Socially-Aware Traffic Management

Michael Seufert, George Darzanos, Ioanna Papafili, Roman Łapacz,
Valentin Burger, and Tobias Hoßfeld

Abstract Socially-aware traffic management utilizes social information to optimize traffic management in the Internet in terms of traffic load, energy consumption, or end user satisfaction. Several use cases can benefit from socially-aware traffic management and the performance of overlay applications can be enhanced. We present existing use cases and their socially-aware approaches and solutions, but also raise discussions on additional benefits from the integration of social information into traffic management as well as practical aspects in this domain.

Key words: Social Information, Social Awareness, Traffic Management, Content Storage, Content Delivery, Service Mobility, Network Security

Michael Seufert
University of Würzburg, Institute of Computer Science, Würzburg, Germany,
e-mail: seufert@informatik.uni-wuerzburg.de

George Darzanos
Athens University of Economics and Business, Department of Informatics, Athens, Greece,
e-mail: ntarzanos@aueb.gr

Ioanna Papafili
Athens University of Economics and Business, Department of Informatics, Athens, Greece,
e-mail: iopapafi@aueb.gr

Roman Łapacz
Institute of Bioorganic Chemistry of the Polish Academy of Sciences, Poznań Supercomputing and Networking Center, Poznań, Poland, e-mail: romradz@man.poznan.pl

Valentin Burger
University of Würzburg, Institute of Computer Science, Würzburg, Germany,
e-mail: valentin.burger@informatik.uni-wuerzburg.de

Tobias Hoßfeld
University of Würzburg, Institute of Computer Science, Würzburg, Germany,
e-mail: hossfeld@informatik.uni-wuerzburg.de

1 Introduction

In online social networks (OSNs) users voluntarily provide information about themselves, their interests, their friends and their activities, especially about their current situation or exceptional events. Nowadays these so called social signals are ubiquitous and can not only be collected from OSNs (e.g., friendships, interests, trust-relevant metadata), but also from applications (e.g., messaging or call patterns) and sensors (e.g., location). Social awareness harvests these signals, extracts useful and re-usable information (e.g., users' social relationships, activity patterns, and interests), and exploits them in order to improve a service.

Recently in the field of traffic management in the Internet, works were conducted which utilize social information, for example, to avoid congestion, increase bandwidth, or reduce latency. In that context, social awareness links social signals and information, such as social network structure, users' preferences and behaviors, etc. to network management mechanisms. This means that such mechanisms exploit the information in order to perform efficient network management, content placement, and traffic optimization to enhance the performance of an overlay application (e.g., video streaming, file sharing). As this promising research field has yet got little attention, this work will provide an insight to this new topic.

In Section 2, we will define socially-aware traffic management and present the involved actors. In the remainder of the paper, we focus on use cases in which the involved stakeholders can benefit from social information. In Section 3, we present three use cases i) content storage and delivery, ii) service mobility, and iii) network security, in which social awareness can benefit involved stakeholders. Finally, Section 4 raises discussions on additional benefits from the integration of social information into traffic management and concludes this work.

2 Terminology, Definitions, and Actors

In order to put the description and discussion of socially-aware traffic management and its use cases on a firm footing, we start with definitions of important terms and present the involved actors.

2.1 Terminology and Definitions

Social signals are any signals which are emitted by persons. In the special case of Internet services, we consider a social signal to be a signal which is emitted in the Internet by an end user of an Internet application. Thus in

fact, any interaction of an end user with an Internet service is a social signal. A signal itself contains no information, but information can be created out of them when evaluated in the right context. As signals range from simple logins to a service to complex service requests which might include interactions with other users or the environment, examples are manifold. In the context of online social networks, these signals are, e.g., friendship requests and confirmations, indications of interest or liking, or postings about external events. Another example are location data which are created by sensors of mobile devices and are transmitted when using an Internet service.

Social information is defined as information about one or more persons, or their relationships. It is deduced from bringing social signals into an appropriate context which allows for the generation of new insights about respective users or relationships between users. For example, evaluating the signals that user A sent a friendship request to user B, and B's confirmation of that request, will generate the social information the A and B are friends. Evaluating the same signals differently (i.e., in a different context), will give the information that A and B were online and used the OSN service at the time the signals were emitted. As another example, evaluating the location data signals of user C in the context that every second Saturday the location data is the same, and that there is a football stadium at that specific location, will create the information that C is a supporter of a certain football team. Thus, it can be seen that the created information depends on the particular evaluation of the social signals, and might require additional (external) information in order to create new social information. Usually, such partial information which requires external information to generate new social information is called meta information.

In general, the term **social awareness** implies the utilization of social information for a specific purpose. In the context of Internet services, we will consider social awareness to be the utilization of social information to improve an Internet service. Social awareness can include the collection of social signals and production of social information, but also a collaboration with a social information provider (see below) is possible. Taking provided or generated social information as an input, social awareness will exploit this information in order to deliver a higher service quality to end users and/or to provide the service more efficiently.

Socially-aware traffic management is a special case of social awareness, in which social information is used to improve traffic management on the Internet. Traffic management are means in order to handle the transportation of data across the networks. As not only link capacities increase in the Internet, but also traffic volumes become larger due to new service levels and new applications (e.g., cloud applications), it is necessary in order to avoid congestion, deliver applications in acceptable quality, and to save energy, resources, and costs. Traffic management can be employed by the service itself, e.g., by service quality selection or scheduling of transmissions, or by the network operators. Their methods typically include, but are not

limited to, prioritization, routing, bandwidth shaping, caching, or offloading. The utilization of social information shall enable the improvement of classical traffic management solutions as well as the development of novel traffic management approaches.

