enhanced COntent distribUtion with Social INformation

www.ict-ecousin.eu

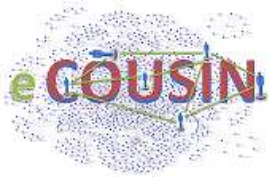

Deliverable D7.3

Final Dissemination, Exploitation and Standardization, Report

Public deliverable, Version 1.0, 30 April 2015

Authors

ORANGE Yannick Le Louédec, Bertrand Mathieu, Patrick Truong

AL-BELL Danny De Vleeschauer, Chris Hawinkel

IMDEA Nicola Bui

TSP Ángel Cuevas Rumín

ALUD Klaus Satzke

TUD David Hausheer

TI Fabio Mondin(Editor), Claudio Venezia

UC3M Rubén Cuevas

Reviewers Nicola Bui

Abstract

Deliverable D7.3 is the final dissemination, exploitation and standardization report of the eCOUSIN project. Deliverable D7.3 presents the results of the dissemination and standardization activities of the consortium partners, as well as an updated version of their exploitation plans.

Project funded by the European Union under the
Information and Communication Technologies FP7 Cooperation Programme
Grant Agreement number 318398

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The objectives of eCOUSIN WP7 have been to ensure the dissemination of the project's results, to coordinate its visibility at an international level through publications, contributions to standards, participation to or arrangement of specific events, and to manage the exploitation plan of the involved partners.

The present Deliverable D7.3 is the final dissemination, exploitation and standardization report of the eCOUSIN project.

Deliverable D7.3 presents the project dissemination results, the actions taken and the channels reached during the whole project life, with as goal to focus on opportunities to collaborate with and impact on the scientific communities.

The consortium promoted the project views and results by means of a considerable dissemination effort. The total number of dissemination actions is close to 100 just taking into account papers, publications, workshops and presentation related to the scientific community. Moreover, a set of papers was promoted over traditional and specific media. More than 10 specific workshops were organized to communicate on the project achievements and to exchange results and ideas with third parties. A final dissemination activity involving all the use-cases is planned for EUCNC 2015 in Paris on June 2015, just a few days before the final review. All the use cases will be showcased at this event by means of posters and live demos.

This report describes also the main contributions of the consortium to standard bodies. The standardisation effort touched three standard bodies mainly: W3C, IRTF ICNRG and also, even if less directly, OMA. In W3C the use case was brought in order to have a direct influence from the project to the standard. And a tangible contribution was produced for IRTF ICN RG.

Finally Deliverable D7.3 includes an updated version of the exploitation plan of all the partners, taking into account the final project outcomes. Each partner has a credible exploitation plan for the project outcomes, with real exploitation perspectives at short time, such as on the personal content sharing cloud use case. Besides Deliverable D2.3 provide complementary insights on exploitation perspectives, by presenting business plans for the selected use cases and the expected benefits for the main concerned stakeholders.

1. INTRODUCTION

Content Distribution Services are booming and they will be responsible for the majority of future Internet traffic. In parallel, Online Social Networks (OSNs) have become today's most popular Internet applications. The widespread adoption of OSNs has drastically changed the way content is consumed in the Internet, as content consumption is nowadays highly impacted by the information shared by users through OSNs and the popularity of a given content is most often dictated by its "social" success.

With such a "social-content revolution", operators need to evolve and optimize their network to avoid being overwhelmed by the ever growing traffic volumes resulting from this paradigm change.

In this context the goal of eCOUSIN has been to design a novel social-aware network architecture with built-in content dissemination functionalities that exploits the social-content interdependencies to improve its efficiency.

This goal translates into four specific objectives:

- The implementation of high performance distributed tools for collecting necessary data to study and model the social-content interdependencies;
- The improvement of the scalability of network infrastructures when handling content by exploiting social information;
- The design of an on-net operational framework that tightly integrates network functionalities and content-related service functionalities;
- The design of algorithms that exploit social information for placing and delivering contents in an optimized manner with a special focus on mobile environments.

eCOUSIN aims to provide a clear added-value in use cases involving social-content interdependencies such as personal content sharing clouds, social-assisted time-unconstrained content delivery, Content-Centric Networking for social-driven content delivery, and enhanced content placement in distribution networks using users' social and coarse-grain location information.

The present Deliverable D7.3 is the final dissemination, exploitation and standardization report of the eCOUSIN project. The dissemination activities are detailed in Section 2. Section 3 provides an updated version of the partners' exploitation plan. Section 3 presents the standardisation activities achieved by the partners. And Section 4 concludes this deliverable.

2. DISSEMINATION ACTIVITIES

Extensive large-spectrum dissemination work, devised to achieve awareness for the project activities and results, has been performed for all the project duration, including the presence in the scientific community through papers, talks, demo at relevant international conferences/events, Production of research material, articles, collaborations etc.

Table 1 summarizes the main dissemination areas taken into account by the project. For each area the main results are presented in the later paragraphs.

Activity / Channel	***Remarks***
eCOUSIN Web Presence	The website is a key channel for dissemination. Its URL can be simply added as a signature on any support: project document, mail, brochures, professional card, etc. Internet users can get a general overview of the project in 2 clicks while interested people can retrieve all its public deliverables easily.
Flyers and Brochures	Flyers and brochures are made available in printed form at conferences and events. They can also be retrieved on eCOUSIN website.
Project Public Deliverables	The project public deliverables are made available on eCOUSIN website, providing full details on most of eCOUSIN's studies and outcomes.
Publications in magazines and journals	The eCOUSIN project intends to publish most of its scientific results in specialized magazines and journals.
Participation to conferences and events: presentations and posters	National and international conferences organised by institutions, universities and research organisations are important opportunities to share project results with other experts in the field. Such participation consists in delivering presentations and/or in having a stand with posters. Presentations allow reaching a large audience whereas posters can be used as a support for more individualized and in-depth exchanges.
Workshops	eCOUSIN arranged workshops to collect requirements, inputs and feedbacks from market stakeholders and the research community, and to communicate on the project findings, notably with the advisory board members. They therefore also help for dissemination.

Table 1. Summary of planned dissemination channels

2.1 eCOUSIN Web presence

The website is the first step in helping dissemination. The project coordinator registered the domain name "*ict-ecousin.eu"* and created the project website [2].

Figure 1: eCOUSIN Web presence

The home page of the web site provides an overview of the project, based on the abstract of the project's description of work:

"Content Distribution Services are booming and they will be responsible for the majority of future Internet traffic. In parallel, Online Social Networks (OSNs) have become today's most popular Internet application. The widespread adoption of OSNs has drastically changed the way content is consumed in the Internet, as content consumption is nowadays highly impacted by the information shared by users through OSNs and the popularity of a given content is most often dictated by its "social" success.

With such a "social-content revolution", operators need to evolve and optimize their network to avoid being overwhelmed by the ever growing traffic volumes resulting from this paradigm change. To this aim, the goal of eCOUSIN is to design a novel social-aware network architecture with built-in content dissemination functionalities that exploits the social-content interdependencies to improve its efficiency. This goal translates into the following specific objectives: (i) the implementation of high performance distributed tools for collecting necessary data to study and model the social-content interdependencies; (ii) the improvement of the scalability of network infrastructures when handling content by exploiting social information; (iii) the design of an on-net operational framework that tightly integrates network functionalities and content-related service functionalities; and (iv) the design of algorithms that exploit social information for placing and delivering contents in an optimized manner with a special focus on mobile environments. We envision that all these will be fundamental components of the future Internet architecture".

The public part of the website is available to any web user to fulfil the following three roles:

1. It delivers the general information about the project, including list of participants, objectives, and work plan;
2. It is a complete repository of all information delivered by the project (public deliverables, scientific publications, etc.);
3. It acts as a unified point of contact for the project and aggregates interest in eCOUSIN.

The new version of the website includes a twitter widget with eCOUSIN's twitter account (@ICT_eCOUSIN).

The website also includes a private area, which contains electronic versions of released private deliverables and quarterly reports in order to facilitate the communication inside the project and the follow-up of the project by the Commission services. The project coordinator is the sole responsible for giving access to this private site.

2.2 Flyers and brochures

The eCOUSIN project has produced flyers and brochures for the purpose of communicating on its objectives and outcomes. Flyers and brochures are made available in printed form at conferences and events, as well as on eCOUSIN website. An overview of the project has also been already provided for the European Commission brochure on Future Internet Cluster [1].

2.3 Public Project Deliverables

The following table lists the public deliverables of the eCOUSIN project.

Deliverable	***Deliverable name***	***Delivery date***
D1.1	eCOUSIN Project WebSite (public area and restrictive area)	M3
D2.1	Initial report on Use Cases and Requirements	M3
D2.2	Initial System Architecture Specification	M9
D2.3	Business and Economic Analysis	M15
D2.4	Final Report on Use Cases, Requirements, System Architecture Specification, Privacy and Regulations	M21
D3.1	Measurement, Modelling, and Prediction of Social-Content Interdependencies (First Version)	M12
D3.2	Initial Release of Measurement and Prediction Software	M18
D3.3	Measurement, Modelling, and Prediction of Social-Content Interdependencies (Final Version)	M26
D4.1	Preliminary Report on the Design of Technical Solutions on Content Placement and Delivery	M12

D4.2	Final Report and Initial Software Release of the Design Extensions and Preliminary Implementation of the Technical Solutions on Content Placement and Delivery	M18
D4.3	Final Implementation and Software Release of the Technical Solutions on Content Placement and Delivery	M26
D5.1	Requirements for Social-Enhanced Content Centric and Mobile Network Infrastructures	M12
D5.2	Modules and Interfaces for Social-Enhanced Content Centric and Mobile Network Infrastructures (Preliminary)	M18
D5.3	Modules and Interfaces for Social-Enhanced Content Centric and Mobile Network Infrastructures (Final Release)	M26
D6.1	Preliminary Plan for System Integration and Assessment (without commitment about the final implementation)	M15
D6.2	Final Plan for System Integration and Assessment	M20
D6.4	Final System Software release	M29
D6.5	Assessment Report	M30
D7.1	Project Presentation Preparation for Publishing at the Website of the Project	M1
D7.2	Intermediate Dissemination, Exploitation and Standardization, Report	M12
D7.3	Final Dissemination and Exploitation Plan Report	M30

Table 2. Public Project Deliverables

2.4 Dissemination of the Scientific and Technical achievements

As the problems arising within eCOUSIN are intellectually challenging, it is desirable to report on both the experience learnt within the project and the solutions proposed to different problems within the scientific community. This is done via the traditional approaches used in this community, namely generating publications as well as organizing scientific events on topics related to eCOUSIN.

eCOUSIN publishes its results in leading, high-level, international conference proceedings and journals. It should be noted that the experience of key personnel (e.g., TPC members of international conferences such as Infocom, Globecom or editors of IEEE journals) in the consortium provides the right expertise to target high level quality output. Additionally, the achieved results will be the basis of books and/or book chapters. Timely publications will improve the project's visibility and enable exploitation of its results.

Additionally, special issues of international journals with eCOUSIN-related research publications will further increase the visibility of eCOUSIN worldwide as well as provide a suitable forum to discuss eCOUSIN ideas with the research community.

2.5 Events and external liaisons

Events and external liaisons include large trade shows; specific industry events focused on Content Distribution networks, EC focused events, academic conferences and larger academic fora. Attendance at these events can take a number of different forms, ranging from a basic presence to interact with people, to making presentations on the project, to engaging in visionary panel discussions to highlight the potential of the technology, to demonstrating the technology in action. It is strongly in the interest of the member organizations to maximize their visibility in these different fora and hence to maximize the visibility of the project.

As well as the above attendance at international events, the consortium aimed at organising specific events, focused on the project's subject in particular and at which it is able to showcase its solutions and technology. The consortium involves key personnel in high visible panels and organized two workshops during its lifetime. Invitations to the most active researchers in the area were issued, including the Advisory Board members.

The project aims also at being proactive with regards to establishing linkages with parties working on related activities both inside and outside Europe. For the former, the objective is to ensure that the work of eCOUSIN complements other work taking place in a European context; for the latter, the objective is to ensure that the work of eCOUSIN remains on the forefront of technology development in this domain, thus strengthening Europe's future potential.

Concertation meetings, where related projects meet to identify synergies, were attended by the Project Coordinator and when necessary technical experts. These are important for dissemination of results, coordination of work, and maximization of project impact.

Furthermore, as it is important for the consortium to showcase its technology with maximum impact, the consortium aims also at investigating opportunities to deploy its solutions for coverage during events such the EuCNC conference planned on June 2015.

2.6 Dissemination results

This section presents the main dissemination results:

- Subsection 2.6.1 contains the list and abstract of papers/publications related to the project already accepted/presented into conferences
- Subsection 2.6.2 contains the list and abstracts of papers/publications which have been submitted but still under review.
- Subsection 2.6.3 contains the list of participation to conferences/presentations
- Subsection 2.6.4 contains the collaboration actions, including the organization of workshops, the participation to EU collaboration activities, and other dissemination actions

2.6.1 Accepted Papers and Publications

Table 3 shows the list of accepted papers and publications. Complementary information, including abstract and detailed list of authors, are provided after the table.

