Future Network & MobileSummit 2013 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2013
ISBN: 978-1-905824-37-3

SmartenIT Cloud Traffic Management Approach and Architectural Considerations

Ioanna Papafili[1], Julius Rückert[2], Patrick Poullie[3],
Sergios Soursos[4], Thomas Bocek[3], Krzysztof Wajda[5],
David Hausheer[2], George D. Stamoulis[1], Burkhard Stiller[3],
[1]*Athens University of Economics and Business,*
Department of Informatics, Athens, Greece
[2] *Technische Universität Darmstadt, Department of Electrical Engineering and Information Technology, Darmstadt, Germany*
[3]*Universität Zürich, Department of Informatics, Zürich, Switzerland*
[4] *Intracom SA Telecom Solutions, Peania, Greece*
[5]*AGH University of Science and Technology,*
Department of Telecommunications, Krakow, Poland

Abstract: The Internet has seen a strong move to support new overlay applications offered over the cloud, which demand for a coherent and integrated control in underlying heterogeneous networks in a scalable, resilient, and cost-efficient manner. A tighter integration of network management and cloud service functionality can lead to cross-layer optimization of operations and management, thus, being a promising approach to offer a large business potential in operational perspectives for all players involved. In this paper the SmartenIT approach is presented, which targets at a cross-layer traffic management by providing incentives to the involved stakeholders, *i.e.*, providers of cloud-based applications, network operators, cloud operators, and end-users, to ensure a QoE-awareness. SmartenIT designs and develops traffic management solutions and mechanisms to meet accordingly special application requirements and to address efficiently load and traffic patterns, exploiting also social information in terms of user relations, interests and behaviour patterns. Additionally, the energy efficiency with respect to both end-user devices and underlying networking and cloud service provisioning infrastructure is tackled to ensure an operationally efficient management. Furthermore, this paper describes three scenarios of major interest to SmartenIT and discusses the possible impact of the SmartenIT approach on metrics of interest associated to each involved stakeholder. Finally, a conceptual architectural framework is introduced and major design axes and describe key functionalities are described; this conceptual architecture constitutes the basis for the development of the SmartenIT architecture.

Keywords: Cloud, applications, network, traffic management, incentives, social-awareness, energy efficiency, distributed management, stakeholders.

1. Introduction

New overlay applications which include content delivery platforms, such as YouTube, online storage systems, *e.g.*, Dropbox, Google Drive, and online social networks such as Facebook and Twitter, become increasingly popular among Internet users and generate traffic volumes that constitute an even larger share of global IP traffic [1]. Additionally, new overlay applications are also highly popular in mobile devices, *e.g.*, smartphones, tablets and phablets, whose number is rapidly increasing; even though these devices are optimized for portability and long battery life, they have shown in the past limited computing and storage resources. Thus, due to their lack of computational and storage resources, as well as the need for high availability in order to serve simultaneously millions of users distributed across the globe, new overlay applications for mobile devices are served over the cloud.

Cloud computing has emerged as a major paradigm for the current and the Future Internet (FI) [2]. Reasons are manifold and include high decentralization of services implying increased resilience and availability [3, 4], increased benefit and economies of scale due to resource sharing among many users, lower operation and maintenance costs including lower energy consumption [5], easier administrative and maintenance task, and eventually ease of access to various applications, content and data from different locations and devices. Consequently, a large variety of new overlay applications, *i.e.*, applications served over the cloud spanning resources that are geographically scattered, already exist and gain increasingly high popularity; the support of those in the wired and wireless access network domain determines a challenge and a commercial must.

However, enormous traffic volumes and traffic patterns generated by popular cloud-based applications greatly influence performance and costs (both OPEX and CAPEX) of involved stakeholders. In particular, network operators are required to handle extremely high and unpredicted traffic volumes passing through their networks, perhaps through expensive transit links, affecting also other types of traffic. Cloud operators are obliged by Service Level Agreements (SLAs) to meet high Quality-of-Service (QoS) requirements for the services they offer, while they are burdened by high energy consumption and severe maintenance tasks in their data center facilities. End-users, both corporate and home ones, may experience increase latency, or service outages; while application and services providers may need to deal with potentially poor Quality-of-Experience (QoE) delivered to their subsequently unsatisfied end-users.

Moreover, cloud-based applications are expected to generate even more traffic in the years to come; this increase will also be driven by the tremendous increase of mobile devices [1]. Additionally, highly diverse QoE requirements of these applications as well as the devices used to access them result in very different communication path quality levels and management tasks. To successfully address such requirements, traffic management mechanisms have to support effectively scalability, agility, stable QoE, and flexibility, in an integrated set of network and overlay service functionality.

