Future Media Networks

Research Challenges 2010

October 2010

Produced by the Future Media Networks cluster of Networked Media Systems FP7 projects, unit D2, DG Information Society and Media, European Commission.

European Commission
Information Society and Media

Europe Direct is a service to help you find answers to your questions about the European Union.

Freephone number (*):

00 800 6 7 8 9 10 11

(*) Certain mobile telephone operators do not allow access to 00 800 numbers or these calls may be billed.

(EN) LEGAL NOTICE

By the Commission of the European Communities, Information Society and Media Directorate-General.

Cataloguing data can be found at the end of this publication.

Luxembourg: Publications Office of the European Union, 2010

ISBN 978-92-79-17393-6

doi:10.2759/37178

PRINTED ON WHITE CHLORINE-FREE PAPER

Future Media Networks
Research Challenges 2010

October 2010

Produced by the Future Media Networks cluster of Networked Media Systems FP7 projects, unit D2, DG Information Society and Media, European Commission.

European Commission
Information Society and Media

Table of **Content**

1. INTRODUCTION

1.1 The Future Media Networks Cluster

The Future Media Networks cluster[1] of the Networked Media Systems FP7 projects is concerned with the algorithms, protocols, techniques, tools and platforms for the support, by future networks, of advanced interactive media applications and the seamless delivery of a range of content including text, images, audio-visual streams and 3D media, virtual and augmented worlds, and games.

The approach of the projects in this cluster is based around the idea that the development of future Internet technologies should not proceed in isolation of the way networks are used. The projects start from the point of view of the demanding network problems posed by distributed media applications and content delivery systems. The cluster is therefore in a privileged position to research novel networking solutions in the context of such demanding usage drivers.

The projects in the cluster have a spectrum of approaches to the networking of such applications and content, ranging from P2P overlays that make no assumptions on the underlying network through to visionary and potentially disruptive novel network architectures that are able to access, discover and route content natively. Topics of common interest that are currently under investigation include:

- Network architectures for supporting advanced media applications, including:
- Evolutionary approaches such as overlay networks,
- Interim solutions through techniques such as deep packet inspection and network virtualisation,
- Visionary networking such as native content-centric networking, and,
- Migration paths and roadmaps towards interim and visionary architectures.
- Adaptation of content to network capabilities.
- P2P applications and their interaction with underlying networks.
- Distributed versus centralised media/content servers and positioning of caching functions for content distribution systems.
- Proposals for new business models investigating alternative relationships between content producers, users, service providers, ISPs and CDNs.
- Personalisation and enrichment of media/content according to user context.
- Development of QoE metrics and accountability means to determine the entities responsible for QoE degradations.

1 See section 13 for the complete list of projects and their contact details.

1.2 Overview of this document

This document identifies a set of research challenges in Future Media Networks from the perspective of the projects in the cluster. The major research challenges can be summarised as follows:

- New network architectures to enable content-centric networking by exploiting the potential merging of routing and server functionalities is key for networked media.
- Open and federated platforms for content storage, distribution and search will increase access to large quantities of media and avoid the problems of fragmentation and monopolisation by a few large content publishers/aggregators.
- Cross-layer cooperation and optimisation: between multimedia applications and wireless networks, between overlay networks and ISPs, between service and network providers, considering the users, their preferences and their expected QoE.
- Quality of Experience remains an open area of investigation. The research issues include: defining the aspects that affect QoE; establishing a set of QoE metrics and their relationship with objective measurements; how the different actors in the chain of delivering media experiences can be accountable for their impact on QoE.
- Integration of interactive multimedia communications with multiplayer online games presents a range of challenges, especially with regard to synchronisation and scalability.
- Inherent privacy-preserving features need to be built into future media networks.
- Techniques for reducing power consumption in media delivery platforms need further investigation to assess their effectiveness.
- New business models are required in the delivery chain of future networked media and the solutions to the above research challenges may potentially have a significant impact.

2. FUTURE NETWORKED MEDIA ENVIRONMENTS

Network-based applications and services have mostly followed the client-server model in the past. To a great extent most content has been centrally created and various communications networks have been involved as a distribution means to get information from the single source of the content to many consumers. This has been the model of the printed press, radio and television, theatre and cinema, music performance and the distribution of recordings and it has been replicated in the Internet.

In recent years there has been a trend for more user participation in the Internet-based services that they use. There has been an explosion of user generated, tailored and recommended content and social networking is beginning to replace traditional communications technologies such as email and websites. However, even though information is being created, modified, edited and consumed by a large number of distributed participants, almost all of these services still rely on servers that house and operate the applications on behalf of service providers in large data centres at strategic locations around the Internet. Typical examples of popular applications that only exist for, by and because of significant user participation are Facebook, YouTube, Flickr, Digg, eBay, Second Life and Wikipedia. All of these depend on the client-server model with servers that need to be adequately dimensioned and carefully positioned to ensure an adequate Quality of Experience for their users.

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, operating, or even being the application itself. The traditional model of media streams being produced centrally and then distributed through a largely passive network to passive consumers may still have a role but more advanced, interesting and attractive future applications will be characterised by having many sources from many viewpoints that need to be navigated, dynamically mixed and tailored to suit the needs, tastes and quirks of individual users according to their context. It will be hard to meet the requirements of these future applications with current hierarchical content distribution networks, as this hierarchical model is not well adapted to serve dynamic content from multiple sources especially considering the needs of low-latency live content streams.

In addition to being more distributed and interactive, future media applications will demand much more from the communications network than they do today. For example, new immersive, collaborative environments will require three-dimensional, multi-sensory user participation which demands high capacity, low latency communications channels to exchange a range of information in a large variety of formats. This ever increasing pressure on the network is amply demonstrated by the following trends:

- The increase in demand for HD (and beyond) quality from consumers on the one hand, and from manufacturers of consumer devices such as plasma screens on the other.
- merging standards for high-quality video content e.g. H.264 level 5, with typical resolutions of 2K x 1K or 4K x 2K, requiring bitrates of up to 240Mbps.
- With the standards for ultra HD coming from Japan and 3DTV following not far behind the amount of bandwidth that the networks will require to deliver will become 10 to 100-fold higher than today's capacity.
- Web 2.0 goes video and live. With the rapid deployment of wired and wireless broadband Internet access for any user, a largely reversed traffic load pattern with a vast number of ingress points for media content has to be handled by the operator networks.
- Modern wireless technologies (currently HSPA, with LTE and WiMAX in the near future) are being deployed widely, increasing mobile uplink and downlink bandwidth available for mass contribution and consumption of high-quality streaming media from anywhere.

3. CONTENT-CENTRIC NETWORKING

Content Centric Networking as one of the major themes driving research challenges in Future Media Networks over the next decade. Traditional communication networks have always been based around establishing communications paths to specified end-points, connected to servers that may in turn contain content, media, services. This model requires users or their applications to know where content is located or even the methods by which the content should be accessed. This model is no longer appropriate for advanced Networked Media Systems. The location of content, servers, services or users should not be of concern. One of the major research challenges is therefore how to make communications network as location-independent as possible. Furthermore, as applications become more distributed and interactive, many of the traditional network functions of naming, addressing, routing and forwarding are strongly influenced by the applications and the content/media in question. There is therefore a strong case for considering whether network layer functions should be separate from application logic. Memory is getting cheaper and will soon approach the same prices as hard disks. Servers and routers may be able to "merge". Routers will be able to store a significant amount of content and this will potentially produce a significant paradigm shift on how routers can be built. Networks will be able to replicate content efficiently at a small cost. In this new "visionary" reality, routing and content signalling protocols will need to be redesigned almost from scratch. This presents a great opportunity for Europe to be at the core of this new reality.

Within this main theme of Content-centric Networking there are several additional research challenges that may be summarised in the following:

- Multi-stream, multi-sensory stream coordination and synchronisation: Media services will be able to use a plethora of types of content with the need for coordination and synchronisation. These streams will potentially be produced and edited by several people simultaneously which puts the real-time requirements greater than anything we have today.
- Distributed Search: For many types of content, where the owners of the meta-data may not want it exposed, search will need to be distributed. Search will not only be by keywords but by images, videos, thought-readings. Results need to be accurate and with a response time not much higher than centralised services. Moreover, it can be foreseen that many specialised search engines will be created and a unified view of the search results will be crucial.
- Location based media: Content will be displayed dependent on user location and context. Augmented reality applications, with many potential uses, will create many business opportunities for the European content industry. Smart cities, car windscreen displays, targeted outdoor publicity are just some examples. This presents substantial challenges from the network delivery point of view. User mobility, for example, needs to be part of the future Internet infra-structure and content needs to be easily found and retrieved using whatever access technologies the user has available.
- Personalisation and Rating: The success of services like Hulu demonstrates that targeted publicity can be a successful business model. With the "end of broadcast media" users will want more and more personalised media (both content and/or publicity) and managing the scale of users is a big challenge. The use of social networks, preferably distributed and open, can be a powerful and efficient tool. This personalisation also has to include rating. With the continuous growth of content available in the Internet, users will have to rely on each other to rate this content (children safe systems are also a big driver for this). Doing this in a safe way, where content producers cannot easily fake the ratings even with crowdsourcing, is a major challenge.

- Privacy protection: With increased personalisation comes a substantial increase in potential loss of privacy. Users will want to limit the entities that know which information they accessed. In some scenarios they may even want to limit this to zero entities. Privacy needs to be at the core of Future Internet protocols and distribution systems.
- New business models: At the end of this decade, the content and media industries are panicking for survival. Newspapers all over the world are closing or having zero profits due to the widespread availability of "free" news. Content producers (musicians, TV stations, movie studios) have also failed to come with solid business models for the digital age where the cost of replicating bits is zero. Micro-payment systems that are open, secure and not owned by a single set of organisations are a potential enabler for future business models. Users will be able to purchase news and content at a smaller cost than today but in a much easier way. This will require research on how to do this efficiently and in a distributed way.
- QoE accountability: As more interactive and high bandwidth content is going to be made available, it will be more difficult to guarantee a high quality of experience. Although network mechanisms to achieve this are well defined (much of it done in previous European projects), mechanisms still need to be built into the Internet fabric than allow users to easily find who is responsible for bad quality of service. Providers need to be accountable in order for markets and user choices to be a reality.