Ref	*Conference / Journal*	*Title*	*Partners*	*Publishing Date*
(01)	Technical Report, University of Cambridge, 2012	Mitigating I/O latency in SSD-based Graph Traversal	UCAM	November 2012
(02)	Journal of Network and Computer Applications, Elsevier, 2012	Evaluating Opportunistic Networks in Disaster Scenarios	UCAM	November 2012
(03)	World Wide Web Conference (WWW' 13)	Google+ or Google-?: Dissecting the evolution of the new OSN in its first year	UC3M, TSP	May 2013
(04)	IEEE International Conference on Computer Communications(IEEE Infocom 2013)	DataSpotting: Exploiting Naturally Clustered Mobile Devices to Offload Cellular Traffic	ALUD	June 2013
(05)	IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM 2013)	Large scale analysis of HTTP adaptive streaming in mobile networks	ORANGE	June 2013
(06)	7th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2013)	Towards a Mobility-Supporting Incentive Scheme for Peer-to-Peer Content Distribution	TUD	June 2013
(07)	ACM Special Interest Group on Management of Data (ACM SIGMOD 2013)	Scale-up Graph Processing: A Storage-centric View	UCAM	June 2013
(08)	Springer Journal of Social Network Analysis and Mining, July, ISSN:	What's in Twitter I Know What Parties are Popular and Who You are	UCAM	July 2013

	1869-5450, 2013	Supporting Now!		
(09)	IEEE COMSOC MMTC E-Letter, Vol.8, No.4	Addressing the Wireless Content Challenge	ALUD	July 2013
(10)	IEEE 21ST International Symposium on Modeling, Analysis And Simulation of Computer And Telecommunications (MASCOTS 2013)	HTTP adaptive streaming in mobile networks: characteristics and caching opportunities	ORANGE	August 2013
(11)	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013)	Centrality and Mode Detection in Dynamic Contact Graphs; a Joint Diagonalisation Approach	UCAM	August 2013
(12)	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM' 13)	Analysis of publicly disclosed information in Facebook profiles	TSP	August2013
(13)	IEEE International Conference on Peer-to-Peer Computing (P2P 2013)	Investigating the Reaction of BitTorrent Content Publishers to Antipiracy Actions.	TSP, UC3M, IMDEA	September 2013
(14)	IEEE International Conference on Peer-to-Peer Computing (P2P 2013)	RASP: Using OpenFlow to Push Overlay Streams into the Underlay	TUD	September 2013
(15)	KSII TRANSACTIONS ON INTERNET AND INFORMATION SYSTEMS	Connectivity Properties of Real BitTorrent Swarms	UC3M, IMDEA	September 2013

(16)	ACM international Conference in Online Social Networks (COSN 2013)	Are Trending Topics useful for marketing? Visibility of Trending Topics vs Traditional Advertisement	UC3M	October 2013
(17)	IEEE/ACM Transactions on Networking	Unveiling the Incentives for Content Publishing in Popular BitTorrent Portals	UC3M, IMDEA, TUD	October 2013
(18)	Joint SmartenIT/eCOUSIN Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM13)	Determining Leaders and Clusters in Video Consumption	A-LBELL, ORANGE	October 2013
(19)	CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013)	Reciprocity with Virtual Nodes: Supporting Mobile Peers in Peer-to-Peer Content Distribution	TUD	October 2013
(20)	CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013)	Volume is not Enough: SVC-aware Server Allocation for Peer-assisted Streaming.	TUD	October 2013
(21)	IEEE International Conference on Network Protocols (ICNP 2013)	Optimizing Energy Consumption and QoE on Mobile Devices	TUD	October 2013
(22)	IEEE Network	Understanding the evolution of multimedia content in the Internet through BitTorrent glasses	TSP, UC3M	November 2013
(23)	ACM High Performance Mobile	Opportunistic Message Routing using Multi-layer Social Networks	UCAM	November 2013

	Opportunistic Systems (ACM HP-MOSys 2013)			
(24)	IEEE Global Communications Conference (GLOBECOM' 13)	"Current City" Prediction for Coarse Location Based Applications on Facebook.	TSP	December 2013
(25)	Communications of the ACM, vol. 56, no. 12	Where's in a name? A survey of mobility in information-centric networks	UC3M, ALUD	December 2013
(26)	Elsevier Computer Networks, Volume 59, Pages 77-90	TorrentGuard: Stopping scam and malware distribution in the BitTorrent ecosystem	UC3M	February 2014
(27)	Personal and Ubiquitous Computing. Vol 18, Issue 2, pages 379-411	Understanding the locality effect in Twitter: Measurements and Analysis	UC3M	February 2014
(28)	Transactions on Emerging Telecommunicatio ns Technologies	On the Tweet arrival process at Twitter: analysis and applications	UC3M	February 2014
(29)	IEEE Transactions on Parallel and Distributed Systems	BitTorrent Locality and Transit Traffic Reduction: When, Why and at What Cost?	UC3M	April 2014
(30)	2nd Workshop on Name Oriented Mobility (NOM) of the 33rd International Conference on Computer Communications (INFOCOM'14)	A CCN-Based Social Network Application Optimising Network Proximity	ORANGE	April 2014
(31)	IEEE/IFIP Network Operations and Management Symposium (NOMS 2014)	Flexible Traffic Management in Broadband Access Networks using Software Defined Networking	TUD	May 2014
(32)	European	A Model for Throughput	IMDEA	May 2014

	Wireless 2014, Special Session on stochastic modeling for protocol design and optimization	Prediction for Mobile Users		
(33)	Elsevier Future Generation Computer Systems.	On Exploiting Social Relationship and Personal Background for Content Discovery in P2P Networks	TSP, UC3M	June 2014
(34)	8th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2014)	Towards Decentralized, Energy- and Privacy-aware Device-to-Device Content Delivery	TUD	June 2014
(35)	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014)	Alike People, Alike Interests? A Large-Scale Study on Interest Similarity in Social Networks	TSP	August 2014
(36)	IEEE International Conference on Peer-to-Peer Computing (P2P 2014)	Clubbing with the Peers: A Measurement Study of BitTorrent Live	TUD	September 2014
(37)	The 1st KuVS Workshop on Anticipatory Networks	Modelling Throughput Prediction Errors as Gaussian Random Walks	IMDEA	September 2014
(38)	IEEE International Conference on Network Protocols (ICNP 2014)	Optimizing Mobile Prefetching by Leveraging Usage Patterns and Social Information	TUD	October 2014
(39)	JNSM Special Issue on Management of Software Defined Networks	Software-Defined Multicast for Over-the-Top and Overlay-based Live Streaming in ISP Networks	TUD	December 2014

(40)	The 19th IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (IEEE CAMAD 2014)	Demo: Mobile Social Prefetcher using Social and Network Information.	IMDEA, TUD	December 2014
(41)	Elsevier Decision Support Systems. Feb. 2015.	Alike people, alike interests? Inferring interest similarity in online social networks	TSP	February 2015
(42)	The 18th International Conference on Intelligence in Next Generation Networks (ICIN 2015)	Networking impact of a local-aware content-based delivery for twitter-like applications	ORANGE	February 2015
(43)	The 6th ACM Multimedia Systems 2015 Conference (MMSys 2015)	Demo: Media Download Optimization through Prefetching and Resource Allocation in Mobile Networks.	IMDEA, TUD	March 2015
(44)	Workshop on Software-Defined Networking and Network Function Virtualization for Flexible Network Management (SDNFlex 2015)	Open Elastic Provisioning of Hardware Acceleration in NFV Environments	TUD	March 2015
(45)	International Conference on Networked Systems (NetSys 2015)	Demonstration of NetworkCoverage – A Mobile Network Performance Measurement App	TUD	March 2015
(46)	International Conference on Networked Systems (NetSys	The Potential of Social-aware Multimedia Prefetching on Mobile Devices	TUD	March 2015

	2015)			
(47)	The 2nd Workshop on Communication and Networking Techniques for Contemporary Video, in conjunction with the 34th IEEE International Conference on Computer Communications (INFOCOM 2015)	Anticipatory Quality-Resource Allocation for Multi-User Mobile Video Streaming	IMDEA	April 2015
(48)	The 14th IFIP Networking 2015 Conference	Lightweight Mobile Bandwidth Availability Measurement	TUD, IMDEA	May 2015
(49)	IFIP Networking 2015	TopT: Supporting Flash Crowd Events in Hybrid Overlay-based Live Streaming	TUD	May 2015

Table 3. Accepted Publications

(01) Roy A, Nilakant K, Dalibard V, Yoneki E. **"Mitigating I/O latency in SSD-based Graph Traversal"**. *Technical Report, University of Cambridge, 2012 (UCAM-CL-TR-823).*

(partially related to eCOUSIN)

Abstract: Mining large graphs has now become an important aspect of many applications. Recent interest in low cost graph traversal on single machines has lead to the construction of systems that use solid state drives (SSDs) to store the graph. An SSD can be accessed with far lower latency than magnetic media, while remaining cheaper than main memory. Unfortunately SSDs are slower than main memory and algorithms running on such systems are hampered by large IO latencies when accessing the SSD. In this paper we present two novel techniques to reduce the impact of SSD IO latency on semi-external memory graph traversal. We introduce a variant of the Compressed Sparse Row (CSR) format that we call Compressed Enumerated Encoded Sparse Offset Row (CEESOR). CEESOR is particularly efficient for graphs with hierarchical structure and can reduce the space required to represent connectivity information by amounts varying from 5 % to as much as 76%. CEESOR allows a larger number of edges to be moved for each unit of IO transfer from the SSD to main memory and more effective use of operating system caches. Our second contribution is a runtime prefetching technique that exploits the ability of solid state drives to service multiple random access requests in parallel. We present a novel run along SSD Prefetcher (RASP). RASP is capable of hiding the effect of IO latency in single threaded graph traversal in breadth-first and shorted path order to the extent that it improves iteration time for large graphs by amounts varying from 2.6X-6X.

(02) Martin-Campillo A, Crowcroft J, Yoneki E, Marti R. **"Evaluating Opportunistic Networks in Disaster Scenarios"**. *Journal of Network and Computer Applications, Elsevier; November 2012.*

(partially related to eCOUSIN)

Abstract: Forwarding data in scenarios where devices have sporadic connectivity is a challenge. An example scenario is a disaster area, where forwarding information generated in the incident location, like victims' medical data, to a coordination point is critical for quick, accurate and coordinated intervention. New applications are being developed based on mobile devices and wireless opportunistic networks as a solution to destroyed or overused communication networks. But the performance of opportunistic routing methods applied to emergency scenarios is unknown today. In this paper, we compare and contrast the efficiency of the most significant opportunistic routing protocols through simulations in realistic disaster scenarios in order to show how the different characteristics of an emergency scenario impact in the behaviour of each one of them.

(03) González R, Cuevas R, Motamedi R, Rejaie R, Cuevas A. **"Google+ or Google-?: Dissecting the evolution of the new OSN in its first year"**. *World Wide Web Conference (WWW' 13); Rio de Janeiro (Brazil); May 2013.*

Abstract: In the era when Facebook and Twitter dominate the market for social media, Google has introduced Google+ (G+) and reported a significant growth in its size while others called it a ghost town. This begs the question "whether G+ can really attract a significant number of connected and active users despite the dominance of Facebook and Twitter". This paper tackles the above question by presenting a detailed characterization of G+ based on large scale measurements. We identify the main components of G+ structure, characterize the key features of their users and their evolution over time. We then conduct detailed analysis on the evolution of connectivity and activity among users in the largest connected component (LCC) of G+ structure, and compare their characteristics with other major OSNs. We show that despite the dramatic growth in the size of G+, the relative size of LCC has been decreasing and its connectivity has become less clustered. While the aggregate user activity has gradually increased, only a very small fraction of users exhibit any type of activity. To our knowledge, our study offers the most comprehensive characterization of G+ based on the largest collected data sets.

(04) Bao X, Lin Y, Lee U, Rimac I, Choudhury R.R. **"DataSpotting: Exploiting Naturally Clustered Mobile Devices to Offload Cellular Traffic"**. *IEEE International Conference on Computer Communications (IEEE Infocom 2013); Turin (Italy); June 2013.*

Abstract: The proliferation of pictures and videos in the Internet is imposing heavy demands on mobile data networks. Though emerging wireless technologies will provide more bandwidth, the increase in demand will easily consume the additional capacity. To alleviate this problem, we explore the possibility of serving user requests from other mobile devices located geographically close to the user. For instance, when Alice reaches areas with high device density - Data Spots - the cellular operator learns Alice's content request, and guides her device to nearby devices that have the requested content. Importantly, communication between the nearby devices can be mediated by servers, avoiding many of the known problems of pure ad hoc communication. This paper argues this viability through systematic prototyping, measurements, and measurement-driven analysis.

(05) Gouta A, Hong C, Hong D, Kermarrec A. M, Le Louédec Y. **"Large scale analysis of HTTP adaptive streaming in mobile networks"**. *IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM 2013); Madrid (Spain); June 2013.*

Abstract: HTTP Adaptive bitrate video Streaming (HAS) is now widely adopted by Content Delivery Network Providers (CDNPs) and Telecom Operators (Telcos) to improve user Quality of Experience (QoE). In HAS, several versions of videos are made available in the network so that the quality of the video can be chosen to better fit the bandwidth capacity of users. These delivery requirements raise new challenges with respect to content caching strategies, since several versions of the content may compete to be cached. In this paper we present analysis of a real HAS dataset collected in France and provided by a mobile telecom operator involving more than 485,000 users requesting adaptive video contents through more than 8 million video sessions over a 6 week measurement period. Firstly, we propose a fine-grained definition of content popularity by exploiting the segmented nature of video streams. We also provide analysis about the behavior of clients when requesting such HAS streams. We propose novel caching policies tailored for chunk-based streaming. Then we study the relationship between the requested video bitrates and radio constraints. Finally, we study the users' patterns when selecting different bitrates of the same video content. Our findings provide useful insights that can be leveraged by the main actors of video content distribution to improve their content caching strategy for adaptive streaming contents as well as to model users' behavior in this context.

(06) Wichtlhuber M, Hausheer D. **"Towards a Mobility-Supporting Incentive Scheme for Peer-to-Peer Content Distribution".** *7th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2013); Barcelona (Spain); June 2013.*

Abstract: The distribution of bulk content such as video streams and system updates is becoming increasingly important on mobile devices. For this kind of data transfers, the application of the Peer-to-Peer (P2P) paradigm is attractive for content providers, as it allows to leverage clients' (peers') resources for the dissemination process. The contribution of resources can be encouraged by applying reciprocal incentives, which prevent a peer from consuming more resources than he contributed to the system. However, reciprocation has drawbacks, e.g., in terms of high energy consumption and low performance, when applied to mobile peers. This work substantiates these problems through measurements and discusses a virtual node concept as a possible solution.

(07) Yoneki E, Roy A. **"Scale-up Graph Processing: A Storage-centric View".** *ACM Special Interest Group on Management of Data(ACM SIGMOD 2013), New York (USA), June 2013.*

Abstract: The determinant of performance in scale-up graph processing on a single system is the speed at which the graph can be fetched from storage: either from disk into memory or from memory into CPU-cache. Algorithms that follow edges perform random accesses to the storage medium for the graph and this can often be the determinant of performance, regardless of the algorithmic complexity or runtime efficiency of the actual algorithm in use. A storage-centric viewpoint would suggest that the solution to this problem lies in recognizing that graphs represent a unique workload and therefore should be treated as such by adopting novel ways to access graph structured data. We approach this problem from two different aspects and this paper details two different efforts in this direction. One approach is specific to graphs stored on SSDs and accelerates random access using a novel prefetcher called RASP. The second approach takes a fresh look at how graphs are accessed and suggests that trading off the low cost of random access for the approach of sequentially streaming a large set of (potentially unrelated) edges can be a winning proposition under certain circumstances: leading to a system for graphs stored on any medium (main-memory,

SSD or magnetic disk) called X-stream. RASP and X-stream therefore take - diametrically opposite - storage centric viewpoints of the graph processing problem. After contrasting the approaches and demonstrating the benefit of each, this paper ends with a description of planned future development of an online algorithm that selects between the two approaches, possibly providing the best of both worlds.

