



Improving end-to-end QoE via close cooperation between applications and ISPs

Bertrand Mathieu, Selim Ellouze, Nico Schwan, David Griffin, Eleni Mykoniati, Toufik Ahmed, Oriol Ribera Prats

► To cite this version:

Bertrand Mathieu, Selim Ellouze, Nico Schwan, David Griffin, Eleni Mykoniati, et al.. Improving end-to-end QoE via close cooperation between applications and ISPs. *IEEE Communications Magazine*, 2011, 49 (3), pp.136 - 143. 10.1109/MCOM.2011.5723811 . hal-00958117

HAL Id: hal-00958117

<https://hal.science/hal-00958117>

Submitted on 12 Mar 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving end-to-end QoE via close cooperation between applications and ISPs

Bertrand Mathieu, Selim Ellouze, Nico Schwan, David Griffin, Eleni Mykoniati, Toufik Ahmed, Oriol Ribera Prats

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 users at the centre of a massively multi-participant communications environment. Furthermore, future networked media environments will be high-quality, multi-sensory, multi-viewpoint and multi-streamed, 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. It would require operators to upgrade the capacity of their infrastructure by several orders of magnitude to ensure end-to-end quality of service. Instead, we advocate the development of intelligent cross-layer techniques that, on the one hand, will mobilise 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. This paper presents an architecture that enables this level of cooperation between the application providers, the users and the communications networks, so that the QoE of the application should be improved and the network traffic optimised.

Index Terms— cooperation, ISP, overlay, media applications, P2P, network services.

I. 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 centres at strategic locations across the Internet to ensure an adequate Quality of Experience 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 the users at the centre of a massively multi-participant 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, multi-sensory, multi-viewpoint and multi-streamed, 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 multi-media streams is pushed to the 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 end points.

Rather than simply throwing bandwidth at the problem, we advocate the development of intelligent cross-layer techniques that, on the one hand, will mobilise 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. This paper presents an architecture that enables this level of cooperation between the actors and elaborates on the related interactions.

Section II presents some example applications to illustrate the benefits of collaboration between the applications and the

Manuscript received September the 01st, 2010

Bertrand Mathieu, Selim Ellouze are with Orange Labs, Lannion, France.

Nico Schwan is with Bell Labs, Stuttgart, Germany.

David Griffin, Eleni Mykoniati are with University College London, UK

Toufik Ahmed is with University of Bordeaux, France.

Oriol Ribera Prats is with Telefonica I&D, Madrid, Spain.

network. Section III presents the big picture of our approach. Section IV elaborates on the cross-layer interactions. The functional architecture which captures the high-level building blocks of our system is detailed in Section V. Section VI of this paper describes research and standardisation initiatives related to our approach. Finally, a summary of the paper and a discussion of future work is presented in Section VII.

II. EXAMPLE 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 the interpolation of multiple audio-visual streams from different viewpoints; exchange of live, multi-sensory and contextual information between participants; and the discovery and navigation of distributed content, information and users. The common denominator being the collaborative production, processing and consumption of a mixture of live, archived and background high quality media from multiple sources with demands than 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 as a bicycle race like the Tour de France (figure 1). In this scenario numerous fixed and mobile sources, such as professional media organisations, 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 to 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 presents problems that can not 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 quantities 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 Quality of Experience (QoE) for the users.



Figure 1: Bicycle use-case

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, avatar gestures as well as share additional multimedia data such as live video, 3D models, text and presentation slides (figure 2). Participants can move around the virtual meeting space, attend presentations, establish special interest discussion groups, socialise in coffee breaks, etc. 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 by the 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 synchronised between many consumers. The efficiency of the system can therefore be greatly improved by organising the overlay with regards to the position of content within the virtual space and making use of ISP-provided network services such as localised in-network content caching or multicast for distributing state changes of common objects to reduce latency in live updates, reduce network load and therefore costs.



Figure 2: 3D virtual conference

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 organised 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 by enabling specific network services such as multicast distribution or caching to be invoked in areas of densely populated receivers.

III. OVERALL APPROACH

Because applications will be more participatory and interactive, today's model of centralised or replicated servers in large data centres 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 a close and strong cooperation between ISPs and overlay applications for an optimised delivery of content to end-users.

This cooperation is achieved via the comprehensive, media-aware and open *ENVISION* interface 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;
- optimising 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 specialised 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 network system is illustrated in figure 3 which in particular highlights the *ENVISION* interface in blue. Further we can see the overlay application network on top of the Internet, consisting of nodes provided by one or more Service Provider (SP), the users themselves and, optionally, nodes of the ISPs. There will be separate overlay instances for each application. The different applications may be more or less dependent on SP nodes. At one extreme, there might be the SP acting simply as the P2P software that is run by the users.