4. CONTENT MEDIATION IN FUTURE MEDIA NETWORKS

This section introduces as new research challenges in content distribution some outstanding problems that need further study regarding this area and the benefits that a mediation framework might bring to content distribution. In particular, two key problems have been identified in that ecosystem:

- The increasing number of intermediaries makes much content accessible only for particular user communities, resulting in global content search and direct access being fragmented. The key issue lies in the lack of a global content naming scheme to access the content, which forces end-users to search the content through relevant intermediaries, maintaining a multiplicity of accounts, front-ends, tools and applications for content discovery and consumption.
- Today's networks are unaware of the content they are transporting. Due to this unawareness, networks cannot apply the most appropriate end-to-end transport strategy to provide the adequate quality of experience for the end users. Besides, flash crowds in live events, which can potentially lead to traffic peaks, cannot be efficiently addressed. Moreover, even when the networks are well prepared, intermediaries acting as Content Distributors (Internet Content Providers such as YouTube or CDN providers such as Akamai, or P2P platforms such as Octoshape) cannot be aware of the network capabilities, traffic conditions, or the transmission requirements for the content. Therefore, the content is delivered far from the most efficient way.

These problems stem from a common issue in the current Internet paradigm that should be highlighted as relevant research challenge: there is not an appropriate linkage between the "worlds" of content and data transmission, hindering the effective coordination between them.

At the present time, the research community is devoting efforts to help to overcome this lack, being one of the most popular approaches allowing the treatment of the content as a primitive itself. Thus the proposed architectures are designed to deal with all aspects of the content natively, facilitating the access and publication of content. Alternative approaches to current computer networks, such as Content Centric Networks [JAC09] or PSIRP [PSIRP], are oriented in this direction.

A key challenge is in how to provide an adequate linkage between both worlds (content and data transmission), allowing Internet Service Providers to act as mediators for content publication and distribution. One such approach is based around a content-oriented Internet that provides fast and simple content access with network-awareness that is gradually deployable over the current and the Future Internet. For this purpose, a unified interface where content is treated as a first citizen is envisaged. In addition, the mediation by Internet Service Providers would make possible an improvement of content delivery in terms of quality and effective bandwidth utilisation.

Figure 1 shows a 2-plane approach for that mediation, including the so-called Content Mediation Plane (CMP), in charge of offering those unified interfaces to Content Consumers and Content Providers, and a Content Forwarding Plane (CFP), in charge of the delivery of the content based on its knowledge of both the network and server status.

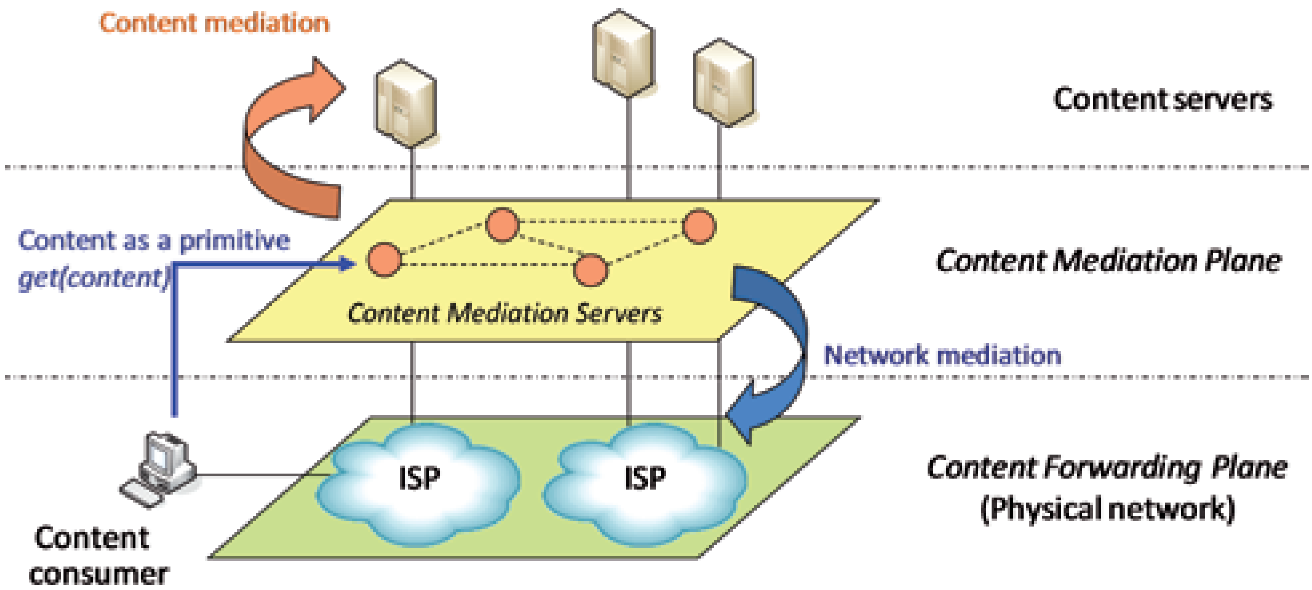

Figure 1 Content Mediation in Future Media Networks

Fundamentally the CMP is a unified, secure and scalable Internet-wide content organisation structure (from global content naming and addressing publication of the content to the discovery, and the final delivery of the content to Content Consumer) which defragments the various content access and distribution avenues regardless the nature and requirements of the content (e.g. real-time vs. elastic, pre-recorded vs. live etc.).

On the other hand, the CFP is basically a content transport system supporting different types of distribution modes (unicast, inter- and intra-domain multicast, anycast, peer-to-peer) that is adaptable to the network and content server conditions.

Thanks to the mediation offered by the CMP, the following key technical advantages can be achieved:

- Unified access to the content whatever its nature and location
- Content delivery with guaranteed QoS
- Point-to-multipoint content delivery capabilities, reducing bandwidth needs for live content
- 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.

Related to the mediation-based approach, there are several research issues that might be developed in the area of the Future Media Networks:

- The design of a global content naming and addressing scheme supported by an infrastructure capable of scalable content search and resolution.
- The creation of a unified interface for the end-users to access the content. In addition, this access must be independent from the content location.
- Hand to hand with this mediation framework, there must exist content forwarding equipment able to perform content-based forwarding at speeds similar to the ones in IP-based forwarding while preserving the coordination with upper layers.

5. OPEN CONTENT DISTRIBUTION NETWORKS

The domain of multimedia content over the Internet is experiencing major breakthroughs and structuring changes. The rise of intensive and personalised consumption models, the multiplication of home terminals providing access to audiovisual content over Internet (e.g. connected TVs and game consoles), the success of social networks, the emergence of High Definition (HD) content and the availability of high-quality content (e.g. Live Olympics) contribute to the strong growth of online multimedia traffic (approximately 50% to 60% per year in core networks the last five years).

This traffic heavily loads on the network infrastructures of the telecom operators, which requires huge investments they state they can (or should) not necessarily compensate for alone. This problem has already had its share of headlines [IPLAYER] [OFCOM]. It crystallises now around the so-called "Network Neutrality" debate [NET] and it is expected to worsen in the coming years with the traffic increase [CVI10]. Another consequence is the key role CDNs (Content Networks Distribution) [CDN] are now acting in the Internet, with as main advantages a decrease of traffic on the Internet core networks and transit links (if the caches of the CDN are outside the core networks, see Figure 2), and also the typical reduction in latency in delivery of the content (by avoiding the need to cross the Internet backbone). It is estimated that approximately a third of all traffic Internet is distributed through a CDN [WAR08].

However, CDNs today have a number of shortcomings that will become more prominent in the near future when the volume of multimedia content to be delivered will be one or two orders of magnitude higher.

Firstly, CDNs completely rely on best effort delivery for the "CDN last mile", the path between the end user and the last server delivering the content to it (See Figure 2). It is important to note that the term "CDN last mile" can be somewhat deceptive as from a network perspective it includes both the last mile access (e.g. xDSL link), and the aggregation network (and sometimes even part of a regional network, see Figure 2). And although the access links are in principle capable of carrying video flows in standard definition (access bit rates have been growing at a compound aggregate growth rate of about 40% per year and mid 2008 a typical bit rate of an access link was between 2 and 10 Mb/s [TJE04]), and will be able to carry high definition video content in a few years from now, the aggregation networks are heavily overbooked. Since until recently the Internet was mainly used for request/response type of services generating on-off type of traffic (where the user requests an object and then spends some time reading the downloaded information), overbooking factors of 20 and more are not uncommon. The evolution towards more streaming (multimedia) services will require sustainable bit rates with much lower overbooking factors. Merely upgrading the aggregation networks to resolve the current overbooking and the expected two order of magnitude increase in multimedia traffic will not be trivial for the network operator, not form a technical point of view, nor economically viable. New solutions for delivery over the aggregation network need to be designed.

