Exhibition Proposal

<u>Proponent's Name</u>

Christian Koch[1], Julius Rückert[1], Nicola Bui[2,3], Foivos Michelinakis[2,3], Guido Fioravantti[2,3], Joerg Widmer[2] and David Hausheer[1]

<u>Proponent's emails</u>

ckoch@ps.tu-darmstadt.de, rueckert@ps.tu-darmstadt.de, nicola.bui@imdea.org, foivos.michelinakis@imdea.org, guido.fioravantti@imdea.org, hausheer@ps.tu-darmstadt.de, joerg.widmer@imdea.org

<u>Proponent's Institution</u>

[1]Peer-to-Peer Systems Engineering, TU Darmstadt, Germany

[2]IMDEA Networks Institute, Spain

[3]Universidad Carlos III de Madrid, Spain

<u>Proponent's CV</u>

Christian Koch is a researcher and first-year PhD student at the Department of Electrical Engineering and Information Technology of the Technische Universität Darmstadt since January 2014. He holds master degrees in computer science and IT security from Technische Universität Darmstadt, where he majored in communications technology and network security. His master thesis, written at the Fraunhofer Institute for Secure Information Technology, examined the migration of huge data volumes in heterogeneous cloud environments and proposes a migration service. This service is included in a solution that helps companies use cloud storage services safely, called *OmniCloud*. His research interest is in the area of mobile communication, where he focuses on offloading mechanisms targeting multimedia content. Christian Koch is a member of the European Unions FP7 project eCOUSIN, which focuses on the optimization of future networks. A special focus is applied on social aspects of communication. Christian Koch plans to finish his PhD thesis in early 2017.

Julius Rückert is a Ph.D. student at the Departement of Electrical Engineering and Information Technology of the Technische Universität Darmstadt since October 2011. He holds a bachelor as well as a master degree in computer science from TU Darmstadt.

Nicola Bui has been the CEO of Patavina Technologies, spin-off of the university of Padova and operating in the ICT field designing and developing embedded systems. At the same time, he has been researcher fellow with Consorzio Ferrara Ricerche (CFR) and with the Department of Information Engineering (DEI) at the University of Padova for seven years. During this period he has been involved in many European and Italian funded projects, such as e-SENSE and SENSEI, on wireless sensor networks, IoT-A, aimed at defining a reference architectural model for the Internet of Things, SWAP, dealing with energy harvesting in the IoT and WISEWAI on the realization of a urban wireless sensor network in Padova. Nicola just joined IMDEA Networks as a Research Engineer.

Foivos Michelinakis received his electrical and computer engineering diploma from the National Technical University of Athens (NTUA) in February 2011 (top 10% of class). During his subsequently army service, he was in the Research and Informatics Corps of the Greek Army, located at the School of Military Informatics for Senior Officers. Since October 2012, he is a PhD student at IMDEA Networks.

Guido Fioravantti is studying Telecommunications at the University Carlos III of Madrid. He is collaborating with IMDEA Networks on the eCOUSIN project dealing with the optimization of mobile networks. Recently, he worked with Heartbeat Records Madrid as a sound engineer and music producer.

Joerg Widmer is Research Professor at IMDEA Networks Institute in Madrid, Spain. He received his M.S. and PhD degrees in computer science from the University of Mannheim, Germany, Germany in 2000 and 2003, respectively. His research focuses primarily on wireless networks, ranging from MAC layer design and interference management to mobile network architectures. From 2005 to 2010, he was manager of the Ubiquitous Networking Research Group at DOCOMO Euro-Labs in Munich, Germany, leading several projects in the area of mobile and cellular networks. Before, he worked as post-doctoral researcher at EPFL, Switzerland on ultra-wide band communication and network coding. He was a visiting researcher at the International Computer Science Institute in Berkeley, CA, USA and University College London, UK. Joerg Widmer authored more than 100 conference and journal papers and three IETF RFCs, holds several patents, serves on the editorial board of IEEE Transactions on Communications, and regularly participates in program committees of several major conferences. Recently, he was awarded an ERC consolidator grant as well as a Spanish Ramon y Cajal grant. He is senior member of IEEE and ACM.