(08) Boutet A, Kim H, Yoneki E.**"What's in Twitter I Know What Parties are Popular and Who You are Supporting Now!".** *Springer Journal of Social Network Analysis and Mining, ISSN: 1869-5450; July 2013.*

Abstract: In modern politics, parties and individual candidates must have an online presence and usually have dedicated social media coordinators. In this context, we study the usefulness of analysing Twitter messages to identify both the characteristics of political parties and the political leaning of users. As a case study, we collected the main stream of Twitter related to the 2010 UK General Election during the associated period—gathering around 1,150,000 messages from about 220,000 users. We examined the characteristics of the three main parties in the election and highlighted the main differences between parties. First, the retweet structure is highly clustered according to political parties. Second, users are more likely to refer to their preferred party and use more positive affect words for the party compared with other parties. Finally, the self-description of users and the List feature can reflect the political orientation of users. From these observations, we develop both an incremental and practical classification method which uses the number of Twitter messages referring to a particular political party or retweets, and a classifier leveraging the valuable semantic content of the List feature to estimate the overall political leaning of users. The experimental results showed that the proposed incremental method achieved an accuracy of 86 % for classifying the users' political leanings and outperforms other classification methods that require expensive costs for tuning classifier parameters and/or knowledge about network topology.

(09)Rimac I, Hilt V. **"Addressing the Wireless Content Challenge"**. *IEEE COMSOC MMTC E-Letter, Vol.8, No.4, July 2013.*

Abstract: The ever increasing demand for high-volume content in the Internet is imposing heavy demands on mobile data networks. The growth in demand is forecasted to far outpace the increase in capacity through advances in wireless technologies. To alleviate this problem, we propose to serve user requests for content from caches in cellular access nodes and from other mobile devices located geographically close to the requesting user."

(10) Gouta A, Hong D, Kemmarec A.M, Le Louédec Y. **"HTTP adaptive streaming in mobile networks: characteristics and caching opportunities".** *IEEE 21ST International Symposium on Modeling, Analysis And Simulation of Computer And Telecommunications (MASCOTS 2013); San Francisco (USA); August 2013*

Abstract: Cellular networks have witnessed the emergence of the HTTP Adaptive Streaming (HAS) as a new video delivery method. In HAS, several qualities of the same videos are made available in the network so that clients can choose the best quality that fits their bandwidth capacity. This has particular implications on caching strategies with respect to the viewing patterns and the switching behavior between video qualities. In this paper we present analysis of a real HAS dataset collected in France and provided by the country's largest mobile phone operator. Firstly, we analyse the viewing patterns of HAS contents and the distribution of the encoding bitrates requested by mobile clients. Secondly, we give an in-depth analysis of the switching pattern between video bitrates during a video session and assess the implication on the caching efficiency. We also model this switching based on

empirical observations. Finally, we propose WA-LRU a new caching algorithm tailored for HAS contents and compare it to the standard LRU. Our evaluations demonstrate that WA-LRU performs better and achieves its goals.

(11) Fay D., Kunegis J., Yoneki E. **"Centrality and Mode Detection in Dynamic Contact Graphs; a Joint Diagonalisation Approach"**. *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013); Niagara Falls (Canada), August 2013.*

Abstract: This paper presents a technique for analysis of dynamic contact networks aimed at extracting periods of time during which the network changes behaviour. The technique is based on tracking the eigenvectors of the contact network in time (efficiently) using a technique called Joint Diagonalisation (JD). Repeated application of JD then shows that real-world networks naturally break into several modes of operation which are time dependent and in one real-world case, even periodic. This shows that a view of real-world contact networks as realisations from a single underlying static graph is mistaken. However, the analysis also shows that a small finite set of underlying static graphs can approximate the dynamic contact graphs studied. We also provide the means by which these underlying approximate graphs can be constructed. Core to the approach is the analysis of spanning trees constructed on the contact network. These trees are the routes a broadcast would take given a random starting location and we find that these propagation paths (in terms of their eigenvector decompositions) cluster into a small subset of modes which surprisingly correspond to clusters in time. The net result is that a dynamic network may be approximated as a (small finite) set of static graphs. Most interestingly the MIT dataset shows a periodic behaviour which allows us to know in advance which mode the network will be in. This has obvious consequences as individuals in the network take differing roles in differing modes. Finally, we demonstrate the technique by constructing a synthetic network with 4 underlying modes of operation; creating synthetics contacts and then used JD to extract the original underlying modes.

(12) Farahbakhsh R, Han X, Cuevas Á, Crespi N.**"Analysis of publicly disclosed information in Facebook profiles".***IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013); Niagara Falls (Canada), August 2013.*

Abstract: Facebook, the most popular online social network is a virtual environment where users share information and are in contact with friends. Apart from many useful aspects, there is a large amount of personal and sensitive information publicly available that is accessible to external entities/users. In this paper we study the public exposure of Facebook profile attributes to understand what type of attributes are considered more sensitive by Facebook users in terms of privacy, and thus are rarely disclosed, and which attributes are available in most Facebook profiles. Furthermore, we also analyze the public exposure of Facebook users by accounting the number of attributes that users make publicly available on average. To complete our analysis we have crawled the profile information of 479K randomly selected Facebook users. Finally, in order to demonstrate the utility of the publicly available information in Facebook profiles we show in this paper three case studies. The first one carries out a gender-based analysis to understand whether men or women share more or less information. The second case study depicts the age distribution of Facebook users. The last case study uses data inferred from Facebook profiles to map the distribution of worldwide population across cities according to its size.

(13) Farahbakhsh R, Cuevas Á, Cuevas R., Rejaie R, Kryczka M, González R, Crespi N. **"Investigating the Reaction of BitTorrent Content Publishers to Antipiracy Actions"**. *IEEE International Conference on Peer-to-Peer Computing (P2P 2013); Trento (Italy); September 2013.*

Abstract: During recent years, a few countries have put in place online antipiracy laws and there has been some major enforcement actions against violators. This raises the question that to what extent antipiracy actions have been effective in deterring online piracy? This is a challenging issue to explore because of the difficulty to capture user behavior, and to identify the subtle effect of various underlying (and potentially opposing) causes. In this paper, we tackle this question by examining the impact of two major antipiracy actions, the closure of Megaupload and the implementation of the French antipiracy law, on publishers in the largest BitTorrent portal who are major providers of copyrighted content online. We capture snapshots of BitTorrent publishers at proper times relative to the targeted antipiracy event and use the trends in the number and the level of activity of these publishers to assess their reaction to these events. Our investigation illustrates the importance of examining the impact of antipiracy events on different groups of publishers and provides valuable insights on the effect of selected major antipiracy actions on publishers' behavior.

(14) Rückert J, Blendin J, Hausheer D. **"RASP: Using OpenFlow to Push Overlay Streams into the Underlay"**. *IEEE International Conference on Peer-to-Peer Computing (P2P 2013); Trento (Italy); September 2013.*

Abstract: Today, an increasing number of live video streaming services are delivered over the Internet. Content Delivery Networks are used for a scalable and cost-efficient delivery, usually ending at the edge of the residential Internet Service Provider (ISP) networks. From here on, the streams have to be transported to the clients using IP unicast. Peer-to-peer (P2P) streaming mechanisms promise to further improve this situation for content providers by shifting some load to the client devices themselves. While this is promising for the content provider, in both cases the costs stay with the ISPs. They have to carry individual streams with the same content through their network. To address this problem, in this demo paper the OpenFlow-based cross-layer approach RASP is proposed that allows ISPs to offer a network layer multicast service with P2P support for traffic originated from outside their own network.

(15) Cuevas R, Kryczka M, Cuevas A, Guerrero C, Azcorra A. **"Connectivity Properties of Real BitTorrent Swarms".** *KSII Transactions on Internet and Information Systems, September 2013*.

Abstract: BitTorrent is one of the most important applications in the current Internet. Despite of its interest, we still have little knowledge regarding the connectivity properties of real BitTorrent swarms. In this paper we leverage a dataset including the connectivity information of 250 real torrents and more than 150k peers to carefully study the connectivity properties of peers. The main topology parameters of the studied swarms suggest that they are significantly less resilient than random graphs. The analysis of the peer level connectivity properties reveals that peers continuously change more than half of their neighbours. Furthermore, we also find that a leecher typically keeps stable connections with a handful of neighbours with which it exchanges most of its traffic whereas seeders do not establish long term connections with any peer so that they can homogeneously distribute chunks among leechers. Finally, we have discovered that a significant portion of the studied peers (45%) have an important locality-biased neighbourhood composition.

(16) Carrascosa J, Gonzalez R, Cuevas R, Azcorra A. **"Are Trending Topics useful for marketing? Visibility of Trending Topics vs Traditional Advertisement".** *ACM Conference on Online Social Networks (COSN 2013); Boston (USA); October 2013.*

Abstract: Trending Topics seem to be a powerful tool to be used in marketing and advertisement contexts, however there is not any rigorous analysis that demonstrates this. In this paper we present

a first effort in this direction. We use a dataset including more than 110K Trending Topics from 35 countries collected over a period of 3 months as basis to characterize the visibility offered by Local Trending Topics. Furthermore, by using metrics that rely on the exposure time of Trending Topics and the penetration of Twitter, we compare the visibility provided by Trending Topics and traditional advertisement channels such as newspapers' ads or radio-stations' commercials for several countries. Our study confirms that Trending Topics offer a comparable visibility to the aforementioned traditional advertisement channels in those countries where we have conducted our comparison study. Then, we conclude that Trending Topics can be useful in marketing and advertisement contexts at least in the analyzed countries.

(17) Cuevas R, Kryczka M, Cuevas A, Kaune S, Guerrero C, Rejaie R. **"Unveiling the Incentives for Content Publishing in Popular BitTorrent Portals"**. *IEEE/ACM Transactions in Networking, October 2013*

Abstract: BitTorrent is the most popular peer-to-peer (P2P) content delivery application where individual users share various types of content with tens of thousands of other users. The growing popularity of BitTorrent is primarily due to the availability of valuable content without any cost for the consumers. However, apart from the required resources, publishing valuable (and often copyrighted) content has serious legal implications for the users who publish the material. This raises the question that whether (at least major) content publishers behave in an altruistic fashion or have other motives such as financial incentives. In this paper, we identify the content publishers of more than 55 K torrents in two major BitTorrent portals and examine their characteristics. We discover that around 100 publishers are responsible for publishing 67% of the content, which corresponds to 75% of the downloads. Our investigation reveals several key insights about major publishers. First, antipiracy agencies and malicious users publish "fake" files to protect copyrighted content and spread malware, respectively. Second, excluding the fake publishers, content publishing in major BitTorrent portals appears to be largely driven by companies that try to attract consumers to their own Web sites for financial gain. Finally, we demonstrate that profit-driven publishers attract more loyal consumers than altruistic top publishers, whereas the latter have a larger fraction of loyal consumers with a higher degree of loyalty than the former.

(18) De Vleeschauwer D, Hawinkel C, Le Louédec Y. **"Determining Leaders and Clusters in Video Consumption"**. *Joint SmartenIT/eCOUSIN Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications(SETM 2013); Zurich (Switzerland); October 2013*

Abstract: In this paper we analyze a trace of a deployed VoD (video on demand) system. Users issue requests for content items in an online VoD catalogue at given moments in time. Based on this information alone we identify communities of users that have similar content preferences, which we refer to as implicit social communities. We find that there is evidence for a limited number of groups of similar users. Next we also determine lead users, i.e., users that consume popular content items consistently before other users do. We show that such users can be identified in the considered data set. We also explain how these two pieces of information could be used to improve recommendation systems and content distribution networks.

(19) Wichtlhuber M, Heise P, Scheurich B, Hausheer D. ***"Reciprocity with Virtual Nodes: Supporting Mobile Peers in Peer-to-Peer Content Distribution"***. *CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013); Zurich (Switzerland); October 2013.*

Abstract: The Peer-to-Peer (P2P) paradigm offers scalable means to perform bulk data distribution, e.g., for small businesses which cannot afford huge upfront investments, by incorporating user's resources in the dissemination process. Due to the proliferation of smartphones with wireless broadband connectivity and the increasing convergence of fixed and mobile platforms, a growing number of users are expected to participate in P2P content distribution networks wirelessly. However, the P2P approach only works if users are willing to contribute resources. A commonly applied incentive scheme is the well-known Tit-for-Tat approach, where each peer is forced to contribute as much bandwidth to the network as he consumes. Nevertheless, reciprocal schemes discriminate resource poor mobile devices in terms of energy and upload bandwidth, as they are device-bound instead of being user-bound. In this work, an incentive scheme featuring virtual nodes is presented, which allows mobile devices to seek help from other devices owned by its user, e.g., the user's home gateway or a supporting cloud instance. Preliminary results are presented in the scope of a P2P streaming scenario.

(20) *Rückert J, Abboud O, Kluge M, Hausheer D.* ***"Volume is not Enough: SVC-aware Server Allocation for Peer-assisted Streaming".*** *CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013); Zurich (Switzerland); October 2013.*

Abstract: Peer-assisted delivery of video content has shown a great potential to reduce upload bandwidth requirements for content providers by exploiting idle client resources in the video dissemination process. As primary content sources, the servers run by content providers play a critical role in such systems, making their adequate provisioning a key part of the streaming mechanism. While dynamic resource provisioning has been studied before, little is known about resource allocation for streaming of scalable media content. Besides the pure amount of resources, here, the quality level of the delivered video content becomes relevant. The spreading of video blocks with the wrong quality can lead to situations where peers are forced to reduce their video qualities, despite them having enough download capacity. To address this problem, in this paper, a new SVC-based adaptation policy and a request-based extension to it are proposed, enabling content providers to manage their streaming services in a video quality-aware manner. Prototypical evaluations show that the mechanisms outperform existing quality-agnostic approaches in terms of delivered SVC video quality.

(21) Kaup F, Hausheer D. **"Optimizing Energy Consumption and QoE on Mobile Devices".** *IEEE International Conference on Network Protocols (ICNP 2013), Göttingen (Germany); October 2013.*

Abstract: The increased availability and data rates of cellular 3G/4G networks combined with the growing use of mobile applications highly affect the Quality of Experience (QoE) perceived by the end-user. The QoE is affected in two ways: First, the data rates in the networks are low when multiple users simultaneously request content; second, the transmission of data over slow connections consumes a considerable amount of energy compared to faster connections. Both can be avoided by better management of the available resources. This paper proposes a new approach, taking the energy efficiency into account as a key QoE aspect. Based on user mobility models, the available connectivity can be predicted, from which estimates for the energy consumption and expected QoE can be derived. An architecture is sketched, which combines QoE prediction for current and future network connections with energy efficiency on mobile devices. The key idea to optimize the QoE is to defer connections, if a better connectivity is predicted in the near future and the Quality of Service (QoS) requirements allow delaying the connection.

(22) Farahbakhsh R, Cuevas A, Cuevas R, Gonzalez R, Crespi N. **"Understanding the evolution of multimedia content in the Internet through BitTorrent glasses"**. *IEEE Network; November 2013*.

Abstract: Today's Internet traffic is mostly dominated by multimedia content and the prediction is that this trend will intensify in the future. Therefore, main Internet players, such as ISPs, content delivery platforms (e.g. Youtube, Bitorrent, Netflix, etc.) or CDN operators, need to understand the evolution of multimedia content availability and popularity in order to adapt their infrastructures and resources to satisfy clients requirements while they minimize their costs. This paper presents a thorough analysis on the evolution of multimedia content available in BitTorrent. Specifically, we analyze the evolution of four relevant metrics across different content categories: content availability, content popularity, content size and user's feedback. To this end we leverage a large-scale dataset formed by 4 snapshots collected from the most popular BitTorrent portal, namely The Pirate Bay, between Nov. 2009 and Feb. 2012. Overall our dataset is formed by more than 160k content that attracted more than 185M of download sessions.