2.2 Actors of Socially-Aware Traffic Management, Their Goals, and Possible Benefits

With socially-aware traffic management, five actors and their goals have to be considered. Note that each actor can be a separate stakeholder, but stakeholders can also have multiple roles.

The **cloud service provider** or **application provider** provides an Internet service to end users. The offered service might be running on own infrastructure or on the infrastucture of a cloud operator. The application provider is interested in monetization of offered services which includes reduction of Internet service provider (ISP) infrastructure and cloud resource consumption costs. Satisfaction of end users is a crucial issue as it is directly related to the number of customers. To ensure this goal the Quality of Service (QoS)/Quality of Experience (QoE) requirements should be met Fiedler et al (2010). If social information is utilized, QoS/QoE parameters for services may be improved and also new services may be offered. Moreover, infrastructure costs can be reduced if social information is exploited to increase the utilization of resources.

The **datacenter operator** or **cloud operator** is operating a datacenter/cloud infrastructure. He offers cloud resources, e.g., storage, computation, to the application provider, while he buys Internet connectivity and inter-connectivity of his sites from an ISP. The cloud provider is mainly interested in monetizing his infrastructure and reduce his costs. Monetization of infrastructure is done by fulfilling service level agreements (SLAs) with the application provider and therefore guaranteeing satisfactory QoS parameters for end users. Reduction of costs, in case of cloud providers, focuses on best possible utilization of hardware – both resource-wise and energy-wise. As it depends on ISPs to provide network access, this stakeholder will seek the best SLA conditions for himself.

The **Internet service provider (ISP)** is operating a communication network infrastructure. His main interest is monetizing his infrastructure. This can be increased by high quality of network services that translates into satisfaction of cloud operators and also of end users (Hoßfeld et al, 2009). Supporting new services, possibly by employing social information, may be attractive for application providers, and simultaneously makes the ISP more competitive towards end users and cloud providers. Such new services can also enable reduction of costs by both more efficient use of own resources and keeping transit link traffic as low as possible.

The **end user**'s main concern is his own QoE (Fiedler et al, 2010), network access cost, and energy consumption (Ickin et al, 2012). This stakeholder is rather not involved in other stakeholders' interactions, being primarily a client of ISP and application provider. It is noteworthy, that costs in case of end user often can be expressed by being exposed to advertisements instead of being involved in the monetary flow.

The **social information provider** wants to benefit from his social information. Therefore, he can provide or sell social information to application providers or ISPs in order for the latters to support optimization decisions, e.g., content placement.

3 Use Cases for Socially-Aware Traffic Management

The exploitation of social information may lead in significant benefits for all involved stakeholders, i.e., ISP, cloud operator, application provider and the end user. Therefore, three indicative use cases are presented. For *content storage and delivery*, three variations of socially-aware traffic management are described, that are centralized, distributed, or hierarchical content delivery platforms. Moving towards practical applications, we investigate information spreading in OSNs and its employment in socially-aware caching solutions for video streaming, and we overview existing traffic management solutions that employ social signals to perform efficiently content placement or prefetching. Next, we describe the *service mobility* use cases, which involves WiFi offloading, content placement for mobile users, and service placement. Finally, we provide some insight to a third use case, i.e., *network security* employing social information to defend against Sybil and DDoS attacks.

3.1 Content Storage and Delivery

Internet traffic has increased manifold in the last few years. Drivers for this increase include inter alia the increased popularity of video streaming applications (e.g., YouTube, NetFlix), the emergence of a multitude of new overlay applications such as online storage (e.g., Dropbox, Google Drive) and online social networks (e.g., Facebook, Twitter), the high increase of mobile devices (e.g., smartphones, tablets) and the upcoming trend of moving both storage and computing capacity to the cloud which allows more, even smaller players to enter the market (Cisco (2012)). Concerning video as a key application contributing largely to the overall IP traffic, video and specifically user-generated content (UGC) sharing (e.g., home-made videos) has evolved to a major trend in OSNs. Three variations of the content storage and delivery use case are described where social information is employed to achieve efficiency in content

delivery, in terms of either content placement or pre-fetching. Moreover, we present work in literature which analyze content spreading in OSNs and show already existing socially-aware caching solutions for video streaming. Finally, we briefly overview related works which employ social awareness in order to handle the huge traffic volumes generated by video sharing over OSNs.

3.1.1 Exploitation of Social Information by a Centralized Content Delivery Platform

We consider a use case inspired by the evaluation scenario described in Traverso et al (2012). Specifically, we consider an OSN having users around the globe who share videos via the OSN which are stored in third-party owned online video streaming platform such as YouTube. This content can be viewed by their online friends, their friends' friends, etc. through the Friend-of-Friend (FoF) relationship.

In order to meet the content demand by users of the video streaming platform, who are located worldwide, the video platform is operated on a geo-diverse system comprising multiple points-of-presence (PoPs) distributed globally. These PoPs are connected to each other by links, which can either be owned by the entity that also owns the PoPs, or be leased from network providers. Each user is assigned and served out of his (geographically) nearest PoP, for all of his requests as depicted in Figure 1.

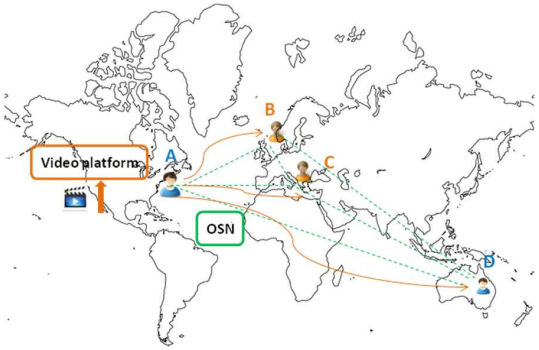

Fig. 1 Content delivery in geographically distributed PoPs.

