

Socially-Aware Traffic Management*

Michael Seufert¹, George Darzanos², Valentin Burger¹, Ioanna Papafili², Tobias Hoßfeld¹

¹ University of Würzburg, Institute of Computer Science, Würzburg, Germany
Email: {seufert | valentin.burger | hossfeld}@informatik.uni-wuerzburg.de

² Athens University of Economics and Business, Department of Informatics, Athens, Greece
Email: {ntarzanos | iopapafi}@aueb.gr

Abstract: Socially-aware traffic management exploits social signals 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. In the talk we show interdisciplinary efforts between communication networks and social network analysis. Specifically, we give an overview on existing use cases and solutions, but also raise discussions at the workshop on additional benefits from the integration of social information into traffic management.

1 Introduction

In online social networks (OSNs) users voluntarily provide information about their lives, 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 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 behaviours, 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 paper will provide an insight to this new topic.

In this extended abstract we will only focus on the use case video streaming being the most prominent example of overlay applications, and present socially-aware caching

*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.

solutions and algorithms. In the talk we show interdisciplinary efforts between communication networks and social network analysis. Specifically, we give an overview on existing use cases and solutions, but also raise discussions at the workshop on additional benefits from the integration of social information into traffic management.

2 Example Use Case: Video Streaming

Concerning video as a key application contributing largely to the overall IP traffic [Cis12], video and specifically user generated content (UGC) sharing (e.g. home-made videos) has evolved to a major trend in OSNs.

We consider an OSN having users around the globe who upload UGC to an online video streaming platform such as YouTube and share the content (URLs) via the OSN where it can be viewed by their online friends and their friends of friends, etc. 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. Each user is assigned and served out of their (geographically) nearest PoP, for all of his requests.

Placing data close to the users is an approach followed by most content delivery networks to improve Quality of Experience for the end user and reduce costs for video platform providers. Therefore all content uploaded by a user A is first uploaded to the nearest respective PoP_A . When content is requested by another user B, the nearest PoP to B, i.e. PoP_B , is contacted and the request is served if the content is available. If the content is not available in PoP_B , a request is made to PoP_A and the content is brought to PoP_B . The scenario considered in this section is inspired by the evaluation scenario described below in [THT⁺12].

3 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 [SYC09] 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 [SMMC11] 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.

In [THT⁺12], the authors propose Tailgate which addresses long-tailed content (photos, video) by deriving and using social information and meta-information from Twitter, such as social relationships, regularities in read access patterns, and time-zone differences so as to predict where and when the content will likely be consumed. The content is then pushed where-ever before it is needed.

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

Social awareness can be used in different ways to improve traffic management. To exploit the information of social networks it is important to understand how information is spread in social networks and how to identify important persons and relationships.

In [BHMW11], 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 [RTT⁺11], 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.

Next, in [WPD⁺10], the authors inferred the network structure of Facebook from crawls, 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.

Finally, in [WSWY12], 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.

5 Discussion

The research field of socially-aware traffic management opens new perspectives for improved service delivery in the Internet, but also new interdisciplinary research challenges. From the network traffic management perspective, it is unclear how to realize 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 OSN provider, Internet service providers, cloud providers offering video services, and finally the end user. Hence, there must be an incentive-compatible network management mechanism which satisfies the requirements of involved stakeholders (e.g. high Quality of Experience for end users, low traffic/congestion in Internet service providers' links, lower energy consumption in data centers where video servers run), and which is based on well-defined open protocols like IETF ALTO. 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 (Nov 2012 - Oct 2015) for a tighter integration of

network management and service functionality to offer a large business potential for all players involved.

From a social network 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. There is a tradeoff between accuracy and costs of social information, which may be adjusted by appropriate (temporal and spatial) sampling methods. Further, the computational complexity of some algorithms, e.g. to identify relevant nodes in the network responsible for video cascades, has to be addressed. Additionally, a major challenge is privacy which has to be ensured and has to be integrated in the solution space.

In our talk we want to deepen this topic, provide additional use-cases, and trigger discussions on current and future interdisciplinary efforts between communication networks and social network analysis to advance socially-aware traffic management in the Internet.

Literatur

- [BHMW11] Eytan Bakshy, Jake M. Hofman, Winter A. Mason und Duncan J. Watts. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the 4th ACM international conference on Web search and data mining, WSDM '11*, Seiten 65–74, New York, NY, USA, 2011. ACM.
- [Cis12] Cisco. Cisco Visual Networking Index: Forecast and Methodology, 2011-2016. Bericht, Cisco, 2012.
- [RTT⁺11] Amit Ruhela, Rudra M. Tripathy, Sipat Triukose, Sebastien Ardon, Amitabha Bagchi und Aaditeshwar Seth. Towards the use of Online Social Networks for Efficient Internet Content Distribution. In *Advanced Networks and Telecommunication Systems (ANTS), 2011 IEEE 5th International Conference on*, Seiten 1–6, 2011.
- [SMMC11] Salvatore Scellato, Cecilia Mascolo, Mirco Musolesi und Jon Crowcroft. Track Globally, Deliver Locally: Improving Content Delivery Networks by Tracking Geographic Social Cascades. In *20th Intl. Conf. on World Wide Web*, 2011.
- [SYC09] Nishanth Sastry, Eiko Yoneki und Jon Crowcroft. Buzztraq: predicting geographical access patterns of social cascades using social networks. In *Proceedings of the 2nd ACM EuroSys Workshop on Social Network Systems*, Seiten 39–45. ACM, 2009.
- [THT⁺12] Stefano Traverso, Kévin Huguenin, Ionut Triestan, Vijay Erramilli, Nikolaos Laoutaris und Kostantina Papagiannaki. Tailgate: handling long-tail content with a little help from friends. In *21st Intl. Conf. on World Wide Web*, 2012.
- [WPD⁺10] Mike P. Wittie, Veljko Pejovic, Lara Deek, Kevin C. Almeroth und Ben Y. Zhao. Exploiting Locality of Interest in Online Social Networks. In *Proceedings of the 6th International Conference, Co-NEXT '10*, Seiten 25:1–25:12, New York, NY, USA, 2010. ACM.
- [WSWY12] Zhi Wang, Lifeng Sun, Chuan Wu und Shiqiang Yang. Guiding internet-scale video service deployment using microblog-based prediction. In *INFOCOM, 2012 Proceedings IEEE*, Seiten 2901–2905, 2012.