The design and implementation of such mechanisms is the target of SmartenIT (Socially-aware MAnagement of new oveRlay application Traffic combined with Energy Efficient in the InterneT). SmartenIT addresses these aforementioned issues by handling the explosion of traffic generated by overlay applications, services, and respective end-points in the wired and wireless domain, cross-cutting multiple operator's domains in support of combined intra-domain cloud services; proposing an agile network management and operation framework, while addressing a better control of heterogeneous networks and their network functionality and integrating overlay service functionality in support of stable QoE; exploiting social information and meta-information on users' relationships, interests and their behavior patterns; and eventually, evaluating the new highly distributed mechanisms, optimized control, and energy efficiency aspects.

In the remainder of this paper, related approaches in literature in Section 2 are qualitatively assessed, while, in Section 3 describes in detail the SmartenIT approach, the methodology to be followed, and major scenarios of interest, where also stakeholders characterization is briefly discussed. Section 4 presents SmartenIT's preliminary architectural considerations and design axis. Finally, Section 5 draws main concluding remarks and discusses future steps.

2. Related Work

Traffic management solutions in the transport or network layer are well known and well described throughout the literature. The two classic approaches are Integrated Services

(IntServ) [23] and Differentiated Services (DiffServ) [24, 25]. While the former allows for fine-grained traffic management, it lacks scalability; while on the other hand, DiffServ provides a scalable mechanism for managing network traffic and providing QoS on IP (Internet Protocol) networks, however it is more coarse-grained. Then, in order to provide QoS-guaranteed differentiated services across Autonomous Systems (AS), Multi-Protocol Label Switching (MPLS) was proposed [26]. MPLS deploys short labels to establish virtual paths and thereby allow rapid packet switching. Thus, the focus is now laid on an overview on related work on traffic management solutions applied in the application layer.

Concerning traffic management in higher layers, SmoothIT [9] focuses on techniques to align overlay and underlay networks. SmoothIT promotes the exchange of information between the physical network and Peer-to-Peer (P2P) overlay networks and end-users, providing them incentives to cooperate so as to achieve an all-win situation. SmoothIT's decentralized and incentive driven approach makes it highly relevant for SmartenIT. Similarly to SmoothIT, IETF Application-Layer Traffic Optimization (ALTO) tries to align overlay and underlay networks as well [10].

The NetStitcher approach [11] targets at an efficient handling of inter-data center traffic. Its aim is to use paid-for but unutilized bandwidth of links inter-connecting data centers by using it for backup, bulky updates propagation, and data migration. Addressing energy efficiency for data centers, the Nano Data centers (NaDa) are introduced at ISPs' gateways [12] and are therefore more economic with respect to heat dissipation, traffic localization, and energy efficiency (main objective for SmartenIT).

Following similar rationale regarding scheduling with the NetStitcher, Tailgate [13] addresses user generated long-tail content exploiting also information and meta-information derived by social networks, so as to efficient content placement. Another approach exploiting social information is SocialTube [14], which addresses content distribution over social networks, *i.e.*, video sharing systems over Facebook. SocialTube is a P2P Video-on-Demand (VoD) system which splits users into followers and non-followers of single video owner based on extensive observations of their behavior, and using a push-based pre-fetching mechanism improves content delivery. Moreover, WebCloud [15] provides a concrete implementation to improve content distribution in social networks. Video sharing performance is improved through a proxy server's assistance.

A different approach involves the implementation of a cloud-based cooperative cache system [16] to make provision of data-intensive processes cheaper and faster, and focuses on highly requested content. In [17], cloud performance is pursued to be increased by caching strategies, considering also the retrieval capacity of edge nodes and servers in hybrid content cloud architectures. Moreover, CloudAngels in [18] aims at reaching the theoretical lower bound of the Minimum Distribution Time (MDT) in peer-assisted content delivery. To increase the aggregate upload speed in the overlay and thereby improve the MDT, Angels, *i.e.*, dedicated helper nodes, are deployed, while complex swarming strategies are also discussed. CloudAngels assumes that all nodes in the cloud are cooperative, *i.e.*, can be actively choreographed by the content provider to reduce MDT.

Finally, Software-Defined Networking (SDN) is considered, which is an emerging technology enabling smart configuration and optimization of physical and virtualized networks. SDN allows for a flexible traffic management in a more dynamic way by software-based algorithms and mechanisms. A popular specification that enabled SDNs is OpenFlow [29], which allows dynamic updates and adaptations in the routing and switching behavior through a software-based controller.

