

Deliverable D7.2

Intermediate Dissemination, Exploitation and Standardization, Report

Public deliverable, Version 1.0, 31 October 2013

Authors

FT Yannick Le Louédec
AL-BELL Danny De Vleeschauer, Chris Hawinkel
IMDEA Joerg Widmer
TSP Ángel Cuevas
ALUD Ivica Rimac, Klaus Satzke
TUD David Hausheer
TI Claudio Venezia (Editor), Fabio Mondin
UCAM Eiko Yoneki
UC3M Ruben Cuevas

Reviewers Luc Le Beller, Yannick Le Louédec (FT)

Abstract

Deliverable D7.2 is the intermediate dissemination, exploitation and standardization report of the eCOUSIN project.

Deliverable D7.2 presents the project dissemination strategy, including the dissemination target audience groups, the actions and channels to reach these groups during the project life, and the first opportunities for external collaborations. Then it provides a preliminary identification of the expected results exploitable by the consortium partners, and the project's initial approach concerning the possible use of and contributions to standards. Finally it reports the dissemination results obtained after one year.

© Copyright 2012 eCOUSIN Consortium



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

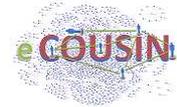
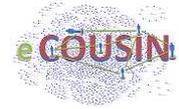


TABLE OF CONTENTS

TABLE OF CONTENTS	2
EXECUTIVE SUMMARY.....	3
1. INTRODUCTION	4
1.1 eCOUSIN background and objectives	4
1.1.1 Context and focus of research.....	4
1.1.2 Objectives and expected results.....	4
1.2 WP7 Scope.....	4
1.3 Scope of Deliverable D7.2.....	5
2. DISSEMINATION STRATEGY	6
2.1 eCOUSIN Web presence	7
2.1.1 Presentation of the project on the Web site.....	8
2.2 Flyers and brochures	8
2.3 Public Project Deliverables	9
2.4 Dissemination of the Scientific and Technical achievements.....	10
2.5 Events and external liaisons	10
2.6 Dissemination actions achieved during the first project period.....	11
3. EXPLOITATION	29
3.1 France Telecom / Orange (FT)	29
3.2 Telecom Italia (TI)	30
3.3 Alcatel-Lucent (A-LBELL and ALUD)	30
3.4 Foundation IMDEA Networks (IMDEA).....	30
3.5 Cambridge University (UCAM).....	31
3.6 Technical University Darmstadt (TUD)	31
3.7 Institut Telecom SudParis (TSP).....	31
3.8 University Carlos III of Madrid (UC3M).....	32
4. STANDARDISATION	33
4.1 Participation to Joint Opensocial+W3C Workshop.....	33
5. CONCLUSION	34
REFERENCES.....	35



EXECUTIVE SUMMARY

Deliverable D7.2 is the first dissemination report of the eCOUSIN project. It provides the initial plan for eCOUSIN dissemination, exploitation, and standardisation activities to be achieved during the project life, as well as the dissemination results obtained after one year.

The objective of this document, which is structured in three main sections, is to describe the planned exploitation, dissemination and standardization activities of the eCOUSIN consortium partners. Besides being advertised by means of many dissemination channels, work in eCOUSIN is bound to be directly useful to the partners operation, with relevance to the ones with industrial activities. This document identifies possible uses for the work developed in eCOUSIN by each of eCOUSIN's partners and lays out plans for it, in a later stage also eligible joint exploitation plans will be taken into account. The consortium will also master the challenge of adhere to and push standardization efforts wherever applicable with particular emphasis on the Social Oriented aspects. Being standard and compliant to already available guidelines is definitely a major goal of eCOUSIN.

Despite eCOUSIN is just in its first stage at the time of writing, the consortium has been very active in these areas and relevant results are already worthy to be mentioned.

1. INTRODUCTION

1.1 eCOUSIN background and objectives

1.1.1 Context and focus of research

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.

1.1.2 Objectives and expected results

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 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 shall 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.

1.2 WP7 Scope

The objectives of WP7 are to ensure the dissemination of the project's results and to coordinate its visibility at an international level through publications, contributions to standards and other industrial forums, participation or arrangement of specific events.