(23) Socievole A, Yoneki E, De Rango E, Crowcroft J ."**Opportunistic Message Routing using Multi-layer Social Networks".** *ACM High Performance Mobile Opportunistic Systems (ACM HP-MOSys 2013); Barcelona (Spain); November 2013.*

Abstract: In opportunistic networks, the nodes usually exploit a contact opportunity to perform hop-by-hop routing, since an end-to-end path between the source node and destination node may not exist. Most social-based routing protocols use social information extracted from real-world encounter networks to select an appropriate message relay. A protocol based on encounter history, however, takes time to build up a knowledge database from which to take routing decisions. An opportunistic routing protocol which extracts social information from multiple social networks, can be an alternative approach to avoid suboptimal paths due to partial information on encounters. While contact information changes constantly and it takes time to identify strong social ties, online social network ties remain rather stable and can be used to augment available partial contact information. In this paper, we propose a novel opportunistic routing approach, called ML-SOR (Multi-layer Social Network based Routing), which extracts social network information from multiple social contexts. To select an effective forwarding node, ML-SOR measures the forwarding capability of a node when compared to an encountered node in terms of node centrality, tie strength and link prediction. These metrics are computed by ML-SOR on different social network layers. Trace driven simulations show that ML-SOR, when compared to other schemes, is able to deliver messages with high probability while keeping overhead ratio very small.

(24) Chanthaweethip W, Han X, Crespi N, Chen Y, Farahbakhsh R, Cuevas Á. **"Current City Prediction for Coarse Location Based Applications on Facebook"**. *IEEE Global Communications Conference (GLOBECOM' 13). Atlanta (USA). December 2013.*

Abstract: Location-Based services with social networks improve users' experience and enrich people's social live. However, location information is often inadequate due to privacy and security concerns. We seek to infer users' 'Current City' on Facebook for coarse location based applications. We first extract users' multiple explicit and implicit location attributes, and analyze correlations of these attributes from two perspectives: user-centric and user-friends. We observe that both user-centric and user-friends location attributes tightly correlate to a user's Current City (e.g., 60% of users stay in their hometown, 60% of users live in the same city as 50% of their friends). Based on extensive analysis and observations on location attributes correlations, we have constructed a Current City Prediction model (CCP) using artificial neural network (ANN) learning frameworks. The experimental results indicate that we achieve accuracy levels of 84% for city-level prediction and 98%

for country-level which are increases of 9% and 18%, respectively than what is possible with Tweecalization.

(25) Tyson G, Sastry N, Cuevas R, Rimac I, Mauthe A. **"Where is in a Name? A Survey of Mobility in Information-Centric Networks".** *Communications of the ACM (CACM); Vol. 56, No. 12; December 2013.*

Abstract: Host mobility has been a long standing challenge in the current Internet architecture. Huge proportions of traffic are now attributed to mobile devices; however, despite this prominence, mobility often remains a badly handled concept. Some have recently argued that the main reason for this lies in its choice of what to name. The Internet Protocol (IP) names hosts based on their topological network location. Through this, it intrinsically binds the what (the name) to the where (the address). Consequently, a mobile host moving its physical location is often required to change its name creating numerous problems. Observations such as this have led to a flurry of research looking at how the future Internet could be re-designed. A prominent example is that of information-centric networks (ICNs). ICNs propose a key paradigm shift, which involves replacing the Internet's existing host-based naming scheme with an information-based one instead. This article therefore chooses to follow Shakespeare's advice and ask \what's in a name?", rather than IP's approach of \where is in a name?". Through this principle, an ICN becomes an infrastructure that revolves around the provision of uniquely identified content to consumers, rather than the routing of data between device pairs. By removing the use of host-centric naming, it is therefore hoped that it will be possible to seamlessly change a host's physical and topological location without needing to perform the types of complex network management that host-centric networks require (e.g. creating forwarding between home and foreign addresses). In this article we aim to explore and review these concepts and ideas. We first explore what an ICN is, before investigating some of the key benefits of designing a network around the concept of information. From this, we then present some prominent ICN proposals before using these to identify important remaining challenges.

(26) Kryczka M, Cuevas R, Cuevas Á, Guerrero C, Azcorra A. **"TorrentGuard: stopping scam and malware distribution in the BitTorrent ecosystem".** *Elsevier Computer Networks, Volume 59, Pages 77-90; February 2014.*

Abstract: In this paper we conduct a large scale measurement study in order to analyse the fake content publishing phenomenon in the BitTorrent Ecosystem. Our results reveal that fake content represents an important portion (35%) of those files shared in BitTorrent and just a few tens of users are responsible for 90% of this content. Furthermore, more than 99% of the analysed fake files are linked to either malware or scam websites. This creates a serious threat for the BitTorrent ecosystem. To address this issue, we present a new tool named TorrentGuard for the early detection of fake content. Based on our evaluation this tool may prevent end users from downloading more than 35 millions of fake files per year. This could help to reduce the number of computer infections and scams suffered by BitTorrent users. TorrentGuard is already available and it can be accessed through both a webpage or a Vuze plugin.

(27) Cuevas R, González R, Cuevas A, Guerrero C. **"Understanding the Locality Effect in Twitter: Measurement and Analysis"**. *Springer Personal and Ubiquitous Computing (Special Issue in Cross-Community Mining). February 2014.*

Abstract: Twitter is one of the most popular applications in the current Internet with more than 500M registered users across the world. In this paper we conduct a comprehensive analysis to

understand the geographical characteristics of Twitter using Cross Community Mining (CCM) techniques. Specifically, we study the locality level shown by the three main elements of Twitter namely, users, relationships and information flow. For this purpose we rely on a dataset including the geolocation information of more than 17M, 100M and 3.5M users, relationships and tweets, respectively. Our main findings are: (1) most of the Twitter users perform their activity from an area of at most few hundred Kms covering few cities within a unique country; (2) the location (i.e., country), and in particular factors such as language or Twitter popularity within a country, dictate the level of locality in the relationships of users and Twitter conversations originated in that country. The combination of these factors reveals the presence of four types of country locality profiles that we carefully analyze and compare in the paper.

(28) González R, Muñoz A, Hernéndez J.A, Cuevas R. **"On the Tweet arrival process at Twitter: analysis and applications".** *Transactions on Emerging Telecommunications Technologies; February 2014.*

Abstract: this work provides a novel measurement-based analysis of the tweet arrival traffic process at Twitter. The analysis considers more than one million total tweets collected at 48 different times of the day (o'clock and half-past every hour). We observe a 3.5-tweet/ms average rate with a valley of 2.5 tweets/ms at 10 am (GMT+1) and a peak at 3 pm (GMT+1) of about 5 tweets/ms. We further model the traffic pattern as a Gaussian process, and we validate such an assumption with multiple normality tests. Finally, we overview a number of applications where such a model may show its utility, namely infrastructure dimensioning and upgrading, the detection of outlier events, energy efficiency and so on.

(29) Cuevas R, Laoutaris N, Xiao Y, Siganos G, Rodriguez P. **"BitTorrent Locality and Transit Traffic Reduction: When, Why and at What Cost?"**. *IEEE Transactions on Parallel and Distributed Systems, April 2014*

Abstract: A substantial amount of work has recently gone into localizing BitTorrent traffic within an ISP in order to avoid excessive and often times unnecessary transit costs. Several architectures and systems have been proposed and the initial results from specific ISPs and a few torrents have been encouraging. In this work we attempt to deepen and scale our understanding of locality and its potential. Looking at specific ISPs, we consider tens of thousands of concurrent torrents, and thus capture ISP-wide implications that cannot be appreciated by looking at only a handful of torrents. Secondly, we go beyond individual case studies and present results for few thousands ISPs represented in our dataset of up to 40K torrents involving more than 3.9M concurrent peers and more than 20M in the course of a day spread in 11K ASes. Finally, we develop scalable methodologies that allow us to process this huge dataset and derive accurate traffic matrices of torrents. Using the previous methods we obtain the following main findings: (1) Although there are a large number of very small ISPs without enough resources for localizing traffic, by analyzing the 100 largest ISPs we show that Locality policies are expected to significantly reduce the transit traffic with respect to the default random overlay construction method in these ISPs; (2) contrary to the popular belief, increasing the access speed of the clients of an ISP does not necessarily help to localize more traffic; (3) by studying several real ISPs, we have shown that soft speed-aware locality policies guarantee win-win situations for ISPs and end users. Furthermore, the maximum transit traffic savings that an ISP can achieve without limiting the number of inter-ISP overlay links is bounded by ``unlocalizable'' torrents with few local clients. The application of restrictions in the number of inter-ISP links leads to a higher transit traffic reduction but the QoS of clients downloading ``unlocalizable'' torrents would be severely harmed.

(30) Mathieu B, Truong P. **"A CCN-Based Social Network Application Optimising Network Proximity"**. *2nd Workshop on Name Oriented Mobility (NOM) of the 33rd International Conference on Computer Communications (INFOCOM'14); Toronto (Canada); April 2014;*

Abstract: Online Social Networking (OSN) applications attracted millions of people in few years and are considered as the success story of current Internet applications. However, how they work is unclear for both end-users and researchers, since the developers keep the system architecture secret and use encryption mechanisms. In this paper, we present the main outcomes of our analysis of one of the most well-known OSNs, Twitter, focusing on the networking behaviour, the involved entities, their location, etc. Detecting that the current network behaviour of those applications is not in line with current end-users behaviours, where locality is important, we propose to adapt the OSNs to a Content-Centric Networking (CCN) approach, which could improve the delivery and reduce the network load and server load. Indeed, CCN is a new networking paradigm focusing on interest based requests of content itself independently on location and so has features very close to OSNs. This paper then introduces our proposal for a CCN-based architecture for OSN applications, with a naming and routing proposal. We also present the demonstrator we have implemented as proof-of-concept and the testbed we set up for functional testing purposes.

(31)Rückert J, Bifulco R, Rizwan-Ul-Haq M, Kolbe Hans-Joerg, Hausheer D. **"Flexible Traffic Management in Broadband Access Networks using Software Defined Networking".***IEEE/IFIP Network Operations and Management Symposium (NOMS 2014); Krakow (Poland); May 2014*.

Abstract: Over the years, the demand for high bandwidth services, such as live and on-demand video streaming, steadily increased. The adequate provisioning of such services is challenging and requires complex network management mechanisms to be implemented by Internet service providers (ISPs). In current broadband network architectures, the traffic of subscribers is tunneled through a single aggregation point, independent of the different service types it belongs to. While having a single aggregation point eases the management of subscribers for the ISP, it implies huge bandwidth requirements for the aggregation point and potentially high end-to-end latency for subscribers. An alternative would be a distributed subscriber management, adding more complexity to the management itself. In this paper, a new traffic management architecture is proposed that uses the concept of Software Defined Networking (SDN) to extend the existing Ethernet-based broadband network architecture, enabling a more efficient traffic management for an ISP. By using SDN-enabled home gateways, the ISP can configure traffic flows more dynamically, optimizing throughput in the network, especially for bandwidth-intensive services. Furthermore, a proof-of-concept implementation of the approach is presented to show the general feasibility and study configuration tradeoffs. Analytic considerations and testbed measurements show that the approach scales well with an increasing number of subscriber sessions.

(32)Bui N, Michelinakis F, Widmer J. **"A Model for Throughput Prediction for Mobile Users".** *European Wireless 2014 - Special Session on stochastic modeling for protocol design and optimization; Barcelona (Spain); May 2014.*

Abstract: In this paper we propose a stochastic model to predict user throughput in mobile networks. In particular, the model accounts for uncertainty such as random phenomena (e.g., fast fading) or inexact information (e.g., user location) to derive the statistical distribution of the user throughput. Such a model is highly useful for aiding scheduling and resource allocation decisions. In addition, we provide a taxonomy of prediction techniques to investigate error sources and the main characteristics of prediction accuracy. Finally, we show the versatility of the model by analyzing LTE

user throughput for the case where knowledge of either the user's actual position or the congestion level in the cell is inexact.

(33)Han X, Cuevas Á, Crespi N, Cuevas R, Huang X. **"On Exploiting Social Relationship and Personal Background for Content Discovery in P2P Networks**". *Elsevier Future Generation Computer Systems, Vol. 40, pp.17-29; June 2014.*

Abstract: Content discovery is a critical issue in unstructured Peer-to-Peer (P2P) networks as nodes maintain only local network information. However, similarly without global information about human networks, one still can find specific persons via his/her friends by using social information. Therefore, in this paper, we investigate the problem of how social information (i.e., friends and background information) could benefit content discovery in P2P networks. We collect social information of 384,494 user profiles from Facebook, and build a social P2P network model based on the empirical analysis. In this model, we enrich nodes in P2P networks with social information and link nodes via their friendships. Each node extracts two types of social features – Knowledge and Similarity – and assigns more weight to the friends that have higher similarity and more knowledge. Furthermore, we present a novel content discovery algorithm which can explore the latent relationships among a node's friends. A node computes stable scores for all its friends regarding their weight and the latent relationships. It then selects the top friends with higher scores to query content. Extensive experiments validate performance of the proposed mechanism. In particular, for personal interests searching, the proposed mechanism can achieve 100% of Search Success Rate by selecting the top 20 friends within two-hop. It also achieves 6.5 Hits on average, which improves 8x the performance of the compared methods.

(34) Nobach L, Hausheer D. **"Towards Decentralized, Energy- and Privacy-aware Device-to-Device Content Delivery**". *8th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2014), Brno (Czech Republic); June 2014.*

Abstract: Device-to-Device (D2D) content delivery is a new approach to directly exchange content between mobile devices, which allows to offload traffic from infrastructure-based networks and thus reduces the risk of congestion. While centralized D2D approaches rely on the mobile operator to discover nearby devices and initiate a content transfer, in decentralized D2D the devices autonomously and opportunistically organize themselves to deliver content to each other. The latter approach is more flexible as it does not depend on a single operator, however, it typically requires more energy for scanning other devices. Another issue is privacy, since the decentralized approach reveals a user's interest for content to all devices in proximity. Therefore, this paper sketches an new approach towards decentralized energy- and privacy-aware D2D content delivery. The proposed approach addresses the energy loss that occurs when constantly and spontaneously scanning multiple unknown devices for content. Furthermore, the paper identifies privacy requirements and proposes first steps towards a privacy-aware D2D solution.