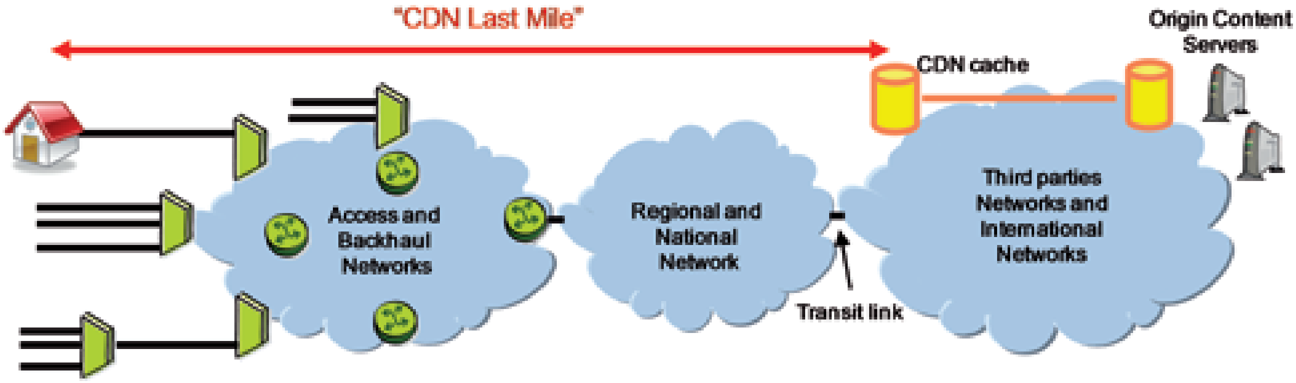

Figure 2 High level CDN view

Secondly, CDNs do not take into account primary network information for delivering content to the end-user. As they mainly are positioned as an application layer overlay network, they may rely on some application layer probing of network conditions, but they do not interface directly with the network elements or management elements of the network operator. Especially for solving the delivery over the aggregation networks, a more network-aware CDN could use valuable information such as specific topological information, pre-congestion status on particular links, or even localised content consumption pattern deviations.

Thirdly, CDNs today are closed systems consisting of an almost worldwide distribution of servers that rely mainly on a set of proprietary mechanisms and interfaces. Multiple CDN systems coexist side by side but without any interoperability in between them. Whereas openness of the CDN system would definitely benefit global CDN systems in terms of multi-vendor solutions, for deep CDNs extending into the aggregation or even access networks, openness is an absolutely mandatory property. It will not make economically sense to deploy multiple servers (proprietary, from different CDN operators) next to each aggregation or access node. Only an open system that can be used for content from any origin can be a viable solution.

Finally, the CDN business model today is limited to offering content producers the prospect of lower latency access for the end-user (and network operator's lower transit costs when deploying the CDN servers). For real quality multimedia delivery over the aggregation and access network, CDNs will have to interact with network level mechanisms from the (access) network provider. Hence also the business model will have to be extended to incorporate this relation, and provide incentives for all parties to invest in future evolution of the network and service infrastructures, while at the same time comply with regulation guidelines.

Several research initiatives have been launched to tackle these issues by investigating several technical solutions in a balanced value chain with benefits for the end-user, content provider, service provider, and network operator. Some of the major research issues are described below (see Figure 3):

- Deployment of caches deeper into the aggregation networks, possibly up to the access nodes: Although content popularity is heavy-tailed (see [AND04]), recent reports [BEN08], [WEB08] have argued that recommendation engines (guiding the user to the next piece of content) and fashion trends (i.e., the word of mouth hyping content such that it is consumed in bulk over a relatively short period of time) considerably steepen the content popularity curve, making the content far more cacheable. Since this implies that the same piece of content is consumed by many users, bringing copies of this content closer to the user can alleviate typical congestion caused by flash crowds (i.e. a larger than anticipated set of users attempting to access a just published piece of content). So, an architecture that allows storing content closer to the user than current CDNs do, will be better equipped to handle the surge in overall bit rate due to the consumption of video over the Internet, all this at a cost much lower than just upgrading the network with more transport capacity.

- Network-controlled adaptive content delivery: Independent of the deepness of the deployed caching architecture, there is always a network path between the place where the content is stored and the end-user that has to be traversed, and that can be congested. Techniques need to be researched where the quality of the video is gracefully degraded as the congestion level increases, exploiting users' tolerance for smooth and temporary quality degradations (rather than access denial or long lasting black screens). It is based on a law of diminishing returns when encoding audiovisual information: the bit rate associated with base quality is usually low enough, but the additional bit rate to gain one unit in quality is higher as the quality gets higher. Consequently, if the network is dimensioned to transport the highest quality under normal circumstances, there is a lot of headroom to reduce the bit rate during congestion times before the quality drops below unacceptable levels. Scalable video coding (SVC) can be used to implement such congestion control that gracefully degrades the video quality.

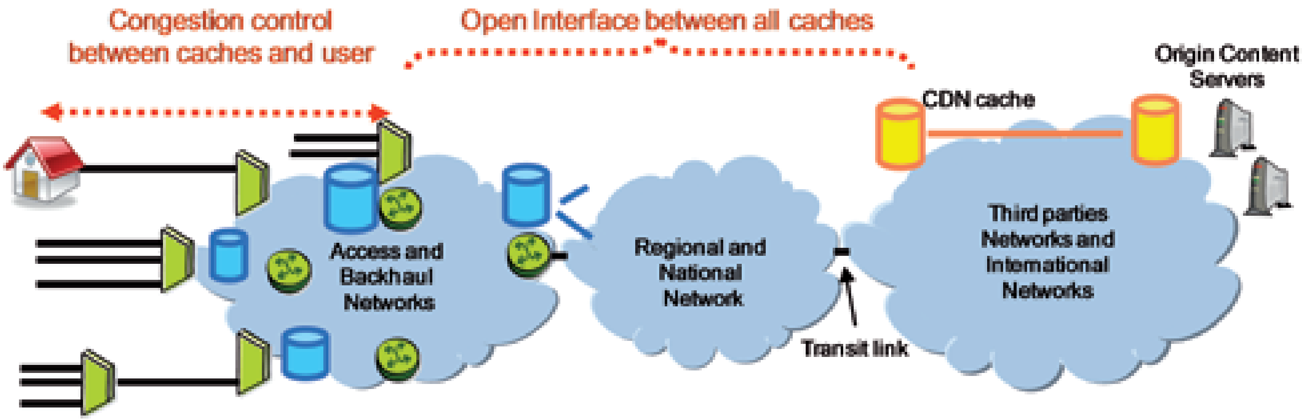

Figure 3 Architecture for Open Content Distribution

Content distribution services deployed over telecom networks are still basically set up on a two-layer model, with connectivity networks at the lower level and application overlays at the upper level. Yet current strategic development trends dictate the necessity for integrated networks with increased real-time interactions between these layers in both fixed and mobile networks.

In the literature the term "Context Awareness Principle" is used to refer in a global manner to such dynamic interactions between the connectivity and applicative layers. Context-awareness is promising in gathering and managing different types of context information related to users, networks, terminals and services; and therefore in enabling: i) Content adaptation: allowing dynamic adaptation of content according to different users and their contexts as well as the network, terminals and services contexts, and ii) Extending the Quality of Service scope to tackle the user satisfaction, the resource level (access network capabilities, terminal(s) capabilities), and the semantic level (content and service metadata, user profile and preferences). Network/service cross-layer systems designate the whole set of mechanisms providing such an increased dynamic interaction between the connectivity and applicative layers.

In the short term the cross-layer systems shall mainly focus on interacting with "classic" network features: device discovery and IP connectivity management, mobility management, network resource monitoring and reservation. Yet these features alone constitute strong drivers for increased cross-layer interactions, as illustrated by two initiatives: "IP Multimedia Subsystem" [IMS] and "Application-Layer Traffic Optimisation" [ALTO].

In the medium term, the evolutions of usages and terminals will bring about a major shift in the content distribution architectures, possibly leading to an in-depth transformation of the current connectivity networks owned by the telcos into feature-rich, convergent and open context-aware network clouds. The main consequence for the cross-layer systems will therefore be an expansion of their functional scope, from "basic" network features to advanced content distribution and application-centric ones (incl. streaming, storage, transcoding processes, etc.).

As regards the long term, numerous ongoing works aim at completely revising existing networks in a 10 year perspective. They converge towards a model relying on disruptive concepts, such as network of information and content centric networking. Research in this field is still at an early stage and there is as yet no common vision of what the network of information will finally be. The first issue and first priority will be therefore to make this vision emerge by working in a collaborative mode, through open innovation projects, and later by influencing the standards. One DARPA project worthy of note is that led by Van Jacobson of PARC (Palo Alto Research Center) on Content-Centric Networking [PARC]. The US activity has stimulated other regions to examine the need to sponsor similar research in collaborative projects such as FP7 4WARD and SAIL projects [4WARD] and [SAIL].

6. MEDIA ECOSYSTEMS FOR USER-SERVICE-NETWORK COOPERATION

The section proposes a novel concept towards the deployment of a networked 'Media Ecosystem'. The proposed solution is based on a flexible cooperation between providers, operators, and end-users, finally enabling every user – first – to access the offered multimedia services in various contexts, and – second – to share and deliver his own audiovisual content dynamically, seamlessly, and transparently to other users. The idea is to provide Content-awareness to the Network Environment, Network- and User Context-awareness to the Service Environment, and adapted services/content to the End-User for his best service experience possible, taking the role of a consumer and/or producer.

Digital multimedia services and networked media content are progressively and rapidly playing a key role in citizens' economic, social and cultural prosperity, bringing breakthrough opportunities in various domains, from industry, communication to education, culture, entertainment, and inevitably business. A new "Media Ecosystem" is hence foreseen to arise, gathering a mass of not only existing but also new potential content creators and media service providers, essentially stemming from customers (active users).

Such "Media Ecosystem", by analogy with the ecology or business counterparts, can be characterised by interworking environments[2], to which the actors belong and through which they collaborate, in the networked media domain. Those environments are:

- The User Environment, to which the End-Users belong;
- The Service[3] Environment, to which the Service and Content Providers belong;
- The Network Environment, to which the Network Providers belong.