WP7 manages the exploitation plan, coordinates and arranges all activities related to the dissemination of the eCOUSIN project results within the research community, standardization bodies and towards markets players.



1.3 Scope of Deliverable D7.2

Deliverable D7.2 is the intermediate dissemination, exploitation and standardization report of the eCOUSIN project.

Deliverable D7.2 presents the project dissemination strategy, including the dissemination target audience groups, the actions and channels to reach these groups during the project life, and the first opportunities for external collaborations. Then it provides a preliminary identification of the expected results exploitable by the consortium partners, and finally the project's initial approach concerning the possible use of and contributions to standards.

2. DISSEMINATION STRATEGY

Extensive large-spectrum dissemination work, devised to achieve awareness for the project activities and results, are to be will be achieved by all partners in several ways, including:

- Presence with technical papers, demonstrations, or talks (e.g. panels) at relevant international conferences, workshops, ICT initiatives, technical events, industrial forums and cooperation with European stakeholders (all partners)
- Production of leading-edge research material suitable for publication in international Journals (IEEE, ACM, Elsevier, Wiley, Kluwer, etc) (all partners)
- Cooperation with other projects in related areas (all partners)
- Regular updates of WP7 activities to be published at the eCOUSIN website providing the latest information on project activities and achievements

Different activities will be started and different channels exploited to communicate on eCOUSIN project's objectives and to disseminate its outcomes. These activities and channels are summed up in Table 2. More information is provided in the following subsections.

Activity / Channel	Remarks
eCOUSIN Website	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. An initial list of candidate journals and magazines has been identified.
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 will arrange workshops and open events to collect requirements, inputs and feedbacks from market stakeholders and the research community, and to communicate on the project

	findings, they will 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].