From now on this section focuses on the overlay level while the network-level will be described in section IV via the cross-layer interactions.

Given that the applications are global in coverage, and require end-to-end traffic optimisation involving multiple hops in different networks, it is necessary to collect information from many underlying networks, via the *ENVISION* interface. 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 harmonisation 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 optimisation of the application. The use of this information will benefit algorithms that are needed for optimising the distribution of the content. These algorithms determine which application resources need to be involved, how to best interconnect the participants and distribute the load and the content to achieve the best QoE given the available resources.

Also content adaption 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 pre-generation strategies on the subset of the most popular content is inefficient and cannot meet the continuously increasing number of users' changing interest and heterogeneous terminals capabilities. For the success of future networked media the adaptation of content to the capabilities of the core and access networks, user devices and user preferences is one of the targeted objectives. The cooperation between the ISPs and the content adaptation services provided within *ENVISION* enables dynamic, distributed and on-the-fly content adaptation.

Content adaptation, therefore, has two dimensions: personalising 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 (application-layer overlay as well as the ISP's layer-3 network). The latter could also allow the upload capacity of participants, especially those that act as content sources, to be

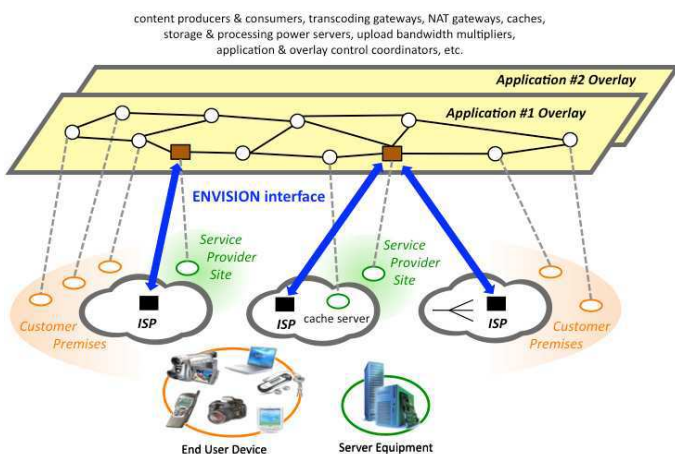


Figure 3: Overview of the system and its relationship with users, ISPs and overlays applications

boosted by making use of parallel connections across several available access networks.

A set of comprehensive application metadata is required for characterising: (1) end-users, including personal information such as name, location, content access rights, user preferences in terms of content type, quality, etc.; (2) device capabilities such as class, network interfaces, codec facility, power/CPU requirements, display capability, adaptation, caching facilities, etc.; (3) content, such as the type, characteristic, angle of view, encoding, cost, etc.

Adaptation is achieved through mechanisms which dynamically adjust the content via the use of scalable video coding, video layer coding, or adaptive streaming, being performed at source, 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.

In order to evaluate the benefit the cooperation for both the overlay and the ISP, we are investigating models capturing the utility-maximising behaviour of overlay applications and the profit-maximising behaviour of ISPs, exploring the resulting Nash equilibriums. This study aims at proving our approach of cooperation, removing tensions between the preferences of the overlay and the ISP, such those that have been analysed from a game-theoretical perspective [4].

IV. CROSS-LAYER INTERACTIONS

The cross-layer interactions we propose will help future applications to provide an adequate QoE by being aware of the underlying networks' capabilities and limitations and by being able to invoke the facilities and services provided by the underlying ISP networks in a dynamic and flexible way. These networks also get information from the overlay so that they can optimise the traffic in their networks, mobilise resources and adapt to the overlay applications, eventually transparently.

Towards this aim, the exchange of rich information between the ISP and the applications is required. The ISP could inform the application about the capabilities of its network: details of access networks (type, link capabilities, coverage, etc.), current network services status (availability of multicast groups, availability of caches, etc.) as well as on other possible metrics (load of routers, bandwidth, delay, etc.). The application, in turn, could inform the ISP about its traffic demand: information related to users (e.g. user location and estimated traffic matrix) or information related to content (quantity of sources, their bitrate, adaptive coding, etc.). This information exchanged between the application and the ISP might be more or less dynamic, might be expressed in absolute or relative terms, and goes further than information reflecting the preferences and policies of the involved business entities as it is currently defined in the IETF ALTO [6] working group by also including information on objectively measured properties of the application and the network.