(35) Han X, Wang L, Park S, Cuevas A, Crespi N. **"Alike People, Alike Interests? A Large-Scale Study on Interest Similarity in Social Networks".** *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining; Beijing (China); August 2014*

Abstract: In this paper, we present a comprehensive empirical study on the correlations between users' interest similarity and various social features across three interest domains (i.e., movie, music and TV). This study relies on a large dataset, containing 479,048 users and 5,263,351 user-generated interests, captured from Facebook. We identify the social features from three types of the users'

information - demographic information (e.g., age, gender, location), social relations (i.e., friendship), and users' interests. The results reveal that the interest similarity follows the homophily principle, which could be further harnessed by various practical applications and services.

(36) Rückert J, Knierim T, Hausheer D. **"Clubbing with the Peers: A Measurement Study of BitTorrent Live".** *IEEE International Conference on Peer-to-Peer Computing (P2P 2014), London (United Kingdom); September 2014.*

Abstract: The peer-to-peer approach can greatly help to cope with highly dynamic live streaming workload by using idle client resources. Yet, P2P streaming typically comes at the cost of increased streaming delays caused by the inevitable multi-hop forwarding of content by peers within the overlay. Various P2P streaming approaches have been proposed aiming at a good tradeoff between flexibility, streaming delay, and costs in terms of traffic overhead for both content providers and clients. Recently, BitTorrent Inc. released a new P2P live streaming system termed BTLive, specifically targeted at low delay and low overhead. For content providers investigating the applicability of BTLive's approach, it is essential to understand its properties as well as its limitations. So far, no publicly available study exists that quantitatively analyzes BTLive's performance. To this end, this paper presents a measurement study of the official beta version of BTLive. The study aims to answer the following key questions: How peer-to-peer is BTLive? How delay optimized is BTLive? What is the overhead of BTLive? To answer these questions, traces of real BTLive traffic between a broadcast server and a number of peers deployed across Europe have been analyzed.

(37) Bui N, Widmer J. **"Modelling Throughput Prediction Errors as Gaussian Random Walks".** *The 1st KuVS Workshop on Anticipatory Networks; Stuttgart (Germany), September 2014*

Abstract: One of the most critical aspects of anticipatory networking is assuming that future system conditions can be estimated. In this paper we address how accurate the current state of the art predictors are in providing a forecast of short term throughput. We propose a simple model for the short term prediction error based on Gaussian Random Walks that allows for mathematical analysis of the impact of imperfect future knowledge on network optimization.

(38) Koch C, Hausheer D. **"Optimizing Mobile Prefetching by Leveraging Usage Patterns and Social Information."***IEEE International Conference on Network Protocols (ICNP 2014); PhD Forum Paper; Raleigh (USA); October 2014.*

Abstract: Real-time entertainment constitutes the majority of traffic in today's mobile networks. The data volume is expected to increase in the near future, whereas the mobile bandwidth capacity is likely to increase significantly slower. Especially peak-hour traffic often leads to overloaded mobile networks and poor user experience. This increases costs for the mobile operator, which has to adapt to the peak demand by capacity overprovisioning. The new approach proposed in this paper aims to leverage the user's context and video meta-information to unleash the potential of video prefetching. Based on observed user interactions with social networks, the videos a user consumes from social neighbours can be predicted. Moreover, the user's daily routine even enables a prediction of the time when videos are consumed as well as the network capabilities available at that point. First results show that partial prefetching based on content categories provides a potential for efficiently offloading mobile networks. Additionally, the user experience can be improved as freezing playbacks of videos can be decreased. Initial results show a high potential for category-based prefetching.

(39) Rückert J, Blendin J, Hausheer D. **"Software-Defined Multicast for Over-the-Top and Overlay-based Live Streaming in ISP Networks".** *JNSM Special Issue on Management of Software Defined Networks, December 2014.*

Abstract: The increasing amount of over-the-top (OTT) live streams and the lack of global network layer multicast support poses challenges for a scalable and efficient streaming over the Internet. Content Delivery Networks (CDNs) help by delivering the streams to the edge of almost every Internet Service Provider (ISP) network of the world but usually also end there. From there on, the streams are to be delivered to the clients using IP unicast, although an IP multicast functionality would be desirable to reduce the load on CDN nodes, transit links, and the ISP infrastructure. IP multicast is usually not available due to missing control and management features of the protocol. Alternatively, Peer-to-Peer (P2P) mechanisms can be applied to extend the overlay multicast functionality of the CDN towards the clients. Unfortunately, P2P only improves the situation for the CDN but makes it more challenging for the ISP as even more unicast flows are generated between clients inside and outside the ISP network. To tackle this problem, a Software-Defined Networking-based cross-layer approach, called Software-Defined Multicast (SDM), is proposed in this paper, enabling ISPs to offer network layer multicast support for OTT and overlay-based live streaming as a service. SDM is specifically tailored towards the needs of P2P-based video stream delivery originating from outside the ISP network and can easily be integrated with existing streaming systems. Prototypical evaluations show significantly improved network layer transmission efficiencies when compared to other overlay streaming mechanisms, down to a level as low as for IP multicast, at linearly bounded costs.

(40) Koch C, Ruckert J, Bui N, Michelinakis F, Fioravantti G, Hausheer D, Widmer J.**"Demo: Mobile Social Prefetcher using Social and Network Information".** *The 19th IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (IEEE CAMAD 2014); Athens (Greece); December 2014.*

Abstract: The load on mobile networks increased over the past years and is predicted to further grow rapidly. Mobile network operators are facing new challenges to deliver mobile data in a user-satisfying way. This situation can be improved, e.g., by offloading traffic to WiFi networks or by shifting traffic to times where good network conditions can be leveraged, e.g. non-peak hours. One way to offload content is speculative prefetching. By knowing which contents will be requested in the future allows to prefetch it in advance. To this end, the users' content consumption has to be predicted. Furthermore, predicting properties of the smartphone's mobile connectivity, e.g., strong signal, allows a better user experience as well as an optimized network resource allocation. The reason for this is that the packet loss, retransmission and congestions can be avoided. To this end, the Mobile Social Prefetcher app aims at relieving the mobile network in a two-fold way. On the one hand, prefetching of promising videos from the users Online Social Networks is performed if WiFi is available. Videos posted on the user's OSN have been shown likely to be watched. On the other hand, if no prefetching opportunity can be used, a specific network optimized streaming is performed. This is provided by a video player which, both, reduces the load for the network operator and decreases stalling events of video playbacks. The proposed approach prefetches videos from video posts on the user's Facebook feed. Furthermore, the app considers current and previously observed cellular network information of the smartphone to optimize the mobile data throughput. This way, both, the operators' and the users' needs are reflected by the approach demonstrated.

(41) Han X, Wang L, Crespi N, Park S, Cuevas Á. "**Alike people, alike interests? Inferring interest similarity in online social networks".** *Elsevier Decision Support Systems; February 2015.*

Abstract: Understanding how much two individuals are alike in their interests (i.e., interest similarity) has become virtually essential for many applications and services in Online Social Networks (OSNs). Since users do not always explicitly elaborate their interests in OSNs like Facebook, how to determine users' interest similarity without fully knowing their interests is a practical problem. In this paper, we investigate how users' interest similarity relates to various social features (e.g. geographic distance); and accordingly infer whether the interests of two users are alike or unalike where one of the users' interests are unknown. Relying on a large Facebook dataset, which contains 479,048 users and 5,263,351 user-generated interests, we present comprehensive empirical studies and verify the homophily of interest similarity across three interest domains (movies, music and TV shows). The homophily reveals that people tend to exhibit more similar tastes if they have similar demographic information (e.g., age, location), or if they are friends. It also shows that the individuals with a higher interest entropy usually share more interests with others. Based on these results, we provide a practical prediction model under a real OSN environment. For a given user with no interest information, this model can select some individuals who not only exhibit many interests but also probably achieve high interest similarities with the given user. Eventually, we illustrate a use case to demonstrate that the proposed prediction model could facilitate decision-making for OSN applications and services.

(42) Truong P, Bertrand M, Peltier J. **"Networking impact of a local-aware content-based delivery for twitter-like applications".** *The 18th International Conference on Intelligence in Next Generation Networks (ICIN 2015); Paris (France); February 2015*

Abstract: Online Social Networks (OSN) have attracted a lot of end-users in the last few years. The success of Facebook, Twitter, Snapchat leads to a huge increase of data transiting in the network of telecom operators. To face to this increase, network operators have to find solutions. The Named Data Networking (NDN) paradigm, focusing on interest-based requests for content, is a good possible networking solution, having native caching features. Furthermore, it has been detected that a large number of people follow people who are in a close vicinity to them, except for very popular people. However, the current networking behaviour of the OSN applications do not take this into consideration. In this paper, we propose a locality-aware NDN based delivery for OSNs in order to improve the quality for end users (reduced delivery time) as well as to reduce the network load (and peering traffic) for network operators and the server load for the application provider. We modelled our solution for evaluating the networking impact and the tests we have performed proved the efficiency of our proposed solution, enabling to largely reduce the data delivery path as well as the peering traffic.

(43) Koch C, Bui N, Ruckert J, Fioravantti G, Michelinakis F, Wilk S, Widmer J, Hausheer D.**"Demo: Media Download Optimization through Prefetching and Resource Allocation in Mobile Networks".** *The 6th ACM Multimedia Systems 2015 Conference (MMSys 2015), Portland (USA); March 2015*

Abstract: Mobile network operators are expected to face significant traffic increase in the upcoming years. One alternative method is to intelligently move transmissions to times of network underutilization, either on 3G/4G or by offloading to WiFi. Video content, predicted by Cisco to constitute 69% of mobile traffic, offers the greatest potential for offloading. To this end, the demonstrated app strives to relieve the mobile network in a two ways. First, long-term prefetching of promising videos based on posts from the user's Online Social Network feed is performed. The knowledge about which video is likely being requested in the near future offers the opportunity to schedule the transmission according to its probability of being watched. Second, the approach is complemented with short-term prefetching, which is used whenever a specific content could not be

downloaded by long-term prefetching. In this case, resources are optimized so as to maximize the communication efficiency while preserving the quality of service. The demonstrated app considers the smartphone's observed cellular network history to optimize the mobile throughput. A customized video player implements both the long-term and short-term prefetching. It reduces both the load on mobile networks, decreases playback pausing events and hereby achieves a high QoE. Thus, the player addresses both the operators' and the users' needs.

(44) Nobach L, Hausheer D.**"Open Elastic Provisioning of Hardware Acceleration in NFV Environments".** *Workshop on Software-Defined Networking and Network Function Virtualization for Flexible Network Management (SDNFlex 2015);Cottbus (Germany); March 2015.*

Abstract: Network Functions Virtualization (NFV) is a new paradigm to move network tasks currently running on dedicated, vendor-specific hardware to elastic, virtualized environments, similar to IaaS cloud computing. A major challenge of NFV is to reach the performance known from dedicated hardware appliances, which often leverage ASIC, FPGA or NPU-based hardware acceleration to increase throughput and reduce delay. In this paper, a framework for elastic provisioning, programming and configuration of acceleration hardware for virtual network functions (VNFs) is proposed. Using this framework, VNFs can offload selected parts of their workload to heterogeneous acceleration processors, which may be shared among all VNF instances in the framework for improved resource utilization.

(45) Kaup F, Jomrich F, Hausheer D. **"Demonstration of NetworkCoverage – A Mobile Network Performance Measurement App".** *International Conference on Networked Systems (NetSys 2015); Cottbus (Germany); March 2015.*

Abstract: Optimizing the Quality of Experience (QoE) of mobile applications over cellular networks requires detailed knowledge of the underlying network and its performance. Parameters of interest are, besides the signal strength and availability of technologies, the Round-trip Time (RTT) and available throughput of individual cells at a given location. This information is generally not readily available. Therefore, an Android application measuring the cellular network performance was developed. This demonstration shows the NetworkCoverage App, being implemented to provide visual feedback of the measured network quality to the users, and sample the cellular network in an efficient manner to later support traffic scheduling improvements based on reliable data.

(46) Wilk S, Rückert J, Thräm T, Koch C, Effelsberg W,Hausheer D. **"The Potential of Social-aware Multimedia Prefetching on Mobile Devices".** *International Conference on Networked Systems (NetSys 2015); Cottbus (Germany); March 2015.*

Abstract: The access to Online Social Networks (OSNs) and to media content shared over these platforms accounts for a large part of today's mobile Internet traffic. Accessing media content and specifically videos is still challenging and costly over cellular networks. Mobile contracts usually have a data cap, the transmission quickly drains the battery of mobile devices, and the connection qualities can vary greatly, depending on the cellular network coverage. Prefetching mechanisms that fetch content items beforehand, in times when the mobile device is connected to a WiFi network, have a high potential to address all these problems. Yet, such a mechanism can only be effective if relevant content can be predicted with a high accuracy. Therefore, in this paper, an analysis of content properties and their potential for prediction are presented. An initial user study with 14 Facebook users running an app on their mobile device was conducted. The results show that video consumption is very diverse across the users. Nevertheless, a small number of attributes turned out

to be promising for the prediction. This work discusses the evaluation setup, the data analysis and describes their potential to define an effective prefetching algorithm.

(47) Bui N and Valentin S and Widmer J. **"Anticipatory Quality-Resource Allocation for Multi-User Mobile Video Streaming".** *The 2nd Workshop on Communication and Networking Techniques for Contemporary Video, in conjunction with the 34th IEEE International Conference on Computer Communications (IEEE INFOCOM 2015); Hong Kong(Hong Kong), April 2015.*

Abstract: Mobile video delivery forms the largest part of the traffic in cellular networks. Thus optimizing the resource allocation to satisfy a user's quality of experience is becoming paramount in modern communications. This paper belongs to the line of research known as anticipatory networking that makes use of prediction of wireless capacity to improve communication performance. In particular, we focus on the problem of optimal resource allocation for steady video delivery under maximum average quality constraints for multiple users. We formulate the problem as a piecewise linear program and provide a heuristic algorithm, which solution is close to optimal. Based on our formulation we are now able to trade off minimum video quality, average quality and offered network capacity.

(48) Michelinakis F, Bui N, Fioravantti G, Widmer J, Kaup F, Hausheer D. **"Lightweight Mobile Bandwidth Availability Measurement"**. *The 14th IFIP Networking 2015 Conference; Toulouse (France); May 2015*

Abstract: Mobile data traffic is increasing rapidly and wireless spectrum is becoming a more and more scarce resource. This makes it highly important to operate the mobile network efficiently. In this paper we are proposing a novel lightweight measurement technique that can be used as a basis for advanced resource optimization algorithms to be run on mobile phones. Our main idea leverages an original packet dispersion based, technique to estimate both per user capacity and asymptotic dispersion rate. This allows passive measurements using only existing mobile traffic. Our technique is able to efficiently filter outliers introduced by mobile network schedulers. In order to verify the feasibility of our measurement technique, we run a week-long measurement campaign spanning two cities in two countries, different radio technologies, and covering all times of the day. The campaign demonstrates that our technique is effective even if it is provided only with a small fraction of the exchanged packets of a flow. The only requirement for the input data is that it should consist of a few consecutive packets that are gathered periodically. This makes the measurement algorithm a good candidate for inclusion in OS libraries to allow for advanced resource optimization and application-level traffic scheduling, based on current and predicted future user capacity.