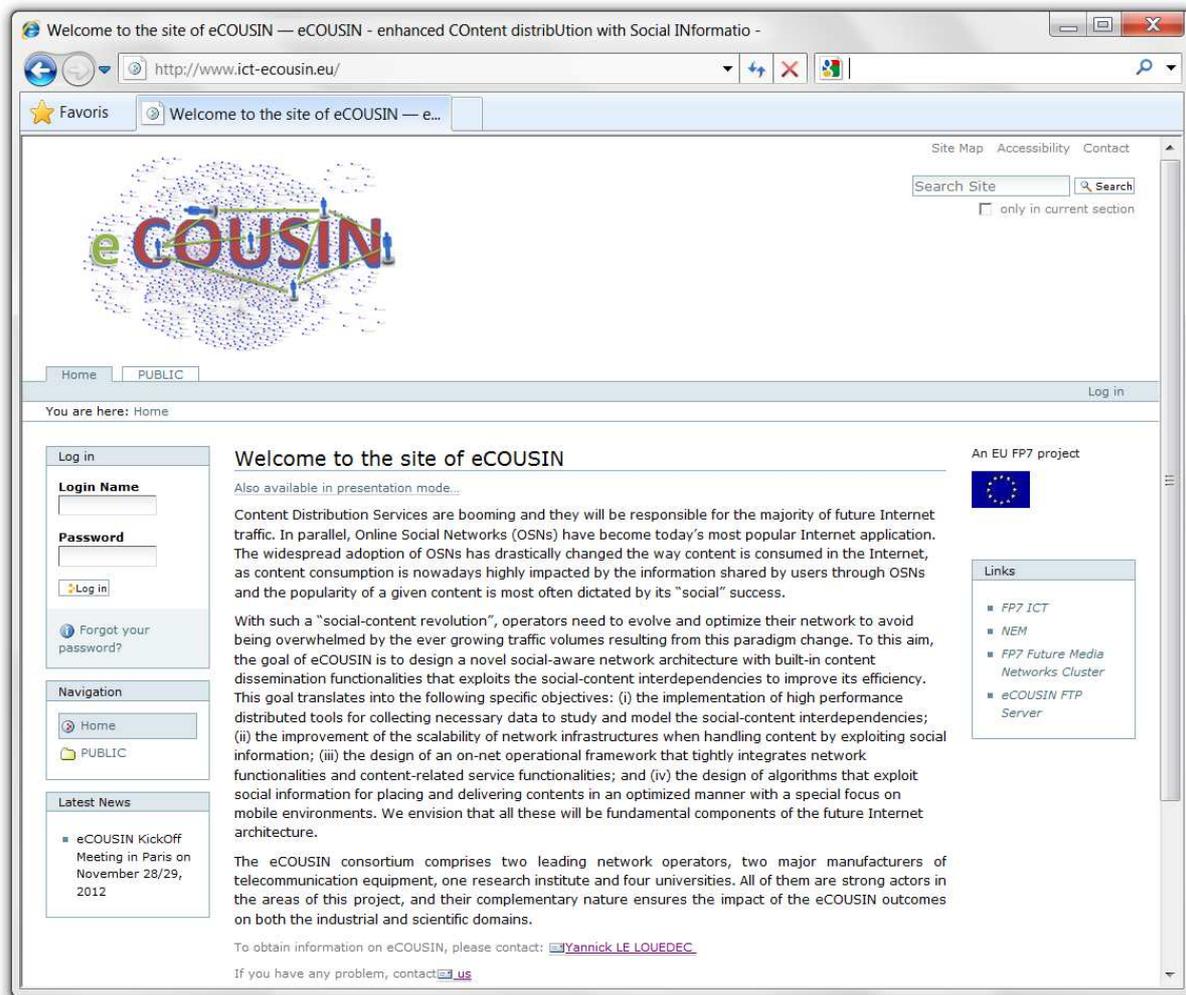


Figure 1: eCOUSIN Web presence

The public part 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, status reports, acknowledge EC contribution and links to other related and relevant sites;
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 will aggregate interest in 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.

The public website is maintained regularly as the project produces results, as papers are published, deliverables released or whenever there is news to report.

The European and FP7 Flags, as well as the caption “An EU FP7 project”, are made highly visible on every page of the web site.

2.1.1 Presentation of the project on the Web site

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”.

In addition, the web site presents the consortium and the work plan. It includes a “contact” section, as well as the Project’s “latest news” and “links” towards the EC web sites “FP7 ICT”, “FP7 Future Internet Cluster”, and “EU NEM”.

2.2 Flyers and brochures

The eCOUSIN project will produce and contribute to flyers and brochures for the purpose of communicating on its objectives and outcomes. Flyers and brochures will be 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	Nature (R - Report; O - Other)	Delivery date
D1.1	eCOUSIN Project WebSite (public area and restrictive area)	O	M3
D2.1	Initial report on Use Cases and Requirements	R	M3
D2.2	Initial System Architecture Specification	R	M9
D2.3	Business and Economic Analysis	R	M15
D2.4	Final Report on Use Cases, Requirements, System Architecture Specification, Privacy and Regulations	R	M21
D3.1	Measurement, Modelling, and Prediction of Social-Content Interdependencies (First Version)	R	M12
D3.2	Initial Release of Measurement and Prediction Software	R	M18
D3.3	Measurement, Modelling, and Prediction of Social-Content Interdependencies (Final Version)	R	M26
D4.1	Preliminary Report on the Design of Technical Solutions on Content Placement and Delivery	R	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	R	M18
D4.3	Final Implementation and Software Release of the Technical Solutions on Content Placement and Delivery	R	M26
D5.1	Requirements for Social-Enhanced Content Centric and Mobile Network Infrastructures	R	M12
D5.2	Modules and Interfaces for Social-Enhanced Content Centric and Mobile Network Infrastructures (Preliminary)	R	M18
D5.3	Modules and Interfaces for Social-Enhanced Content Centric and Mobile Network Infrastructures (Final Release)	R	M26
D6.1	Preliminary Plan for System Integration and Assessment (without commitment about the final implementation)	R	M15
D6.2	Final Plan for System Integration and Assessment	R	M20
D6.4	Final System Software release	R	M29