Although traffic management and economic traffic management are well researched areas across all network layers, additional factors such as energy efficiency and social-aware aspects do play also an important role in traffic management. In case of energy efficiency, the traffic should not only be optimized to find the fastest/cheapest path, but it should also be optimized to use the most energy efficient one, *e.g.*, use smart caching based on social aspects to reduce the traffic.

3. SmartenIT Approach and Application Scenarios

The SmartenIT project aims at the design of incentive compatible network management mechanisms for overlay providers, network providers, and users to ensure QoE-awareness. To achieve these goals, the SmartenIT consortium builds on the concept of incentives and economic-awareness, social-awareness, and energy efficiency; these are the *key concepts* and *design axes* for the SmartenIT's functionality and architecture, which are to be analyzed below in Section 4 as part of the conceptual architecture of the project.

The methodology to be followed within SmartenIT starts from conducting the classification of cloud-based services and investigation and determination of roles of the key stakeholders in delivering cloud-based services to end-users. Next, the identification of business relationships between stakeholders and their complying and/or conflicting interests and incentives for cooperation (such as QoE, monetary incentives, investments, performance, innovation) is to be carried. Then, the combination of advanced economic modeling and specialized network management mechanisms in literature will lead the design of new traffic management mechanisms, the development of a novel architecture and the implementation of a prototype. Finally, theoretical proposals will be evaluated by extensive simulations and test-bed trials in certain scenarios.

Thus, as a starting point, the SmartenIT consortium has envisioned three scenarios of interest that will be employed to validate the SmartenIT architecture, management approach, and functionality and to show the key potential of the SmartenIT approach for the involved stakeholders. These three scenarios and the potential impact of SmartenIT's mechanisms when applied in each one of them are described below.

Inter cloud/inter data center communication and collaboration for energy efficiency. Inter cloud/inter data center communication is considered for data replication to a backup facility, such as another data center located at a remote location, so as to increase fault tolerance and enhance customers' QoE, *e.g.*, the scenario of [11]. It should be noted that data centers may be in general served by multiple ISPs; their interconnection is possible due to the transit ISPs that sell global connectivity network service. This transit service is seamless to both the data centers and the end customers who are purchasing Internet access from their edge ISPs. Alternative scenarios are also possible such as the communication of two data centers for the exchange of data needed to perform computing operations over massive datasets, *e.g.*, using a medical data database in the cloud.

Moreover, the collaboration of data centers for energy efficiency is also addressed within SmartenIT. Data centers may either create a decentralized energy ecosystem, or a centrally-orchestrated one, *e.g.* by a third-party as an Energy broker, where the workload can be exchanged/migrate between data centers to address energy shaping requests from the energy providers, economical demands (lower costs of energy in certain locations) or to prioritize access to stable renewable energy sources. In this scenario, depicted in Figure 1, the information about energy consumption and related implications can be included in SLAs between service users and service providers, while Operational Level Agreements (OLAs) need to be established between data centers for workload migration and access to distributed energy sources.

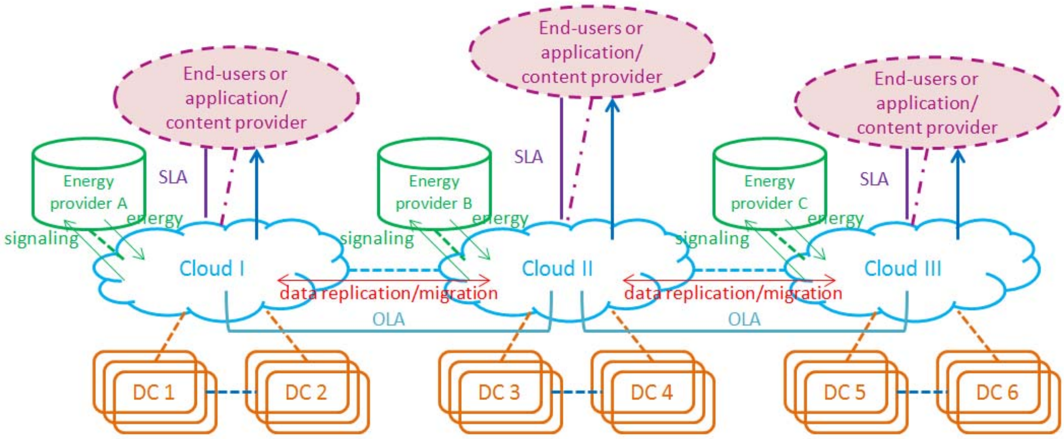

Figure 1: Data center collaboration for energy-efficiency.