(49) Rückert J, Richerzhagen B, Lidanski E, Steinmetz R, Hausheer D. **"TopT: Supporting Flash Crowd Events in Hybrid Overlay-based Live Streaming".** *IFIP Networking 2015, Toulouse (France); May 2015.*

Abstract: Recent studies show that an increasing number of over-the-top live streams is delivered over the Internet. For the delivery of those streams, the dynamically changing and potentially large number of users imposes a major challenge. Flash crowds, where the number of users multiplies or significantly drops in a very small time frame, can cause serious degradations in the streaming performance. Due to the missing support for global network-layer multicast, overlay-based approaches have been broadly studied, showing that, with relaxed time constraints, they can scale well with the number of users. Yet, to support flash crowds, scaling has to happen quickly to keep up also with rapidly changing populations. Only a few approaches exist that focus on this aspect by

influencing the streaming topology and, so far, it is not clear if and how these mechanisms can be applied to state-of-the-art hybrid streaming systems. Therefore, in this paper, TOPT is proposed, integrating new as well as existing mechanisms in a common framework. The evaluation shows that the streaming topology, indeed, plays a major role during flash crowds. The lightweight and decentralized tree-forming and topology optimization mechanisms of TOPT, combined with tracker extensions to attach new peers in batches, greatly help improving the streaming performance in terms of reduced playback interruptions by more than 60% and slight reduction in communication overhead at an acceptable increase in average start-up delays by 24%.

2.6.2 Papers and Publications under review

Besides this list of accepted papers, there is also a number of papers submitted related to the project, for whom the presentation/acceptance day is after the end of the project. Such papers are listed in the table here below, followed by their abstract and detailed info.

Ref	*Conference / Journal*	*Title*	*Partners*
(Sub01)	Internet Teletraffic conference	Behind the NAT - A Measurement Study of Cellular Service Quality	IMDEA,TUD
(Sub02)	Internet Teletraffic conference	Anticipatory Admission Control and Resource Allocation for Media Streaming in Mobile Networks	IMDEA
(Sub03)	ACM Transactions on Intelligent Systems and Technology	Is Your Hidden Location Undercover? Predicting Current City from Profile and Social Relationship	TSP
(Sub04)	Elsevier Expert Systems with Applications	Community Similarity Degree: Finding Similarity to Improve Recommendations in On-line Social Networks	TSP, UC3M
(Sub05)	IEEE Communications Magazine	How Far is Facebook from Me? Facebook Network Infrastructure Analysis	TSP
(Sub06)	Communications of ACM	Characterization of Professional Users Strategies in major OSNs	TSP, UC3M
(Sub07)	PlosOne	Quantifying the Economic and Cultural Biases of Social Media through Trending Topics	UC3M
(Sub08)	ACM/IEEE Transactions on Networking	Assessing the Evolution of Google+ in its First Two Years	UC3M
(Sub09)	MASCOTS 2015	CPSys: A system for mobile video prefetching	Orange, INRIA, TUD

Table 4 – Papers and Publications Under Review

(Sub01) Kaup F, Michelinakis F, Bui N, Widmer J, Wac K, Hausheer D.**"Behind the NAT - A Measurement Study of Cellular Service Quality"**. *Internet Teletraffic conference.*

Abstract: Mobile applications such as VoIP, (live) gaming, or video streaming have diverse QoS requirements ranging from low delay to high throughput. The optimization of the network quality experienced by end-users requires detailed knowledge of the expected network performance. Yet, the achieved service quality is affected by a number of factors, including network operator and available technologies. However, most studies focusing on measuring the cellular network do not consider the performance implications of network configuration and management. To this end, this paper reports about an extensive data set of cellular network measurements, focused on analyzing root causes of mobile network performance variability. Measurements conducted in a 4G cellular network in Germany show that management and configuration decisions can indeed have a substantial impact on the performance. Specifically, it could be observed that the association of mobile devices to a Point of Presence (PoP) within the operator's network can influence the end-to-end RTT by a large extent. Given the collected data a model predicting the PoP assignment and its resulting RTT leveraging Markov Chain and machine learning approaches is developed. Overheads of 58% to 73% compared to the optimum performance have been observed in more than 57% of the measurements.

(Sub02) Bui N, Malanchini, I, Widmer J. **"Anticipatory Admission Control and Resource Allocation for Media Streaming in Mobile Networks"**. *Internet Teletraffic conference.*

Abstract: The exponential growth of media streaming traffic will have a strong impact on the bandwidth consumption of the future wireless infrastructure. One key challenge is to deliver services taking into account the stringent requirements of mobile video streaming, e.g., the users expected Quality-of-Service. Admission control and resource allocation can strongly benefit from the use of anticipatory information such as the prediction of future users demand and expected channel gain. In this paper, we use this information to formulate an optimal admission control scheme that maximizes the number of accepted users into the system with the constraint that not only the current but also the expected demand of all users must be satisfied. Together with the optimal set of accepted users, the optimal resource scheduling is derived. In order to have a solution that can be computed in a reasonable time, we propose an low complexity heuristic. Numerical results show the performance of the proposed scheme with respect to the state of the art.

(Sub03) Han X, Wang J, Wen A, Cuevas A, Crespi N. **"Is Your Hidden Location Undercover? Predicting Current City from Profile and Social Relationship".** *ACM Transactions on Intelligent Systems and Technology.*

Abstract: Privacy has become a major concern in Online Social Networks (OSNs) due to the threats such as advertising spam, online stalking and identity theft. To protect privacy, plenty of users hide or do not fill their privacy-sensitive attributes in OSNs. Existing studies try to infer users' hidden attributes through some other information exposed by the users themselves, which implies a potential disclosure of the hidden attributes. However, these studies do not quantify the exposure risk for a user based on his self-exposed information. Thus, an individual user still cannot understand the exposure probability of his privacy-sensitive attributes, let alone take the effective countermeasures. In this paper, we attempt to study the exposure probability of a user's hidden attributes via his self-exposed information, with a representative privacy-sensitive attribute - current city - in Facebook. To this end, we first design a novel current city prediction approach to disclose a user's hidden current city from his self-exposed information. Based on 371.913 users' information crawled from Facebook, we verify that our proposed prediction approach can predict a user's current

city more accurately than state-of-the-art approaches. Furthermore, relying on the proposed prediction approach, we model the exposure probability that a user's current city can be correctly predicted via some measurable characteristics of the self-exposed information. Eventually, we construct an exposure estimator to assess the current city exposure risk for an individual user, given his self-exposed information. Some case studies are illustrated to show how to use our proposed estimator to protect users' privacy.

(Sub04) Han X, Farahbakhsh R, Cuevas A, Cuevas R, Crespi N. **"Community Similarity Degree: Finding Similarity to Improve Recommendations in On-line Social Networks".** *Elsevier Expert System with Applications.*

Abstract: The irruption of OSNs in the Internet arena brings the possibility of accessing a huge number of diverse users' communities. This opens a new issue on how to select the best communities (among all the available options) to address a particular objective. Therefore, it is necessary to provide simple tools to perform such a selection in an easy way. In this paper we present the Community Similarity Degree (CSD) that is a metric to compute the degree of similarity among the users within a community. To evaluate the utility of our metric we rely on a dataset that includes more than 200K Facebook users. Using this dataset we define four different types of community: Friend-based (group of friends of a user), Interest-based (group of users sharing a common interest), Location- based (group of users from a city) and Random-based (group of users selected at random). We use the CSD to quantify users' similarity based on the interests they share within a community for five well-defined Facebook profile attributes: television, books, music, movies and games. Surprisingly, our results reveal that Interest-based communities are the ones showing a larger similarity degree, with a CSD between 1.5× and 4.5× larger than Friend-based communities. We use this outcome to demonstrate that communities with a larger similarity degree increase the efficiency of recommendation systems. We have emulated an OSN recommendation system in which an Interest-based recommendation strategy outperforms in 52% the efficiency shown by a Friend-based recommendation approach. This result demonstrates the practical usefulness of the proposed CSD metric.

(Sub05) Farahbakhsh R, Cuevas A, Ortiz A, Han X, Crespi N. **"How Far is Facebook from Me? Facebook Network Infrastructure Analysis".** *IEEE Communications Magazine.*

Abstract: Facebook (FB) is today the most popular social network with more than one billion subscribers worldwide. To provide good quality of service (e.g., low access delay) to their clients, FB relies on Akamai which provides a worldwide content distribution network with a large number of edge servers that are much closer to FB subscribers. In this paper we aim at depicting a global picture of the current FB network infrastructure deployment taking into account both native FB servers and Akamai nodes. Towards this end, we have performed a measurement based analysis during a period of two weeks using 463 PlanetLab nodes distributed across 41 different countries. Based on the obtained data we compare the average access delay that nodes in different countries experience accessing both native FB servers and Akamai nodes. In addition, we obtain a wide view of the deployment of Akamai nodes serving FB users worldwide. Finally, we analyze the geographical coverage of those nodes, and demonstrate that in most of the cases Akamai nodes located in a particular country not only service local FB subscribers, but also FB users located in nearby countries.

(Sub06) Farahbakhsh R, Cuevas A, Cuevas R, Crespi N. **"Characterization of Professional Users Strategies in major OSNs".** *Communications of ACM.*

Abstract: On-line Social Networks (OSNs) are being used intensively by countless professional players (e.g., large companies, politicians, athletes, celebrities, etc.) as a means of interacting with a huge amount of regular OSN users. In this paper we study the global strategy of professional users by sector (e.g., Cars companies, Clothing companies, Politician, etc.) across three most popular OSNs, Facebook, Twitter, and Google+.

(Sub07) Carrascosa J, Cuevas R, Gonzalez R, Azcorra A, Garcia D. **"Quantifying the Economic and Cultural Biases of Social Media through Trending Topics".** *PlosOne.*

Abstract: Online social media has recently irrupted as the last major venue for the propagation of news and cultural content, competing with traditional mass media and allowing citizens to access new sources of information. In this paper, we study collectively filtered news and popular content in Twitter, known as Trending Topics (TTs), to quantify the extent to which they show similar biases known for mass media. We use two datasets collected in 2013 and 2014, including more than 300.000 TTs from 62 countries. The existing patterns of leader-follower relationships among countries reveal systemic biases known for mass media: Countries concentrate their attention to small groups of other countries, generating a pattern of centralization in which TTs follow the gradient of wealth across countries. At the same time, we find subjective biases within language communities linked to the cultural similarity of countries, in which countries with closer cultures and shared languages tend to follow each others' TTs. Moreover, using a novel methodology based on the Google News service, we study the influence of mass media in TTs for four countries. We find that roughly half of the TTs in Twitter overlap with news reported by mass media, and that the rest of TTs are more likely to spread internationally within Twitter. Our results confirm that online social media have the power to independently spread content beyond mass media, but at the same time social media content follows economic incentives and is subject to cultural factors and language barriers.

(Sub08) Gonzalez R, Cuevas R, Motamedi R, Rejaie R, Cuevas A. **"Assessing the Evolution of Google+ in its First Two Years".** *ACM/IEEE Transactions on Networking.*

Abstract: In the era when Facebook and Twitter dominate the market for social media, Google has introduced Google+ (G+) and reported a significant growth in its size while others called it a ghost town. This begs the question that "whether G+ can really attract a significant number of connected and active users despite the dominance of Facebook and Twitter? This paper presents a detailed longitudinal characterization of G+ based on large scale measurements. We identify the main components of G+ structure, characterize the key feature of their users and their evolution over time. We then conduct detailed analysis on the evolution of connectivity and activity among users in the largest connected component (LCC) of G+ structure, and compare their characteristics with other major OSNs. We show that despite the dramatic growth in the size of G+, the relative size of LCC has been decreasing and its connectivity has become less clustered. While the aggregate user activity has gradually increased, only a very small fraction of users exhibit any type of activity and an even smaller fraction of these users attract any reaction. The identity of users with most followers and reactions reveal that most of them are related to high tech industry. To our knowledge, this study offers the most comprehensive characterization of G+ based on the largest collected data sets.

(Sub09) Gouta A, Hausheer D, Cuevas R, Kermarrec A, Koch C, Le Louédec Y, Rückert J. **"CPSys: A system for mobile video prefetching".** *MASCOTS 2015*

Abstract: Online media services are reshaping the way video content is watched. People with similar interests tend to request same content. This provides enormous potential to predict which content users are interested in. Besides, mobile devices are commonly used to watch videos which popularity is largely driven by its social success. In this paper, we design CPSys a Central Predictor System to prefetch relevant videos for each user. To fine tune our prefetching system, we rely on a large dataset collected from a large mobile carrier in Europe. The rationale of our prefetching strategy is first to form a graph and build implicit or explicit ties between similar users. On top of this graph, we propose the Most Popular and Most Recent (MPMR) policy to predict relevant videos for each user. We show that CPSys can achieve high performance with respect to the correct prediction ratio and by significantly reducing the traffic overhead. We further show that CPSys outperforms other prefetching schemes that have been presented and studied in the state of the art. At the end, we provide a proof-of-concept implementation of our prefetching system.

2.6.3 Participation to Conferences/Presentations

The Next table shows the list of accepted papers and publications. After the table, more information including abstract and detailed list of authors can be found.