Each of the above environments is nowadays present in real deployments, but there is a profound lack of collaboration between them. The User context is not taken into consideration by the Service (or Content) Provider delivering the service (content), which, as a consequence, is not capable of adapting the service (content) to the capabilities of the user. Likewise, current architectures do not include any exchange of content- and network-based information between the Network layers and upper layers, implying that Content-Aware Networks and Network-Aware services as well as applications have difficulties to emerge. This network service neutrality considered many years as a good principle, proves nowadays to be a weak solution for network services, especially considering the new multimedia communications and their increasing importance in the future Internet.

Consequently, the real challenge and actual objective is to find the appropriate means for efficient cooperation between entities of various environments so as to provide the End-User with the best service experience. Additionally, the "to be deployed Media Ecosystem" must enable users to ubiquitously access any service, over any media and via different devices in a customised way, including QoS/QoE attributes, while also providing the means for them to become engaged and assume the role of content/service creator. Thus, the "Media Ecosystem" must also allow Service Providers to offer complex services composed with multiple elements coming from different originators, and offering valuable opportunities for service exploitation.

2 "Environment", is a generic and comprehensive name to emphasise a grouping of functions defined around the same functional goal and possibly spanning vertically one or more several architectural (sub-) layers. This name is used to characterise its broader scope.

3 "Service", is a high level service, seen at application/service architectural layer.

6.1 Challenges beyond the state-of-the-art

Several new features are required inside the Network/Service/User Environments in order to enable the future generic "Media Ecosystem" that:

- At network level, realises and offers to upper layers, a rich and virtualised networked multimedia space, through an all-IP prototype for environment, which can be customised and exploited by all sectors for delivering networked media content. Challenges in this area are:
    - Applying Content-Aware Network (CAN) concepts [SUB01], to perform network/transport intelligent content-aware processing (routing, forwarding, dynamic adaptation, security, etc.) [CAR03] [AGA07] for existing and future emerging applications in a scalable, open and optimised way. To this end, a new virtual CAN layer is proposed on top of traditional network/transport layer;
    - Realising Distributed Management and Control for customising the CANs to respond to the upper layer needs, including 1-1, 1-n and n-m communications, and also allow efficient network resource exploitation at network provider level;
    - Performing cross layer optimisations between the virtual CAN layer and upper layers, including but not limited to P2P approach [HEL09]. This optimisation will take benefit from network awareness capabilities of the upper layers;
    - Extending CAN functionalities for achieving an efficient collaboration with elements in the Service Environment, enabling content-awareness and network-awareness;
- At service/content level, delivers enriched networked media services and content, which can be efficiently exploited by End-Users. Areas for study in this area include:
    - Elaborating a new approach for the delivery of services which includes the Home-Box as a new essential element in the service distribution chain, capable of advanced functionalities, such as handling service management and adaptation, user mobility, security;
    - Creating a new virtual Home-Box layer, composed of virtually interconnected Home-Boxes (in traditional distributed or P2P mode), capable of advanced ways of service/content provisioning;
    - Clearly identifying and dissociating the Service/Content Providers' roles and capabilities and the Home-Box layer role and capabilities in terms of service/content exploitation and delivery, which will lead to the vision of their efficient cooperation;
        - Enhancing the types of services to be made available:
        - From the delivery point of view, through the servers/Home-Boxes, in various modes;
        - From the discovery point of view, with the introduction of a new type of Service Registry;
        - From the management point of view, with an efficient management process;
        - From the composition point of view, with the realisation of a Service Composition Engine supporting streaming media applications;
    - Achieving an efficient collaboration with the User Environment and with the CAN Network Environment, enabling User Context-awareness, Network-awareness and Content-awareness;

- At user level, allows the users to consume and/or generate content and exploit services delivered by components of the Service Environment. Issues for investigation are:
    - Adding new dimensions to the user by giving him the possibility to have several roles, such as:
        - Content and Service Consumer;
        - Content and Service Provider;
        - Content and Service Manager;
    - Elaborating a User Profile for characterising the static and dynamic parameters of the user and his context, in order to be exploited by Service Environment elements (HB) for the delivery of adapted services;
    - Permitting any user to access/deliver/manage any service/content on any device from anywhere and at any time, thanks to a specific User/Service interface and QoE monitoring tool at the user's terminal side;
    - Achieving efficient collaboration with the Service Environment, enabling User Context-Awareness, for the best End-User experience.
- At all levels, performs a constant monitoring in several points of the service distribution chain and regulates a twofold adaptation action, at the virtual Home-Box Layer and at the virtual CAN Layer [HEL09].

6.2 Outstanding problems

Some of the major problems to be solved are:

- Scalable and robust service-aware networking architectures definition, including:
    - Connectivity-to-network, network-to-network services, network service-to-service;
    - Cross-domain interoperability and deployment;
    - Optimal orchestration of available resources and systems; Interrelation and unification of the communication, storage, content and computation substrata;
- End-to-end delivery adapted to User Profile, heterogeneity on consumption and transport;
- Services availability, mobility and security;
- Flexible, hierarchical, autonomic and recursive management.

In term of business actors, solutions are needed that would allow:

- Network Providers to preserve all their deployed and invested infrastructures, but doing - at wish- incremental or radical further development;
- Content/Service Creators to keep the way they create new added-value; services/content, but having new capabilities offered;
- Content/Service Providers to enhance their offering with respect to user needs;
- Consumers to smoothly adapt to a new way of using content/service;
- Transparent process in terms of difficulties, apparent process in terms of improvements;
- No necessity for drastic investments, new vision for more revenues;
- Propose new elements and tools to improve their everyday work and life.

6.3 Possible approaches

An evolutionary Future Internet architecture can be foreseen that focuses on content and the associated service(s). The goal is to transform corporate, academic, ISP, etc. networks into a service-centric environment. The existing network environment is upgraded, while maintaining full backwards compatibility, in order to:

- Adapt network dynamics to inherently recognise and support services with special needs for treatment, resources and functionalities: service awareness;
- Improve existing and incorporate new network services: multicast, anycast, security, mobility, QoS, e.g.:
    - Provide true E2E service awareness while covering core, aggregation/metro, access and local networks;
    - Create end-to-end flow based multi-level security;
    - Create multiple virtual private networks for smart grid, public safety on the same infrastructure;
- Provide an efficient compatible environment for the transport of existing and new types of end-users network applications data – e.g. for multimedia streaming, teleconferencing, VoIP, cloud computing, gaming, social networks, sensor networks, naming [JAC09] etc.;
- The solution will be flexible, covering all aspects related to network-neutral or non-neutral behaviour.

The key points of the solution are:

- Network Resources Virtualisation:
    - Slicing techniques;
    - Service-centric overlay approach;
    - Efficient flexible and distributed FI architecture;
- Awareness:
    - Service awareness at the network level;
    - Network awareness at the service level;
- To facilitate Service/Network coupling, thus offering a dramatic improvement of Quality of Experience for a variety of services;
- Resource management: hierarchical, autonomic, adaptable and virtualisation-capable providing:
    - Resource aggregation and optimal allocation;
    - Differentiated service handling;
    - Flexible service management;
    - Trust and security;
- The management systems accommodates different degrees of local autonomy of the sub-networks/sub-systems depending on the business model applied;
- Inclusion of new **business model** features, capable to satisfy the Network Infrastructure Providers, Operators, Service Providers, End-Users, and/or third parties as stakeholders.

7. COOPERATION AND CROSS-LAYER OPTIMISATION BETWEEN NETWORK AND APPLICATION OVERLAYS

Future networked media environments will differ significantly from today's applications in two important dimensions. They will be high-quality, multi-sensory, multi-viewpoint and multi-streamed, relying on HD and 3D video which will place unprecedented demands on networks for high capacity, low-latency, low-loss communications paths. Advanced media applications will also 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 user and provider resources, to provide and access a seamless mixture of live, archived and background material. In addition to meeting the capacity and quality challenges, future networks also need to provide mechanisms for the highly dynamic discovery of distributed content and other participants and to support the communications between unpredictable and arbitrarily large meshes of network endpoints, distributed around the entire globe.

High-definition, highly interactive networked media applications pose major challenges to network operators. Multi-sourced content means higher quantities of data throughout the network, putting additional pressure at the network edge for unprecedented 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. Rather than simply throwing bandwidth at the problem an alternative, more scalable approach is to develop 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. Applications need to be able to accommodate unpredictably large numbers of participants in a cost-effective way, while still maintaining high responsiveness to deliver a high Quality of Experience to the participants. Content, which itself is changing dynamically in scale and context according to user participation and behaviour, needs to be adapted to network capacity and capabilities, and networks need to be aware of the nature and needs of the content it is transporting.

The challenge of cross-layer optimisation can be addressed by:

- increasing the degree of cooperation between ISPs and the networked applications they are conveying;
- optimising application overlay networks to make best use of the capabilities of the underlying networks and the participant end users; in particular trying to localise the traffic generated by the application as much as possible, as well as avoiding network bottlenecks;
- providing the means by which service providers can access and mobilise specialised network resources to achieve efficient distribution of highly demanding content streams;
- enabling dynamic adaptation of the content to meet the abilities of the underlying networks.

Because applications will be more participatory and interactive, their fundamental way of working is heavily dependent on user behaviour. This implies that it isn't just the content that is housed in diverse locations but that the logic of the application itself will also be distributed. This means that 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 and with the providers' servers to form the participatory, multi-sourced communications applications envisioned. The application level interaction between participants, the discovery and processing of information/media sources and the distribution of tailored content are all performed by the same network of participants, acting as peers at both the application logic and the content distribution levels.

The following sections outline the research challenges and issues that need to be addressed to enable the future media applications to run over the Internet providing adequate Quality of Experience (QoE) for the users and cost-efficiency for the involved business parties.

