

A game-theoretic evaluation of an ISP business model in caching

Jörn Künsemöller¹ · Nan Zhang² · Kimmo Berg³ · João Soares⁴

Published online: 9 January 2016
© Springer Science+Business Media New York 2016

Abstract Internet traffic volume is increasing and this causes scalability issues in content delivery. This problem can be addressed with different types of caching solutions. The incentives of different stakeholders to pay for these solutions are not known. However, it has been identified

that Internet service providers (ISPs) need to be involved in the process of cache deployment due to their ownership of the network. This work evaluates a new business model where ISPs charge content providers (CPs) for a caching service because CPs benefit from more efficient content distribution. We provide conditions for sustainable paid in-network caching and their numerical evaluation in order to aid strategic decision-making by CPs, ISPs, and Cloud storage providers (CSPs). Although ISP caching as a paid service may not be an equilibrium, it turns out to be Pareto optimal at the right pricing. This encourages cooperation between CPs and ISPs. CSPs may choose cache friendly physical locations for their facilities in order to provide the necessary capacity to the ISPs. However, the required amounts are in all likelihood too small to be an incentive for the CSPs. ISP caching as a paid service can be an equilibrium when future benefits are considered and when the ISPs terminate caching-related improvements of service quality for clients who do not pay for caching.

The work of J. Künsemöller is funded by the German Research Foundation (DFG) within the Research Training Group Automatism and the Collaborative Research Center On-The-Fly Computing (SFB 901). The work of N. Zhang has been performed in the framework of the CELTIC-Plus project C2012/2-5 SIGMONA. The work of J. Soares is funded by the Portuguese Foundation of Science and Technology (No SFRH/BDE/51102/2010). K. Berg acknowledges the funding from Emil Aaltosen Säätiö through Post doc -pooli.

✉ Jörn Künsemöller
jkuensem@mail.uni-paderborn.de

Nan Zhang
nan.zhang@aalto.fi

Kimmo Berg
kimmo.berg@aalto.fi

João Soares
joao-m-soares@telecom.pt

Keywords Caching · Business model · Cloud storage · Content distribution

1 Introduction

Internet traffic has increased over the last years not only because of a growing user base, but also because data intensive services, such as video streaming, have become more common. This development will presumably continue, e.g. global Internet traffic is expected to triple between 2014 and 2019, where 80 percent of consumer Internet traffic will be video (Cisco 2015). The delivery of such content causes a lot of data transfer in the Internet backbone. This traffic can be reduced by caching technologies, which prevent

¹ Computer Networks Group, University of Paderborn, EIM-I, Warburger Str. 100, Paderborn, Germany

² Department of Communications and Networking, Aalto University, Otakaari 5A, Espoo, Finland

³ Department of Mathematics and Systems Analysis, Aalto University, Otakaari 1F, Espoo, Finland

⁴ Portugal Telecom Inovação e Sistemas, Rua Eng. José Ferreira Pinto Basto, Aveiro, Portugal

repeated transport of the same data over long distances and also provide a better user experience due to lower latencies. While caching is of benefit to both the Internet service provider (ISP) and the content provider (CP), they utilize caching independently. ISP's may cache for internal optimization and many CPs make use of specialized content delivery networks (CDNs) to improve user experience.

This paper analyzes a business model for a cooperative caching solution, where ISP and CPs share the caching costs and, in consequence, the ISP caches more data. This reduces traffic and, hence, cost of the ISP's network and increases the distribution of content at the same time (since additional copies are deployed in the network). The question is whether or not such a business model can prevail in the market and under what conditions it may exist.

We provide general observations and conditions for sustainable paid caching in content-ISP networks as well as an evaluation of these conditions based on real-world numbers. This provides knowledge about market interaction, which is important for strategic decisions by CPs, ISPs, and cloud storage providers (CSPs), while the numerical examples help to assess the impact of changing magnitudes like service prices and costs. Paid ISP-side caching can serve as an alternative to a costly continuous expansion of ISP network capacities or questionable methods like bandwidth throttling (Grove et al. 2013), and may be combined with other traffic reduction approaches like content optimization (Wong 2013). Further, our studies are of high importance in a cloud context. While today cloud computing relies on the power of big data centers, studies have already pointed out the benefits of having distributed clouds, e.g. (Soares et al. 2012). In such environments, ISP caching and (CDNs) can be flexibly managed according to demand by making use of cloud storage facilities.

Game theory is used for the analysis, because of its ability to consider behavior and interaction of several actors. We propose game-theoretic models for both the feasibility of the business models as well as the potential resource allocation in the cloud. For cloud resources, we investigate how ISP caching may affect the location of cloud storage facilities. Further, we study the long-term incentives of such a paid ISP caching service in a repeated game.

This paper extends and improves the presentation of a previously published conference paper on the new business models of ISP caching (Künsemöller et al. 2013). We further contribute original research on resource allocation and on the long-term incentives for ISP caching, which were not discussed in Künsemöller et al. (2013). Additionally, quantitative evaluations are performed for the business model game and the resource allocation game.

The remainder of this paper is organized as follows. Related research is presented in Section 2. Section 3 gives important background information about caching today

and describes the considered business models. Two game-theoretic models are set up in Section 4 in order to analyze the feasibility of the business models as well as the resource allocation in the cloud. Section 5 investigates Nash equilibrium and Pareto optimality conditions in these two games. An evaluation of these conditions is presented in Section 6. The long-term incentive of CPs to pay for ISP caching is discussed in Section 7 and Section 8 concludes the paper.

2 Related work

In the past, wholesale market failure, due to the end-to-end quality of service coordination failures at the IP layer, the information asymmetries between ISPs and CPs and the significant costs related to contracting with different players, have resulted in the entry of CDN providers into the content delivery market (Faratin 2007). Today, the importance of the ISPs' involvement in cache deployment is becoming evident as CPs (e.g. Netflix), who have been using CDNs for a long time, are now deploying their own caches within the ISP's network (Netflix 2015). Some state-of-the-art works study the ISP's involvement in the caching process either as ISP-operated CDNs (Wulf et al. 2010; Cho et al. 2011), close cooperation of CDN and ISP (Frank et al. 2013), or ISP transparent caching (Katsaros et al. 2011; Psaras et al. 2012; Carofiglio et al. 2013; Kimmerlin et al. 2014; Wang et al. 2014). However, most neglect the business aspects and target the study more from a resource perspective. For example, Cho et al. (2011) presents a technical solution for ISP-operated CDN, where the efficient use of network resources is analyzed.

In Pham (2015), similar to our analysis, ISP caching is investigated in a game-theoretic approach with regard to ISPs' and CPs' economic interests. However, it focuses on a comparison of caching quantities in equilibrium in the Internet versus an information-centric networking scenario, while we investigate ISP caching as an optional enriched Internet service.

Several papers study possible cooperation between ISPs and CPs at a control level by utilizing game theory. The authors in Jiang et al. (2009) look into different approaches that an ISP can take in managing traffic engineering and server selection, ranging from running the two systems independently to designing a joint system. The surprising conclusion from this work is that in the case of two independent systems, extra visibility between the two systems results in a less efficient outcome. Server selection and traffic engineering is also studied in DiPalantino and Johari (2012). Although these works study the cooperation between ISPs and CPs, it is important to highlight that their focus is on cooperation from a control perspective, while ours is on cooperation in costs of the caching system.