For these scenarios, SmartenIT will elaborate on traffic management mechanisms, *e.g.*, alternative route selection, to efficiently perform exchange of data between storage and computational resources at low operational cost, in terms of both inter-connection costs and low energy consumption, simultaneously assuring high users' QoE. Furthermore, the stakeholders analysis performed in [28] indicates that the control, market offering, and orchestration of the end-to-end inter data center communication and/or orchestration for energy efficiency could be performed either by the data center owners, which would integrate a SmartenIT functionality, or alternatively by a third party entity as an Energy broker or a SmartenIT Over-The-Top (OTT) provider.

Social meta-/information exploitation for content delivery. In this scenario, users of a social network, *e.g.*, Facebook, are located around the globe and upload User-Generated Content (UGC) such as home-made videos to a video streaming platform, *e.g.*, YouTube, as illustrated in Figure 2. This content can be viewed by these users' online friends or their friends' friends. Content delivery is performed by the video platform, which is operated on a geo-diverse system comprising multiple Points-of-Presence (PoPs) distributed globally.

In this setup, the exploitation of social information can lead to more sophisticated traffic management, reduction of the traffic load on the inter-domain links, reduction of the operating costs for the ISPs, and significant performance of users' QoE, *i.e.*, decrease latency. Questions that need to be answered by SmartenIT is *what* meta-information can be used, under *which* circumstances can it be acquired, *e.g.*, collaboration with OSN (On-line Social Network) providers, and *how* to utilize it in a way mutually beneficial for all parties.

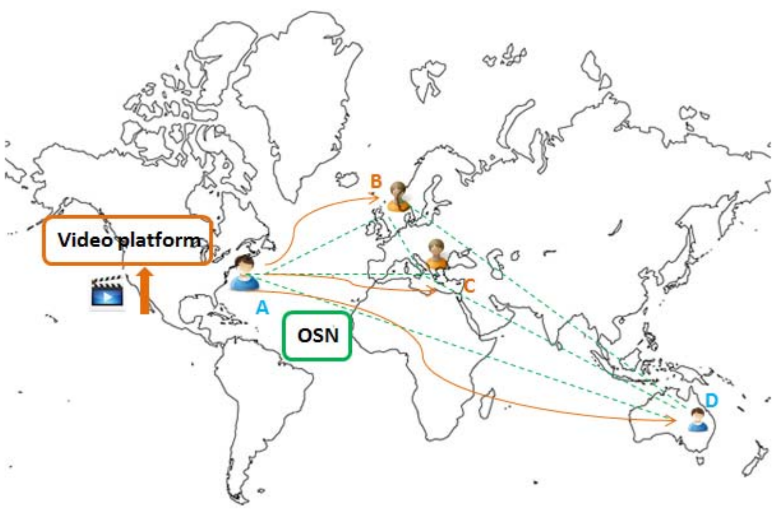

Figure 2: Content dissemination of UGC over an OSN.

The stakeholder analysis in [28] suggests that a SmartenIT mechanism could be seen either as a set of architectural modules embodied within video delivery platform, or alternatively, such a module could be employed by a new entity, *e.g.*, an OTT Provider, and placed among the existing stakeholders, so as to collect social information and provide it to the CDN or the ISP, or the End-Users themselves in the peer-assisted setup.

Global service mobility. Global service mobility is related to the services' ability to follow user's location. In more detail, the global service mobility goal is to serve content and applications to end-users with high availability and high QoE levels, for example by the use of coordinated CDNs, different cloud providers, or distributed private data centers.

As depicted in Figure 3, there exist two alternative scenarios for global service mobility; these are the *one-time data move*, and the *gradual data move*. Even though both scenarios define different strategies, they can co-exist as different options to enable service mobility. In the one-time data move, the service should fully move its resources to the new location at once, while on the other hand, in the gradual data move, the service should observe (considering a time-frame) whether the user remains in the new location. The gradual data move can be considered a set of smaller one-time data moves but depending on *policies*, *e.g.*, based on most accessed data or most triggered actions in the service.

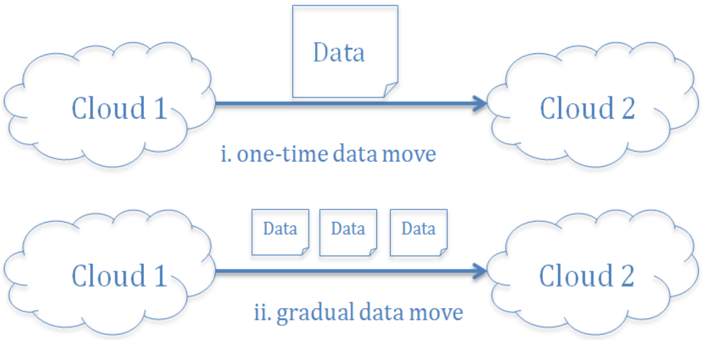

Figure 3: Global service mobility.