In addition to information exchange our approach enables future networked media applications to make use of advanced network services in a dynamic and flexible way to

achieve a cost-efficient delivery of high QoE for their users. Examples of such network services include: (1) 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; (2) caching: via the use of specialised nodes, provided either by ISPs or third party entities, to optimise the delivery and save bandwidth in the network; (3) bandwidth on demand: to enable the delivery towards end-users over multiple access networks simultaneously, and provide bandwidth on demand over aggregated access networks; (4) dynamic QoS mapping: invocation and mapping of application QoS requirements network capabilities, end-users devices and access networks; (5) ads/text insertion: in order to offer added-value services that might be monetised by network operators, (6) content adaptation: the presence of heterogeneous end-users devices and network infrastructures will require multiple versions of the same resource that can be efficiently generated using content adaptation. In our solution, the network services can be dynamically requested by overlay applications and instantiated by ISPs.

Overlay nodes are mobile and can potentially join different networks, meaning that a static configuration of the address and access details of the ENVISION interface of a single access network is insufficient. Instead we advocate the use of a dynamic protocol that allows nodes or their surrogates to discover the right ENVISION interface with respect to their current location.

V. OVERALL FUNCTIONAL ARCHITECTURE

In this section, we describe the overall functional architecture of our system (figure 4). As it is a functional architecture, we do not specify whether the building blocks are centralised or distributed or located on specific network or service provider equipment or end-user devices, nor we give recommendations regarding their implementation, such as co-locating functional blocks in the same server or distributing them over many nodes. Rather, the architecture is defined to identify the functionality required to achieve the necessary interaction between application overlays and underlying ISPs and to identify the interfaces and interactions required between the functional blocks to facilitate cross-layer optimisation.

The underlying network is composed of several Autonomous Systems (AS). A managing functionality in the network level is required in each ISP which is willing to cooperate by providing an ENVISION interface. As the aimed cooperation will allow layers to provide information and invoke services reciprocally, appropriate security mechanisms for authentication, authorisation and accounting facilities are needed. At the overlay level, this desired cooperation requires a functionality that allows information sharing. New interactive media applications offering high QoE for their users require functions such as content adaptation that could be provided by various entities including third party service providers. The interfaces were defined after studying the interactions between these functionalities delimited into blocks in the architecture.

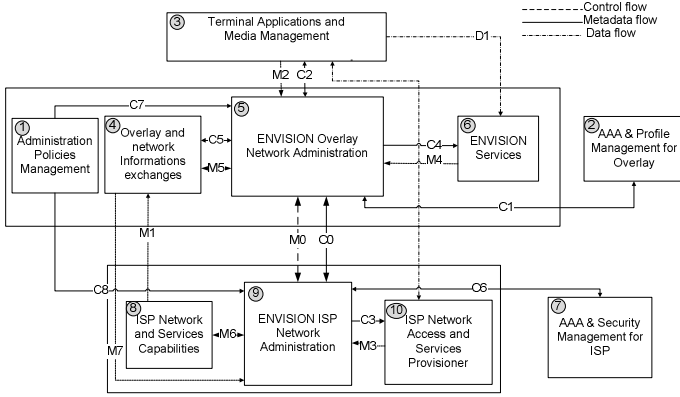


Figure 4: ENVISION Functional Architecture

In the following paragraphs the roles and the interfaces of the individual blocks are described. From the ten main blocks that have been identified, the lower blocks, named 7 to 10, are the network part of our system (e.g. network-level, managed by ISPs), while the others blocks are at the overlay level.

At both levels two blocks are assigned for AAA (Authentication, Authorisation, and Accounting) purposes. Block 2, “*AAA & Profile Management for Overlay*”, enables authentication of users joining the overlay, ISPs cooperating with the overlay and potentially third party services providers. Block 7, “*AAA & Security Management for ISP*”, enables the authentication of overlays cooperating with the ISP. Both blocks offer standard AAA functions, such as accounting facilities, access authorisation, profile management, with security mechanisms and can be accessed respectively by the blocks 5 “*ENVISION Overlay Network Administration*” and 9 “*ENVISION ISP Network Administration*”.

At the network level, the block 8, “*ISP Network & Services Capabilities*” processes, stores and provides information about the network such as the network description, the offered network services, the recommended routing, the ISP policies and preferences, the level of cooperation with the overlay, etc. Via the metadata interface M1, part of that information can be given to the block 4, “*Overlay and Network information exchange*”.