There is also important research that investigates other interactions in the market. Dán (2011) studies how the cooperation between ISPs can influence transit traffic costs with respect to the cached content in a scenario, where ISPs have caching capabilities. Two game-theoretic models for cooperative caching are put forward: one, where the ISPs follow a selfish strategy and another, where the interests of the neighboring ISPs are also taken into account. The results show that by cooperating, ISPs can achieve considerable gains, even if they follow a selfish strategy. The gains can further increase, when also taking the neighboring ISPs' interests into consideration. Khare and Zhang (2011) also aims at a more efficient routing with a combination of a non-uniform bandwidth pricing by the ISP and a CDN-side cost-aware routing. Other works focus on the self-interaction within ISPs or CPs themselves, e.g. Shakkottai and Srikant (2006), Lee et al. (2008), and Shrimali et al. (2010).

3 Background

The main benefit of caching lies in the reduction of redundant traffic, which is caused by repeated requests and delivery of the same data. Especially for the ISPs, caching means less transit or peering costs due to reduced traffic volumes flowing outside their networks. In this paper, traffic flowing outside an ISP's network is defined as distant traffic, whereas traffic within the ISP's network is considered as local traffic. In addition, caching may also reduce local traffic, however, due to its minor significance compared to distant traffic costs, this paper ignores this effect.

From a CP's perspective, caching may reduce latency for its end users due to the proximity of the cache servers to the end users. Similarly, the end user perceives the ISP's service quality to improve with caching.

As a consequence of the increasing importance of caching, different existing caching technologies are operating in parallel. In addition, new technologies that utilize caching are being developed. This section gives a brief introduction to the basic concepts, such as the Internet topology and caching technologies, and explains the assumptions adopted in this paper.

3.1 Internet topology

ISPs typically operate in tiers (Labovitz et al. 2010), where the Tier-1 ISPs form the Internet backbone and offer transit services to lower tier ISPs. The lower tier ISPs can be divided into content heavy networks (content ISPs) and eyeball heavy networks (eyeball ISPs), based on the customers they serve (Faratin et al. 2008). Content ISPs serve mostly content providers and generate outbound traffic, whereas eyeball ISPs have end users as customers and incur more

inbound traffic. As this paper analyzes the business relationship between content providers and ISPs, the focus is on content ISPs and referred to simply as ISPs for the rest of the paper.

For full connectivity, the lower tier ISPs can either buy transit from upper tier ISPs or utilize peering agreements (Norton 2014). The transit provider sells access to its entire routing table and usually charges the ISPs based on usage by using the 95th percentile measurement method. On the other hand, peering agreements provide access only to the peer's network and have traditionally been settlement free. However, if the peers do not derive equal value from the peering relationship, paid peering can be advocated.

Due to the importance of content delivery, the logical topology of Internet is changing and large content providers (e.g. Google) and CDNs (e.g. Akamai) are increasingly forming their own backbone networks (Labovitz et al. 2010). The ISPs can then directly peer with the large CPs or CDNs for better quality of service and cost savings from reduced transit traffic volume.

3.2 Caching technologies

3.2.1 Web caching

Web caching can be considered as the first caching technology in the market. The demand for web caching became evident over a decade ago (Barish and Obraczka 2000) when the usage of the World Wide Web (WWW) increased dramatically. The idea is to temporarily cache web sites in the proxy servers or in end users' browsers to more efficiently serve the subsequent requests. In addition to the ISP controlled proxy caching, web caches can also be deployed by the content providers closer to the content servers to offer origin server load balancing. The difference compared to CDNs and cloud storage lies in web caching being limited to HTTP traffic as well as the temporary nature of the cache storage. Furthermore, web caching is done in a transparent manner, which means that the ISPs cache web pages without any agreements with the CPs.

3.2.2 Content delivery networks

A CDN (Dilley et al. 2002) operates as an overlay to the basic Internet and divides the end-to-end connection into two pieces: one between the CP and the CDN servers, the other between the CDN servers and the end users. The CDN provider co-locates its data centers into the ISP's network and caches the CP's content based on the contract type: either the content is cached after it is requested for the first time or the CP can push certain content into the cache before it is requested (Dilley et al. 2002). In addition, this paper

assumes that when CPs use CDNs, all content from these CPs are served from the CDN servers.

Traditionally, CDNs (Vakali and Pallis 2003) are operated by third-party CDN providers, which are here called pure-play CDNs. In addition, CDNs used to have settlement-free peering agreements with smaller ISPs for co-locating the data centers (Faratin 2007). However, the relationships are changing and ISPs are increasingly charging CDN providers for the co-location service (Level3 2010; der Veen 2011). Other changes are also taking place today: for example, ISPs and CPs are increasingly building their own CDN networks (Telefonica 2014; AT&T 2015; Google 2015; Fitchard 2012). As a response, the pure-play CDN providers are offering CDN licenses to ISPs (Akamai 2015a). In addition, the CDN providers are working towards interconnectivity between themselves through initiatives such as CDNi (Peterson and Davie 2014).

3.2.3 Cloud storage

Cloud computing (Vaquero et al. 2009) is a paradigm for better and easier hardware and software management. Clouds are pools of virtualized resources, such as software, hardware and services, that can be easily accessed. The idea of the cloud is to move the infrastructure to the network, which reduces the costs of resource management and offers better scalability and flexibility.

The cloud paradigm offers mainly three service categories: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS). In the caching case, only IaaS is relevant, where a CSP virtualizes its resources so that they can be split and assigned dynamically to the customers. Typically, the customer is charged only for the actual used storage and the service level agreements (SLAs) guarantee the quality of service.

3.2.4 In-network caching

Furthermore, in-network caching schemes, such as ISP-driven caching, information-centric networking (Ahlgren

et al. 2012) and SDN optimized in-network caching (Costa-Requena et al. 2014), are widely researched. With in-network caching, the content is cached in the network elements, e.g. routers and servers, when it passes through.

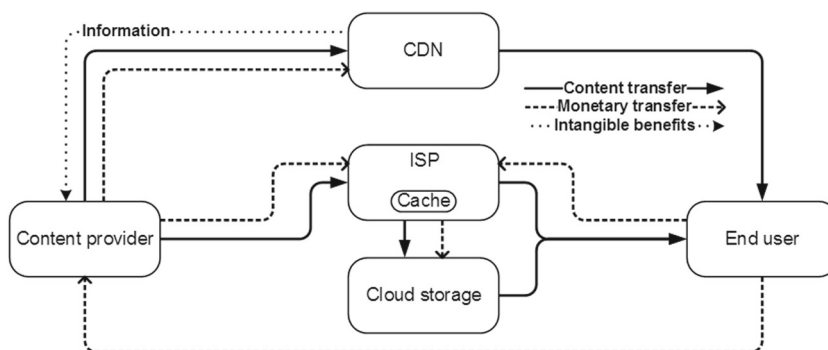
In ISP-driven caching, the ISPs place cache servers or caching enabled routers into their own network and cache the content either transparently or according to agreements with CPs. In addition, the ISP can choose to utilize a third-party storage provider (e.g. a CSP) or build their own caching infrastructure. Information-centric networking and SDN optimized caching operates with a similar concept. However, in information-centric networking, routing is done based on content names instead of host addresses (Ahlgren et al. 2012). Additionally, SDN optimized in-network caching utilizes a cache controller to relocate cache locations depending on the end users' requests (Costa-Requena et al. 2014).

3.3 Assumptions

This paper assumes a simplified content delivery ecosystem with only CPs, ISPs, CDNs, CSPs and end users, the value network of which is illustrated in Fig. 1. The value network shows the exchanges between each of the stakeholders divided into 1) content transfer, 2) monetary transfer and 3) intangible benefits, where the intangible benefits include, for example, information on the end users' preferences, customer loyalty and brand recognition. The two caching schemes considered in this paper are ISP-driven caching and pure-play CDNs due to their high impact in the current content delivery market.