Addressing global service mobility proves to be highly challenging, as one may need to decide which resources in the cloud, or even which cloud, to use to provide the service to the end-user taking into account numerous factors, *e.g.*, user's location, costs of traffic transfer, cost of access and cost of usage of a given cloud, cost of migration of data, resources or virtual machines between clouds used to realize the service. The objective of this scenario is to develop and design mechanisms that will achieve seamless mobility of services (*i.e.*, move services close to the end-user) and, most significant, high QoE for end-users in terms of high download throughput and short download times, along with reduction of traffic replicated among data centers and energy consumed thereby. The initial stakeholders' characterization [28] suggests that collaboration should be pursued among the cloud operator and the application provider (*i.e.*, cloud operator's customer) to apply data moving policies in a way mutually beneficial for all. Moreover, global service mobility could also be achieved involving edge ISPs and end-users by a mechanism like NaDas [12].

4. SmartenIT's Conceptual Architecture

As introduced earlier, the SmartenIT approach aims at incentive-compatible cross-layer network management mechanisms involving a number of different entities and spans across several research areas of traditional network and traffic management. Thus, the conceptual architecture for the SmartenIT approach is presented here as the basis for the initial software component architecture. Figure 4 gives an overview of the different involved domains, the three design axes, and the basic concepts and functionality of the SmartenIT approach.

There are three domains, where the SmartenIT solution will apply; plus a forth domain, the application one. All domains are depicted as grey boxes in Figure 4. The application domain includes the cloud-based applications as primary sources of the traffic to be managed. Prominent examples for such applications are cloud storage and video streaming services as well as online social networks. All three impose high requirements for the management mechanisms to be applied in the other three domains:

Data center/Cloud. This domain covers entities and concerns that target the overlay and cloud management. Here, data centers for hosting cloud services but also other overlay-based services are relevant.

Core and fixed/mobile access network. Serving as bridge between data centers and the end user devices, the core and access network exhibits a special role for the management of traffic. It is an aggregation point of traffic from a large number of different services and applications. It unites a huge number of independently managed autonomous systems. The optimization of content delivery, the management of traffic flows, as well as the monitoring of the traffic are relevant aspects covered by this domain. Focus in mobile data delivery and mobile backbone will be also placed here.

End user. End user devices are considered as well as possibilities to manage traffic to obtain a certain service quality for the user.