7.1 Fostering Cooperation between Applications and Networks

Application-layer networks are global overlays running on top of the Internet and it is essential that the participants have a high QoE considering the highly demanding nature of interactive multi-participatory communications including HD and 3D video. Today's media overlay applications are mainly in the field of file transfer, because live media streaming applications offer a limited QoE. This means that future overlays need to be aware of the underlying networks' capabilities (and weaknesses) and to be able to influence how their data is transported across the application-layer network using the facilities of the underlying ISP networks. This is a challenging task considering the inter-domain nature of applications with participation of users around the entire world.

The P4P/ALTO initiative investigates how overlay networks and ISPs can cooperate to optimise file-based traffic being generated by P2P applications and transported over the ISP's infrastructure. The ISP is able to indicate a preference for which peers should exchange data to avoid over-utilisation of its network or the unnecessary loading of high-cost resources such as inter-provider links. The P2P network benefits by avoiding congested parts of the network, resulting in higher average throughput.

The overlay-ISP interaction in ALTO is of limited functionality for future media services. The ISP provides information to the overlay on its preference for the ranking of peers according to the list/traffic matrix supplied by the overlay. However, live and on-demand multimedia services cannot be deployed efficiently without relying on network services. Examples of valuable services to effectively increase the capacity are multicasting, caching, fallback to lower bitrates by transcoding and quality adaptation. Increase in quality requires access to bandwidth reservation, traffic differentiation, mobility management, identity authentication, authorisation and geo-location. These network capabilities – today hidden in the walled gardens of network providers – are required to implement efficient network-aware services. To achieve efficient cross-layer integration, these network capabilities need to be made available by network providers to service developers and integrators. A major challenge, therefore, is to expand and enhance the overlay-ISP interaction, by developing a comprehensive, media-aware, open and standardised interface between the ISPs and application overlays.

Besides explicit interaction with the ISPs, the application should also continuously monitor the network status, this may be possible by integrating in it a fully distributed monitoring platform, which is in charge to continuously monitor both application layer performance indices, such as the quality of the individual peer-to-peer connections; and 2) Quality of Experience, expressing in a concise way the end-user QoE by aggregating individual peer-to-peer connection qualities. Network metrics, such as network latencies, packet losses, path capacity and available bandwidth, etc. Measurements collected through network monitoring provide complementary information with respect to ALTO that can be promptly exploited to trigger overlay reconfigurations when it is needed (i.e. when some congestion arise) [NAPA].

7.2 Optimising Overlay Applications

The scope of the overlay-ISP interaction in ALTO is limited to the viewpoint of a single ISP and the peers located on its domain. Given that our applications are global in coverage, and require end-to-end traffic optimisation involving several hops in different networks, it is necessary to collect information from many underlying networks. There are several problems associated with collecting and using this information: data from one network may conflict with that provided by another; the quantity and quality of the information may differ from ISP to ISP and some may not offer any information at all. The harmonisation of the information gleaned from the ISPs, the aggregation of the information collected from different ISPs, its auditing and augmentation with additional data collected by the overlay and its subsequent use for the global optimisation of the application is a major research challenge.

Interaction with the underlying network provides essential information to a number of application functions that need to be optimised in order to support the future networked media environments.

In high volume applications where users are interested only in part of the available content, and this interest changes quickly over time (e.g. the micro-journalism use case), the application must provide the interest management techniques to determine which of the content sources need to be distributed and to which groups of users, so that the QoE remains high and within the capabilities of the given application and network resources. Selecting the appropriate subset of sources, needs to be optimised to best match, not only the interest of the participants, but also the capabilities of their access means and the underlying network conditions, while minimising the impact on the overlay topology.

The distribution of the content from a given source to a large and dynamic group of participants is a challenge in itself. Intelligent algorithms need to be in place to determine which application resources (e.g. caching servers, NAT traversal gateways etc.) 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. The related topology creation and data scheduling algorithms need to be optimised depending on the requirements of the particular application for throughput, delay, loss and responsiveness to application-layer interactions. Resiliency is also an important objective, in particular in overlay applications where users participate in the distribution of the content but may dynamically change their participation to a particular stream, or the application altogether and leave the overlay with high churn. Caching techniques for content that remains popular over a period of time are also of particular interest as they can significantly reduce the amount of resources required to distribute this content.

Finally, robust and efficient control data management techniques are required to enable the coordination of the media application through a control layer operating across a widely distributed set of nodes. Dynamically changing information regarding users and content sources participating in the application needs to be distributed efficiently and consistently to a large number of nodes. This is further complicated by the dynamics of participants' changing interest in content sources that are constantly joining and leaving as well as new network and application resources becoming available. Data management may be distributed uniformly to a number of overlay nodes, or the application space may split to distinct areas of responsibility. The protocols electing the nodes to participate in the infrastructure, assigning data to nodes or forming and inter-connecting the areas of responsibility, replicating the data and handling churn are all outstanding research areas.

7.3 Supporting Heterogeneous Access Means

Until now digital coding and encoding have been designed following the client/server paradigm but now applications will have to deal with the fact that the content may come from several sources and terminal devices with different capabilities, residing in networks that offer different services. Applications will have to adapt and select "quality layers" with a brand new set of constraints and circumstances. Another major challenge, therefore, the adaptation of content to the capabilities of the core and access networks, user devices and user preferences. This includes the encoding of media streams and other forms of content for transmission from the originator to the application; processing and mixing of media sources in the distributed application to produce derived content - adding value, customising and tailoring the data/media; encoding and distributing the derived content to the consumer(s). 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).

Content adaptation for network-aware multi-participatory interactive applications presents several challenges. Intelligent techniques are required to dynamically select the content sources and corresponding quality levels to be transmitted to each recipient, following the application layer interactions and responding to the changes in network conditions. These techniques need to maximise the delivered QoE, be responsive to changes in the environment while minimising the impact on the overlay topology and reconfiguration. The content distribution algorithms need to be enhanced to treat differently the data packets depending on their layered codec priority.

7.4 Mobilising Network and Service Infrastructure Resources

The distribution of high-volume content like 3D HD video to a large number of recipients requires a huge amount of bandwidth, storage, processing and other service infrastructure resources. Optimisation techniques at the application and network layers will reduce the required resources to a minimum but this is still a significant amount of resources to accommodate the needs of advanced media applications if they rely only on traditional means.

Pre-provisioning and advance payment for bandwidth or other resources is considered highly inefficient for new applications that may need to accommodate an arbitrarily large set of end users, with unpredictable topological distributions and traffic profiles. It is therefore essential that the applications take advantage of the resources of the end users themselves wherever and to the extent that this is possible; user bandwidth, processing power and storage resources differ significantly depending on their terminal device (STB, PC, mobile phone, etc.) and their physical access network (ADSL, FTTH, Wi-Fi, WiMAX, 3G, etc.). To achieve this, applications need to provide incentives to their end users. Even assuming the full cooperation of end users the demand for resources may still exceed supply. This is very likely to be the case in mobile environments where upload capacity, storage and processing power of end devices might be very limited.

An alternative to be investigated is the active participation of the ISP in the application overlay. An ISP may directly benefit from strategically contributing resources to the overlay. Providing a node for caching content, for example, may result in reducing the load in the ISP's inter-domain links. It is expected, however, that the resources required by the application may exceed those that the ISP finds beneficial to contribute voluntarily. An ISP may also offer service infrastructure nodes and associated bandwidth resources to the overlay, for a price. The ISP is in an advantageous position to offer such resources compared to other third party service infrastructure providers, for several reasons. Firstly, the ISP can provide service infrastructure with associated network level guarantees. Secondly, provided that it has an appropriate prior agreement with its customers, the ISP can also mobilise the resources of its customers, through controlling their set-top boxes, for example, and offer them to the application to act as content caches or bandwidth multipliers. In the latter case a multiplication effect is provided by the users' STBs downloading only part of the content and uploading it as many times as possible.

7.5 Alignment of Incentives between Actors

Another important challenge for future media networks is the alignment of incentives between the diverse set of economic actors that participate in the creation, post-production, distribution, consumption and reproduction of digital media works.

If one considers the traffic matrix of greatest benefit to a given content distribution overlay, it is clear that it will depend on its preferences regarding cost, QoS, resource availability, data caching and replication. On the other hand, if one considers the traffic matrix of greatest benefit to the ISP providing it with network connectivity, it will depend on its infrastructure and transmission costs, the background traffic that it carries and its traffic engineering policies. Thus, a tension arises between the preferences of the overlay and the ISP and this can be analysed from a game-theoretical perspective [ABEA06].

Similarly, issues surrounding content popularity and geographical correlation with content demand [Nef05] will create incentives tensions between media consumers and ISPs [CJMW05]. Content delivery systems might not be equipped to distribute content in the way that would be most attractive to content producers [BKR06], and content distributors might find it difficult to provide incentives for high-quality user-generated content [YAA08]. Finally, the capacity for regulators and the legal system to modify this complex technical environment is not completely understood [Yoo02, Sid06].

The tensions between the main actors in the production and distribution of digital content has been very visible in the context of net neutrality [Cro07], and it has been accompanied by many measurement tools that can be used to detect the blocking or throttling behaviour that can emerge as a rational strategy for ISPs [WCL08].

Proposals to shift this equilibrium away from throttling and into more overlay-friendly states include bandwidth auctions [Yan08] and infrastructure caching, if the ISP can attract a sufficiently large number of subscribers [GFGS07].

Regarding the analysis of cooperative equilibrium outcomes between ISPs and managed overlays, work has been done through the use of the Shapley value of coalitional games [MmCL07], and there is a long tradition of work regarding incentives issues in routing. Finally, it is important to note that there is currently a large interest in the standardisation of interface to guide the equilibrium outcomes of these overlay-ISP interactions [PMG09], and an increased interest in the research community in exploring the issues behind the interaction of routing and content distribution [DJ09].