Placing data close to the users is an approach followed by most content delivery networks (CDN). Therefore all content uploaded by a user A is first uploaded to the nearest PoP, i.e., PoP_A. When content is requested by another user B, the nearest PoP to B, i.e., PoP_B, is contacted and if the content is available there, the request is served. The content can be already present at PoP_B, if content was first uploaded there or was brought there by

an earlier request. If the content is not available in PoP_B, then a request is made to PoP_A and the content is brought to PoP_B.

In such as setup, social relationships between the users of the OSN can be taken into account to predict where a piece of content will be requested next, i.e., by which PoP. For instance, it is expected that due to their social relationship, users of the online social network will request a piece of content that a friend of them, e.g., user A, has uploaded to the video platform with higher probability than users that have no social relationship with A.

The so-called social awareness involves the exploitation of such social information (e.g., social graph, social relationships among users, behavior patterns, login time, time spent in the OSN) in order to predict where and by whom an uploaded video will be consumed. Such predictions can be employed to develop socially-aware mechanisms such as TailGate proposed in Traverso et al (2012) that will enable pre-fetching of the uploaded video in various locations (e.g., PoPs).

3.1.2 Exploitation of Social Information by a Distributed Content Delivery Platform

An alternative use case involves the dissemination of video content in a peer-to-peer (P2P) fashion among an OSN's users, taken from Li et al (2012). We consider again an OSN, whose users are scattered around the globe and upload videos, i.e., UGC to an online video streaming platform like YouTube. This content, similarly to the scenario described in the previous section, can be viewed by their friends, their friends' friends, etc.

End users, also called peers, download parts of the file, e.g., chunks or blocks, and are considered to be able to re-upload them to another peer. Additionally, a proxy server is considered to orchestrate the content dissemination as a P2P tracker or to participate in it. In the latter case, the proxy server is connected to the content provider, which is an end user in case of UGC. Moreover, multiple proxy servers are considered to be also distributed globally and each one of them to have a specific domain of influence, e.g., an ISP's domain, an Autonomous System (AS).

The initial content provider uploading a video to the proxy server, the proxy server itself, and the peers participating in the dissemination of that particular video are considered a swarm. Furthermore, the video parts exchange among peers is performed based on some specific peer and chunk selection policy. As mentioned before, placing video chunks close to the end users is an approach followed by most CDNs as it leads to lower latency and stall time, and thus high QoE for end users. Therefore, social information can be extracted from OSN by the video platform owner, so as to predict by whom a video uploaded to the proxy server will be viewed. These users can be preferably included in the dissemination swarm. Once they want to access

the video, they have lower delay (thus, a better QoE) because part of the file is already on their device.

According to Traverso et al (2012) and Li et al (2012), direct friends (1-hop friends) and friends of friends (FoF or 2-hops friends) of a user C have high probability (more than 80%), to watch a video uploaded or posted by C. Social information, such as users' interests, e.g., sports or music, prove to be also important, as users, which have a FoF relationship with C and share the same interests, are highly likely to watch a video uploaded by C.

3.1.3 Exploitation of Social Information by a Hierarchical Storage Platform

Another interesting use case of applying the knowledge derived from OSNs is improving the internal decision making algorithms in advanced distributed hierarchical storage management systems.

Hierarchical Storage Management (HSM) is an approach to manage high volume of data in the way that data are categorized and moved between storage types to reduce the storage cost as well as to optimize an access and the energy consumption of data storage management. A hierarchy level is assigned to a storage media. The first level is represented by high-speed high-cost devices destined for data set that is frequently accessed by applications. Other data, for example older and thus less popular, can be automatically moved to a slower low-cost storage media.

Usually, three levels of storage hierarchy are defined, as illustrated in Figure 2. The first level is a high-speed systems, such as hard disk drive array, the second one is slower, such as optical storage, and the last one may be implemented as magnetic tape drives. As the technology of the first level is the most expensive the size of it is smaller than the storage sizes of other levels.

HSM can be also interpreted as a tier storage technique, although sometimes storage specialists see differences between them (Yoshida (2011)). The basic difference seems to be the way how datasets are accessed. In case of HSM, inactive data are moved to the levels of slower storages and can be accessed directly again only after migrated back to the first high-speed level. On the contrary, the tier approach allows fetching data from any tier any time.

Nowadays, when the amount of data is rapidly growing, HSM offers a substantial benefit from managing storage devices efficiently, especially in large-scale networks, storage and computational environments, such as clouds. In particular, a common deployment scenario involves resources of a cloud residing in remote geographical locations, while end users perceive its resources as a consistent pool available for allocation (e.g., IaaS model (Mell and Grance (2011))). This operation is highly related to the so-called service mobility discussed in Section 3.2.

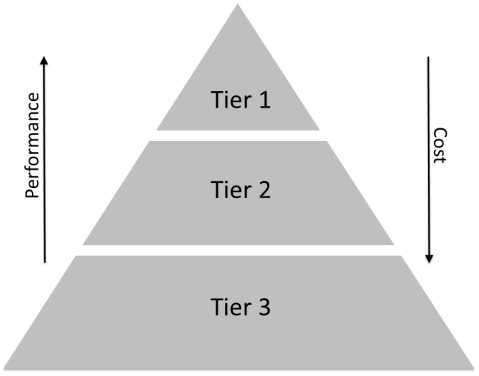

Fig. 2 Storage tiers in HSM along with performance and cost trends. (Source: inspired by Ganley (2013))

Moreover, one cloud operator may utilize storage resources or assign specific works to another cloud operator, e.g., in the context of a cloud federation, to achieve e.g., load balancing, reduction of his individual energy consumption, etc. In order to optimize these operations of data migration, social signals can be exploited by the cloud operators so as to predict not only the amount but also the where and when of future demand. As a result, end users will experience better QoE, i.e., faster access to data, while the cloud operators will achieve more accurate utilization of their storage hierarchies (tiers) and in consequence lower energy consumption.