David Hausheer has been assistant professor at the Department of Electrical Engineering and Information Technology of the Technische Universitaet Darmstadt since May 2011. He holds a diploma degree in electrical engineering and a Ph.D. degree in technical sciences from ETH Zurich. From 2005–2011 he has been employed as a senior researcher and lecturer at University of Zurich, Switzerland, while being on leave as a visiting scholar at EECS, UC Berkeley from October 2009 to April 2011 under an SNSF fellowship for advanced researchers. David has been a co-applicant for several national and international research projects, including the German CRC MAKI, SNSF DaSAHIT, and the EU projects SESERV, SmartenIT, and eCOUSIN, as well as the Cisco URP project SCRIPT and the Deutsche Telekom project D-Nets. His research interests include several networked systems research areas, including software-defined networking, peer-to-peer and overlay networks, energy-efficient networking, and network economics. David has co-authored more than 70 peer-reviewed publications (e.g., IEEE CNSM, IFIP/IEEE IM, IEEE ICC, IEEE LCN, IEEE/IFIP NOMS, IFIP Networking, IEEE P2P, IEEE ICNP, IEEE CCNC, Springer JNSM, Elsevier ComNet). He is also executive committee member of the IEEE Computer Society TCCC, steering committee member of IFIP/ACM AIMS, and associate editor of Wiley IJNM. Moreover, he has been co-editor of the book "Towards the Future Internet – A European Research Perspective" published by IOS Press and guest editor of the IEEE Computer special issue on "Software Defined Networking" and the Wiley IJNM special issue on "Economic Traffic Management".

Project

This work has been supported in parts by the EU (FP7/#317846 , SmartenIT and FP7/#318398, eCOUSIN).

Exhibition Title

Demo: Mobile Social Prefetcher using Social and Network Information

Objective

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.

Description of the Demo

Section I. INTRODUCTION

The mobile Internet is one of the biggest enrichments of the 21th century. Several services appeared in the last years which are solely offered as smartphone apps. Additionally many services, e.g. Facebook, are mainly used as mobile apps. Worldwide the mobile data traffic grew by 81% in the year 2013 to 1.5 exabytes per month [1]. This represents about an 18-fold increase compared to the year 2000 [2]. The per-month mobile data volume is expected to increase tenfold, whereas the mobile bandwidth capacity is expected to only increase twofold till 2018. This raises new challenges for the mobile operators and endangers a stable service for the customers. Mobile carriers already reacted to this by introducing data caps to LTE data plans [3].

The main idea of this demo is to show the user and provider advantages obtained by the eCOUSIN solutions from two different angles: on the one hand we will visualize the benefit brought by the technical contributions on the network performance by means of controlled experiments; on the other hand, the user experience will be evaluated through a mobile application for smartphones using the prefetching and bandwidth optimization functionalities. This can be done by shifting a part of the data to be transferred to other access technologies, such as WiFi, e.g. when the user is at home or at work. This is especially interesting for data-intensive content that is not time-critical, for example videos which are accessed on demand. Mobile data plans are usually limited in their monthly transfer volume or based on a pay-per-use tariff. Prefetching efficiency highly depends on the possibility to predict content that is likely to be requested in the future and, consequently, can be prefetched. Prediction can happen at different time scales: short-term to optimize the bandwidth usage of the mobile connection (e.g., to smooth traffic peaks in mobile networks and to reduce the load of congested cells) or long-term to actively load contents in advance. For the latter case, it has been shown that OSNs are predictive for the video consumption of single users O^2SM [4].

In this demonstration, we will present an Adroid app called Mobile Social Prefetcher which is able to prefetch videos if WiFi is available and stream videos with a high QoE, if not. In particular, when the user requests the smartphone to play a video, the app opens an own video player instead of the regular one, which enables this benefits. This player takes over the playback of the prefetched video rather than, e.g. the YouTube app. If the video is not prefetched it will be streamed from remote by dynamically adopting the transfer rate depending on the playback buffer and the network conditions. In the following, we describe the application architecture in Section II and the details of the demonstration in Section III.

Section II. MOBILE SOCIAL PREFETCHER ARCHITECTURE

The architecture of the Mobile Social Prefetcher app is shown in Fig. 1. The main components are the Social Data Manager, the Content Prefetcher, and the Network component. All components run on the user's mobile device.