We assume the existence of peering or transit agreements between multiple ISPs, which allow one ISP to offer caching services that go beyond its network. Thus, the CP has a business relation with only one ISP, which we can assume to be its local one. The transit ISP and the ISP serving the end users are not illustrated in Fig. 1 for simplicity reasons. In addition, this paper assumes that the ISPs do not have existing caching infrastructure and need to build the caching service before offering it to the CPs.

Fig. 1 Value network of a simplified Internet content delivery ecosystem



From an ISP's perspective, it has two strategic decisions regarding caching. First, it has to decide whether to cache or not. If the ISP decides to cache, it has to decide how to price the caching service. The third decision relates to whether to buy caching services from a third party (e.g. CSP) or build their own caching infrastructure. The first two decisions combine into three potential business models: 1) Basic service, 2) ISP internal network optimization and 3) ISP-driven caching service. The three business models serve as the basis for the business model game introduced in Section 4.1. and are briefly explained below. In addition, the resource allocation game presented in Section 4.2. discusses the third decision of the ISP, i.e. leasing vs. building the caching system.

3.3.1 Business model 1: basic service

The first business model represents a situation in which the ISP decides not to deploy caching and stays in its traditional market: access provision and traffic transmission. In this business model, the ISP charges the CPs only for the network access and offers a best-effort service.

If the CP wishes to improve the quality of experience (QoE) to its end users, it can either deploy its own caching system or buy the service from a CDN. For example, Google is a content provider that has its own caching system (Google 2015) and MTV Networks uses Akamai's services (Akamai 2015b).

3.3.2 Business model 2: ISP-internal network optimization

Business Model 1 does not fully comply with the current situation of the network, because most ISPs employ caching at some level, e.g. web caching. Thus, the second business model explains a situation in which the ISP caches content, but does not charge CPs for the caching service. The main incentive for the ISP is in reducing costs through optimizing its own network and reducing transit traffic volume. In addition, the content providers may see an increase in the QoE perceived by their end users and, thus, more CPs might use the service offered by the ISP that does caching.

In this situation, the CP pays the ISP for the network access and the traffic volume in the traditional way. The difference compared to Business Model 1 is that the CP does not have a direct relationship with the CDNs, though the ISPs could also outsource the actual caching to CDNs or CSPs. Whether the ISP builds its own caching network or outsources the service depends on the cost efficiencies of the solutions, since the CP does not pay for the caching service in either case.

3.3.3 Business model 3: ISP-driven caching service

In the third business model, ISPs are offering caching services to the CPs for an extra fee. The CPs can be charged based on the amount of cached data, the traffic volume generated by the caches or the combination of the two. In all three charging models, the CP contracts only with the ISP. However, the actual caching could be done by a third party as explained above. We assume that the ISP charges by bandwidth, since the ISP pays for the bandwidth required for transit as well and ISPs are trying to save in transit costs by caching the content. The acceptable range of caching related payment for the CP and the profitable range of prices charged by the ISPs are evaluated in Section 6.

4 Game-theoretic setup

This section provides two separate game models that represent the interaction of ISPs and CPs regarding the different business models and the resource allocation interaction between ISPs and CSPs.

4.1 Business model game

In the following, we want to identify conditions under which the business models that are presented in Section 3 can exist as an equilibrium in a market situation. We make use of game theory as it can model market dynamics by considering the incentives of all actors.

A simple two-player game with ISP and CP as the players is set up to compare their different action options with respect to the options of the other one. The ISP can choose to either route all data requests to the CP or to install caches and meet requests from there. We neglect the time that is needed to deploy the caching service for now and address this issue in Section 7. The CP can choose between a traditional Internet service (Business Models 1 and 2), a service that involves a payment for caching (Business Model 3) and the utilization of a CDN (ISP's competition for caching). The resulting situation depends on the decisions of both parties. Each decision combination features a specific utility for each player as a measure of how valuable the resulting situation is (higher is better).

Table 1 depicts the game setup in the normal form. The upper left corner represents Business Model 1, the upper right represents Business Model 2. In the middle left, the CP is willing to pay for caching but the ISP does not decide to cache. Business Model 3 can be found in the middle right. U and V denote CP's and ISP's utilities, respectively.

All utilities depend on billing: the ISP prefers a higher payment from the CP, while for the CP a lower payment

Table 1 Utilities in the payment model game

		ISP	
		DON'T CACHE	CACHE
CP	TRADITIONAL	U_1 / V_1	U_2 / V_2
	PAY FOR CACHING	U_3 / V_3	U_4 / V_4
	CDN	U_5 / V_5	U_6 / V_6

Each decision combination results in the situation of specific value of CP and ISP

is more valuable. The ISP's utility is reduced by any operational expenses. Service fees that a player charges for its service are denoted as prices. Other financial factors like operational expenses are denoted as costs. When one player pays the other, the charged price affects the utility of one player negatively and the other player's utility positively. The same price variable is used with different sign to reflect that; the important difference to cost variables is that the players can arbitrarily set the prices, while the costs are factors that are determined by the environment. In the following, the utility functions for the different outcomes of the game model are presented.

The traditional Internet service payment depends on the required bandwidth b . We assume that this bandwidth consists of a short routing distance component b_{local} and a long routing distance component b_{distant} ($b = b_{\text{local}} + b_{\text{distant}}$). Both b_{local} and b_{distant} cause costs in the local network (c_{local}), while b_{distant} also accumulates additional transit costs (c_{distant}).

When the ISP is caching, all distant demand can be met without transit traffic by a nearby cache. In reality, some amount of data is most likely delivered over the full distance until it is cached. We consider this amount of traffic as negligible compared to the overall traffic. Hence, in the case of caching, the requests that are otherwise associated with b_{distant} never reach the CP but are entirely served from a local cache.

The ISP charges p_{isp} as a price for bandwidth. In Business Model 1, this utility is decreased by costs in the local network and long-distance transfer (1). In Business Model 2, the ISP has to pay for storage instead of long-distance transfer, and c_{storage} is the cost of data hosting (2). The CP only pays for b_{local} and benefits from better QoE. QoE^+ is the value of QoE improvement perceived by the CP.

$$\begin{aligned} U_1 &= -b \cdot p_{\text{isp}} \\ V_1 &= b \cdot (p_{\text{isp}} - c_{\text{local}}) - b_{\text{distant}} \cdot c_{\text{distant}} \end{aligned} \quad (1)$$

$$\begin{aligned} U_2 &= -b_{\text{local}} \cdot p_{\text{isp}} + \text{QoE}^+ \\ V_2 &= b_{\text{local}} \cdot p_{\text{isp}} - b \cdot c_{\text{local}} - c_{\text{storage}} \end{aligned} \quad (2)$$

The CDN option comes with a service fee for bandwidth p_{cdn} and storage $p_{\text{cdn-storage}}$. We assume that all data requests (including local requests) are met by the CDN and the CP does not obtain any bandwidth from the ISP directly. Similar to ISP caching, all demand can be met from a local CDN server without the long-distance transfer. For example, Akamai (Akamai 2015c) and Netflix (Netflix 2015) peer directly with ISPs, and CPs only deal with Akamai or Netflix. This provides a good user experience irrespective of the ISP's action. We further assume the QoE to be the same as with ISP caching. The ISP saves long-distance transfer costs and we assume that it charges the same bandwidth price from the CDN as it usually charges from the CP (3). Caching when the content is already distributed over the network obviously only causes costs to the ISP without any benefit for either parties (4).

$$\begin{aligned} U_5 &= -b \cdot p_{\text{cdn}} - p_{\text{cdn-storage}} + \text{QoE}^+ \\ V_5 &= b \cdot (p_{\text{isp}} - c_{\text{local}}) \end{aligned} \quad (3)$$

$$\begin{aligned} U_6 &= U_5 \\ V_6 &= V_5 - c_{\text{storage}} \end{aligned} \quad (4)$$

When the service stipulates a caching payment, this does not imply that the ISP actually decides to cache. As the payment is usage-based, no caching fees have to be paid when the ISP is not caching. In this case, the utilities for the service are the same as those of the traditional service when we assume the same bandwidth price (5). In Business Model 3, where caching actually takes place, utilities are based on the service fee for cache bandwidth p_{caching} that the ISP charges from the CP.

$$\begin{aligned} U_3 &= U_1 \\ V_3 &= V_1 \end{aligned} \quad (5)$$

$$\begin{aligned} U_4 &= U_2 - b_{\text{distant}} \cdot p_{\text{caching}} \\ V_4 &= V_2 + b_{\text{distant}} \cdot p_{\text{caching}} \end{aligned} \quad (6)$$

4.2 Resource allocation game

When an ISP adopts caching, it can either build its own caching infrastructure or buy the caching capability from third parties. These third parties include traditional hosting service providers and CSPs. This section investigates using game theory whether the utilization of third-party hosting or the operation of own caching facilities is more profitable to the ISP.