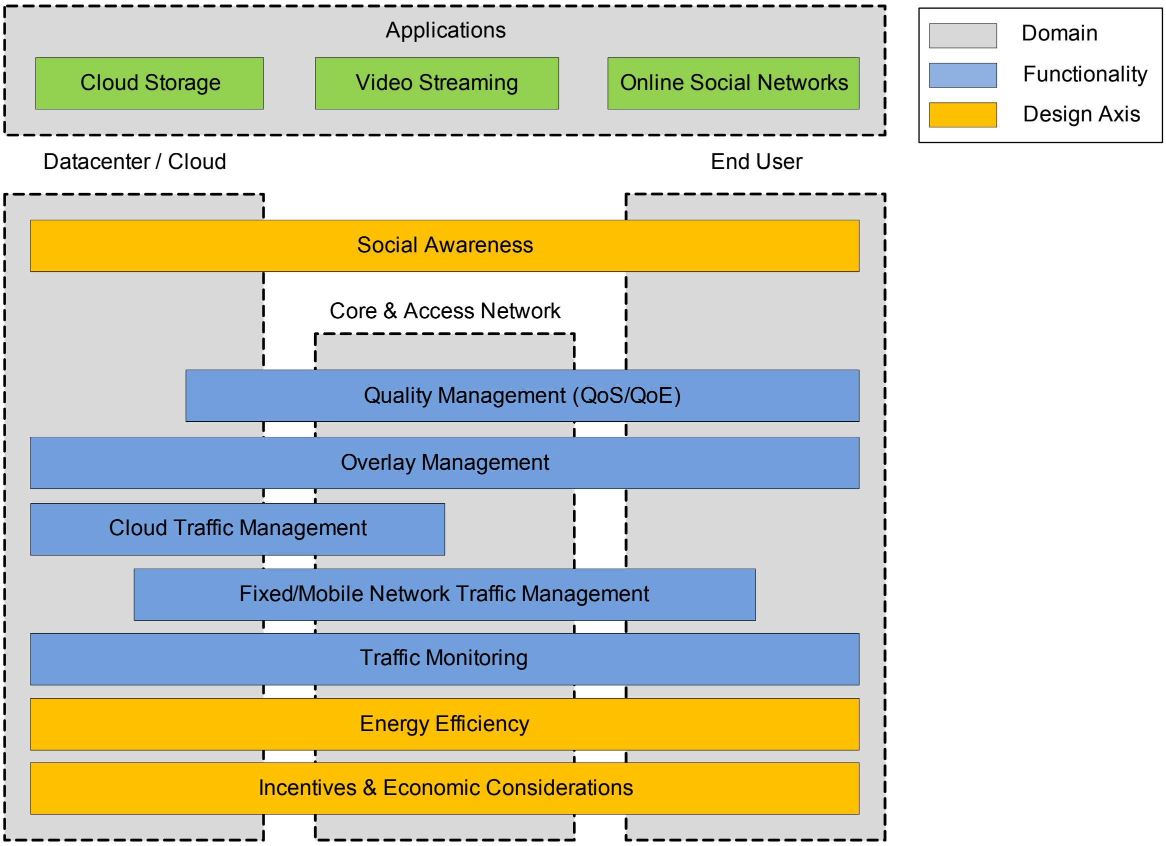

Figure 4: The initial conceptual SmartenIT architecture.

Across these domains the SmartenIT approach follows three design axes, depicted by orange boxes in Figure 4. They are considered for the design and implementation of the newly developed traffic management mechanisms:

Energy efficiency. This axis spans all three main domains of network and traffic management. In all of them there is a need to provide not only operational mechanisms, but also to cause as low energy consumption as possible. At the user side, especially focusing on mobile scenarios, this requires applications that carefully use the resources of the mobile devices, leaving small energy footprints and preserving a long battery lifetime. For the core network, energy considerations are important as a large number of different infrastructure elements are involved. In this context, *e.g.*, mechanisms such as Software Defined Networking (SDN) promise new, especially energy efficient solutions, potentially reducing the energy consumption at, *e.g.*, routers and switches. Finally, energy efficiency for the application provisioning infrastructure, the cloud, and thereby within data centers and its operation are relevant. Altogether, the SmartenIT approach promises to combine the potentials for energy savings for an integrated network and traffic management solution.

Incentives and economic considerations. The SmartenIT architecture will be built on the concept of Economic Traffic Management (ETM) [19] introduced by SmoothIT. This principle suggests that traffic management technologies and solutions need to fulfill incentive compatibility in order to be successfully adopted in practice, namely they need to provide incentives for all stakeholders to cooperate in a joint cross-layer optimization approach, *e.g.*, expose information by providing advice for more efficient handling of resources [20], reward incentive-compatible behavior of players [21], and invest in and deploy infrastructures that imply benefit for other stakeholders, too [22].

Social awareness. SmartenIT will derive and exploit social information (*e.g.*, social graphs and relationships among users) and meta-information (*e.g.*, behavior patterns) in order to predict the generation of traffic, the destination of traffic flows, and possible dissemination patterns, either between end users and the cloud, or between clouds/data centers. To derive such social information, crawling of OSNs constitutes an option, though this approach is inadequate as only limited information can be extracted. To overcome this obstacle, incentives should be provided to OSNs to *expose* such information; the latter is addressed by the second axis.

Within these domains, respecting the three design axes, different *concepts* exist that constitute the core of the SmartenIT approach; they are depicted by blue boxes in Figure 4. The described components provide the basis upon which SmartenIT will derive the component architecture in a later step. Below, major concepts are introduced:

Overlay management. This concept targets at the coordination of overlay nodes running at client devices. This covers overlay nodes to, *e.g.*, coordinate ad-hoc communication between mobile devices, but also includes nano data centers and other communication between clients. A desirable outcome of this concern, for example, is to align the flexible overlay topology with the underlying network infrastructure and to allow for incentive-compatible and energy-efficient communication.

Cloud traffic management. The goal of cloud traffic management is to reduce costs and to achieve an optimized performance of cloud services. Therefore, it needs to consider many aspects ranging from the choice for specific cloud providers in an inter-cloud scenario (Data center/Cloud domain) to the reduction of redundant traffic in the network itself using techniques such as multicast (Core and Fixed/Mobile Access Network Domain).

Network traffic management. Network management is an important concern especially for the core network and, thereby, within ISPs' network. Decisions on optimal traffic engineering, the avoidance of congested links and redundant traffic, and resilient transportation of traffic are key aspects. In this context, the concept of SDN provides a promising approach to design flexible and cost-efficient network and traffic management mechanisms. For traffic management, furthermore, caching is an important topic as it aims at efficient delivery of content. The nature of content originated from the cloud may differ from other traffic, being more dynamic and complex. It is composed of various pieces of data, possibly residing in different data centers around the globe.

