments is underway.

ACKNOWLEDGMENTS

This work was supported by the ENVISION project (<http://www.envision-project.org>), a research project partially funded by the European Union's 7th Framework Program (contract no. 248565). The authors wish to thank all project participants for their valuable comments and contributions to the work described in this article.

REFERENCES

- [1] 3GPP TS 23.228, "IP Multimedia Subsystem (IMS); Stage 2 (Release 7)"; <http://www.3gpp.org>.
- [2] OSA/Parlay X, "3GPP TS 29.199 Release 7 Specifications," Sept. 2007.
- [3] V. Gurbani et al., "A Survey of Research on the Application-Layer Traffic Optimization Problem and the Need for Layer Cooperation," *IEEE Commun. Mag.*, vol. 47, no. 8, 2009, pp. 107–12.
- [4] R. Keralapura et al., "Can ISPs Take the Heat from Overlay Networks?," *ACM HotNets '04*, Nov. 15–16, 2004, San Diego, CA.
- [5] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?," *Internet Measurement Conf.*, Oct. 19–21, 2005, Berkeley, CA.
- [6] H. Xie et al., "P4P: Explicit Communications for Cooperative Control Between P2P and Network Providers," http://www.dcia.info/documents/P4P_Overview.pdf.
- [7] R. Fortuna et al., "QoE in Pull Based P2P-TV Systems: Overlay Topology Design Tradeoffs," *Proc. 10th Int'l. Conf. Peer-to-Peer Comp.*, Delft, The Netherlands, Aug. 2010.
- [8] R. Jimenez, L. E. Eriksson, and B. Knutsson, "P2P-Next: Technical and Legal Challenges"; tslab.ssvl.kth.se
- [9] D. Frey et al., "Solipsis: A Decentralized Architecture for Virtual Environments," *1st Int'l. Wksp. Massively Multiuser Virtual Environments*, 2008.
- [10] S.-Y. Hu et al., "FloD: A Framework for Peer-to-Peer 3D Streaming," *INFOCOM 2008*.

BIOGRAPHIES

BERTRAND MATHIEU [SM] (bertrand2.mathieu@orange-ftgroup.com) is a senior researcher at France Telecom, Orange Labs since 1994. He received a Diploma of Engineering in Toulon, an M.Sc. degree from the University of Marseille, and a Ph.D. degree from the University Pierre et Marie Curie, Paris. His research activities are related to dynamic overlay networks, P2P networks, and information-centric networking. He has contributed to several national and European projects, and published more than 30 papers in international conferences, journals, and books. He is a member of several conferences' Technical Program Committees and an IEEE Senior Member.

SELIM ELLOUZE received his M. Eng degree in telecommunication from Ecole Nationale Supérieure d'Electronique Informatique et Radiocommunications de Bordeaux, Université Bordeaux 1. He is currently working toward a Ph.D. degree as a research engineer in Orange Labs Lannion, France. His research interests include the evolution of the control protocols for IP networks.

NICO SCHWAN is a research engineer with Bell Labs Service

Infrastructure Research Team in Stuttgart, Germany. He is also an academic at the University of Cooperative Education Stuttgart. His current research interests are in the areas of content delivery networks, multimedia Internet overlay applications, and content-centric networking. He has a B.Sc. degree in applied computer science from the University of Cooperative Education Stuttgart.

DAVID GRIFFIN is a principal research associate in the Department of Electronic and Electrical Engineering, University College London (UCL). He has a B.Sc. from Loughborough University and a Ph.D. from UCL, both in electrical engineering. His research interests are in the planning, management, and dynamic control for providing QoS in multiservice networks, P2P networking, and novel routing paradigms for the future Internet.

ELENI MYKONIATI (e.mykoniati@ee.ucl.ac.uk) received a B.Sc. in computer science from Piraeus University, Greece, in 1996 and a Ph.D. degree from the National Technical University of Athens (NTUA), Greece in 2003. She has worked as a research associate for Telscom S.A. Switzerland, the NTUA DB Lab and Telecom Lab, and Algonet S.A. Greece. Since 2007 she is a senior research associate in the Department of Electronic and Electrical Engineering at University College London. Her research interests include business-driven traffic engineering in IP networks, QoS, and peer-to-peer networking.

TOUFIK AHMED is a professor at ENSEIRB-MATMECA School of Engineers in Institut Polytechnique de Bordeaux (IPB) and performing research activities in CNRS-LaBRI Lab-UMR 5800 at University Bordeaux 1. His main research activities concern QoS management and provisioning for multimedia wired and wireless networks, media streaming over P2P networks, cross-layer optimization, and end-to-end QoS signaling protocols. He has also worked on a number of national and international projects. He is serving as TPC member for international conferences including IEEE ICC, IEEE GLOBECOM, and IEEE WCNC. He is currently team leader of the COMET group at LaBRI Lab.

ORIOL RIBERA PRATS is a business consultant specialized in quantitative market analysis at the Telefonica R&D Center in Barcelona, Spain. He received his Telecommunications Engineer degree from the Universitat Politècnica de Catalunya and has had a consolidated career of ten years as product manager in the mobile Internet industry beginning at Nokia, and continuing with co-founding Genaker (a technology start-up) and now at Telefonica.

Ongoing work is focused on the internal functions of each block, including the development of overlay optimization algorithms, source selection algorithms according to context information, dynamic activation logic for network services, such as multicast (and multi-ISPs multicast) and caching/adaptation functions.