D6.5	Assessment Report		M30
D7.1	Project Presentation Preparation for Publishing at the Website of the Project	R	M1
D7.2	Intermediate Dissemination, Exploitation and Standardization, Report	R	M12
D7.3	Final Dissemination and Exploitation Plan Report	R	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 will be done via the traditional approaches used in this community, namely generating publications as well as organizing scientific events on topics related to eCOUSIN.

eCOUSIN will aim at publishing its results in leading, high-level, international conference proceedings and journals including, but not limited to, IEEE JSAC, IEEE Communications Magazine, Elsevier Computer Networks, IEEE Network, Springer Wireless Personal Communications, and IEEE Wireless Communications. 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 will provide 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

As the topic of the project is a hot one at present, it is likely that there will be many opportunities for eCOUSIN to make itself visible. These 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. While it is unrealistic to envisage that eCOUSIN will be represented at all of these events, 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. A particularly relevant event is the annual IST Conference, to which the project commits to attend and contribute.

As well as the above attendance at international events, the consortium will organise specific events which will focus on the project's subject in particular and at which it will be able to showcase its solutions and technology. The consortium will involve key personnel in high visible panels and will organize two workshops during its lifetime. Invitations to the most active researchers in the area will

be issued, including the Advisory Board members, and it is anticipated that this forum will provide a focal point for **Social-aware Content Distribution Network** technology in Europe.

The project will also be proactive with regards to establishing linkages with parties working on related activities both inside and outside Europe. For the former, the objective will be to ensure that the work of eCOUSIN complements other work taking place in a European context; for the latter, the objective will be to ensure that the work of eCOUSIN remains on the forefront of technology development in this domain, thus strengthening Europe's future potential. Members of the consortium already have good linkages with parties in the US and far east working in this area.

Concertation meetings, where related projects meet to identify synergies will be attended by the Project Coordinator and any 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 will also investigate opportunities to deploy its solutions for coverage during events such as conferences.

Finally, the project will explore potential collaborations with *public* domain (mobile) Internet data analysis frameworks as for instance the Cooperative Association for Internet Data Analysis (CAIDA), Measurement Lab (M-Lab) and the Community Resource for Archiving Wireless Data At Dartmouth (CRAWDAD). The interaction with such communities is expected to provide a further positive effect on the project success, especially if this may turn out into the proactive involvement of public domain developers in the implementation of eCOUSIN solutions.

2.6 Dissemination actions achieved during the first project period

Table 3, Table 4 and Table 5 present an overview of the dissemination actions achieved during the first project period:

- Table 3 lists the publications
- Table 4, the participation to conferences and presentations
- Table 5, the collaboration actions, including the organization of workshops and the participation to EU collaboration activities, and any other dissemination actions.

<i>Ref</i>	<i>Conference / Journal</i>	<i>Title</i>	<i>Countries addressed / Size of audience</i>	<i>Partners</i>	<i>Publishing Date</i>
(1)	World Wide Web Conference (WWW' 13)	Google+ or Google-?: Dissecting the evolution of the new OSN in its first year	Worldwide	UC3M, TSP	May 2013
(2)	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM' 13)	Analysis of publicly disclosed information in Facebook profiles	Worldwide	TSP	Sep 2013
(3)	IEEE International Conference on Peer-to-Peer Computing (P2P' 13)	Investigating the Reaction of BitTorrent Content Publishers to Antipiracy Actions.	Worldwide	TSP, UC3M, IMDEA	Sep 2013
(4)	IEEE Network	Understanding the evolution of multimedia content in the Internet through BitTorrent glasses	Worldwide	TSP, UC3M	Nov 2013
(5)	IEEE Global Communications Conference (GLOBECOM' 13)	"Current City" Prediction for Coarse Location Based Applications on Facebook.	Worldwide	TSP	Dec 2013
(6)	Springer Personal and Ubiquitous Computing (Special Issue in Cross-Community Mining).	Understanding the Locality Effect in Twitter: Measurement and Analysis	Worldwide	UC3M	Apr 2013
(7)	IEEE Transactions on Parallel and Distributed Systems	BitTorrent Locality and Transit Traffic Reduction: When, Why and at What Cost?	Worldwide	UC3M	March 2013
(8)	ACM international Conference in Online Social	Are Trending Topics useful for marketing? Visibility of Trending	Worldwide	UC3M	Oct 2013