Besides the management of the core and fixed access network traffic, mobile access networks impose specific requirements for traffic management. Other infrastructure elements, requirements, and limitations are to be considered in this case. It is to be seen how approaches from the core and access network and traffic management can be extended and integrated with those within this field to achieve a *converged* Fixed/Mobile Network Traffic Management solution.

Traffic monitoring. To take network and traffic management decisions, the monitoring of traffic is an essential prerequisite. An accurate assessment of the network condition and its available resources is important to identify the need for and potential of new management approaches. Therefore, it is important to consider this concern in the designing the management solution following of the SmartenIT approach.

Quality management. The quality management concept estimates objective, *i.e.*, QoS, and subjective, *i.e.*, QoE, aggregates and normalizes it and makes it available to the traffic management mechanisms so as to be considered during the optimization process. Thereby, it considers the anticipated levels of user acceptance, and also improvements when introducing a novel approach. Using the later described Quality Monitor, the estimation of latency and throughput can be done from the viewpoint of the end-user, *e.g.*, by sending messages and using locally available information such as the RTT (Round Trip Time).

The described concepts will constitute major ingredients of the SmartenIT architecture. Some of them can be grouped in more generic categories, *e.g.*, monitoring the effects of a specific traffic management mechanism may include measuring QoS/QoE, traffic and energy consumption simultaneously. Additionally, mobile traffic management extends network traffic management in the sense that aspects related to the devices generating this traffic and their requirements need to be considered.

5. Summary and Next Steps

SmartenIT has identified major issues appearing due to the rapidly increasing adoption of cloud computing to offer services and the high popularity of new applications, *e.g.*, enormous traffic volumes generated and unpredicted traffic patterns, high consumption of energy in data centers, and increasingly higher need for power in end-users' mobile devices. These issues, impacting both performance and costs of involved players, namely network operators, cloud operators, application and content providers, and eventually end-users, are addressed by several approaches in the literature in a fragmented manner only.

Thus, this paper briefly presented the EU SmartenIT project, which has identified the need for a seamlessly integrated, cross-layer, incentive-compatible, social-aware, and energy-efficient network management, and addresses these issues in three representative scenarios. SmartenIT provides a high-level description of these scenarios of interest that cover major types of services offered over the cloud, *e.g.*, inter-cloud communication and energy efficiency, global service mobility and social-awareness, and SmartenIT qualitatively predicts a potential optimization to be achieved, while the potential for a collaboration among stakeholders involved is outlined.

Finally, this paper provides a conceptual description of the SmartenIT architecture and a "blueprint" of its main functionality and design axes. The specification of various modules comprising the architecture and their implementation, prototyping, and evaluation in a test-bed and actual deployments constitute the next steps of SmartenIT's work.

Acknowledgements

This work has been performed in the framework of the EU ICT Project SmartenIT (FP7-2012-ICT-317846). Furthermore, all authors would like to thank all SmartenIT project members for providing insights and conducting useful discussions on the topic of this paper.

References

[1] Cisco, Cisco Global Cloud Index: Forecast and Methodology (2011-2016), 2012.

[2] M. P. Papazoglou, W. van den Heuvel: Blueprinting the Cloud, IEEE Internet Computing, Vol. 15, No. 6, November 2011.

[3] Cisco Secure Network Container: Multi-Tenant Cloud Computing, 2010.

[4] I. Poese, B. Frank, B. Ager, G. Smaragdakis, A. Feldmann: Improving Content Delivery Using Provider-aided Distance Information; Internet Measurement Conference (IMC 2010), Melbourne, Australia, November 2010.

[5] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, F. Zhao, Energy-aware Server Provisioning and Load Dispatching for Connection-intensive Internet Services, 5th USENIX Symposium on Networked Systems Design and Implementation, San Francisco, California, U.S.A.., April 2008, pp 337-350.

[6] M. Welzl, M. Mühlhäuser: Scalability and Quality-of-Service: A Trade-off? IEEE Communications Magazine, Vol. 41, No. 6, June 2003, pp 32-36.

[7] S. Sundaram, A. Abdurakhmanov, Y.-T. Kim: WBEM-based Inter-AS Traffic Engineering for QoS-guaranteed DiffServ Provisioning, 1st International Workshop on Broadband Convergence Networks (BCN 2006), April 2006 pp 1-9.