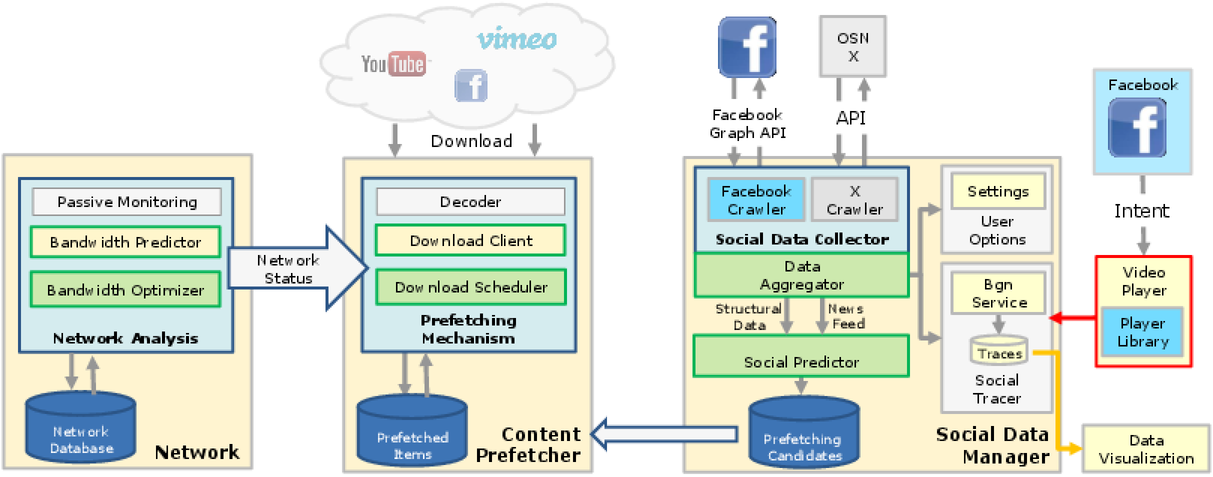

Figure 1 - Architecture of the Mobile Social Prefetcher

The **Social Data Manager** collects and analyses social information as a key functionality of the demonstrator. As the envisioned prefetching application is running on an user's smartphone, it accesses the users OSN with the permissions of this particular user. It is planned to retrieve structural information on the user's one-hop friendship graph, meta-data about the user and his friends, as well as the ongoing interactions of the user with his friends and the interaction history. Besides, the app can access the news feed of the user that aggregates the most recent activities that are likely to be relevant for the user. Currently, the app supports Facebook, but is extensible to include other social networks as well.

The prediction functionality itself might consider rather static social properties of content items, e.g. whether they are part of the user's news feed due to group memberships or shared by direct friends. Besides, also rather dynamic properties could be used to decide on an item's relevance, such as the temporal importance of a relationship to a friend. The latter could be determined using the user's interaction history, e.g. how often a user liked/shared/commented posts by a certain friend.

The **Content Prefetcher** includes the functionalities which are necessary to download the content. Here, prefetching candidates are scheduled to be downloaded by the Download Scheduler. This allows postponing the download until certain conditions are met, e.g. a WiFi connection is available. The Download Client performs the actual download and takes care of connection abortions and also chooses the quality of the video if multiple qualities are available. The Decoder service helps the Download Client to find the URI of the mp4 video file. Therefore it gets a video id as an input and outputs a URI or a list of URIs to the mp4 representations of the video hosted at the content provider or a CDN. The Social Data Collector & Manager continuously monitors the OSNs used, which is at this point in time Facebook only. For every OSN used, a special crawler is used. The Social Data Collector hands the information acquired by the crawler module to the Data Aggregator Module. Here, the date is brought into a common representation which is stored in a local database by the Social Tracer. To determine which items should be prefetched first the Social Predictor stores the candidates with a priority in the local Prefetching Candidates database, which is used by the Content Prefetcher.

The **Network** component consists of the following three sub-modules. The Bandwidth optimization module is intended to dynamically vary the bitrate of a multimedia server in order to adapt to capacity variations that the user mobile experiences during the playback of video streamed from remote. It will take as input: the predicted throughput of a device, QoE constraints specific to the content, and feedback from the mobile client application. In order to be able to feed the Bandwidth Optimization module with predictions and their confidence, a bandwidth prediction module is used. This module will consider statistical information about user mobility and the available bandwidth in a given cell. In order to combine all the statistical information, different predictors will be jointly used. In particular, we use different solution for the short (e.g., tens of seconds, a few minutes) and the medium-long (e.g., tens of minutes, hours) term predictions.