	Networks	Topics vs Traditional Advertisement			
(9)	IEEE/ACM Transactions on Networkings	Unveiling the Incentives for Content Publishing in Popular BitTorrent Portals	Worldwide	UC3M, IMDEA, TUD	Oct 2013
(10)	KSII TRANSACTIONS ON INTERNET AND INFORMATION SYSTEMS	Connectivity Properties of Real BitTorrent Swarms	Worldwide	UC3M, IMDEA	Sep 2013
(11)	IEEE Infocom 2013	DataSpotting: Exploiting Naturally Clustered Mobile Devices to Offload Cellular Traffic	Worldwide	ALUD	June 2013
(12)	Joint SmartenIT/eCOUS IN Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM13)	Determining Leaders and Clusters in Video Consumption	Worldwide	A-LBELL, FT	Oct 2013
(13)	IEEE COMSOC MMTC E-Letter, Vol.8, No.4	Addressing the Wireless Content Challenge	Worldwide	ALUD	July 2013.
(14)	WoWMoM 2013 (Madrid)	Large scale analysis of HTTP adaptive streaming in mobile networks	Worldwide	FT	June 4-7, 2013
(15)	MASCOTS 2013 (August 14-16, 2013, San Francisco)	HTTP adaptive streaming in mobile networks: characteristics and caching opportunities	Worldwide	FT	August 14-16, 2013
(16)	CNSM 2013	Volume is not Enough: SVC-aware Server Allocation for Peer-assisted	Worldwide	TUD	October 2013

		Streaming.			
(17)	SETM 2013	Reciprocity with Virtual Nodes: Supporting Mobile Peers in Peer-to-Peer Content Distribution	Worldwide	TUD	October 2013
(18)	IEEE P2P 2013	RASP: Using OpenFlow to Push Overlay Streams into the Underlay	Worldwide	TUD	September 2013
(19)	AIMS 2013	Towards a Mobility-Supporting Incentive Scheme for Peer-to-Peer Content Distribution	Worldwide	TUD	June 2013
(20)	ICNP 2013	Optimizing Energy Consumption and QoE on Mobile Devices	Worldwide	TUD	October 2013
(21)	ACM HP-MOSys 2013	Opportunistic Message Routing using Multi-layer Social Networks	Worldwide	UCAM	November, 2013
(22)	SOSP poster 2013	SAKYOMI: SSD Prefetcher for Large-Scale Graph Traversal	Worldwide	UCAM	November, 2013
(23)	ACM SIGMOD 2013	Scale-up Graph Processing: A Storage-centric View	Worldwide	UCAM	June, 2013
(24)	Journal of Social Network Analysis and Mining, July, Springer, ISSN: 1869-5450, 2013	What's in Twitter I Know What Parties are Popular and Who You are Supporting Now!	Worldwide	UCAM	July, 2013
(25)	IEEE/ACM ASONAM 2013	Centrality and Mode Detection in Dynamic Contact Graphs; a Joint Diagonalisation Approach	Worldwide	UCAM	August, 2013
(26)	Technical Report, University of Cambridge, 2012	Mitigating I/O latency in SSD-based Graph Traversal	Worldwide	UCAM	2012

(27)	Journal of Network and Computer Applications, Elsevier, 2012	Evaluating Opportunistic Networks in Disaster Scenarios	Worldwide	UCAM	2012
------	--	---	-----------	------	------

Table 3. Publications

(1) *"Google+ or Google-?: Dissecting the evolution of the new OSN in its first year"*. R. González, R. Cuevas, R. Motamedi, R. Rejaie, Á. Cuevas. *World Wide Web Conference (WWW' 13)*. May 13-17, 2013. Rio de Janeiro (Brazil).

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.

(2) *"Analysis of publicly disclosed information in Facebook profiles"*. R. Farahbakhsh, X. Han, Á. Cuevas, N. Crespi. *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM' 13)*. August 25-27, 2013. Niagara Falls (Canada).