3.1.4 Understanding Information Spreading in OSN for Utilization in Traffic Management Algorithms

Social awareness can be used in different ways to improve content delivery. To exploit the information within OSNs it is important to understand how information is spread, and how to identify important nodes in the social graph and the relationships with their friends.

In Bakshy et al (2011), the authors investigate the influence of posts by tracking the diffusion of URLs in Twitter and show that content that is connected with good feeling and interesting content is more likely to be propagated. They also find that the users that have most influence are also the most cost-effective. Hence, influential users post relative rarely, but if they do, the content is of high interest.

In Ruhela et al (2011), the authors collected data from five different sources and investigated the temporal growth and decay of topics in the network and the geographical and social spread of the topics. Besides identifying different classes of temporal growth patterns and time zone differences

in popularity, they find that the social cohesion of users interested in specific content is greater for niche topics. Hence, they propose to use semantic information about the topic to assess the temporal growth, use time-zone information to predict the breakout of popular topics in a specific region, and use social network predictors for niche content. To distribute the content and use cache capacities effectively we need good replica placement algorithms.

Next, in Wittie et al (2010), the authors inferred the network structure of Facebook performing crawling, packet captures, and network measurements. Due to high locality of interests they state that service providers could profit a lot from locality to save traffic on intercontinental paths. Proposed solutions are regional caches or a CDN that connects a global network of server farms at different ISPs to bring the content close to users.

Finally, in Wang et al (2012), the authors explore how patterns of video link propagation in a microblogging system are correlated with video popularity on the video sharing site, at different times and in different geographic regions. Then, they design neural network-based learning frameworks to predict the number and geographic distribution of viewers, in order to deploy a proactive video sharing system. The evaluations show that their frameworks achieve better prediction accuracy compared to a classical approach that relies on historical numbers of views.

3.1.5 Existing Socially-Aware Caching Solutions for Video Streaming

Socially-aware caching tries to predict future access to user generated content (e.g., videos) based on information from OSNs. Hints shall be generated for replica placement and/or cache replacement.

In Sastry et al (2009) the classical approach of placing replicas based on access history is improved. Therefore social cascades are identified in an OSN, and declared affiliations of potential future users (i.e., OSN friends of previous users) are added. In Scellato et al (2011) standard cache replacement strategies are augmented with geo-social information from OSNs. Again social cascades are analyzed to recognize locally popular content which should be kept longer in the cache.

Apart from the increasing popularity of video sharing over OSN, another significant characteristic of content dissemination on top of OSNs that need to be taken into consideration is the long-tailed nature of content, i.e., UGC such as home-made funny videos, etc. Below, some solutions are briefly presented and discussed that focus and address the long-tailed nature of video delivery over OSNs.

In Traverso et al (2012), the authors propose TailGate which derives and uses social information derived from OSNs, such as social relationships, regularities in read access patterns, and time-zone differences for predicting where and when the content will likely be consumed, in order to push the content

where-ever before it is needed. Thus, exploiting the derived social information, long-tail content is selectively distributed across globally spread PoPs, while lowering bandwidth costs and improving QoE. In particular, bandwidth costs are minimized under peak based pricing schemes (95th percentile), but the approach is also beneficial for flat rate schemes.

For the analysis of TailGate, the authors considered the scenario depicted in Figure 3. Specifically, they consider an online video delivery service with users across the world, operated on a geo-diverse system comprising multiple PoPs distributed globally. Each of these interconnected PoPs handles content for geographically close users. In particular, when UGC is created, it is first uploaded to the geographically closest PoP, and then it can be distributed to other PoPs.

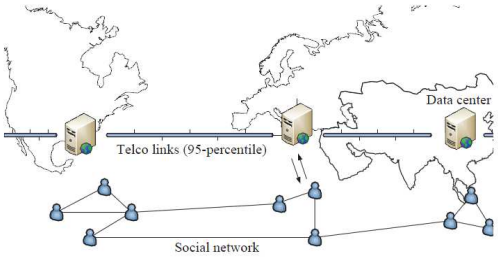

Fig. 3 TailGate's generic distributed architecture for video delivery. (Source: Traverso et al (2012))

In Li et al (2012), the authors identify the fact that the deployment of traditional video sharing systems in OSNs is costly and not scalable. Thus, they propose SocialTube, a peer-assisted video sharing system that explores social relationships, interest similarity, and physical location between peers in OSNs. Specifically, SocialTube incorporates three algorithms: an OSN-based P2P overlay construction algorithm that clusters peers based on their social relationships and interests, an OSN-based chunk pre-fetch algorithm to increase the video prefetch accuracy to minimize video playback startup delay, and a buffer management algorithm. The social network-based P2P overlay has a hierarchical structure that connects a source node with its followers, and connects the followers with other non-followers.

Moreover, in order to reduce the video startup latency, the social network-based pre-fetching algorithm is employed. This algorithm dictates that when a source node uploads a new video to a centralized video server, the source also pushes the prefix, i.e., the first chunk, of the video to its followers. Additionally, it is pushed to the peers in the interest clusters matching the content of the video, because there is a high probability that it will be requested to be watched, since followers watch almost all videos of the source.