Event	Location	Date	Title	Authors
Invited talk	Alcatel Lucent Bell Labs, US	March 11, 2013	Google+ or Google-*?: Dissecting the evolution of the new OSN in its first year*	Rubén Cuevas (UC3M)
Invited talk	AT&T Research Lab, US	March 12, 2013	Google+ or Google-? *: Dissecting the evolution of the new OSN in its first year*	Rubén Cuevas (UC3M)
Dagstuhl seminar on Future Internet [3]	Dagstuhl, Germany	March 24-27, 2013	Future Internet	Markus Hofman, Bell Labs
Dagstuhl Seminar on Future Internet [3]	Dagstuhl, Germany	March 24-27, 2013	Software-Defined Networking - Challenges and Opportunities	David Hausheer (TUD)
Talk	University of Cambridge, UK	March 2013	Large-Scale Data Processing	Eiko Yoneki (UCAM)
Invited Talk	University of Exeter, UK	May 2013	Empirical Approach for Modelling Dynamic Contact Networks	Eiko Yoneki (UCAM)
Invited Talk	ISI Torino, Italia	May 2013	Influential Neighbours Selection for Information Diffusion	Eiko Yoneki (UCAM)

Talk	MSR Cambridge, UK	May 2013	RASP: Large-Scale Graph Traversal with SSD Prefetching	Eiko Yoneki (UCAM)
Invited Talk	SIGMOD, USA	Jun 2013	Scale-up Graph Processing: A Storage-centric View	Eiko Yoneki (UCAM)
AIMS 2013 Conference (Tutorial)	Barcelona, Spain	June 2013	EmanicsLab: A European Research Network tailored to Network and Service Management	David Hausheer (TUD)
Invited Talk	Graphlab workshop, USA	July 2013	Scale-up Graph Processing: A Storage-centric View	Eiko Yoneki (UCAM)
Future Networks and Mobile Summit event (http://www.futurenetworksummit.eu/2013/)	Lisboa, Portugal	July 03-05, 2013	Participation of eCOUSIN to the poster session ofthe conference	All consortium partners
Invited Talk	University of Saskatchewan, Canada	August 2013	Empirical Approach for Modelling Dynamic Contact Networks	Eiko Yoneki (UCAM)
Invited Talk	University Massachusets Amherst, USA	August 8, 2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	Boston University, USA	August 9, 2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited talk	McGill University, Montreal, Canada	August 18, 2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	INRS, Montreal, Canada	September 4, 2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	University of Ottawa, Canada	September 6, 2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	University of Nottingham, UK	September 2013	Fluphone to Epi-Pi: Towards Quantifying Contact Networks for	Eiko Yoneki (UCAM)

			Digital Epidemiology	
Invited Talks	Northeaster University	October 9, 2013	Google+ or Google-? Characterizing User Engagement in Major OSNs	Ruben Cuevas (UC3M)
Invited Talks	Boston University	October 10, 2013	Are Trending Topics Useful for Marketing? Visibility of Trending Topics vs. Traditional Advertisements	Rubén Cuevas (UC3M)
eCOUSIN 2-day Workshop with eCOUSIN's Advisory Board, collocated with IMC 2013	Barcelona, Spain	October 21-22, 2013	eCOUSIN 2-day Workshop with eCOUSIN's Advisory Board, collocated with IMC 2013	All consortium partners
SOSP Poster 2013	Farmington, USA	November, 3/6, 2013	SAKYOMI: SSD Prefetcher for Large-Scale Graph Traversal	Eiko Yoneki (UCAM)
Future Internet Assembly (FIA)	Athens, Greece	March 18/20, 2014	Presented Poster:"eCOUSIN: enhanced COntent distribUtion with Social INformation"	F. Mondin, C. Venezia (TI)
Workshop "Digital Word, Large Scale and Complexity" at Institut Mines-Télécom (TSP)	Paris, France	March 26/27, 2014	Presented Poster:"enhanced COntent distribUtion with Social Information"	R. Farahbakhsh (TSP)
Workshop "Digital Word, Large Scale and Complexity" at Institut Mines-Télécom (TSP)	Paris, France	March 26/27, 2014	Presented Poster:"Understanding the Evolution of Multimedia Content in the Internet through BitTorrent glasses"	R. Farahbakhsh, A. Cuevas, R. Gonzalez, N. Crespi (TSP), R. Cuevas (UC3M)
Organizer and co-chair of the Workshop on Name Oriented Mobility (NOM) at Infocom 2014	Toronto, Canada	April 28, 2014	Workshop on Name Oriented Mobility (NOM)	B. Mathieu (Orange)
European Conference on Networks and Communications	Bologna, Italia	June23/26, 2014	Prepared Poster "eCOUSIN: enhanced COntent distribUtion	F. Mondin, C. Venezia (TI), Y. Le

(EuCNC'2014)			with Social INformation"	Louédec (Orange)
European Conference on Network and Communications (EuCNC'2015)	Paris, France	June 29, 2015	Prepared Show case of Project results	All consortium Partners

Table 5. Conferences and Presentations

2.6.3.1 Focus on EUCNC 2015

The eCOUSIN consortium proposed and got accepted a 12 square meter exhibition booth at EuCNC 2015, the European Conference on Networks and Communications. Being very close to the end of the project, we could propose to show all the main results of eCOUSIN during the event. In particular our exhibition includes examples of all of the four demonstrators; and three of them are shown live at the event, while the remaining one is broadcasted from the testbed location.

In addition, all the partners contribute to the booth by providing posters, flyers, presentations and videos. We expect our exhibition to have a great impact on the audience thanks to the interactive demos and eye catching visuals. Finally, it is going to be a unique occasions to showcase all of the project results in a single venue.

2.6.4 Collaboration Activities and Workshops

The Next table shows the list of collaboration activities, workshops and any other dissemination action not falling in the previous sections.

Collaboration activity	Place	Date
Workshop: "Optimization of Network Resources for Content Access and Delivery" with FP7 ENVISION, OCEAN, ALICANTE, ETICS, FUSION and eCOUSIN projects (http://www.envision-project.org/workshops/oncad12/index.html)	Lannion, France	September 6, 2012
Participation to the EC's plenary Concertation meeting on October 11th 2012 in Brussels with a presentation of the eCOUSIN project at the agenda.	Brussels, Belgium	October 11, 2012
Project Flyer for the project website and for the European Commission brochure on Future Internet Cluster [1].	NA	February 2013
Participation to the 11th FP7 Future Networks Concertation Meeting	Brussels	February 27-28, 2013
Participation of the consortium to the call for the Fifth Future Internet Award (http://www.fi-dublin.eu/fi-award/)	NA	April 1, 2013
Submission of the proposal "Social Aware Smart Content Delivery Networks" at EIT ICT Labs Call 2014	NA	May 15, 2013
Publication "Large scale analysis of HTTP adaptive streaming in mobile networks" (Ali Gouta, Yannick Le Louédec (Orange)) in the issue of July 2013 of COMSOC's MMTC E-Letter (Multimedia Communications Technical Committee of IEEE COMSOC, http://committees.comsoc.org/mmc/eletters.asp)	NA	June 10, 2013
Joint SmartenIT/eCOUSIN Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM), co-located with CNSM 2013 (http://www.cnsm-conf.org/2013/workshops.html). Organized by TUD and University of Zurich (Burkhard Stiller, Coordinator of SmartenIT)	Zurich, Switzerland	October 18, 2013
Participation to the 12th FP7 Future Networks Concertation Meeting Participation of eCOUSIN with a presentation entitled "Mobile Traffic Scheduling Architecture Improving Energy Efficiency and QoE" (prepared by TUD and presented by UC3M) to the "Green and Energy-efficient Networking" Workshop at the Future Internet Cluster meeting on October 22, 2013.	Brussels	October 22-23, 2013
Joint meeting between the EU FP7 eCOUSIN and Social Sensor projects to identify common interests	Audobridge conference	February 13, 2014

eCOUSIN Workshop 2014, co-located with the Conference on Online Social Networks (COSN 2014)	Dublin, Ireland	September 29-30, 2014
Participation to the 12th FP7 Future Networks Concertation Meeting Participation to the Future Internet Cluster Position paper on Smart Networks and Novel Architectures. Dissemination of an updated version of eCOUSIN factsheet.	Brussels	October 23, 2014
Participation to the 12th FP7 Future Networks Concertation Meeting Participation with a presentation of eCOUSIN achievements as finishing project, and a presentation from eCOUSIN partners for the open discussion on "SDN/NFV and Network Apps – a view from the cluster".	Brussels	March 25, 2015
The software tools and data traces made available in their public version in the frame of deliverable D3.2 are considered to be used in other EU projects, e.g. EU FP7 SmartenIT.	NA	NA

Table 6. Collaboration Activities

2.6.4.1 Focus on the Joint SmartenIT/eCOUSIN Workshop held on October 18, 2013

The CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications “SETM 2013” aims to bridge a gap, by addressing social-aware and incentive-compatible network management mechanisms in support of overlay and cloud applications in an integrated and thus efficient manner. SETM 2013, which was organized by David Hausheer and Tobias Hossfeld as a joint workshop between SmartenIT and eCOUSIN in conjunction with this year’s CNSM 2013 conference, is a continuation of three successful workshops entitled “Workshop on Economic Traffic Management (ETM)”. The last event organized in 2010 was collocated with ITC22. The workshop title has been changed to better reflect the scope of the workshop that basing on economic perspectives in network management extends its scope towards social awareness, cloud computing, content awareness and CDNs. That way the workshop is following current trends in that area.

Following these intentions, SETM 2013 provided a single-track and one-day program, including keynotes, three technical sessions, and a poster session. The keynotes were given by eCOUSIN Advisory Board member Bruce Maggs (Professor, Department of Computer Science, Duke University and Vice President, Research Akamai Technologies, USA) and Volker Hilt (Head of Networked Services Research, Alcatel-Lucent Bell Labs, Germany). The SETM 2013 workshop was financially sponsored by Orange and Intracom Telecom. Overall, a total of 15 papers were registered. Out of those, 14 papers were finally submitted and one paper was withdrawn. Before the papers entered the review process, the TPC chairs fast-rejected four papers, including plagiarized papers and papers which were completely out of scope.

Each of the remaining 10 papers received in total 3-4 reviews. Based on the thorough reviews, seven papers were finally selected for the workshop program, with two papers being accepted as full papers and 5 papers being accepted as short papers, which are presented during SETM 2013.

Additionally, the poster session that is organized jointly with the SVM workshop includes seven invited posters from SETM.

With respect to the technical sessions, one focus lies on video services in the Internet responsible for the majority of Internet traffic, as discussed in Session 1. The proposed approaches for improving video delivery ranges from concrete algorithmic solutions (*Multi-source Cooperative Adaptation for QoE-aware Video Multicast Rate control*) to social-aware concepts for recommendation system and content distribution networks (*Determining leaders and clusters in video consumption*), while also economic aspects and business interests are analyzed (*Video Delivery over Next Generation Cellular Networks*). Network traffic management solutions are considered in Session 2 with a focus on the federation of clouds and available bandwidth provisioning systems for offering reliable network transfer services to the cloud systems (*Networking solutions in the federation of clouds*) and the problem of cost-optimal distribution of inevitable inter-domain traffic on multiple links dynamic solutions (*Dynamic Traffic Management mechanism for active optimization of ISP costs*). Session 3 considers user involved approaches to realize a socially-aware traffic management solution which targets three popular use cases: data offloading to WiFi, content caching/prefetching, and content delivery (*HORST - Home Router Sharing based on Trust*) but also to optimize energy and upload bandwidth of smartphones by proper incentive and collaboration schemes (*Reciprocity with Virtual Nodes: Supporting Mobile Peers in Peer-to-Peer Content Distribution*).

3. EXPLOITATION

The technology developed in the framework of the eCOUSIN project is a key component in the strategy of the industrial partners since it will allow them to evolve current Content Distribution Network architectures based on Social Networks information exploitation. The resulting new technological solutions will be exploited for:

- CDN architecture product evolution roadmap
- Trialling social networks-aware CDN technology enhancements
- Development of new networking hardware functionality
- Product customization and differentiation
- Trialling of new services
- Developing new services.

The prototypes developed in the project will provide a technological edge to the partners involved which, based on the know-how acquired, will be able to contribute to a successful technology transfer for product development activities.

In the following, an indicative exploitation plan per partner is provided:

3.1 Orange

ORANGE is both a network service provider and a service provider.

As a network provider ORANGE is coming up against major issues in carrying Internet traffic to its end-users. This traffic is amazingly picking up, mainly driven by the growing success of social networks and audiovisual applications over Internet. As end-user pricing are flat rate and IP transit prices are declining slowly, both cannot compensate this traffic growth. ORANGE needs technical solutions which ease the pressure of this traffic, reduce the network cost or at least differ investment, while improving the quality of service for end-users. Besides, in a wider and longer term perspective, one may observe that the transfer of content has become predominant in the Internet- rather than the connection of end-points -, while the infrastructure is still based on connections. The foreseen incompatibility between the infrastructure and expected evolution of usages calls for breakthrough solutions. Social Networks being destined to be one, if not THE, major Internet application, ORANGE contributed in eCOUSIN to investigate the relationship between OSN applications and networking delivery systems, including ICN. eCOUSIN project has thus been a great opportunity to contribute to searches on the design of the Future Internet, in the continuation of previous projects, e.g., FP7 projects 4WARD and SAIL, as well as French ANR Connect project. The outcomes from eCOUSIN are already exploited internally inside Orange for studies and searches to define the requirements and specifications for the naming, addressing and content placement in future 5G and convergent networks.

As a service provider, ORANGE aims to leverage the outputs of eCOUSIN to provide the best services with best quality of experience to its end users. At the time of the project launch, ORANGE was expanding into the social networks and sharing applications markets. One typical illustration of this was the participation of ORANGE in Dailymotion. The outputs of eCOUSIN, especially the monitoring and modelling tools from WP3 (content consumption patterns, user behaviour dynamics) and the social aware network systems from WP4 (content placement, look-up and delivery), aimed at being used to prepare the evolution of the recommendation system and the social network features in the portal.But Orange strategy has evolved during the project lifetime, illustrated by the on-going negotiations in year 2015 to withdraw from Dailymotion. Yet the initial objective remains the same for Orange Labs: leverage the outputs of eCOUSIN to provide the best services with best quality of

experience to its end users. In this vein the contributors from Orange have been strongly active in the last year of the project to foster the dissemination and exploitation of eCOUSIN's outcomes towards Orange operational units and services, with a specific focus on Use Case 3: mobile content prefetching. The standard dissemination and exploitation process of innovations inside Orange Group consists in going through business analysis on the one hand and through demonstrations of prototype implementations on the other hand. Both actions are still on-going inside Orange at the time this report is produced. As an illustration, the developments achieved on use case 3 have been integrated in a platform from Orange that covers a larger spectrum of applications (including content delivery but also conversational services and email), and whose purpose is to provide seamless continuity services on any network. The demonstration of this platform was selected for the last yearly research and development show of Orange on December 2014.

3.2 Telecom Italia

Telecom Italia expects eCOUSIN results will provide a valid foundation for novel highly distributed Social Networking scenarios where the central aggregators won't play a crucial role. Telecom Italia is investigating mechanisms to enable customers to share their multimedia contents directly to and from their devices. This will go through the definition of intelligent content diffusion and discovery mechanisms. Another important feature which will be undertaken is the mechanism for empowering social content sharing by means of context awareness. In general, eCOUSIN findings in the opportunities of leveraging on network infrastructures will improve the effectiveness of this kind of services.

Concretely, the objective is to exploit the technical results achieved by eCOUSIN to extend Telecom Italia's commercial product ranges. Telecom Italia actually wants to be an operator, not just a carrier. This is the reason why TI's marketing is pushing both on cloud services and on media-centre devices. The "CuboVision" device is already on the market, working as an UPnP decoder/media centre, as well as a Smart Hard Disk from IoMega in partnership with Telecom Italia. Telecom Italia intents to implement eCOUSIN use cases on such devices and to flash out the value added services offered to its end users, the ultimate goal being to stimulate them to use the bandwidth sold by Telecom Italia as core business, in particular to encourage them to move to fibre connections. The target market size is being analysed in the frame of Work Package 2.