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.

(3) *"Investigating the Reaction of BitTorrent Content Publishers to Antipiracy Actions"*. R. Farahbakhsh, Á. Cuevas, R. Cuevas, R. Rejaie, M. Kryczka, R. González, N. Crespi. *IEEE International Conference on Peer-to-Peer Computing (P2P ' 13)*. September 9-11, 2013. Trento (Italy).

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.

(4) *"Understanding the evolution of multimedia content in the Internet through BitTorrent glasses"* R. Farahbakhsh, Á. Cuevas, R. Cuevas, R. Gonzalez, N. Crespi . *IEEE Network*. To be published issue Nov. 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, BitTorrent, 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.

(5) *"Current City Prediction for Coarse Location Based Applications on Facebook"*. W. Chanthaweethip, X. Han, N. Crespi, Y. Chen, R. Farahbakhsh, Á. Cuevas. *IEEE Global Communications Conference (GLOBECOM' 13)*. December 9-13, 2013. Atlanta (USA).

Abstract: Location-Based services with social networks improve users' experience and enrich people's social life. 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.

(6) R. Cuevas, R. González, A. Cuevas, C. Guerrero. "Understanding the Locality Effect in Twitter: Measurement and Analysis". *Springer Personal and Ubiquitous Computing (Special Issue in Cross-Community Mining)*. April 2013.

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.

(7) R. Cuevas, N. Laoutaris, Y. Xiao, G. Siganos, P. Rodriguez. "BitTorrent Locality and Transit Traffic Reduction: When, Why and at What Cost?". *IEEE Transactions on Parallel and Distributed Systems*, Accepted March 2013.

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.

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

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

Abstract: 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.

(10) R. Cuevas, M. Kryczka, A. Cuevas, C. Guerrero, A. Azcorra “Connectivity Properties of Real BitTorrent Swarms”, *KSII Transactions on Internet and Information Systems*.

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.

(11) “DataSpotting: Exploiting Naturally Clustered Mobile Devices to Offload Cellular Traffic”. X. Bao, Y. Lin, U. Lee, I. Rimal, R.R. Choudhury. *IEEE International Conference on Computer Communications*, 2013. Turin (Italy)

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.

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

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.

(13) *Addressing the Wireless Content Challenge"*. Ivica Rimac and Volker Hilt., *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."

(14) *"Large scale analysis of HTTP adaptive streaming in mobile networks"*. Ali Gouta, Charles Hong, Dohy Hong, Anne-Marie Kermarrec, Yannick Le Louédec. *WoWMoM 2013 (IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, Madrid, Spain, June 4-7, 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.

(15) *"HTTP adaptive streaming in mobile networks: characteristics and caching opportunities". Ali Gouta (Orange), Dohy Hong (N2NSoft), Anne-Marie Kermarrec (Inria), Yannick Le Louédec (Orange). MASCOTS 2013 (August 14-16, 2013, San Francisco)*

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.

(16) *Julius Rückert, Osama Abboud, Martin Kluge, David Hausheer: Volume is not Enough: SVC-aware Server Allocation for Peer-assisted Streaming. In: 9th International Conference on Network and Service Management (CNSM 2013), Short Paper, 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.

(17) *Matthias Wichtlhuber, Peter Heise, Björn Scheurich, David Hausheer: Reciprocity with Virtual Nodes: Supporting Mobile Peers in Peer-to-Peer Content Distribution. In: CNSM Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013), Short Paper, 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.

(18) *Julius Rückert, Jeremias Blendin, David Hausheer: RASP: Using OpenFlow to Push Overlay Streams into the Underlay. In: 13th IEEE International Conference on Peer-to-Peer Computing, Demo Paper, 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.

(19) *Matthias Wichtlhuber, David Hausheer: Towards a Mobility-Supporting Incentive Scheme for Peer-to-Peer Content Distribution. In: 7th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2013), 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.

(20) *Fabian Kaup, David Hausheer: Optimizing Energy Consumption and QoE on Mobile Devices. In: IEEE International Conference on Network Protocols (ICNP 2013), 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.