[8] L. Li, M. M. Buddhikot, C. Chekuri, K. Guo: Routing Bandwidth Guaranteed Paths with Local Restoration in Label Switched Networks, IEEE International Conference on Network Protocols (ICNP 2002), Paris, France, November 2002.

[9] S. Soursos, M. Ángeles Callejo Rodriguez, K. Pussep, P. Racz, S. Spirou, G. D. Stamoulis, B. Stiller: ETMS: A System for Economic Management of Overlay Traffic, In: "Towards the Future Internet - Emerging Trends from European Research", IOS Press, FIA 2010, Valencia, Spain, April 2010, pp 1-10.

[10] RFC5693: Application-Layer Traffic Optimization, Problem Statement, October 2009.

[11] N. Laoutaris, M. Sirivianos, X. Yang, P. Rodriguez: Inter-Datacenter Bulk Transfers with NetStitcher, ACM SIGCOMM 2011, Toronto, Canada, August 2011.

[12] V. Valancius, N. Laoutaris, L. Massoulie, C. Diot, P. Rodriguez: Greening the Internet with Nano Data Centers. 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT 2009), Rome, Italy, 2009.

[13] S. Traverso, K. Huguenin, I. Trestian, V. Erramilli, N. Laoutaris, K. Papagiannaki: Handling Long-Tail Content with a Little Help from Friends, 21st International Conference on World Wide Web (WWW 2012), New York, U.S.A., April 2012.

[14] Z. Li, H. Shen, H. Wang, G. Liu, J. Li: SocialTube: P2P-assisted Video Sharing in Online Social Networks, INFOCOM, 2012 Orlando, Florida, U.S.A., pp 2886-2890.

[15] F. Zhou, L. Zhang, E. Franco, A. Mislove, R. Revis, R. Sundaram: WebCloud: Recruiting Social Network Users to Assist in Content Distribution, 11th IEEE International Symposium on Network Computing and Applications, Cambridge, Massachusetts, August 2012, U.S.A., pp 10-19.

[16] D. Chiu, A. Shetty, G. Agrawal, Elastic Cloud Caches for Accelerating Service-Oriented Computations. ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis. Washington, DC, U.S.A., November 2010.

[17] M. Bjorkqvist, L.Y. Chen, M. Vukolic, Xi Zhang: Minimizing Retrieval Latency for Content Cloud, INFOCOM, 2011 Shanghai, China, April 2011, pp 1080-1088.

[18] R. Sweha, V. Ishakian, A. Bestavros: Angels in the Cloud: A Peer-assisted Bulk-synchronous Content Distribution Service, IEEE International Conference on Cloud Computing (CLOUD), Washington DC, U.S.A., July 2011, pp 97-104.

[19] T. Hoßfeld, D. Hausheer, F. Hecht, F. Lehrieder, S. Oechsner, I. Papafili, P. Racz, S. Soursos, D. Staehle, G. D. Stamoulis, P. Tran-Gia, B. Stiller. An Economic Traffic Management Approach to Enable the TripleWin for Users, ISPs, and Overlay Providers. In: "Towards the Future Internet – A European Research Perspective", FIA Prague Book, IOS Press, FIA 2009, Prague, Czech Republic, May 2009, pp 24-34..

[20] P. Racz, S. Oechsner, F. Lehrieder: BGP-based Locality Promotion for P2P Applications. 19th IEEE International Conference on Computer Communications and Networks (ICCCN 2010), Zürich, Switzerland, August 2010.

[21] K. Pussep, S. Kuleshov, C. Groß, S. Soursos: An Incentive-based Approach to Traffic Management for Peer-to-Peer Overlays. 3rd Workshop on Economic Traffic Management (ETM 2010), Amsterdam, The Netherlands, September 2010, LNCS Vol. 6236.

[22] I. Papafili, S. Soursos, G. D. Stamoulis: Improvement of BitTorrent Performance and Inter-domain Traffic by Inserting ISP-owned Peers. 6th International Workshop on Internet Charging and QoS Technologies (ICQT 2009), Aachen, Germany, May 2009, LNCS Vol. 5539.

[23] RFC 1633: Integrated Services in the Internet Architecture: An Overview, June 1994
[24] RFC 2475: Architecture for Differentiated Services, December 1998
[25] RFC 3260: New Terminology and Clarifications for DiffServ. April 2002
[26] RFC 2702: MPLS Traffic Engineering September 1999
[27] RFC 3469: Framework for MPLS-based Recovery February 2003
[28] SmartenIT Deliverable D1.1: Report on Stakeholders Characterization and Traffic Characteristics, April 2013, http://www.smartenit.eu/publications.
[29] OpenFlow: The OpenFlow Switch Specification 1.1.0, http://openflow.org, April 2013.