As regards the Personal Content Sharing use case, the dissemination and exploitation plan consists in identifying a set of ultra-broadband subscribers interested in trialling the service. These users actually won't have to make new subscriptions to any specific service but just installing the eCOUSIN package on their PCs or Network Access Server and start sharing content. They will identify potential friends or acquaintances that are already connected by their preferred social networks and invite them to install the eCOUSIN client. The installation is straightforward and a user-friendly user interface will drive them through the sharing process. The service enables to share private contents without having to upload them anywhere. Of course the more uplink bandwidth available the more the user experience will be pleasant up to be able to access friend's contents via standard UPnP devices like if they were stored locally.

Telecom Italia believes that this could lead some customers to evaluate and finally perform the Internet connectivity upgrade to Fibre. Providing end users with improved connectivity without introducing new services that truly benefit from such improved connectivity wouldn't be an effective way to market it, in so far customers generally expect to perceive an added value.

In the very last few days Telecom Italia started pushing commercially a brand new set top box called TIM-VISION, which strongly leverages on fibre to provide video content of high quality to the users.

As a proof of this, Telecom Italia just signed a deal with the main TV satellite provider in Italy, so to offer interesting first quality content. The Tim Vision Set top box is then evolving day by day and when ready to be a product the personal sharing clouds application will be one of the main candidates for inclusion as an app.

3.3 Alcatel-Lucent

Alcatel-Lucent's solutions and product portfolio enable carriers, Internet service providers, enterprises and governments worldwide, to deliver voice, video, multimedia, and data communication services to end-users. Alcatel-Lucent is a market leader in broadband access systems and triple-play solutions.

Alcatel-Lucent recognizes that the upcoming wave of Internet applications where content will be discussed on online social networks will have a huge effect on the content distribution patterns. This will very likely put a high strain on all segments of the network (operated by Alcatel-Lucent's customers). The results of the eCOUSIN project will be extremely useful to decide which features will be required in Alcatel-Lucent's new products and in future evolutions of its current products, with which the operators can ensure the end-user's quality of experience while at the same time allowing operators to exploit their network efficiently.

Bell Labs, the research organization of Alcatel-Lucent, mainly studies the technical aspects of the innovative eCOUSIN architecture and collaborates with the eCOUSIN partners to assess and enhance the performance of the content delivery via online social networks. Bell Labs will continue to disseminate novel research results at leading international conferences (e.g., ITC, Globecom, Infocom), has contributed to ensuring the openness of the model by proposing contributions (jointly with eCOUSIN partner FT) to relevant standardization bodies (such as IRTF) and will proceed to disseminate project results internally to ensure that Alcatel-Lucent's business units incorporate the right features into the roadmaps of their products. Project results are expected to be of particular interest to improve content delivery efficiency in ALU's Velocix transparent cache and content delivery platform and ALU's CloudBand solution:

- The Velocix Transparent Cache allows network service providers to intercept Internet traffic and cache content from popular websites to reduce transport, peering and transit expenses while improving the user experience. The eCOUSIN results related to caching and pre-fetching may help to further improve the caching algorithm of this transparent cache.
- The Velocix Digital Media Delivery Platform comprises a range of appliances and services that allows Network Service Providers to build and operate their own content delivery platform for rich digital media taking into account subscriber data to deliver more personalized, more dynamic and more equitable video service experience to TVs, tablets, smartphones, game consoles and PCs. The eCOUSIN results related to data mining of social relations may prove helpful in making the subscriber data even more relevant.
- The CloudBand solution enables service providers to accelerate adoption of Network Function Virtualization (NFV), providing a fully integrated solution that orchestrates infrastructure, applications, and network in a single virtualized platform. It consists of two building blocks, i.e., the CloudBand Management System and the CloudBand Node. The former orchestrates, automates, and optimizes virtual network functions across the service provider's distributed network and data centres, providing a coherent view of the entire NFV infrastructure as a single carrier grade pool. The latter is an all-in-one compute, storage, and network node system including hardware and software designed for efficient remote operation of distributed clouds. Some of the use cases studied in eCOUSIN could benefit from this flexible architecture to be rapidly deployed.

3.4 Foundation IMDEA Networks

IMDEA Networks was created to help reducing the gap between academia and industry. The outcome of eCOUSIN will result in the following benefits for IMDEA Networks: First, it represents an opportunity to increase IMDEA's Intellectual Property assets, helping to supplement their income from the public and private sector, while at the same time allowing them to perform technology transfer to industry. Second, it will provide the Ph.D. students that work at the Institute the knowledge of innovative solutions and technologies that will prepare them to work in technologically advanced jobs. Third, it will help maintain and even increase IMDEA Network's technological leadership as a research institution in Content Distribution and new Social-inspired technologies, and fourth, it will create opportunities for new start-up companies that spin off from IMDEA Networks.

3.5 Technical University Darmstadt

As an academic partner TUD's primary focus is to disseminate the results developed in eCOUSIN to the scientific community by publishing them in highly ranked journals and conferences which show a close relationship to the planned research. Additionally, TUD will exploit the outcomes of the project in lectures and seminars held at TUD, which will benefit directly with first-hand research results and interesting new seminar topics. In addition, a number of semester-, diploma-, and master theses with topics related to the planned research will benefit from the project. Finally, 1-2 doctoral theses have performed directly on the project and have, thus, highly benefited from the project's outcome, research results, and publications.

3.6 Institut Telecom SudParis

During the first year of the project, TSP included two seminars in 2012/2013 academic year covering two topics close related to eCOUSIN project: On-line Social networks (OSNs) and P2P networks. The first one covers fundamental aspects of OSNs as well as more used techniques to retrieve data, which mostly relates to the activities carried out in WP3. The second seminar covers the functionality and data collection techniques of one of the system distributing a large volume of content over the world. This seminar was schedule in the MSc of Communications Network and Services (COMSNETS) and taught but members of TSP actively participating in eCOUSIN project.

Since beginning of year 2014, TSP and UC3M have been managing their exploitation plans in a frame of a close collaboration, which include three main types of actions:

1- Education: UC3M has added into the two subjects of the Master in Telematics Engineering (Internet Measurements and Content Distribution Architectures) teaching material associated to the different technologies designed and developed in eCOUSIN and plan to extend this material with the results obtained during the third year. In addition, UC3M and TSP have started in the academic year 2014/2015 a MSc Dual-Degree Program in the area of Internet Communications and Services that includes. This Dual-Degree involves the UC3M MSc in Telematics Engineering referred above and the TSP MSc in Communication Networks & Services. Both masters will incorporate teaching material associated to the technologies developed in eCOUSIN.

2- Commercial Exploitation of the OSN measurement tools: UC3M and TSP are in contact with several companies that are interested in obtaining different social media data analysis for marketing or public image purposes. The OSN measurement tools designed and

implemented in eCOUSIN as well as the know-how on OSN data processing and analysis acquired during the project will be fundamental for such exploitation activity.

3- Bring the developed know-how for OSN measurements tool implementation into new EU proposals within the H2020 program: UC3M and TSP are using their know-how on developing measurement tools in the OSN domain as a valuable contribution to new EU proposals in the context of the H2020 program. This know how has been included in 3 H2020 EU that are currently under evaluation.

3.7 University Carlos III of Madrid

During the first year of the project, the master of Telematic Engineering at UC3M included the following subjects in the 2012/2013 academic year in which the technologies and topics covered by the eCOUSIN project are taught: “Architecture of the Internet”, “Content Delivery Networks” and “Internet Measurements”. The two former subjects are related to WP4 since they discuss the current Architecture of Internet and Content Distribution infrastructures. Furthermore, they present the research conducted in this area including specifically the solutions under consideration in eCOUSIN. The latter subject covers (among others) the measurement techniques developed in the WP3 of eCOUSIN. For more details on the Master of Telematic Engineering we refer the reader to the corresponding Web site [4].

Since beginning of year 2014, TSP and UC3M have been managing their exploitation plans in a frame of a close collaboration, which include three main types of actions, reported in Section 3.7

4. STANDARDISATION

This section focuses in the early achievements of eCOUSIN in the standardization area. eCOUSIN has two objectives here: exploit the latest developments from the standardisation bodies that are relevant for eCOUSIN project, and contribute actively to standards to promote eCOUSIN's technical outcomes. The Social aspects are the first to be targeted in the standardization context; and as the technical activity of eCOUSIN matures the consortium extends and defines a larger list of feasible targets.

4.1 Participation to Joint Opensocial+W3C Workshop

A project representative had the possibility to participate to the Opensocial+W3C Joint Workshop about "The future of business", which was held in San Francisco (CA, USA) on 2013, August, 8th and 9th.

Telecom Italia tried to take on the eCOUSIN topics and points of view to W3C which was the actual organizer of the meeting by participating actively in the program committee.

The eCOUSIN project was invited to write and present a position paper which had its own slot and was discussed in a "panel-like" session by all the audience together with other three paper close in terms of subjects.

W3C is the place where the web standards are developed and decided while opensocial on the other side is taking on the view of distributed network architecture, with focus on mobile applications and their use of network resources. Even if the project did not have the possibility to actively participate in writing down standards specification in W3C, with activities such as writing down RFCs, taking the project views in terms of efficient use of network resources can really make the difference.

Standards should be aware of network efficiency issues, in order to keep in mind that both on mobile and fixed networks the new social orientation of content distribution is an issue that must be kept in mind. By showing the project's use cases and solutions to W3C audience we pointed out the attention on the development of a new social-aware standard.

After the meeting a new technical working group opensocial+w3c was created and we plan to keep an eye on that.

On the other hand, we had the possibility to meet people and get feedbacks from a group of the most active web experts in the web technologies panorama. The discussion panel in which the position paper was placed was held together with another position paper about Open Mobile Alliance's Federated Social Network Enabler (named SNEW, for Social NEtwork Web). The standardisation activity in OMA is being taken on by Telecom Italia and other partners in OMA and this technology will be included in the eCOUSIN's Personal Content Sharing Cloud use case.

Moreover, we met people from other EU funded projects, such as the openI project, with whom we started some informal talks in order to identify synergies between the projects, evaluate collaborations, joint sessions etc.

The project plans in terms of standardisation is to keep the standardisation fora as aware as possible of the project concerns, which we try to address by means of the project's use cases.

4.2 Contribution to W3C Social WG (TI)

The exploitation of social data is a critical foundation of eCOUSIN's architectural framework. But today's social APIs and protocols do not allow an easy transfer of social data between existing social networks. Moreover the lack of standard APIs makes it difficult for Web application developers to embed social functionality from third-party sites into their apps. Therefore it is very important to foster and contribute to standardization in this area.

eCOUSIN has been part of the Social Workshop reported in the above chapter, and beyond directly contributed to the efforts that have finally led to the creation of W3C working groups in this area.

eCOUSIN planned a contribution to the W3C Social Interest Group having as objective to develop relevant use cases, a unified social architecture, and the associated terminology.

eCOUSIN planned also to contribute to the newly created Social Web Working Group to standardize:

- an extensible transfer syntax for activities like status updates, with the ActivityStreams 2.0 data-format as a possible starting point
- an API for third-party social content embedding, with the OpenSocial 2.5.1 Activity Streams and Embedded Experiences APIs Member Submission being as a possible starting point
- a protocol for social network federation, with the Web Mention protocol and the Linked Data Platform as possible starting points.

The start-up phase of the two working group was a bit slow due to some issues related to previous patent which actually blocked the advances for a while. After this slow start, the two working group finally started their activity. In these terms, as anticipated in deliverable D7.2, eCOUSIN was present at the W3C Technical Plenary + Advisory Committee (TPAC), held in Santa Clara in November 2014.

The Project itself was presented as a use case to take into account in setting up the social standards, with a specific focus on the Personal Sharing Clouds use case. This was an active contribution to the working groups by the project, which previously pushed in order to build up those working groups.

In the very last period of the project the eCOUSIN project pushed some requirements in the W3C Social Working Group Related with specific reference to the forms ActivityPump, Micropub, SoLID, ActivityStreams 2.0 and Vocabularies

4.3 Contribution to IRTF ICNRG

As suggested during the mid-term review by the European Commission, the project had an active contribution to IRTF-ICNRG standardisation group.

IRTF ICN RG aims at addressing the research challenges on ICN, which include:

- Naming schemes for ICN, including scalable name resolution for flat names
- Scalable routing schemes
- Congestion control, QoS approaches, and caching strategies
- Metrics that make it possible to evaluate ICN implementations in a consistent manner
- Security and privacy, including scoping of information objects and access control to them

The objective of eCOUSIN was to disseminate towards ICN RG the outcomes of the works achieved on eCOUSIN's Use Case 6, "Information-Centric and Social-Driven Content Delivery".

The following Internet draft, result of the joint effort of Alcatel Lucent and Orange, has been submitted to the IRTF ICNRG Research Group, and has been presented during IETF92 meeting in Dallas, Texas on March 24, 2015:

Truong P, Mathieu B, Stephan E, K. Satzke.**"Named data networking for social network content delivery".** *[draft-truong-icnrg-ndn-osn-00.txt].*

Abstract:

Online Social Networking (OSN) applications have attracted millions of people over the last few years. Their traffic represents a large part of the traffic of the Internet. For instance, Facebook represents near 25 percent of the Internet traffic [5][6], and a part of this traffic is exchanged amongst groups of end-users which are located in the same geographic area. In this document, we introduce a Named Data Networking (NDN) architecture to improve the delivery of OSNs contents requested by end-users in the neighbourhood of the publishers: Having the knowledge of the social network graph and the end-users network location, a SDN-based NDN controller dynamically configures the NDN routers to route the interest requests directly between the end-users.

5. CONCLUSION

This final report contains all the dissemination and standardisation activities achieved by the consortium during the project lifetime, as well as an updated version of each partner's exploitation plan.

The consortium promoted the project views and results by means of a considerable dissemination effort. The total number of dissemination actions is close to 100 just taking into account papers, publications, workshops and presentation related to the scientific community. Moreover, a set of papers was promoted over traditional and specific media. More than 10 specific workshops were organized to communicate on the project achievements and to exchange results and ideas with third parties. A final dissemination activity involving all the use-cases is planned for EUCNC 2015 in Paris on June 2015, just a few days before the final review. All the use cases will be showcased at this event by means of posters and live demos.

The standardisation effort touched three standard bodies mainly: W3C, IRTF ICNRG and also, even if less directly, OMA. In W3C the use case was brought in order to have a direct influence from the project to the standard. And a tangible contribution was produced for IRTF ICN RG.

Finally each partner has a credible exploitation plan for the project outcomes, with real exploitation perspectives at short time, such as on the personal content sharing cloud use case. Besides Deliverable D2.3 provide complementary insights on exploitation perspectives, by presenting business plans for the selected use cases and the expected benefits for the main concerned stakeholders.

REFERENCES

[1] http://www.fines-cluster.eu/fines/jm/Publications/Download-document/409-FInES-Horizon-2020_Position-Paper-v2.0_final.html

[2] http://www.ict-ecousin.eu

[3] http://www.dagstuhl.de/mat/index.en.phtml?13131

[4] http://www.uc3m.es/portal/page/portal/postgraduate_studies/masters/Master_in_Telematic_Engineering/Program