8. PEER-TO-PEER CONTENT DELIVERY

The current infrastructure of the Internet is not suited to simultaneous transmission of live events to millions of people (i.e. broadcasting). The problem is that a dedicated stream of data must be sent to every single user. With millions of potential users, the simultaneous streams of data will easily congest the Internet. In addition this solution requires to the content provide a bandwidth cost that scales linearly with the number of users. For several years, we have been told that the answer to this problem is "multicasting", whereby the data stream is distributed to many local servers that subsequently "re-broadcast" the content to local users. However, most IP routers of the Internet cannot support multicasting–and there seems to be no financial incentive for the ISPs to introduce multicasting.

Also, the use of Audiovisual Media is moving from a collective and passive approach to personal active behaviour, at home and in mobile situations outside the home. At the same time use patterns are shifting towards non-linear usages, moving away from the classic model of linear broadcast TV. The TV set no longer has the monopoly of delivery of audiovisual content; the PC and related media centres, mobile phones, and potentially initiatives from new stakeholders are all becoming increasingly important.

In such heterogeneous environments, efficient content delivery needs optimised unicast, multicast, broadcast, and also support for new mechanisms that have been made possible by the recent advances in P2P grids. This situation has important consequences for the existing business models and institutions, as well as for content production, content distribution, and end user experience on various terminals. This particular holds for stakeholders that propose P2P services based on heterogeneous terminals and networks, together with the demand from users of transparent service continuity. This makes Peer-to-Peer -based technologies that can provide efficient and low-cost delivery of professional and user created content essential for the technologically-competitive future Europe.

A current research challenge is how is to move forward the technical enablers to facilitate new business scenarios for the complete value chain in the content domain from a linear unidirectional push mode to a user centric, time and place independent platform paradigm. A platform approach allows modular development and modular applications, enables knowledge sharing and facilitates technology integration, code and skill re-use. This translates to fast development of new content delivery applications that build value for service and content providers.

The core business of the main European commercial broadcasters is being eroded in today's multi channel era. All players are looking at diversified revenues coming from digital repurposed channels, the Internet, interaction and transaction income.

The TV and advertising landscape is changing dramatically, especially:

- The changing nature and importance of linear broadcast TV vis-à-vis the developing TV spheres of IPTV, broadband TV (classical streaming P-2-p) and Mobile TV (DVBH and 3G);
- The growing ineffectiveness of interruptive awareness building 30 seconds TV spots and the trend towards Internet-like ad production, ad insertion and ad delivery systems and strategies;
- The growing popularity of time-shift TV induced by the massive proliferation of PVRs in the home;
- The rise in broadband deployment worldwide;
- The need for broadcasters to generate more diversified revenues and thus to generate interaction and transaction income;

- The rising popularity of user generated video content and the opportunities offered by Web 2.0-like developments in the mobile and TV domains;
- The changing landscape of selling TV formats worldwide and the production of advertising films induced by the changes described above.

The changing role of the broadcaster is characterised by:

- End users that become more demanding in terms of wanting content delivered at a time and in a place that is convenient to them;
- Content owners from the Hollywood studios to the music labels and football clubs that begin to experiment with new forms of online distribution, experimenting with partners and business models;
- The rise in user generated content including photo sharing, blogs, short films or shared websites and personally created TV like video channels;
- The rise of fourth generation P2P networks that enable video and audio delivery in SDTV and potentially HDTV quality, solving the bandwidth problem of classical AV streaming applications and fully introducing the economics of the Internet into the TV world.

The result will be a move:

- from large audiences that are so appealing to advertisers to fragmented and segmented audiences;
- from programming scheduled by broadcasters to content scheduled by consumers;
- from an environment when hits have the monopoly and make the money to an environment in which niche and older content are being monetised.

The result is an environment, in which telcos offer triple play solutions and move into (IP)TV to effectively stop cable operators from converting accustomed DSL customers into customers enjoying broadband in the home via cable and thus effectively also compete with traditional broadcasters. The result is also an environment in which large DTH broadcasters such as BSkyB buy social communities and ISPs (operators), established online brands such as Yahoo, AOL, and Google move into the video domain and new online video brands suddenly emerge (YouTube) – with the right advertising and promotional content strategy being key to success (see recent purchase of DoubleClick by Google).

Finally, the one-to-many broadcasting models may soon be replaced or at least challenged by one-to-one and P2P unicast models with a return path leading to the end of the unchallenged reign of the classical linear 24/7 TV channel, as we know it, and its partial substitution by huge on-demand servers offering any time, anywhere entertainment content to customers with web like or Google like advertising in-between, and leaving huge live broadcasts (sports and breaking news) as the main sphere of classical free TV.

9. PEER-TO-PEER COLLABORATIVE MEDIA PLATFORMS AND SOCIAL NETWORKS

Building on the P2P research issues identified in the previous section a further set of research challenges are related to P2P multimedia content distribution mechanisms making use of social networks. The challenges include investigation of scalable media coding techniques including both standard and state of the art research methods (wavelets, multiple description coding), combined with new transport and real time streaming protocols deployed over peer to peer architectures.

Additional novelty is introduced by investigating mechanisms for the discovery of media resources and the selection of peer nodes based on information collected, maintained and provided by social networks such as Facebook and twitter. Such data is created by user communities over the Internet, and the incorporation of extra information in streams such as social graphs, text and metadata generated according to end user profiles can be supported and exploited by P2P media distribution platforms.

An associated research challenge is concerned with the introduction of social networking information to media production and media distribution schemes, making them content aware, and adapted to the personal needs and requirements of end users.

Information from social networks can be used to enhance content search and locality discovery using data and derived information concerning common interest amongst neighbours in social graphs to find peers who may be participating in related live streams or may have stored/cached content of interest to a user. Similar techniques could be used to build personalised and tailored recommender systems that are constructed and refined according to the likes, dislikes and developing trends within a user's social neighbourhood.

Many media applications today are embedding social network features for comment, chat, and creating shared media experiences between friends or other groups of users with a common interest. There are several related research challenges: Firstly, how to exploit knowledge of social network relationships to predict how media consumption may be correlated between groups of users. This information can be used to dimension media servers and network resources to avoid congestion and improve QoE. Secondly, a major hindrance to exploitation of social network data is the fragmentation of the population of social network users into numerous proprietary and closed social networks. This issue is compounded by the fact that each new game or media application tends to build its own social network around it rather then building upon the rich data available about existing social relationships. Also applications are often restricted to execute within the confines of specific social network platform. A major research challenge, therefore, that would benefit the exploitation of social network graphs for future media networking is in finding solutions to open up social network platforms to allow cross-platform information exchange and usage. Of course, reliable mechanisms to preserve privacy are an essential prerequisite.

A final research challenge is to integrate these areas of research into a common target: the creation of a novel, state of the art platform for real time media distribution over the Internet that takes advantage of user profiles and information gleaned from social networks for media creation and distribution, with the purpose of enabling high-quality, content-aware, personalised, network-adaptable content streaming to end users.

10. CROSS-LAYER OPTIMISATION OF MULTIMEDIA OVER WIRELESS IP LINKS

IPTV provides telecommunication operators the opportunity to better serve the video market and better compete with the industry transition to HD (High Definition) and unicast video (e.g. Video on Demand), both creating challenges to the existing CATV (Cable TV) and SAT (Satellite) providers. One of the main challenges in IPTV is the line rate of the access lines, which is predominantly ADSL.

While preserving existing IPTV networks models and leaving existing infrastructures intact, a research challenge is how to optimise the bandwidth of IPTV for the delivery of multiple HD streams over a single ADSL line, and by that enabling multiple HD channels per household.

Currently, the DSL access network creates a bandwidth bottleneck which does not allow for adequate provision of personalised High Definition video content to the subscriber.

Algorithms are required for data dropping that are based on the nature of the video content "Content aware data dropping algorithm". The full bandwidth IPTV video streams will flow through the transport access to the algorithm that is located at the IPTV transport network to decide which data to drop and which to let flow to the STB and TV. This decision will control the bandwidth of the video stream that will flow through the "Bandwidth Bottleneck". The narrower the bottleneck is more data shall be dropped. However, as more data is dropped, there is an impact on the quality of experience (QoE) of the viewer. Therefore, it is mandatory that the data dropping will be done in a manner that will cause minimal effect on the QoE according to measurable metrics.

Video transmission over communication networks results in packet loss due to limited resources, limited bandwidth and bit errors, which results in severe degradation in perceived video quality and annoying watching experience. Most of the existing video codecs employs various VLC and predictive coding techniques to achieve a high compression efficiency, which makes the compressed bit-stream vulnerable to transmission errors. A single bit error in the video stream may render the correct decoding of future codewords impossible. Furthermore, owing to predictive coding, the effects of packet loss are likely to be propagated to the neighbouring video blocks both temporally as well as spatially. Additionally, the effected surface region due to packet losses tends to increase in the succeeding frames, because if those corrupted pixels were used as reference for prediction of other pixels, then the error will also propagate to predicted pixels. Due to the use of motion compensation techniques, the predicted pixels may not be located at the same position as the reference pixels, which lead to spread in corruption region in succeeding frames. It has been observed, that due to these dependencies in the encoded video bit stream only 3% packet loss incurred in the transmission of MPEG video affects 30% of the frames.

More recently, several factors have been identified that may have significant impact on the perceptual quality of transmitted videos over IP networks, namely:

- the number of erroneous frames caused by a packet loss and subsequent error propagation (i.e. the error length);
- the severity of a loss (i.e. the difficulty in concealing a lost frame), which can be measured by the peak signal-to-noise ratio (PSNR) drop after a loss;
- the loss position, which is measured by the time to the end of the sequence (also known as the "forgiveness effect");
- the number of losses in a sequence; and
- the loss pattern (e.g. clustered or spread).