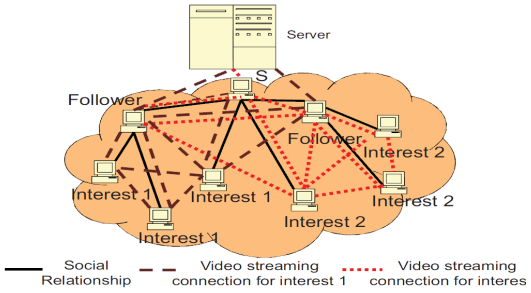

Fig. 4 SocialTube's P2P overlay structure based on social relationships and interests. (Source: Li et al (2012))

In Zhou et al (2012), the authors examined crawled data from Facebook and observed that a significant fraction of Internet traffic contains content that is created at the edge of the network, i.e., UGC. Moreover, they observed that users are in general significantly more interested in the content that is uploaded by their friends and friends-of-friends, while traffic local to a region is produced and consumed mostly in the same region, which is contrary to the case of traditional web content. Furthermore, they argue that while caching the most popular 10% of traditional content would allow to satisfy at least half of all requests, this caching technique would perform significantly worst for content with a more even popularity distribution.

Therefore, they propose WebCloud, a content distribution system for OSNs that works by re-purposing client web browsers to help serve content to others, and which tries to serve the request from one of that user's friends' browsers, instead of from the OSN directly. WebCloud is designed to be deployed by a web site, such as the provider of an OSN, to be compatible with the web browsers (no plug-ins) of today and to serve as a cache for popular content.

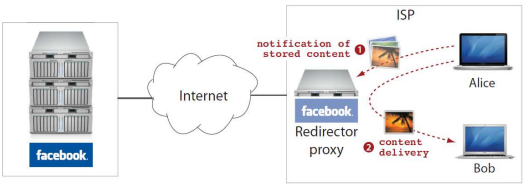

Fig. 5 Content sharing in WebCloud, where Alice first informs the proxy of locally stored content. When Bob requests content from the proxy, the proxy fetches it from Alice and delivers it to Bob, thereby keeping the content exchange local. (Source: Zhou et al (2012))

The authors claim that WebCloud trying to keep the content exchange between two users within the same ISP and geographic region, to reduce both

OSN's and the ISP's costs. WebCloud emulates direct browser-to-browser communication by introducing middleboxes, which are called redirector proxies (Figure 5). The proxy determines if any other online local user has the requested content and if so, fetches the content from that user's browser and transmits it to the requestor. Should no local user have the content, the browser fetches the content from the OSN.

3.2 Global Service Mobility

In today's Internet, services and applications have to be seamlessly available to end users at any time and location. Especially mobile users, i.e., users who access services by means of mobile devices from any location, pose severe challenges for mobile service providers, cloud operators and application providers.

3.2.1 Exploitation of Social Information for WiFi Offloading and Service Placement

New services have ever increasing network requirements which strain current mobile networks, and feed the desire of network operators to offload traffic to WiFi networks. On the other hand, application providers rely more and more on the cloud concept, which allows moving services within or among datacenters worldwide, and thereby foster the mobility of services. Both approaches can benefit from the utilization of social information, as well as data related to the location of the end user.

Location information can be retrieved easily either from the mobile network provider, from services (e.g., Ruhela et al (2011)), or from end users. Location data are often used by applications, and even shared by end users, e.g., as meta-information of postings, or as explicit postings of locations in OSNs or in specialized services like Foursquare[1].

This allows for the monitoring of such social signals and facilitates the creation of mobility patterns for different users. One step beyond, these patterns can be used to predict the location of a user in the future. These predictions can be also used to improve the content delivery described above in Section 3.1 for mobile users. Both pre-fetching and caching algorithms will achieve a higher accuracy which improves the cache hit ratio for mobile users, and thus, improves users' QoE. In the following we will present other aspects of global service mobility, such as WiFi offloading and service placement.

Offloading data to WiFi networks has already been in the focus for some years, and can be considered as providing fast and energy-efficient Internet

[1] https://foursquare.com

access for mobile end users. Offloading traffic using WiFi networks can save between 75% and 90% of the energy for network transmissions compared to 3G connectivity only (Gautam et al (2013)). At the same time, the risk of network caused stallings may be reduced by increased data rates on WiFi, improving the QoE of the end user.

Service placement is a generalization of content placement, i.e., instead of only placing content at the appropriate caching locations (cf. Section 3.1), whole services are placed. This includes the creation, termination, and migration of virtual machines which are running the service. Especially cloud services which are based on the elasticity of clouds can benefit from such placement. The description is mainly based on Biancani and Cruschelli (2013).

Service placement is interesting both from ISPs' and application providers' perspectives. Application providers are interested in maintaining a good ratio of revenue and costs. An optimal placement of a service among a number of cloud providers or own datacenters can help optimizing costs as well as meeting end users' QoE requirements. However, the optimal placement of services is also in the interest of the ISP to reduce his operational costs, as a disadvantageous service placement can increase traffic from outside the provider's AS. Thus, both stakeholders could collaborate, e.g., by using an ALTO (Alimi et al (2013)) style approach.

The placement of service will be optimized by taking into account social information, e.g., where services may be popular in specific regions or for which specific groups. Such information can be aggregated from different sources, i.e., from a direct cooperation with OSN applications and end users, or from the ISP who might exploit its aggregated knowledge on users interests and mobility patterns. A side effect of these new possibilities is an expected improvement of the perceived network quality on end user devices. This is achieved by locating services closer to the end user, reducing the delay, and improving the network throughput.

3.2.2 Existing Socially-Aware Solutions for Service Mobility

Fon[2] started a WiFi sharing community in 2006 by offering a home router device with a separate shared WiFi network which could be accessed by every community member. Similar approaches are the hotspot databases Boingo[3] and WeFi[4], and Karma[5] which adds social reciprocity to WiFi sharing. Also the research community investigated incentives and algorithms for broadband access sharing (Mamatas et al (2010)), and architectures for ubiquitous WiFi access in city areas (Sastry et al (2007); Vidales et al (2009)).