The basic idea is that third parties can offer storage cheaper due to better economies of scale, especially compared with smaller ISPs, but the ISPs may have to compromise on the cache location. Storage providers, on the other hand, might consider this possibility in their data center site selection in order to gain ISPs as their customers. However, in this paper, despite the reduced flexibility in

location choices, the caching system from a third party is assumed to offer the same QoE for end users as an ISP’s own caching system. (Section 5.6 discusses the impact on the results when the assumption is that QoE suffers from outsourcing to a third party.) In addition, if the third party is cheap enough and savings over an own caching facility exceed extra traffic costs, the ISP has an incentive to use the third-party storage provider. Both the ISP and the third-party provider have to benefit from a situation, where the third party is involved in ISP caching or this is not likely to happen.

Another two-player game with ISP and a third party CSP as the players is set up. As discussed in Section 3, the ISP might operate own storage equipment or utilize third party facilities (e.g. a cloud storage service). The CSP can either optimize economies of scale or partition its facilities in order to place them in several locations within the ISP’s network that are more appropriate for caching. Table 2 shows the normal form of the game. W denotes CSP’s utility.

The utilities depend on the price that is asked for some amount of storage, storage costs and the amount of stored data. The model distinguishes the data that the CSP stores for other clients d_{tp} and the amount of data that the ISP caches d_{isp} . Storage in its own facilities costs the ISP an amount of $c_{storage}$. The CSP asks $p_{tp-storage}$ for storage and the production cost of CSP storage is denoted by $c_{tp-storage}$. When the location is chosen for best size, the production costs of the CSP are reduced due to economies of scale. The cost reduction compared to the location that is best for caching is represented by the coefficient EoS. For instance, $EoS = 0.9$ means a 10 percent decrease of production cost, $EoS = 1$ means that no economies of scale apply. Accordingly, the CSP’s utility varies with the location of its facilities, while the ISP’s utility is not affected as long as the ISP uses its own storage facilities (7 and 8).

$$\begin{aligned} W_7 &= d_{tp} \cdot (p_{tp-storage} - c_{tp-storage} \cdot EoS) \\ V_7 &= -d_{isp} \cdot c_{storage} \end{aligned} \tag{7}$$

$$\begin{aligned} W_9 &= d_{tp} \cdot (p_{tp-storage} - c_{tp-storage}) \\ V_9 &= -d_{isp} \cdot c_{storage} \end{aligned} \tag{8}$$

When the CSP serves the ISP for caching, the additional demand increases the CSP’s revenue (9 and 10). Instead of the cost for own equipment, the ISP pays a service fee. The ISP has additional network transfer cost $c_{tp-transfer}$, when

the CSP chooses to optimize economies of scale and places facilities in a relatively remote location. This affects all distant data requests, which is $b_{distant}$. We assume that the CSP’s facility is still close enough, so that this does not have a significant impact on the quality of experience to the end user. We also assume that the ISP does not charge the CSP for cache related traffic, since the ISP is in the user role in that case and would be charged back by the CSP anyway.

$$\begin{aligned} W_8 &= (d_{tp} + d_{isp}) \cdot (p_{tp-storage} - c_{tp-storage} \cdot EoS) \\ V_8 &= -d_{isp} \cdot p_{tp-storage} - b_{distant} \cdot c_{tp-transfer} \end{aligned} \tag{9}$$

$$\begin{aligned} W_{10} &= (d_{tp} + d_{isp}) \cdot (p_{tp-storage} - c_{tp-storage}) \\ V_{10} &= -d_{isp} \cdot p_{tp-storage} \end{aligned} \tag{10}$$

5 Game analysis

Section 4 presented two game-theoretic models regarding ISP-driven caching. This section identifies the conditions for Nash equilibria and Pareto-efficient outcomes in these games. For the business game outcomes, a notation (CP’S DECISION, ISP’S DECISION) is used and in the resource allocation game, the outcomes are denoted as (CSP’S DECISION, ISP’S DECISION).

Note that the games are not completely independent from each other and decisions in one game can affect the other. For example, the ISP’s choice in resource allocation affects the cost of storage in the business model game and the ISP’s choice whether or not to cache affects the amount of resources to allocate. However, the effect of one game’s outcome only shows in the values of the other game’s constants. This means that the formal analysis of the two games can be conducted independently, since it holds true irrespective of any specific values. Nevertheless, keep in mind that it is possible that the outcome of one game can affect whether the required condition for e.g. a Pareto-efficient outcome of the other game can be met. Refer to Section 6 for numerical evaluation of the formal conditions that result from the following analysis.

5.1 Equilibrium conditions in the business model game

A Nash equilibrium describes a situation where no player can unilaterally deviate to any better outcome (Fudenberg and Tirole 1991). Now, we examine which outcomes can be equilibria in the business model game and under which conditions.

(CDN, DON’T CACHE) is an equilibrium, if $U_5 \geq U_1 = U_3$ (CP should not deviate) and $V_5 \geq V_6$ (ISP should not deviate). The second is always true because caching facilities do have a cost, i.e. $c_{storage} \geq 0$. The first condition requires that the CP values the QoE improvement more than

Table 2 Utilities in the cache hosting game

		ISP	
		OWN INFRASTRUCTURE	THIRD PARTY
CSP	BEST SIZE	W_7 / V_7	W_8 / V_8
	BEST LOCATION	W_9 / V_9	W_{10} / V_{10}

the additional costs incurred to the CDN compared to the traditional service:

$$\Leftrightarrow \begin{aligned} U_5 &\geq U_1 \\ \text{QoE}^+ &\geq b \cdot (p_{\text{cdn}} - p_{\text{isp}}) + p_{\text{cdn-storage}}. \end{aligned} \quad (11)$$

Business Model 1 (TRADITIONAL, DON'T CACHE) is an equilibrium, if $U_1 = U_3 \geq U_5$ and $V_1 \geq V_2$. The first condition is exactly the opposite of the previous one; it holds when Inequality (11) is false. Second, the caching of a data object by the ISP is more expensive than the difference between transfer costs of associated bandwidth and the loss of sales of this bandwidth:

$$\Leftrightarrow \begin{aligned} V_1 &\geq V_2 \\ c_{\text{storage}} &\geq b_{\text{distant}} \cdot (c_{\text{distant}} - p_{\text{isp}}). \end{aligned} \quad (12)$$

Business Model 2 (TRADITIONAL, CACHE) is an equilibrium, if $U_2 \geq U_4$, $U_2 \geq U_6$ and $V_2 \geq V_1$. The last condition is again the opposite of what we had before, and it is satisfied when Inequality (12) does not hold. The first condition is satisfied trivially, since the price of caching is positive, i.e. $p_{\text{caching}} \geq 0$. The second condition requires that the ISP's price for local bandwidth is smaller than the overall price for the CDN:

$$\Leftrightarrow \begin{aligned} U_2 &\geq U_6 \\ b \cdot p_{\text{cdn}} + p_{\text{cdn-storage}} &\geq b_{\text{local}} \cdot p_{\text{isp}}. \end{aligned} \quad (13)$$

We expect this to hold always, since we assumed that the CDN provider has to pay the ISP's bandwidth price. If the CDN provider were to set the prices so low that Inequality (13) holds, then its service would not be profitable.