In addition to losses, delays and available bandwidth may also deeply impact the received video quality.

Innovation in the area of sophisticated multimedia source coding schemes aiming to satisfy design criteria and trade-offs in terms of source representation quality, bitrate, delay, encoding/decoding complexity, etc. is a key issue in the modern world where users demand for content “anywhere and anytime”. Today’s approach, relying on traditional separation approaches and focussing on services delivered over homogeneous networks, does not allow to meet the on-going demands to maintain the required Quality of Service (QoS) for each of the users, who have different needs and requirements.

It is necessary to consider the whole transmission chain and all protocol layers to improve the quality of video data transmissions over wireless and mobile IP links. The required areas of improvements can be classified in the following in four categories: i) applications for content delivery, considering limitations of existing solutions at the application layer; ii) networking, considering existing solutions at session, transport and network layer; iii) radio access, where areas of improvement are identified at the data link and physical layer; iv) cross-layer solutions.

As video is being streaming over wireless channels, whose conditions in terms maximum bit rate and communication quality may deeply vary due, e.g., to limited bandwidth availability and high number of clients or receivers’ mobility, the source coding scheme needs to rely on a rate control algorithm more reactive (in the sense of quickly evolving to meet the target bitrate) than the ones traditionally used in long term rate control for video compression. As a matter of fact, the channel conditions may vary too quickly when compared to the convergence time of the rate control algorithm, thus introducing great misbalances and several inadequately settings (as it could happen by using the long term adaptation rate control algorithms proposed in the H.264/AVC JM software). Moreover, in order to improve the quality of a video received over imperfect channels, protection schemes introducing redundancy and thus consuming bandwidth have to be taken into account when selecting coding parameters so to not exceed the maximum available bandwidth.

In addition to the above further research challenges include: firstly, wireless channel status, bandwidth availability and received video quality have to be known at the encoder side: an efficient tool for end-to-end delivery of feedbacks is not available today. Second, feedback information from the multiple receivers has to be aggregated and managed by the transmitter and optimisation criteria have to be adapted to user diversity. Finally, the IP networking to multimedia transmissions has to be optimised for video streaming and point-to-multipoint transmissions.

To achieve this goal, a scheme is required including all elements of major importance in a point to multi-point video streaming chain, in particular video coding, networking modules, MAC layer and physical layer, efficiently communicating together.

The research challenge can be summarised as the maximisation the video quality perceived by the users of the system and the impact of any proposed solutions should be evaluated by considering different quality metrics, based on subjective tests and objective measurements, and different video sources:

- real-time data received from a video camera and a microphone: to evaluate the delay and the video quality perceived by the receivers,
- raw data stored in the server: to show the ability to encode by following user feedbacks and reacting to changing conditions,
- adapt pre-coded streams: to show the ability of adaptation at the server and at the BSs.

Efficient solutions will also minimise the overhead of the communication, which is an inevitable drawback in case of additional protection schemes and end-to-end signalling. The effectiveness of the designed signalling framework and of header compression mechanisms in terms of overhead minimisation will be proved by several measurements.

11. MERGING MULTIPLAYER GAMES AND INTERACTIVE VIDEO

Massively multiplayer online games (MMOGs) allow a large number of online users to inhabit the same virtual world and interact with each other in a variety of collaborative and competing scenarios. Gamers within an MMOG typically become members of active communities with mutual interests, shared adventures, and common objectives. A significant research challenge is in enhancing collaborative activities between MMOG players. This requires new tools for the generation, distribution and insertion of user-generated content (UGC) into existing MMOGs. This UGC may include items (e.g., textures, 3D objects) to be added to the game, live video captured from the game screen and streamed to other players, and videos showing walkthroughs or game tutorials to be watched on demand.

Two proposed technologies to be studied are the in-game graphical insertion technology (IGIT) and a peer-to-peer (P2P) system for the distribution of live video. IGIT is an innovative technology of replacing or inserting content into a game in real time without the need to change the game code in the client or server. For example, billboards can be inserted, tattoos can be added to in-game characters, an area on the screen can be assigned to display user information, and any type of window (Web browser, instant messaging, audio/video chat, etc.) can be inserted floating on or outside the game area. The technology can be implemented on multiple games, making it possible to create a community that is not limited to a specific game or publisher. Enabling thousands of users to communicate UGC represents a significant challenge to networks already occupied by the MMOG client-server data. New techniques are needed for UGC distribution that are "friendly" (supportive and not disruptive) to the MMOG client-server traffic. The key innovation is a P2P system that will allow MMOG gamers to stream live video of the game without interrupting the MMOG data flow and the need to upload the video data to a central server. The outstanding challenges in this area can therefore be summarised as:

- New techniques of replacing an existing texture with a new one, including in-game spot identification methods, in-game asset replacement and insertion.
- P2P video delivery for MMOGs. This requires new application-layer mechanisms to allow players to diffuse captured game video over a P2P overlay without harming the game experience (MMOG-friendly P2P).

Meta data extraction from a game and UGC presentation inside the game: Currently, there are no generic technologies for the extraction of meta-data (e.g., score, level, badges) from software, and in particular games. Most existing solutions rely on source code integration with the game, which requires predefining what exactly should be tracked and how it should be tracked. Current solutions for presenting information back in the application/game assume that the information type is known in advance (in most cases – textual information), the place in which the information is to be displayed is defined in advance, and the type of interaction with the information is defined in advance.

In-game user communication: Commercial solutions such as Xfire provide a crude implementation of in-game communication tools (typically instant messaging), which are overlaid on top of a limited number of existing games.

Personal authoring tools: There are currently no personal authoring tools for in-game solutions. There are some tools that allow screen captures or video captures from within video games, but these result in flat files holding the information, without any real context.

In- and around game advertising: Existing methods are expensive because they need to be integrated into the source code for the MMOG.

P2P streaming for MMOGs: In traditional P2P live streaming systems, one source provides live video to a large number of interested peers, which are organised into a P2P overlay. However, there are new applications where many users can be a source of a different video, and any user can potentially be interested in all available videos. For example, MMOG players may generate video data, stream video live to other users, and watch live video generated by other users. The following three scenarios illustrate possible applications.

- Scenario 1: A player broadcasts live screen-captured video of its game to any other player.
- Scenario 2: A player streams live screen-captured video of its game to a restricted group (guild).
- Scenario 3: A player streams animated virtual 3D objects. The "clients" are players whose virtual position is close to the virtual position of the object.

In these scenarios, a peer cannot participate in all P2P overlays because of constraints on its physical resources, including upload and download bandwidth. Moreover, a player has additional connections with the MMOG game server, which must be given highest priority. The challenge for a user is to adequately allocate its physical resources. One of the main functionalities of TCP is to ensure a fair sharing of the bandwidth among all connections. However, this model for network resource sharing is not efficient when one application uses multiple connections as in P2P streaming. Some technologies such as pcube have been implemented, but they act on the backbone although the congestion occurs in the access network. Integrated services (IntServ) and differentiated services (DiffServ) can also be used. Unfortunately they require the cooperation of all routers. Flow-aware networking (FAN) [OuRo05] provides per-flow differentiation to active flows through implicit admission control and per-flow scheduling. FAN requires the association of the end user and its access router, but the replacement and update of current access routers is costly. TCP Nice, TCP-LP and 4CP make the background transfer more sensitive to network congestion. Though they do not need any support from routers, they require modifying the TCP protocol, which is not realistic.

Therefore developing a practical P2P system that can automatically manage the set of P2P overlays is an important challenge. For Scenario 3, one possible solution is to use the concept of Area of Interest (AoI) [JiHuHu09]. An AoI is defined as the part of the virtual world around a user that generates content. When a peer is within the AoI of a user that is generating high-quality content, or when it belongs to many AoIs, it may experience congestion. The challenge is to design a mechanism for determining the best size of these AoIs. The management of the AoI must take into account the popularity of the virtual place and the capacity of the devices of the players that are located there. Such a management has been shown to be hard in wireless sensor networks [HaLeSi09], but some heuristics can perform well. For Scenario 1 and 2, the concept of AoI needs to be extended beyond the notion of virtual position.

A related challenge is to maximise the number of peers receiving content in the general context where not all peers can be reached because too few resources are available. If we assume tree overlays and consider only one video stream, the problem is to build a tree that spans the maximum number of peers with the constraint that every peer can only serve a limited number of other peers. In the context of many concurrent video streams, the problem becomes even harder with a constrained forest. The building of degree-constrained trees is an NP-hard problem [KoRa03]. Related works deal with the computation of the streaming capacity of a P2P overlay [LiChSe10].

Most P2P live video systems rely on TCP as the transport protocol. However TCP may introduce significant delays. An alternative to TCP is to use UDP and apply application-layer error control. This includes UDP without error control, UDP with FEC, ARQ, and Multiple Description Coding (MDC). However, in most proposed schemes, the error resilience technique lacks network adaptivity. One powerful way to offer network adaptivity is to protect the video data with a rateless code and keep on sending the channel symbols until an acknowledgment is received or a timeout occurs. But as the acknowledgment needs time to reach the sender, the sender may transmit redundant channel symbols. The problem of how to determine packet scheduling algorithms that minimise this overhead has been addressed for a client-server model [AhHaAl10]. A major challenge is to devise a solution for the more challenging context of a P2P network.

Due to the interdependencies of encoded video frames, not all source symbols are equal in importance. Therefore, more important symbols should have stronger protection than less important ones. Despite some recent progress, the topic of unequal error protection with rateless codes remains an open challenge.

In joint source channel coding, an optimal tradeoff between the source coding rate and the channel coding rate is determined. Joint source-channel coding has been studied extensively for traditional erasure codes. However, there is no work on joint source channel coding with rateless codes.