[2] http://www.fon.com

[3] http://www.boingo.com

[4] http://wefi.com

[5] https://yourkarma.com

Valancius et al (2009) propose a distributed service platform, called Nano Data Centers or NaDa, based on tiny (i.e., nano) managed "servers" located at the edges of the network, i.e., in users' premises. With NaDas, both the nano servers and access bandwidth to those servers are controlled and managed by a single entity, typically an ISP. The role of the nano servers can be played by ISP-owned devices like Triple-Play gateways and DSL/cable modems that sit behind standard broadband accesses. Such gateways form the core of the NaDa platform and can host many of the Internet services today operated in datacenters. ISPs can easily employ NaDas by providing new customers with slightly over-dimensioned gateways, whose extra computation, storage, and bandwidth resources are used to host services, all of which will be totally isolated from the end user via virtualization technologies.

Home router sharing based on trust (HORST) (Seufert et al (2013)) is a mechanism which addresses the data offloading use case and combines it with mechanisms for content caching/pre-fetching and content delivery. HORST establishes a user-owned Nano Data Center (uNaDa) on the home router and sets up two WiFi networks (SSIDs) - one for private usage and one for sharing. The owner of the home router shares the WiFi credentials with trusted users via an OSN application and can also request access to other shared WiFis. As HORST knows the location of the users and the WiFis, it can recommend near shared WiFi networks, and automatically request access and connect the users for data offloading. HORST combines content placement for mobile users with data offloading and uses social information in order to predict which content will be requested by which user. As HORST also knows about the current and predicts future users of each shared WiFi from location data, the uNaDa on the home router can be used to cache or pre-fetch delay-tolerant content which will be delivered when the user is connected to the WiFi.

QoE and Energy Aware Mobile Traffic Management (QoEnA) (Kaup and Hausheer (2013)) is a mechanism focused on the improvement of QoE, at the same time reducing the energy consumption on mobile devices by intelligent scheduling of network traffic which is generated on the mobile device. It is based on QoS maps, user mobility prediction, energy models, and QoE models. Thereby, QoEnA schedules traffic which is generated on the mobile device to different connections or locations in order to improve the QoE of the end users while reducing the energy consumption of the mobile device.

Social information can also be used for routing and content placement in mobile ad hoc networks. In Costa et al (2008), a routing framework for publish-subscribe services is described. In such a service, messages (or content items) are tagged with topics and shall be routed to users that are interested in these topics. In the presented framework, predictions of co-locations are based on metrics of social interactions, because socially bound hosts are likely to be co-located regularly. For each message, a best carrier is selected based on interests, mobility, and co-location prediction, to whom the message is for-

warded. The presented socially-aware approach is shown to have advantages in terms of message delivery, delay, and overhead.

Dinh et al (2009) work towards socially-aware routing for mobile ad hoc networks. They present an algorithm to identify modular structures in dynamic network topologies based on interactions, and merge them to a compact representation of the network. This compact representation is well suited for dynamic networks and allows for a faster computation of routing strategies compared to state-of-the-art algorithms.

3.3 Network Security

OSNs can provide valuable social information that can be employed to come up against malicious users and their behavior which leads to large scale attacks, e.g., sybil or DDoS attacks which are described below.

Social information about trust between users can be used both for the self-protection of the OSN and for the protection of other services or applications.This information can be either extracted from the OSN itself, or by creating a graph of trust among the end users of a service or application, i.e., by employing a system in which each user has the ability to create relationships of trust with other users.

In both approaches, users can be represented as nodes of a social graph where an edge between two nodes implies a both-way relationship of trust.

A sybil attack (Douceur (2002)) occurs when a malicious user takes on a large number of identities and pretends to be multiple, distinct users/nodes. When these sybil nodes collude together and comprise a large fraction of systems identities, the attacker gains significant advantage in a distributed system. For example, sybil nodes can work together to distort reputation values, out-vote legitimate nodes in consensus systems, or corrupt data in distributed storage systems.

In order to avoid sybil attacks, the social activity of various nodes as well as their social relationships can be examined in order to verify fake profiles and identify potential malicious nodes in a system. Based on the fact that a social network is fast mixing (Nagaraja (2007)) the social graph can help to reveal malicious users, while this becomes easier as the number of fake (malicious) identities increases. This is due to the fact that it is difficult for a malicious user to establish multiple social relationships between the sybil nodes and real users.

According to Yu et al (2006), sybil nodes form a well-connected subgraph that has only a small number of edges connected to honest users, as depicted in Figure 6. These edges are also called attack edges to the honest network.

As a counter-measure, SybilGuard (Yu et al (2006)) exploits this property of the social graph to identify sybil nodes by finding this small cut and by bounding the number of sybil nodes a malicious user can create. Sybil-

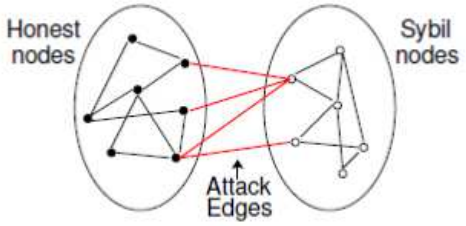

Fig. 6 The complete social graph consisting of the honest network, the sybil network, and the attack edges connecting these two networks. (Source: Yu et al (2006))

Guard relies on a special kind of verifiable random walk in the graph and intersections between such walks. These walks are designed so that the small cut between the sybil region and the honest region can be recognized and used to identify malicious users.