Business Model 3 (PAY FOR CACHING, CACHE) is an equilibrium, if $U_4 \geq U_2$, $U_4 \geq U_6$ and $V_4 \geq V_3$. The first condition never holds and, thus, Business Model 3 cannot be an equilibrium. However, the last condition is satisfied, when the profit that the ISP might have from charging the CP directly for the bandwidth is smaller than the profit from caching. This is the case, when the cost for storage is lower than the increase in revenue:

$$\Leftrightarrow \begin{aligned} V_3 = V_1 &\leq V_4 \\ c_{\text{storage}} &\leq b_{\text{distant}} \cdot (c_{\text{distant}} - p_{\text{isp}} + p_{\text{caching}}). \end{aligned} \quad (14)$$

Finally, (PAY FOR CACHING, DON'T CACHE) can be an equilibrium, if $U_3 \geq U_5$ (Condition 11 does not hold) and $V_3 \geq V_4$ (Condition 14 does not hold). This happens, when the ISP prefers not to cache over Business Model 3 and the CP prefers the ISP services over the CDN. This outcome is equivalent to Business Model 1, when considering the utilities and the caching situation. Though, it can be regarded as a failed Business Model 3, since the CP apparently intends to benefit from caching, but caching does not take place.

We can make the following observations. Whenever the ISP prefers Business Model 2 over 1 (Inequality 12 is not met), it also prefers Business Model 3 over (PAY FOR CACHING, DON'T CACHE) (Inequality 14). Whenever the

ISP prefers (PAY FOR CACHING, DON'T CACHE) over Business Model 3 (Inequality 14 is not met), it prefers Business Model 1 over 2 (Inequality 12). Business Model 1 and 3 can be preferred by the ISP at the same time with the different CP actions, which may be the most interesting case; see Section 5.5.

5.2 Pareto optimality in the business model game

An outcome is called Pareto-optimal or Pareto-efficient, when no other outcome can be found that would improve one player's utility without making the others worse off (Pardalos et al. 2008). We examine under which conditions Business Model 3 is Pareto-optimal.

Pareto optimality requires that there is no (U_i, V_i) such that $U_i \geq U_4$ and $V_i \geq V_4$ and at least one of the inequalities should be strict. A sufficient condition is that $U_i < U_4$ or $V_i < V_4$ for all $i \neq 4$. Since $V_2 < V_4$ and (CDN, DON'T CACHE) Pareto dominates (CDN, CACHE) (since $U_5 = U_6$ and $V_5 \geq V_6$), we do not need to consider cases when $i = 2$ or $i = 6$. We also do not need to consider case $i = 3$, since the utilities are the same as when $i = 1$. Thus, Business Model 3 is Pareto optimal, when Business Model 1 and (CDN, DON'T CACHE) offer a lower utility to either party:

$$(U_1 < U_4 \text{ or } V_1 < V_4) \text{ and } (U_5 < U_4 \text{ or } V_5 < V_4).$$

The condition for $V_1 < V_4$ was presented in Eq. 14. $U_1 < U_4$ holds, when the improved user experience is more valuable to the CP than the difference of caching and Internet service price (15). This depends on the individual case and at a given price level, some CPs probably meet this condition while others do not.

$$\begin{aligned} U_1 &< U_4 \\ \text{QoE}^+ &> b_{\text{distant}} \cdot (p_{\text{caching}} - p_{\text{isp}}) \end{aligned} \quad (15)$$

For $V_5 < V_4$, the difference in revenue from caching and Internet service has to exceed the caching cost:

$$\Leftrightarrow \begin{aligned} V_5 &< V_4 \\ c_{\text{storage}} &< b_{\text{distant}} \cdot (p_{\text{caching}} - p_{\text{isp}}) \end{aligned} \quad (16)$$

$U_5 < U_4$ holds, when the charges by the ISP are lower than those of the CDN (Inequality 17). With the bandwidth payment model, this requires the CDN's price for storage to be at least as high as the difference between the bandwidth prices of the ISP and the CDN.

$$\Leftrightarrow \begin{aligned} U_5 &< U_4 \\ b_{\text{distant}} \cdot (p_{\text{caching}} - p_{\text{cdn}}) - b_{\text{local}} \cdot (p_{\text{cdn}} - p_{\text{isp}}) &< p_{\text{cdn-storage}} \end{aligned} \quad (17)$$

Business Model 3 is Pareto-optimal, when either Condition (14) or Condition (15) holds and, additionally, either Condition 16 or Condition (17) is met. This can be achieved

with an appropriate pricing for ISP caching. For instance, under the assumption that the CDN charges more for bandwidth than the ISP charges for its Internet service ($p_{isp} \leq p_{cdn}$), Condition (17) always holds, when the ISP charges no more for cache bandwidth than the CDN ($p_{caching} \leq p_{cdn}$) and at that price Condition (14) holds for CPs that require a distant bandwidth sufficiently high. Some CPs with a lower bandwidth may instead value user experience enough to meet Condition (15). Furthermore, note that when Condition (16) is met, Condition (14) is also met.

5.3 Equilibria in the resource allocation game

Now, we determine equilibrium outcomes and their conditions in the resource allocation game.

(BEST SIZE, OWN INFRASTRUCTURE) is an equilibrium, when $W_7 \geq W_9$ and $V_7 \geq V_8$. The first condition holds, when there are economies of scale compared to the location that is best for caching (Inequality 18). The second condition is met when the additional transfer cost of cached data is higher than the savings from cheaper storage (Inequality 19).

$$\begin{aligned} W_7 &\geq W_9 \\ \Leftrightarrow \text{EoS} &\leq 1 \end{aligned} \tag{18}$$

$$\begin{aligned} V_7 &\geq V_8 \\ \Leftrightarrow b_{\text{distant}} \cdot c_{\text{tp-transfer}} &\geq d_{\text{isp}} \cdot (c_{\text{storage}} - p_{\text{tp-storage}}) \end{aligned} \tag{19}$$

An equilibrium in (BEST LOCATION, THIRD PARTY) requires $W_{10} \geq W_8$ and $V_{10} \geq V_9$. For the first condition, there must be no economies of scale to the CSP (Inequality 20). Second, the price for CSP’s storage has to be cheaper than own storage facilities (Inequality 21).

$$\begin{aligned} W_{10} &\geq W_8 \\ \Leftrightarrow \text{EoS} &\geq 1 \end{aligned} \tag{20}$$

$$\begin{aligned} V_{10} &\geq V_9 \\ \Leftrightarrow c_{\text{storage}} &\geq p_{\text{tp-storage}} \end{aligned} \tag{21}$$

(BEST SIZE, THIRD PARTY) is an equilibrium, when $V_8 \geq V_7$ and $W_8 \geq W_{10}$. These conditions are the opposites of Conditions (19) and (20).

(BEST LOCATION, OWN INFRASTRUCTURE) is an equilibrium, when $V_9 \geq V_{10}$ and $W_9 \geq W_7$, which are the opposites of Conditions (18) and (21).