One limitation of the UDP protocol is the lack of a congestion control mechanism. Congestion control with UDP can be realised with the Datagram Congestion Control Protocol (DCCP) [RFC4340]. DCCP uses an Explicit Congestion Notification (ECN) bit, which is set on by a congested router. A major problem of DCCP is that all routers on the path from the sender to the receiver should be ECN-aware.

12. REFERENCES

[4WARD] http://www.4ward-project.eu/

[ABEA+06] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter. A survey on networking games in telecommunications. Comput. Oper. Res., 33(2):286–311, 2006.

[AGA07] Agarwal, S. "A case study of large scale P2P video multicast", IP Multimedia Subsystem Architecture and Applications, 2007 International Conference on, pp. 1—5 2007.

[AhHaAl10] S. Ahmad, R. Hamzaoui, and M. Al-Akaidi, "Adaptive unicast video streaming with rateless codes and feedback," IEEE Transactions on Circuits and Systems for Video Technology, vol. 20, pp. 275-285, Feb. 2010.

[ALTO] https://datatracker.ietf.org/wg/alto/charter/

[AND04] C. Anderson "The long tail", Wired magazine 12.10, October 2004.

[BEN08] R.A. Bentley, P. Ormerod, M.E. Madsen, "Shelf space strategy in long-tail markets", http://arxiv.org/abs/0808.1655v2, August 2008.

[BKR06] Jesse C. Bockstedt, Robert J. Kauffman, and Frederick J. Riggins. The move to artist-led on-line music distribution: A theory-based assessment and prospects for structural changes in the digital music market. Int. J. Electron. Commerce, 10(3):7–38, 2006.

[CAR03] A.Carzaniga, A. L. Wolf, "Forwarding in a Content-Based Network", http://www.inf.usi.ch/carzaniga/papers/cw_sigcomm03.pdf

[CDN] http://en.wikipedia.org/wiki/Content_delivery_network

[CJMW05] Hyunseok Chang, Sugih Jamin, Z. Morley Mao, and Walter Willinger. An empirical approach to modeling inter-AS traffic matrices. In IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement, pages 12–12, Berkeley, CA, USA, 2005. USENIX Association.

[COMET] http://www.comet-project.org/

[Cro07] Jon Crowcroft. Net neutrality: the technical side of the debate: a white paper. SIGCOMM Comput. Commun. Rev., 37(1):49–56, 2007.

[CVI10] Cisco Visual Networking Index: Forecast and Methodology, 2009-2014. June 2, 2010. Cisco. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html

[DJ09] D. DiPalantino and R. Johari. Traffic engineering versus content distribution: A game theoretic perspective. In The 28th IEEE International Conference on Computer Communications (INFOCOM 2009), Rio de Janeiro, Brazil, Brazil, 4 2009.

[ENVI] http://www.envision-project.org/index.html#

[GFGS07] M. Garetto, D. R. Figueiredo, R. Gaeta, and M. Sereno. A modeling framework to understand the tussle between ISPs and peer-to-peer file-sharing users. Perform. Eval., 64(9-12):819–837, 2007.

[HaLeSi09 B. Han, J. Leblet, and G. Simon, ``Query range problem in wireless sensor networks," IEEE Comm. Letters, vol. 13, no. 1, pp. 55-57, 2009.

[HEL09] H. Hellwagner et al., "Efficient In-Network Adaptation of Encrypted H.264/SVC Content", Image Communication, vol. 24, 2009, p. 740-758.

[IMS] http://en.wikipedia.org/wiki/IP_Multimedia_Subsystem

[IPLAYER] http://www.bbc.co.uk/iPlayer

[JAC09] Jacobson V., Smetters D., Thornton J., Plass M., Briggs N., Braynard R (2009), "Networking Named Content", Proceeding of ACM CoNEXT 2009. Rome, Italy, December 2009

[JiHuHu09] J.-R. Jiang, Y.-L. Huang, and S.-Y. Hu, ``Scalable AOI-cast for peer-to-peer networked virtual environments," Journal of Internet Technology, vol. 10, no. 2, pp. 119-126, 2009.

[JZSRC08] Wenjie Jiang, Rui Zhang-Shen, Jennifer Rexford, and Mung Chiang. Cooperative content distribution and traffic engineering. In NetEcon '08: Proceedings of the 3rd international workshop on Economics of networked systems, pages 7–12. ACM, 2008.

[KMK+09] N. Kamiyama, T. Mori, R. Kawahara, S. Harada, and H. Hasegawa. ISP-Operated CDN. In 12th IEEE Global Internet Symposium (INFOCOM workshops), Rio de Janeiro, Brazil, April 2009.

[KoRa03] J. Könemann and R. Ravi, ``Primal-dual meets local search: approximating MST's with nonuniform degree bounds," in: Proc. 35th Annual ACM Symposium on Theory of Computing, San Diego, CA, pp. 389-395, 2003.

[LiChSe10] S. Liu, M. Chen, S. Sengupta, M. Chiang, J. Li and P.A. Chou, "P2P Streaming capacity under node degree bound", in Proc. of IEEE International Conference on Distributed Computing Systems (ICDCS), 2010

[MmCL+07] Richard T. B. Ma, Dah ming Chiu, John C. S. Lui, Vishal Misra, and Dan Rubenstein. Internet economics: the use of Shapley value for ISP settlement. In CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference, pages 1–12, New York, NY, USA, 2007. ACM.

[NAPA] http://www.napa-wine.eu

[Nef05] Gina Neff. The changing place of cultural production: The location of social networks in a digital media industry. Annals of the American Academy of Political and Social Science, 597:134–152, 2005.

[NET] http://en.wikipedia.org/wiki/Network_neutrality

[OFCOM] http://www.telecoms.com/itmgcontent/tcoms/features/articles/20017562353.html

[OuRo05] S. Oueslati and J. Roberts, "A new direction for quality of service: flow-aware networking," Next Generation Internet Networks, pp. 226-232, April 2005.

[PARC] http://www.parc.com/work/focus-area/networking/

[PMG09] Jon Peterson, Enrico Marocco, and Vijay Gurbani. Application-Layer Traffic Optimization (ALTO) working group. http://www.ietf.org/dyn/wg/charter/alto-charter.html, 2009.

[PSIRP] http://www.psirp.org/

[RFC4340] http://tools.ietf.org/html/rfc4340.

[SAIL] FP7 project SAIL: Scalable & Adaptive Internet soLutions, Project Reference: 257448, http://tinyurl.com/2ee34ko

[Sid06] J. Gregory Sidak. A Consumer-Welfare approach to Network Neutrality regulation of the Internet. Journal of Competition Law and Economics, 2(3):349–474, 2006.

[SUB01] Barani Subbiah Zartash Afzal Uzmi "Content Aware Networking in the Internet: Issues and Challenges", suraj.lums.edu.pk/~zartash/publications/2001-06-ICC-Content.pdf

[TJE04] T. Tjelta, L.E. Braten, C. Howson, T. Montalant, "Developments in Broadband Wireless Access", Broadband Europe, Brugge (Belgium), 8-10 December 2004.

[WAR08] G. Warrell, C. Lagerling, A. Dafferner, "Live Internet TV – A reality as p2p technologies are being deployed", GPBullhound Independent Technology Research, March 2008.

[WCL08] Jessie Hui Wang, Dah Ming Chiu, and John C.S. Lui. A game-theoretic analysis of the implications of overlay network traffic on ISP peering. Computer Networks, 52(15):2961–2974, 2008. Complex Computer and Communication Networks.

[WEB08] R. Webb, "Online shopping and the Harry Potter effect", http://www.newscientist.com/article/mg20026873.300-online-shopping-and-the-harry-potter-effect.html, December 2008.

[YAA08] Jiang Yang, Lada A. Adamic, and Mark S. Ackerman. Crowdsourcing and knowledge sharing: strategic user behaviour on taskcn. In EC '08: Proceedings of the 9th ACM conference on Electronic commerce, pages 246–255, New York, NY, USA, 2008. ACM.

[Yan08] Xiaowei Yang. Auction, but don't block. In NetEcon '08: Proceedings of the 3rd international workshop on Economics of networked systems, pages 19–24, New York, NY, USA, 2008. ACM.

[Yoo02] Christopher S. Yoo. Vertical Integration and Media Regulation in the New Economy. Vanderbilt Law and Economics Research Paper No. 02-01; Vanderbilt Public Law Research Paper No. 02-01, 2002.

13. CONTRIBUTING PROJECTS

ALICANTE
http://www.ict-alicante.eu/

nextMEDIA
http://www.fi-nextmedia.eu/

CNG
http://www.cng-project.eu/

OCEAN
http://www.ict-ocean.eu/

ALICANTE
http://www.ict-alicante.eu/

nextMEDIA
http://www.fi-nextmedia.eu/

COMET
http://www.comet-project.org/

OPTIMIX
http://www.ict-optimix.eu/

ENVISION
http://www.envision-project.org/

P2P-NEXT
http://www.p2p-next.org/

FutureNEM
http://www.nem-initiative.org/

SARACEN
http://www.saracen-p2p.eu/

NAPA-WINE
http://www.napa-wine.eu/

FMN-cluster
http://jefferson.ee.ucl.ac.uk/fmn-wiki/

European Commission

Luxembourg: Publications Office of the European Union

2010 — 52 pp. — 21 x 29,7 cm

ISBN 978-92-79-17393-6
doi:10.2759/37178

KK-32-10-490-EN-C

Future Media Internet

Research Challenges and the Road Ahead

DG Information Society and Media of the European Commission.

http://cordis.europa.eu/fp7/ict/netmedia/home_en.html

European Commission
Information Society and Media

Publications Office

doi:10.2759/37178

ISBN 978-92-79-17393-6