The short term prediction is achieved by means of a simple ARMA filter [5] whose order and coefficients have been previously tuned according to user's past information. Depending on the user's movement speed, the prediction validity can vary from a few tens off seconds (fast movements) to some minutes (quasi-static scenarios). During the validity time, the ARMA filter predicts the future mobile throughput. The medium-long term prediction is performed by means of statistical models [6]. These models accounts for the degradation on the accuracy of both the user position and the network cell congestion while the prediction is made farther in the future. By combining the two prediction techniques this component is able to decide when it is best to prefetch a content and, if the content has to be streamed from remote, what is the best way to allocate resource in order not optimize the cost.

Section III. DEMONSTRATION

The demo involves an Android smartphone and two WiFi hotspots. One hotspot will be used to simulate a cellular network during the demo (SSID: Cellular). A second hotspot is considered to be an ordinary WiFi hotspot (SSID: WiFi). One part of the demo is shown on the smartphone where the app is installed. The behaviour of the app will be shown by connecting the smartphone to one of the two hotspots. As soon as the user starts a video playback, our video player can be selected and chosen for playback. This allows to interfere the normal behaviour of Android and use the demonstrated app which also includes a custom video player. The demonstration will focus on two scenario, each of them makes use of both the smartphone and an additional laptop to illustrate background processes and additional informations.

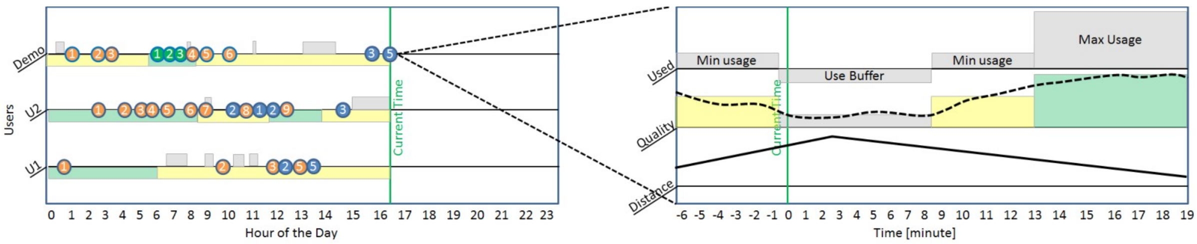

Figure 2 - Left side: circles represent video events on the time line of a given user. These events are color coded so that orange stands for publish time, green for prefetch time and blue for watch time. The area behind the dots is colored yellow or green if the mobile has mobile or WiFi access respectively. Rightside: video 5 is streamed from remote as it has not been prefetched before. The figure illustrates the bandwidth usage policies depending on the measured and forecast signal quality.

Fig. 2 shows an example of a typical usage diagram of our application. The dots on the left indicate video posts on the corresponding users Facebook feed. For example, the first orange circle of the demo user, marked with 1, indicates a video on this users Facebook feed. As soon as WiFi gets available, the app begins to prefetch the videos, in this case video 1, 2, and 3. Due to legal restrictions, prefetching will only be available at the demo device. Prefetched videos are indicated by a green circle. If a user watches a video, this is indicated by a blue circle, e.g., blue circle number 5. In case of this video, there was no possibility to prefetch it in advance, so it has to be streamed. The buffer filling strategy and the network connectivity prediction is shown on the right side. During the demo, this network conditions and the predictions are simulated using real traces. The reason for this is that the prediction model is based on a ARMA model, which is based on regular movement and connectivity patterns of the app user.

The app will be used in a multi-national study, especially in Germany, Spain, and Singapore. In the demonstration it is planned to visualize first results of the data collected of all participating users in a privacy-sensitive way. Fig. 2 shows an example view which we plan to demonstrate. Here the data for a subset of ca. 10-20 users and the demo device will be visualized live. Posts occurring on the users feed are indicated by a dot on the users time-line. The vertical green line indicates the current time, which moves during the demo. Prefetched videos are indicated by highlighting the post-circles green, which applies for the demo device only. Furthermore above the user time-line it is indicated if the screen of the user was enabled and for how long. This information provides us with meaningful data when and for how long users use their smartphone and enables the development of individual prediction models. Specific models are planned as further research. Underneath the user time-line, the network type the user is connected to, is shown. Here three different colors indicate the connection type.

WiFi is indicated as green, 3G as grey, and LTE as yellow. The reason for this is that prefetching performed over WiFi is about 10-times more energy efficient and often faster than over cellular networks [7].