Under the assumption that there are economies of scale to the CSP, Condition (18) always holds and Condition (20) never holds, which means that the best size strategy dominates the best cache location strategy for the CSP. Accordingly, which of the two resulting outcomes is an equilibrium solely depends on whether the CSP service is cheap enough to outweigh the additional transfer costs (whether Condition 19 holds or not).

5.4 Pareto optimality in the resource allocation game

Section 5.3 presented that there cannot be an equilibrium, when the CSP chooses cache-appropriate facility locations. We now investigate whether (BEST LOCATION, THIRD PARTY) can be Pareto-optimal.

Pareto optimality requires $W_i < W_{10}$ or $V_i < V_{10}$ for all $i \neq 10$. Unilateral changes away from (BEST LOCATION, THIRD PARTY) cause lower utilities to either the ISP or the CSP by definition: $W_8 < W_{10}$ holds, since the ISP’s utility is reduced by additional network transfer cost and $V_9 < V_{10}$ holds, because the CSP has less revenue, when the ISP does not use the CSP’s service.

Accordingly, (BEST LOCATION, THIRD PARTY) is Pareto-optimal, when $W_7 < W_{10}$ or $V_7 < V_{10}$. For the first condition, the profit from the service usage by the ISP has to be higher than the additional costs due to the sacrificed economies of scale (22). The second condition is met, when the price for CSP storage is cheaper than own storage facilities (23).

$$\begin{aligned} W_7 &< W_{10} \\ \Leftrightarrow d_{\text{isp}} \cdot (p_{\text{tp-storage}} - c_{\text{tp-storage}}) &> d_{\text{tp}} \cdot c_{\text{tp-storage}} \cdot (\text{EoS} - 1) \end{aligned} \tag{22}$$

$$\begin{aligned} V_7 &< V_{10} \\ \Leftrightarrow c_{\text{storage}} &> p_{\text{tp-storage}} \end{aligned} \tag{23}$$

Since the CSP’s service price is the only common factor of these conditions, the outcome can be Pareto-optimal, whenever the other factors allow a price that fulfills both conditions. Such a price exists, whenever caching data in ISP facilities over CSP facilities causes extra cost that is larger than the extra cost the CSP has, when it stores all data of the other customers at the lower economies of scale of a caching-friendly location:

$$d_{\text{isp}} \cdot (c_{\text{storage}} - c_{\text{tp-storage}}) > d_{\text{tp}} \cdot c_{\text{tp-storage}} \cdot (\text{EoS} - 1). \tag{24}$$

This condition holds, if the amount of cached data d_{isp} is sufficiently high and $c_{\text{tp-storage}}$ is cheaper than c_{storage} .

5.5 Discussion on the business model game

The analysis of the business model game shows that ISP caching, where CPs participate in costs (Business Model 3) cannot be an equilibrium, which discourages the business model of selling ISP-operated caching. On the other hand, there are possible equilibria in most other outcomes with different conditions for costs, prices and the ratio of local-to-distant bandwidth. When these conditions are met, Business Model 3 appears even less likely. An evaluation of these conditions is presented in Section 6.

Suppose that the CP prefers Business Model 1 over the CDN. When the ISPs would offer a lower bandwidth price

p_{isp} for the service that includes a cache payment, this gives CPs an incentive to switch to this service. Accordingly, Business Model 1 cannot be an equilibrium. Suppose that the ISP prefers Business Model 3 over not to cache and Business Model 1 over Business Model 2 (caching is only reasonable when CP pays). Then either the CP or the ISP has an incentive to switch strategies in the four upper market outcomes in Table 1 (counter-clockwise loop), which means that none of these states is an equilibrium.

Another possibility to encourage CPs to pay for caching could be the complete abolition of the traditional service. CPs would have to change their ISP in order to make a contract, where they are not charged for caching. Since the CPs would thereby lose the benefits from the caching service, this would also prevent the incentive to deviate from Business Model 3. On the other hand, not offering the traditional service option could drive away many CPs. Section 7 investigates how Business Model 3 can be supported while holding onto both contract options.

Unstable market conditions, where the ISP is not constantly caching, create fluctuating QoE and cache payment. Assuming that CPs and ISPs want to avoid such discontinuities, they will likely agree on a convenient outcome. When it is Pareto-optimal, Business Model 3 can be reasonable despite not being an equilibrium since it prevents any temporary market situations at the expense of either party. Even in case of other equilibria, e.g. in favor of the CDN, the improved utility that the Pareto-optimal outcome might offer to both ISPs and CPs would encourage such a cooperation. Note, however, that it may be the case that not all CPs meet the Pareto-optimality conditions at the same time.

Although, we mentioned that the paper focuses on content ISPs in Section 3, the same model could in principle also be applied on eyeball ISPs, which primarily serve end users. If we assume that the CPs pay for access to the eyeball network, we can consider the content ISP to be in the same role as the CP in the presented model (The content ISP passes on the price for the eyeball ISP to its customer CPs). If we, on the other hand, assume that the eyeball ISP pays the content network for all the traffic that the eyeballs users create in the content network, p_{isp} would be negative (zero in case of peering), which would encourage caching a lot.

5.6 Discussion on the resource allocation game

Although (BEST LOCATION, THIRD PARTY) cannot be an equilibrium of the resource allocation game, it can be Pareto-optimal (Section 5.4). Therefore, the use of third-party storage for ISP-operated caching is an option that might influence the placement of these storage facilities within the ISP network. Whether this can be the case in the current market situation is evaluated in Section 6.

The CSP has no incentive to switch its facility location, when it expects the ISP to anyway utilize a remote CSP's service, i.e. (BEST SIZE, THIRD PARTY). In (BEST LOCATION, OWN INFRASTRUCTURE), on the other hand, the ISP possibly has an incentive to switch to the CSP's service. Accordingly, the CSP is in charge to induce the desired situation. Since storage facilities cause huge capital expenditures and actually switching back and forth is probably very expensive, agreements foregoing accomplished facts are advisable.

QoE is not regarded in the model, but the end user experience possibly suffers to some extent, when a remote size-optimized CSP's facility is used for caching. A noticeable QoE decrease could negatively influence the ISP's utility in Eq. 9, because it weakens the ISPs selling point of a service with caching. This changes the equilibrium conditions for the situations, where the CSP chooses a location for best economies of scale. The Pareto optimality of the cache-friendly location strategy is not, however, influenced by such a utility change.

The investment in storage facilities is an especially important factor regarding the assumption that ISPs build the caching service on-demand. In the case of own infrastructure the deployment and modifications of, e.g., capacity take some time during which the CPs would have to wait for the service to be operational. This is further discussed in Section 7.

6 Evaluation

The previous section determined the condition sets for each of the potential Nash equilibrium and Pareto optimal outcomes. However, it is not clear which conditions hold in reality and what states can actually be regarded as stable. The main interest lies within the question whether the Pareto optimal conditions for Business Model 3 can be met. Thus, this section evaluates the feasibility of the conditions by applying real market pricing and cost information for ISPs, CSPs and CPs.

Due to the differences in pricing at different geographical locations and the dominance of the U.S. based storage service providers, the data is adopted from the U.S. market in 2015. In addition, all the data is normalized to a period of one month. The costs of storage could not be found, thus, they are left as variables in the calculations. Furthermore, price of caching and the increase in QoE are left to be calculated by the equations. Table 3 summarizes the data and the sources.