Block 10, “*ISP Network Access and Services Provisioner*”, manages the available network services, e.g., multicast, caches, ads insertion, etc. These services are invoked by block 9 via the command interface C3 to optimise the data distribution to answer an overlay request or independently. Results are reported back to block 9 such as the execution feedback or the status of the services. User data (data flows) passes via this block (data interface D0) for being processed and distributed.

Block 9 is responsible for the overlay management at the ISP domain (intra-ISP). Therefore it receives: (1) information about the data distribution from block 4 via M7; (2) specific requests for further actions or specific service mobilisation from block 5 via C0; (3) network information

from block 8; and (4) feedback about services execution and network status from block 10. This information is compiled and its the optimisation algorithm is executed in order to: (1) determine the appropriate actions to undertake for this distribution context; (2) mobilise network resources and services if needed; (3) provide feedback and recommendations to block 5; and (4) update the network information within block 8 via M6.

At the overlay level, block 4 manages the information provided by the ISPs and the overlay. It allows the overlay accessing the information provided by ISPs via interface C5, and additionally it allows ISPs via interface M7 to access information provided by the overlay and eventually any other information shared between them. It contains data related to the distribution context; being data relative to the overlay application (e.g., video codec, bit rate, number of sources, number of users, necessary services, etc.), relative to information provided by the ISPs, via the interface M1.

Block 6, “*ENVISION Services*” manages the availability and execution of the services provided at the overlay level, e.g. content adaptation, caching or multi-view aggregation (i.e. when multiple streams are retrieved and combined into one single stream and distributed again). These services could be provided by the end-users or by third-party entities and that can be added to the overlay swarm. These services are acknowledgeable to block 9 via interface M4 and commanded via interface C4. The data to be processed and distributed is received via D1.

At the upper level, block 3, “*Terminal Application and Media Management*” represents the end-users side and the application related features. It implements functions for: (1) content generation, consumption, search, etc; (2) data flow handling e.g. transmission, reception, synchronisation; (3) providing information about the user (e.g. interest, profile, etc.) and the terminal via interface C2; and (4) providing feedback about QoE to block 5 via interface M2.

Being the major block of the architecture, block 5, “*ENVISION Overlay Network Administration*” is in charge of managing the cross-domain application overlay. It therefore implements the overlay optimisation algorithms. This block is aware of the complete content distribution context, i.e.: (1) content characteristics and processing possibilities, e.g. adaptations; (2) end-user context, e.g. access network, terminal capabilities; (3) underlying network’s capacities and available services, e.g. multicast and (4) available overlay services, e.g. multi-view, transcoding, caching etc. Executing the overlay optimisation algorithms based upon the content distribution context collected from the different blocks, block 9 sets up a distribution strategy by: (1) defining a distribution topology, e.g. selecting appropriate set of peers, requesting a multicast tree from an ISP, etc.; (2) defining the set of actions that will enhance the QoE, e.g. performing content adaptation for a set of end-users, activating a set of cache servers, etc.; (3) defining appropriate actions to answer eventual requests made by underlying networks, e.g. reducing the data rate for a specified region under congestion risks; and (4) sharing the valuable information about the content distribution strategy

with underlying networks. Obviously some of these actions are subject to security checks performed by block 2 via interface C1.

Finally block 1, “*Administration Policies Management*”, is a configuration function defining and updating the policies for the overlay-ISPs cooperation and its execution parameters allowing more flexibility for these operations

VI. RELATED WORK

Overlay applications are currently totally agnostic from the underlying network infrastructure and thus perform end-to-end measurements to gain some knowledge [1], but this is not in cooperation with the ISP and can lead to undermine routing policies of ISPs [2][3]. To avoid this, several initiatives or research projects have been initiated in the last few years related to promote a cooperation between overlay applications and underlying networks.

The P4P initiative [5] and later the IETF ALTO working group [6] investigate how overlay networks and ISPs can cooperate to optimise traffic being generated by P2P applications and transported over the ISP’s infrastructure. In their approach, the ISP is able to indicate preferences which peers should exchange data. Our solution is somehow related to ALTO, however ALTO is of limited functionality for future media services. Therefore our system proposes a much richer interface that will allow true cross-layer cooperation, both in terms of information exchange and via the possibility for overlays to dynamically request ISP’s network services.