(21) A. Socievole, E. Yoneki, F. De Rango and J. Crowcroft "Opportunistic Message Routing using Multi-layer Social Networks". *ACM HP-MOSys, Barcelona, 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.

(22) E. Yoneki, K. Nilakant, V. Dalibard, and A. Roy SAKYOMI: SSD Prefetcher for Large-Scale Graph Traversal, *SOSP poster 2013.*

(23) E. Yoneki and A. Roy "Scale-up Graph Processing: A Storage-centric View". *ACM SIGMOD - GRADES, 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.

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

(25) D. Fay, J. Kunegis, and E. Yoneki "Centrality and Mode Detection in Dynamic Contact Graphs; a Joint Diagonalisation Approach". *IEEE/ACM ASONAM, Niagara Falls, Canada, 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 synthetic contacts and then used JD to extract the original underlying modes.

(26) A. Roy, K. Nilakant, V. Dalibard, and Eiko Yoneki "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 led 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.

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

Event	Location	Date	Title	Authors
Invited Talk	University Massachus sets Ahmerst, USA	08/08/2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	Boston University, USA	09/08/2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited Talk	INRS, Monreal, Canada	04/09/2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)

Invited Talk	University of Ottawa, Canada	06/09/2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited talk	McGill University, Montreal, Canada	18/08/2013	Characterizing user engagement in major OSNs	Ángel Cuevas (TSP)
Invited talk	Alcatel Lucent Bell Labs, US	11/03/2013	Google+ or Google-?: <i>Dissecting the evolution of the new OSN in its first year</i>	Rubén Cuevas (UC3M)
Invited talk	AT&T Research Lab, US	12/03/2013	Google+ or Google-? : <i>Dissecting the evolution of the new OSN in its first year</i>	Rubén Cuevas (UC3M)
Invited Talks	Northeastern University	09/10/2013	Google+ or Google-? Characterizing User Engagement in Major OSNs	Rubén Cuevas (UC3M)
Invited Talks	Boston University	10/10/2013	Are Trending Topics Useful for Marketing? Visibility of Trending Topics vs Traditional Advertisements	Rubén Cuevas (UC3M)
Future Networks and Mobile Summit event (http://www.futurenetworksummit.eu/2013/)	Lisboa, Portugal	July.03-05, 2013	Participation of eCOUSIN to the poster session of the conference	All consortium partners
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
Dagstuhl seminar on Future Internet [3]	Dagstuhl, Germany	24.03.13-27.03.13	Future Internet	Markus Hofman, Bell Labs
Dagstuhl Seminar on Future Internet [3]	Dagstuhl, Germany	24.03.13-27.03.13	Software-Defined Networking - Challenges and Opportunities	David Hausheer (TUD)
AIMS 2013 Conference (Tutorial)	Barcelona, Spain	June 2013.	EmanicsLab: A European Research Network tailored to	David Hausheer (TUD)

			Network and Service Management	
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)
Invited Talk	Graphlab workshop, USA	July 2013	Scale-up Graph Processing: A Storage-centric View	Eiko Yoneki (UCAM)
Invited Talk	University of Saskatchewan, Canada	August 2013	Empirical Approach for Modelling Dynamic Contact Networks	Eiko Yoneki (UCAM)
Invited Talk	University of Nottingham, UK	September 2013	Fluphone to Epi-Pi: Towards Quantifying Contact Networks for Digital Epidemiology	Eiko Yoneki (UCAM)

Table 4. Conferences and Presentations

Collaboration activity	Place	Date
Workshop: "Optimisation 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 plenary's Concertation meeting on October 11 th 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

Table 5. Collaboration Activities

About 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. In addition, specific extensions for mobile networks as well as Information/Content Centric Networking paradigms will be explored. 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 France Telecom / Orange (FT)

FT is both a network service provider and a service provider. As a network provider FT 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. eCOUSIN shall provide FT with technical solutions which will ease the pressure of this traffic, reducing the network cost or at least differing 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. FT has been deeply involved in this domain, contributing to searches on the design of the Future Internet (e.g., FP7 projects 4WARD and SAIL, as well as French ANR Connect project). Social Networks being destined to be one, if not THE, major Internet application, FT will investigate in eCOUSIN the relationship between OSN applications and networking delivery systems, including CDN and ICN, as well as the migration path to the target architecture.