Bandwidths vary largely depending on the demands of the customers. Thus, two values are chosen based on the bandwidths of CDN's demand. Additionally, the bandwidths, 12 and 30 Gbps, can generate a maximum of 4 and

Table 3 Numerical values of the variables

Variable	Value	Source
b	12 Gbps and 30 Gbps	Rayburn (2015)
b_{distant}	$0.45 \cdot b$	Cisco (2015)
c_{distant}	\$630 per Gbps	Norton (2010)
c_{storage}	c_{storage}	–
$c_{\text{tp-storage}}$	$c_{\text{tp-storage}}$	–
$c_{\text{tp-transfer}}$	\$0.05 per GB	AWS (2015)
d_{isp}	$0.1 \cdot d_{\text{tp}}, 0.5 \cdot d_{\text{tp}}$ and $0.8 \cdot d_{\text{tp}}$	–
d_{tp}	540 000 000 GB	Cisco (2014)
EoS	cloud = <i>yes</i> , otherwise = <i>maybe</i>	Armbrust et al. (2010)
p_{caching}	–	–
p_{cdn}	For 10 PB per month, \$0.01 per GB. For 4 PB per month, \$0.015 per GB.	Rayburn (2015)
$p_{\text{cdn-storage}}$	\$0 per GB	Rayburn (2015)
p_{isp}	\$1500 per Gbps	Comcast: Plans & pricing (2015)
$p_{\text{tp-storage}}$	For 5 PB a month, \$0.0275 per GB	AWS (2015)
QoE^+	–	–

10 PB of traffic per month, respectively. Thus, the third party storage pricing for 5 PB is used.

6.1 Business model game

This section restates and discusses the conditions from the business model game (Section 5.1). Real market data are substituted into the conditions and the calculation results are presented. The conditions, which are assumed to hold always are not repeated in this section.

For the state (CDN, DON’T CACHE) to be a Nash equilibrium, Condition (25) was identified. With the higher CDN traffic rate, i.e. 30 Gbps, Condition (25) holds more easily and the probability of the outcome (CDN, DON’T CACHE) being stable is higher.

$$\begin{aligned}
 QoE^+ &\geq b \cdot (p_{\text{cdn}} - p_{\text{isp}}) + p_{\text{cdn-storage}} \\
 &\geq \begin{cases} 3360\$/\text{Gbps} & \text{if } b = 12 \text{ Gbps per month} \\ 1740\$/\text{Gbps} & \text{if } b = 30 \text{ Gbps per month} \end{cases}
 \end{aligned}
 \tag{25}$$

Conditions (26) and (27) are defined for the outcome (TRADITIONAL, DON’T CACHE) to be an equilibrium. Their equivalent calculations show that for smaller traffic rates, Condition (26) is more likely to hold and Condition (27) should always hold. Thus, an outcome, where the CP chooses the traditional service and the ISP does not cache is likely. In addition, the resemblance to the real

market situation shows that the game models are realistically constructed.

$$\begin{aligned}
 QoE^+ &< b \cdot (p_{\text{cdn}} - p_{\text{isp}}) + p_{\text{cdn-storage}} \\
 &< \begin{cases} 3360\$/\text{Gbps} & \text{if } b = 12 \text{ Gbps per month} \\ 1740\$/\text{Gbps} & \text{if } b = 30 \text{ Gbps per month} \end{cases}
 \end{aligned}
 \tag{26}$$

$$\begin{aligned}
 c_{\text{storage}} &\geq b_{\text{distant}} \cdot (c_{\text{distant}} - p_{\text{isp}}) \\
 &\geq \begin{cases} -\$4698 & \text{if } b = 12 \text{ Gbps per month} \\ -\$11745 & \text{if } b = 30 \text{ Gbps per month} \end{cases}
 \end{aligned}
 \tag{27}$$

For (TRADITIONAL, CACHE), Condition (27) is reversed. Accordingly, such an equilibrium requires that the storage cost is smaller than zero. This seems reasonable as investments in caching have to be compensated from either new revenue sources or lower operational costs. Since the ISP prices the caching service at zero in this outcome, no new revenues are generated. The bandwidth savings, on the other hand, are not sufficient to make Business Model 2 more attractive than Business Model 1 at the assessed values. Thus, the requirement for negative storage cost reflects the necessity for cost savings in a cache-enabled network or for new business models, such as Business Model 3 in this paper.

Conditions (28) and (29) determine a range of prices for the ISP’s caching service, in which the outcome (PAY FOR CACHING, CACHE) is Pareto optimal (for at least all the clients that use a CDN today, i. e. they meet Condition 25). As can be seen, the price has to be more than \$0.0019 per GB per month for it to be profitable for the ISP. In addition, the acceptable upper bound for the CPs is 0.0277 per GB per month if the traffic is low, and \$0.0166 per GB per month for higher traffic volumes. Comparing this pricing to the prices of third party storage services, which are on average \$0.0275 per GB per month, the outcome (PAY FOR CACHING, CACHE) seems feasible. However, the outcome is not in equilibrium and CPs may stop paying after the caches are installed. Section 7 investigates how this outcome can be supported as a long-term solution.

$$\begin{aligned}
 p_{\text{caching}} &> c_{\text{storage}}/b_{\text{distant}} + p_{\text{isp}} \\
 &> 0.0019\$/\text{GB}
 \end{aligned}
 \tag{28}$$

$$\begin{aligned}
 &b_{\text{distant}} \cdot (p_{\text{caching}} - p_{\text{cdn}}) - b_{\text{local}} \cdot (p_{\text{cdn}} - p_{\text{isp}}) \\
 &< p_{\text{cdn-storage}} \\
 \Rightarrow p_{\text{caching}} &< \begin{cases} 0.0277\$/\text{GB} & \text{if } b = 12 \text{ Gbps per month} \\ 0.0166\$/\text{GB} & \text{if } b = 30 \text{ Gbps per month} \end{cases}
 \end{aligned}
 \tag{29}$$

6.2 Resource allocation game

This section elaborates on the conditions of the resource allocation game and their real market values. For an equilibrium, where the CSP chooses the best size and the ISP builds its own infrastructure, i.e. (BEST SIZE, OWN INFRASTRUCTURE), Conditions (30) and (31) have to be met. Condition (30) is met, since cloud providers have significant economies of scale in high capacity data storage (Hamilton 2010). Condition (31) shows that the cost of storage has to be quite low for the investment to be feasible for an ISP. Note that the required storage costs are almost the same under the different assumptions, except for the very low cache-related traffic volumes (i.e. d_{isp} is 10 % of third party storage). This means that the impact of the traffic volume is negligible even at relatively small cache-related traffic volumes. On the other hand, this means that in order for the ISP to have an incentive to cache (i.e. caching cost is smaller than long distant transfer cost), the necessary amount of storage would be so small that it is currently not relevant enough for the CSP to give up economies of scale and choose a cache friendly location.

$$EoS < 1 \tag{30}$$

$$\begin{aligned} & b_{distant} \cdot c_{tp-transfer} \geq d_{isp} \cdot (c_{storage} - p_{tp-storage}) \\ \Rightarrow & c_{storage} \leq \begin{cases} 0.0291\$/GB & \text{if } d_{isp} = 0.1 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0316\$/GB & \text{if } d_{isp} = 0.1 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \\ 0.0278\$/GB & \text{if } d_{isp} = 0.5 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0283\$/GB & \text{if } d_{isp} = 0.5 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \\ 0.0277\$/GB & \text{if } d_{isp} = 0.8 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0280\$/GB & \text{if } d_{isp} = 0.8 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \end{cases} \end{aligned} \tag{31}$$

An equilibrium outcome in (BEST SIZE, THIRD PARTY) also requires Condition (30) and additionally Condition (32), which is just the opposite to Condition (31). This means that depending on the cost of storage to the ISP, either (BEST SIZE, OWN INFRASTRUCTURE) or (BEST SIZE, THIRD PARTY) is an equilibrium.

$$\begin{aligned} & b_{distant} \cdot c_{tp-transfer} < d_{isp} \cdot (c_{storage} - p_{tp-storage}) \\ \Rightarrow & c_{storage} > \begin{cases} 0.0291\$/GB & \text{if } d_{isp} = 0.1 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0316\$/GB & \text{if } d_{isp} = 0.1 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \\ 0.0278\$/GB & \text{if } d_{isp} = 0.5 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0283\$/GB & \text{if } d_{isp} = 0.5 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \\ 0.0277\$/GB & \text{if } d_{isp} = 0.8 \cdot d_{tp} \text{ \& } b = 12 \text{ Gbps per month} \\ 0.0280\$/GB & \text{if } d_{isp} = 0.8 \cdot d_{tp} \text{ \& } b = 30 \text{ Gbps per month} \end{cases} \end{aligned} \tag{32}$$

Since both (BEST LOCATION, OWN INFRASTRUCTURE) and (BEST LOCATION, THIRD PARTY) require the exact

opposite of Condition (30) and it is known that CSPs have economies of scale, the outcomes cannot be equilibria.