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

Research work studied P2P support in the context of massively multi-participant virtual environments such as [9] for games, [10] for virtual environments and [11] 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 addressed by our approach. Also dynamic activation of network services via the cooperation between overlay and ISPs for improving the QoE is not taken into consideration.

For the definition of interfaces between the underlying network and the application or control level the well known IMS (IP Multimedia System) [12] designs an interface between the control entity and the network entities. However, this interface is still closed and only usable by the ISPs. Also the possible network services that might be dynamically activated are limited. In another initiative, the Parlay/OSA framework has been 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 [13]. The objectives of this group were similar as 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 the two directions) and the supported services were limited to more traditional telecoms applications rather than media applications we aim to support. Furthermore, those two initiatives do not really address issues for massive multi-participant distributed applications as we envision.

Finally, open interfaces are promoted by the OpenFlow initiative [14], but here the interface is rather related to network configurations or functions, and not to activation of network services by overlay applications.

VII. CONCLUSIONS & FUTURE WORK

In this paper, we have presented a new architecture, fostering cooperation between overlay applications and ISPs for an optimised delivery of services to end-users. Overlay algorithms are optimised thanks to information provided by ISPs and the delivery QoE is further improved via the activation of network services, provided by ISPs. The functional architecture and associated interfaces has been defined to support this level of cooperation between applications and ISPs.

While the approach presented in this article has been developed to support interactive, multi-party, high capacity media applications such as those presented in section II 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 could also make use of other network services provided by the ENVISION interface, such as capacity reservation for background distribution of content to CDN nodes.

The ongoing work is focused on the internal functions of each block, including the development of overlay optimisation 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. Finally evaluation through both simulation and testbed experiments is planned.

ACKNOWLEDGEMENTS

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] Gurbani, V.; Hilt, V.; Rimac, I.; Tomsu, M.; Marocco, E.; "A survey of research on the application-layer traffic optimization problem and the need for layer cooperation", Communications Magazine, IEEE, Vol. 47, N° 8, pp 107 – 112, 2009
- [2] Keralapura, C., C., R., Taft, N., Iannacone, G., "Can ISPs take the heat from overlay networks?", In ACM Workshop on Hot Topics in Networks (HotNets’04), November 15-16, 2004, San Diego, CA, USA

- [3] Karagiannis, T., Rodriguez, P., Papagiannaki, K., "Should Internet service providers fear peer-assisted content distribution?", Internet Measurement Conference (IMC'05), October 19-21, 2005, Berkeley, CA, USA.
- [4] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications", *Comput. Oper. Res.* 33, no. 2, 286–311, 2006
- [5] H. Xie, A. Krishnamurthy, A. Silberschatz, and Y. Richard Yang, "P4P: Explicit Communications for Cooperative Control Between P2P and Network Providers", http://www.dcia.info/documents/P4P_Overview.pdf
- [6] IETF ALTO group status page: <http://tools.ietf.org/wg/alto/>
- [7] R. Fortuna, E. Leonardi, M. Mellia, M. Meo, and S. Traverso, "QoE in Pull Based P2P-TV Systems: Overlay Topology Design Tradeoffs," in *Proceedings of the 10th International Conference on Peer-to-Peer Computing 2010 (P2P'10)*, Delft, The Netherlands, August 2010
- [8] R Jimenez, LE Eriksson, B Knutsson, "P2P-Next: Technical and Legal Challenges", tslab.ssvl.kth.se
- [9] A. Bharambe, J.R. Douceur, J.R. Lorch, T. Moscibroda, J. Pang, S. Seshan, X. Zhuang, *Donnybrook: Enabling Large-Scale, High-Speed, Peer-to-Peer Games*, Sigcomm 2008
- [10] D. Frey, J. Royan, R. Piegay, A.-M. Kermarrec, E. Anceaume, F. Le Fessant, *Solipsis: A Decentralized Architecture for Virtual Environments*, 1st International Workshop on Massively Multiuser Virtual Environments, 2008
- [11] S.-Y. Hu, T.-H. Huang, S.-C. Chang, W.-L. Sung, J.-R. Jiang, B.-Y. Chen, *FloD: A Framework for Peer-to-Peer 3D Streaming*, Infocom 2008.
- [12] 3GPP TS 23.228 IP Multimedia Subsystem (IMS); Stage 2 (Release 7), www.3gpp.org.
- [13] OSA/Parlay X, 3GPP TS 29.199 Release 7 specifications, September 2007.
- [14] OpenFlow Switch Specification, version 1.0, December 31st, 2009; <http://www.openflowswitch.org/documents/openflow-spec-v1.0.0.pdf>