As a service provider, FT expects to leverage the outputs of eCOUSIN to provide the best services with best quality of experience to its end users. FT expands into social networks and sharing applications markets. As an illustration FT has taken recently a 49% share 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), will be used to prepare the evolution of the recommendation system and the social network features in the portal.

3.2 Telecom Italia (TI)

Within Telecom Italia eCOUSIN results are expected to 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 idea is to exploit the technical results achieved by eCOUSIN to by extending our commercial products. Telecom Italia actually wants to be an operator, not just a carrier, that is the reason why the marketing is pushing both on cloud services and on media-center devices. The "CuboVision" device is already on market, working as a Upnp decoder/media centre, while in the next few weeks, a Smart Hard Disk by IoMega with Telecom Italia is going to be put on market. Our Idea is to implement the eCOUSIN use cases on such devices increasing the number of value added services to the final user, stimulating them to use the bandwidth we sell as core business. The target market size has been analysed on work package 2.

3.3 Alcatel-Lucent (A-LBELL and ALUD)

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 disseminate novel research results at leading international conferences (e.g., ITC, Globecom, Infocom), will ensure the openness of the model by proposing contributions (possibly jointly with eCOUSIN partners) to relevant standardization bodies (e.g., IETF, Broadband Forum, ITU-T and ETSI) and will 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 content delivery platform and ALU's CloudBand solution.

3.4 Foundation IMDEA Networks (IMDEA)

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 Cambridge University (UCAM)

UCAM will make the results of this project visible to a wide audience by publications in leading peer-reviewed research conferences and journals, to generate awareness and induce constructive comments and feedback from the scientific and industrial-research community. UCAM will participate in program committees and editorial boards defining key areas of interest for the research community including organization of Magazines and Journal special issues. UCAM will also participate in high-visibility events concerning emerging paradigms (e.g., CCN, NDN, and ICN). These include academic and industry oriented events including national and international bodies. UCAM's participation to this project will increase the number of students trained in the areas covered by the project and M.Sc. or Ph.D. thesis work in the context of the eCOUSIN project will carry their ideas into their future workplaces and academic institutions.

3.6 Technical University Darmstadt (TUD)

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 will be performed directly on the project and will, thus, highly benefit from the project's outcome, research results, and publications.

3.7 Institut Telecom SudParis (TSP)

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. TSP plans to keep these seminars in the academic year 2014/2015 and update some of the content with some of the contributions of TSP in eCOUSIN disseminated through scientific publications.

3.8 University Carlos III of Madrid (UC3M)

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”, “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. In the next academic year UC3M will include again these subjects as part of the academic program. This activity implies a clear transfer knowledge of those technologies developed in eCOUSIN to the future engineers.

For more details on the Master of Telematic Engineering we refer the reader to the corresponding Web site [4].

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 will mature we will extend and define 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.

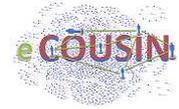
Standards should be aware of network efficiency issues, in order to keep in mind that both on mobile and fixed networks the new social orientation of content distribution is an issue that must be kept in mind. By showing the project's use cases and solutions to W3C audience we pointed out the attention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



5. CONCLUSION

Deliverable D7.2 is the first dissemination report of the eCOUSIN project.

Deliverable D7.2 presents the project dissemination strategy, including the dissemination target audience groups, the actions and channels to reach these groups during the project life, and the first opportunities for external collaborations. Then it provides a preliminary identification of the expected results exploitable by the consortium partners, and finally the project's initial approach concerning the possible use of and contributions to standards.

All the dissemination activities will be reported on a regular basis in the management report and in Deliverable D7.3 (Final Dissemination and Exploitation Plan Report) at Month M30.

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 Telema
tic_Engineering/Program](http://www.uc3m.es/portal/page/portal/postgraduate_studies/masters/Master_in_Telema%20tic_Engineering/Program)