A Distributed Denial of Service attack (DDoS) is another use case, which occurs when multiple systems usually controlled by one malicious entity (e.g., botnets) flood the bandwidth or resources of a targeted system, e.g., a server, in order to make the system unavailable to its intended users.

A similar solution to SybilGuard can be developed against DDoS attacks. The social graph can be used by OSNs or third-parties (e.g., video streaming platforms, or banking institutions) to reveal fake profiles and identify potentially malicious users. In a DDoS attack, a malicious entity may be in control of multiple systems belonging to real users, e.g., by means of trojan horses, and therefore making it unable to detect these problematic profiles through the social graph.

We can overcome this obstacle by observing/monitoring the social behavior/activity of end users; whenever significant changes appear in their activity (e.g., high increase of requests for video viewing), the user should be added in a "suspects' list" implying that some or all of his requests are being denied. Additionally, the same solution can be used in applications or systems which do not contain social information. This can be achieved by asking users to sign up with a social network account or by encouraging them to declare other users in the system that they trust or are socially connected to. Maintaining such a suspects' list, we may avoid or relieve the impact of a DDoS attack.

Nonetheless, there are open issues to be addressed by future research. For example, there is the possibility that also honest users are included in the suspects' list, and thus get denied of the service. Moreover, there is a trade-off between the efficiency of the monitoring of the social behavior of users against the high monitoring effort.

4 Conclusions and Discussions

The research field of socially-aware traffic management opens new perspectives for improved service delivery in the Internet like the discussed use cases content storage, global service mobility, or network security. Nevertheless, the utilization of social information introduces new interdisciplinary research challenges, such as the monitoring of social data, the processing and storage of these data, as well as the integration into existing systems. From the network traffic management perspective, it is unclear how to realize and deploy socially-aware traffic management solutions. In particular, the design and implementation of socially-aware networking functionalities involves several stakeholders of the service delivery chain, like the social information provider, Internet service providers, cloud operators, application providers, and finally the end user. Hence, there must be an incentive-compatible network management mechanism (Hoßfeld et al, 2009) which satisfies the requirements of involved stakeholders (e.g., high QoE for end users, low traffic/congestion in ISPs' links, lower energy consumption in datacenters where services run), and which is based on well-defined open protocols as currently defined by the IETF ALTO "Application-Layer Traffic Optimization" Working Group in Alimi et al (2013). Furthermore, a seamless integration of those socially-aware mechanisms into today's Internet applications and network management is desired. Such architectural and conceptual challenges are currently developed in the FP7 SmartenIT[6] for a tighter integration of network management and service functionality to offer a large business potential for all players involved. An initial architecture with incentives as integral part is presented in Hausheer and Rückert (2013) which follows a modular concept and existing standards and proposals.

From a social data analysis perspective, there are also some practical aspects addressed in order to obtain, maintain, and update social signals from existing platforms due to the huge amount of existing data. While social data can be utilized for various use cases, the monitoring of the data differs in terms of temporal and spatial scale. Different sources for social information can be taken into account and it has to be decided which information to be retrieved from which source and when. Depending on the actual use case, it may be necessary to consider aggregated information, single or selected users, or even the entire OSN topology. Moreover, the monitoring frequency may address different timescales (hours, days, months). In general, there is a trade-off between accuracy and costs of social information, which may be adjusted by appropriate (temporal and spatial) sampling methods. This means that the social monitoring has to be customized for a specific data source from which relevant data is fetched. Then, preprocessing, aggregation, and analysis of

[6] The FP7 project SmartenIT (FP7-2012-ICT-317846) "Socially-aware Management of New Overlay Application Traffic combined with Energy Efficiency in the Internet" is running from Nov 2012 - Oct 2015. More information is available at http://www.smartenit.eu.

the data is necessary before sending the resulting social information to traffic management elements in a system. Especially the design of such algorithms, e.g., to identify relevant nodes in the network responsible for video cascades, and the computational complexity have to be addressed, since scalability is one of the key issues of socially-aware traffic management. Finally, besides those technical challenges, privacy is another major challenge which has to be ensured and integrated in the solution space.

For overcoming the emerging challenges, the tight coupling between social data analysis and the resulting traffic management solutions is required, while socio-informatics is foreseen as a driver to establish an interdisciplinary research community in that interesting domain.

Acknowledgements This work was partly funded by Deutsche Forschungsgemeinschaft (DFG) under grants HO 4770/1-1 and TR257/31-1, and in the framework of the EU ICT Project SmartenIT (FP7-2012-ICT-317846). The authors alone are responsible for the content.

References

Alimi R, Penno R, Yang Y (2013) ALTO Protocol. Tech. rep., Internet Engineering Task Force Application-Layer Traffic Optimization Working Group, URL http://tools.ietf.org/wg/alto/

Bakshy E, Hofman JM, Mason WA, Watts DJ (2011) Everyone's an Influencer: Quantifying Influence on Twitter. In: Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM '11), New York, NY, USA

Biancani M, Cruschelli P (eds) (2013) Deliverable D1.2 Report on Cloud Service Classifications and Scenarios, SmartenIT Consortium (European FP7 STREP No. 317846)

Cisco (2012) Cisco Visual Networking Index: Forecast and Methodology, 2011-2016. Tech. rep., Cisco

Costa P, Mascolo C, Musolesi M, Picco GP (2008) Socially-aware Routing for Publish-Subscribe in Delay-tolerant Mobile Ad Hoc Networks. IEEE Journal on Selected Areas in Communications 26(5):748–760

Dinh TN, Xuan Y, Thai MT (2009) Towards Social-aware Routing in Dynamic Communication Networks. In: Proceedings of the 28th IEEE International Performance Computing and Communications Conference (IPCCC), Phoenix, AZ, USA

Douceur JR (2002) The Sybil Attack. In: Peer-to-peer Systems, Springer, pp 251–260