A. Prefetching

The first scenario to demonstrate is the prefetching scenario. Prefetching will only be performed if the smartphone is connected to WiFi. The video content, on which the app relies, is retrieved from a Facebook feed created for the demo. This Facebook profile is assigned to a couple of channels which post videos regularly. Also, the option to post videos on the Facebook feed manually, for instance from the demo’s audience is possible. To illustrate the app’s behaviour we also plan to implement a Discard button which will delete all prefetched videos the app is aware of. This allows to enforce a re-download or allows to show the second scenario, where a video is streamed in a network-friendly manner. In case the app is connected to WiFi the prefetching of the latest videos becomes visible on the smartphone immediately. For the demo, a live view on the content becoming available at the feed is visualized with orange circles, see Fig. 2.

The app we provide for the demo will have full functionality, especially prefetching is enabled. Furthermore, the events collected from other devices, from users participating in a user study are shown. These events include the network connectivity, like shown in Fig. 2 on the left side. The green area under the user time lines indicate that WiFi is available and the yellow areas indicate cellular network connectivity. On top of the user time lines, we show the time during which the users have turned on the displays of their smartphones. This information gives valuable information on the regularity of smartphone usage. Based on this information, predictions about when the user is likely to use his smartphone in the future can be made. This information is not used in algorithms of the demo app but will be visualized as well and will be used to enhance the app in the future.

B. Streaming

The second scenario to demonstrate is the streaming scenario. If the connection is established to Cellular, the Mobile Social Prefetcher will perform a local playback if the content was prefetched during a previous WiFi connection. All playback functionality is implemented in a custom video player. In case the video was not prefetched, the Bandwidth Optimizer module is used to optimize the video streaming. To demonstrate this during the demo, the bandwidth of Cellular will be shaped to illustrate the app’s behaviour under realistic conditions, where the connectivity is likely to change. The decisions made by the app’s Bandwidth Optimizer module is demonstrated live like shown in Fig. 2 on the right side. At the same time we show the bandwidth of the connection during the playback on the laptop. Furthermore, the buffer filling on the smartphone is shown in parallel. The video playback will be smooth without pausing or stalling. To clarify the difference of our app, we can use Android's default player with the same traffic pattern, which will suffer from stalling events and will not be able to offer a smooth playback.

Section IV. CONCLUSION

The Mobile Social Prefetcher app demonstrated has the potential to enhance the user experience and at the same time to cope with varying network conditions. Relying on the videos posted on the users personal social network feed, the app prefetched videos if WiFi is available. The playback of these videos will be performed locally if they are requested from the user. If a video could not be prefetched because it was posted quite recently or no WiFi was available, a network-friendly streaming is performed. Here, the video player buffer is filled depending on the current network conditions and, therefore, reduces stalling events even under bad network conditions. This is demonstrated in a comparison of our app against the native Android video player under realistic network conditions.

Since we use a video player which can be handled as any other video player our approach minimizes the effort for the user. Concluding, this demonstration offer an interactive experience using Facebook and video playbacks. The video post publishing, prefetching, playback and streaming are shown on a live view during the demo. Additionally, data collected from other users using the app in a user study are also shown live.

ACKNOWLEDGMENTS

The authors would like to thank Bjorn Scheurich, Mohammad Hosseinian, Stefan Wilk, and Timo Thram for their contribution of the mobile social prefetcher implementation.

Requirements

Internet access (broadband would be preferable) and an LED/LCD screen (32 inches or bigger).

References

[1] Cisco Systems, Inc., "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 20132018," 2014.

[2] Sandvine Incorporated ULC, "Global Internet Phenomena 1H 2014 Report," 2014.

[3] Z. Epstein, “Brace yourself: The death of unlimited mobile data is coming,” BGR Media LLC, 2014.

[4] Y. Zhao et al., “O2SM: Enabling Efficient Offline Access to Online Social Media and Social Networks,” in Middleware. Springer LNCS, 2013, Vol. 8275, pp. 445–465.

[5] N. Bui and J. Widmer, “Modelling throughput prediction errors as gaussian random walks,” in 1st KuVS Workshop on Anticipatory Networks, 2014.

[6] N. Bui, F. Michelinakis, and J. Widmer, “A model for throughput prediction for mobile users,” in European Wireless, 2014.

[7] C. Gross et al., “EnerSim: An Energy Consumption Model for Large- scale Overlay Simulators,” in IEEE LCN, 2013.