Fiedler M, Hossfeld T, Tran-Gia P (2010) A Generic Quantitative Relationship Between Quality of Experience and Quality of Service. IEEE Network 24(2):36–41

Ganley B (2013) Optimize the Virtual Desktop Experience Through Strong Back-end Design. Tech. rep., Dell Power Solutions, URL http://i.dell.com/sites/doccontent/business/solutions/power/en/Documents/ps4q13-20130371-ganley.pdf

Gautam N, Petander H, Noel J (2013) A Comparison of the Cost and Energy Efficiency of Prefetching and Streaming of Mobile Video. In: Proceedings of the 5th Workshop on Mobile Video (MoVid '13), New York, NY, USA

Hausheer D, Rückert J (eds) (2013) Deliverable D3.1 Report on Initial System Architecture, SmartenIT Consortium (European FP7 STREP No. 317846)

Hoßfeld T, Hausheer D, Hecht F, Lehrieder F, Oechsner S, Papafili I, Racz P, Soursos S, Staehle D, Stamoulis GD, Tran-Gia P, Stiller B (2009) An Economic Traffic Management Approach to Enable the TripleWin for Users, ISPs, and Overlay Providers. In: Towards the Future Internet - A European Research Perspective, Future Internet Assembly, pp 24–34

Ickin S, Wac K, Fiedler M, Janowski L, Hong JH, Dey AK (2012) Factors Influencing Quality of Experience of Commonly Used Mobile Applications. IEEE Communications Magazine 50(4):48–56

Kaup F, Hausheer D (2013) Optimizing Energy Consumption and QoE on Mobile Devices. In: Proceedings of the IEEE International Conference on Network Protocols (ICNP 2013), Göttingen, Germany

Li Z, Shen H, Wang H, Liu G, Li J (2012) SocialTube: P2P-assisted Video Sharing in Online Social Networks. In: Proceedings of the IEEE INFOCOM, Orlando, FL, USA

Mamatas L, Psaras I, Pavlou G (2010) Incentives and Algorithms for Broadband Access Sharing. In: Proceedings of the ACM SIGCOMM Workshop on Home Networks, New Delhi, India

Mell P, Grance T (2011) The NIST Definition of Cloud Computing. Tech. rep., Recommendations of the National Institute of Standards and Technology

Nagaraja S (2007) Anonymity in the Wild: Mixes on Unstructured Networks. In: Privacy Enhancing Technologies, Springer, pp 254–271

Ruhela A, Tripathy RM, Triukose S, Ardon S, Bagchi A, Seth A (2011) Towards the Use of Online Social Networks for Efficient Internet Content Distribution. In: Proceedings of the IEEE 5th International Conference on Advanced Networks and Telecommunication Systems (ANTS), Bangalore, India

Sastry N, Crowcroft J, Sollins K (2007) Architecting Citywide Ubiquitous Wi-Fi Access. In: Proceedings of the 6th Workshop on Hot Topics in Networks (HotNets), Atlanta, GA, USA

Sastry N, Yoneki E, Crowcroft J (2009) Buzztraq: Predicting Geographical Access Patterns of Social Cascades Using Social Networks. In: Proceedings of the 2nd ACM EuroSys Workshop on Social Network Systems (SocialNets), Nuremberg, Germany

Scellato S, Mascolo C, Musolesi M, Crowcroft J (2011) Track Globally, Deliver Locally: Improving Content Delivery Networks by Tracking Ge-

ographic Social Cascades. In: Proceedings of the 20th International Conference on World Wide Web (WWW2011), Hyderabad, India

Seufert M, Burger V, Hoßfeld T (2013) HORST - Home Router Sharing based on Trust. In: Proceedings of the Workshop on Social-aware Economic Traffic Management for Overlay and Cloud Applications (SETM 2013), Zurich, Switzerland

Traverso S, Huguenin K, Triestan I, Erramilli V, Laoutaris N, Papagiannaki K (2012) TailGate: Handling Long-Tail Content with a Little Help from Friends. In: Proceedings of the 21st International Conference on World Wide Web (WWW2012), Lyon, France

Valancius V, Laoutaris N, Massoulié L, Diot C, Rodriguez P (2009) Greening the Internet with Nano Data Centers. In: Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (Co-NEXT '09), Rome, Italy

Vidales P, Manecke A, Solarski M (2009) Metropolitan Public WiFi Access Based on Broadband Sharing. In: Proceedings of the Mexican International Conference on Computer Science (ENC 2009), Mexico City, Mexico

Wang Z, Sun L, Wu C, Yang S (2012) Guiding Internet-scale Video Service Deployment Using Microblog-based Prediction. In: Proceedings of the IEEE INFOCOM, Orlando, FL, USA

Wittie MP, Pejovic V, Deek L, Almeroth KC, Zhao BY (2010) Exploiting Locality of Interest in Online Social Networks. In: Proceedings of the 6th International Conference on Emerging Networking Experiments and Technologies (Co-NEXT '10), Philadelphia, PA, USA

Yoshida H (2011) The Differences Between Tiering and HSM. URL http://blogs.hds.com/hu/2011/05/the-differences-between-tiering-and-hsm-hierarchical-storage-management.html

Yu H, Kaminsky M, Gibbons PB, Flaxman A (2006) SybilGuard: Defending Against Sybil Attacks via Social Networks. ACM SIGCOMM Computer Communication Review 36(4):267–278

Zhou F, Zhang L, Franco E, Mislove A, Revis R, Sundaram R (2012) WebCloud: Recruiting Social Network Users to Assist in Content Distribution. In: Proceedings of the 11th IEEE International Symposium on Network Computing and Applications (NCA), Cambridge, MA, USA