The outcome (BEST LOCATION, THIRD PARTY) is Pareto optimal, when the storage cost of an ISP-built cloud is higher than \$0.0275 per GB per month (Condition 33) and additionally Condition (34) is met, which sets an upper bound to the storage cost of a third party cloud in terms of EoS. In a real market, CSPs are typically large enough to benefit from economies of scale, whereas ISP-owned caching infrastructure might be too limited in scale, because ISPs are often geographically bounded due to regulatory and policy constraints. This means that either Condition (33) holds or we can at least assume that CSPs are able to offer a price to meet the condition. Further, since *EoS* is defined between 0 and 1, the probability that Condition (34) is met increases with the volume of cached data.

$$\begin{aligned} c_{storage} & \geq p_{tp-storage} \\ & \geq 0.0275\$/GB \end{aligned} \tag{33}$$

$$\begin{aligned} & d_{isp} \cdot (p_{tp-storage} - c_{tp-storage}) \geq d_{tp} \cdot c_{tp-storage} \cdot (EoS - 1) \\ \Rightarrow & c_{tp-storage} \leq \begin{cases} 0.0028/(-0.9 + EoS)\$/GB & \text{if } d_{isp} = 0.1 \cdot d_{tp} \\ 0.0138/(-0.5 + EoS)\$/GB & \text{if } d_{isp} = 0.5 \cdot d_{tp} \\ 0.0220/(-0.2 + EoS)\$/GB & \text{if } d_{isp} = 0.8 \cdot d_{tp} \end{cases} \end{aligned} \tag{34}$$

7 Long term incentives

7.1 Storage elasticity and problem description

The resource allocation game introduced in Section 4.2 investigates whether an ISP should use third-party storage or build its own facilities for caching. These two possibilities not only differ in price and location, but also in flexibility. This has some important consequences once the caches are installed, which are investigated in this section.

When the ISP uses elastic cloud storage, the amount of cache storage can be scaled according to demand. Most importantly, caches can be abolished once a CP stops paying for them. The third party can then use the free capacity for other clients. On the other hand, when the ISP invests in its own caching facilities, the ISP has an incentive to use all the installed capacity irrespective of the payment by the CP, since caching reduces network costs. This becomes apparent, when we change the business model game in Section 4 to cover a setting, where the ISP has already set up its own caching infrastructure. In that case, the ISP also has caching costs when the CP chooses not to pay for caching ($c_{storage}$ applies in V_1), which causes the ISP to prefer to cache as long as distant transfer costs are more than it can charge from the CP for traffic in the traditional contract ($V_2 > V_1$

when $c_{\text{distant}} > p_{\text{isp}}$). Hence, the threat that the ISP gives up caching when the CP stops paying for it is not credible anymore, because the ISP already operates its own caching infrastructure and, in consequence, Business Model 2 is a Nash equilibrium.

This change of the setting after own infrastructure is deployed is not represented in the game. However, it could eventually provide free caching to CPs, when there are not enough other paying CPs present. In consequence, the ISP cannot safely build up a working service and then offer it to the CP. Instead, the CP would have to give an incentive to the ISP to invest in caching equipment by committing to a service that is not yet working and by agreeing to a long-term payment. Although this would be preferable over the status quo, it might be unrealistic, since the CPs very likely cannot or do not want to wait for the service to be deployed after the agreement is made – or to make such a binding contract for a service of unknown real-life performance in the first place. In the following, we investigate how the ISP can have an incentive to invest while at the same time the CP has an incentive to keep paying for the provided caching.

7.2 Incentives in a repeated game

The results of Section 5 suggest that Business Model 3 is not an equilibrium, since the CP has no incentive to make the payment, if the ISP deploys caching. However, this outcome is Pareto-efficient, if the conditions in Section 5.2 are met. Now, it is an interesting question whether there is a way to support Business Model 3 as an equilibrium by changing the game setting such that the CP has an incentive to pay for the caching solution after the system is set up. This would support the ISP’s decision of deploying caching.

The problem with the current game model is that the ISP and the CP make their decisions independently at the same time and these are one-time choices. We propose the following modified game model together with a solution that gives an incentive to the CP to pay for caching in the long term.¹ In the first stage, the ISP decides whether or not to deploy caching. If caching is deployed then they play a repeated game, shown in Table 4, where the CP decides whether or not to pay and the ISP decides whether or not to punish the CP (for not paying). We assume that the ISP can reduce the quality of service by QoE^- such that the CP’s utility becomes lower, but the ISP still gets the

Table 4 Utilities in the repeated game

		ISP	
		PUNISH	NO PUNISHMENT
CP	NO PAYMENT	$U_2 - QoE^-$	U_2
	PAY	$U_4 - QoE^-$	U_4

benefits of caching. We also assume that the quality of service cannot go below the quality level without caching, i.e., $QoE^- \leq QoE^+$, because a worse service might not be acceptable for the CPs and also out of net neutrality considerations (although any punishment at all could turn out problematic in this regard). The other utilities are as before, e.g., $V_4 = V_2 + b_{\text{distant}} \cdot p_{\text{caching}}$.

Now, the outcome with caching and paying can be sustained as an equilibrium, if the players interact for several periods and they are patient enough, i.e., they value future utilities enough. The supporting mechanism could, for example, be a simple trigger strategy. The players are supposed to choose paying and not punishing unless the other party deviates, which triggers a punishment. The ISP’s punishment is to degrade the quality of service for the following period, if the CP did not pay, and similarly the CP can punish the ISP by not paying, if the ISP did not offer a good caching service. We note that the CP has no incentive to make the cache payment, if the payment is higher than the value of the increase in service quality. Therefore, we need to assume that $b_{\text{distant}} \cdot p_{\text{caching}} \leq QoE^+$.

Let us now calculate the required level of patience for the CP. For simplicity, we assume that the game is repeated infinitely many times and the players discount the future utilities with a discount factor δ , where $0 < \delta < 1$. Playing (PAY, NO PUNISHMENT) infinitely many times is a subgame-perfect equilibrium, if the players should not deviate from the path of play when a deviation is followed by the extreme punishment, i.e., playing (NO PAYMENT, PUNISH) infinitely many times (Abreu 1988; Berg and Kitti 2014; 2013). So, the path of play should give higher utility than the best possible deviation, i.e., $U_4 \geq (1 - \delta) \cdot U_2 + \delta \cdot (U_2 - QoE^-)$. The right term means that utility U_2 is received one time and then $U_2 - QoE^-$ after that. From this condition, we can solve the required discount factor δ_{req} (Condition 35). Thus, the CP should pay for caching, if $\delta \geq \delta_{req}$; i.e. the CP is patient enough and the future events have an accordingly high enough weight in the CP’s decision.

$$\delta_{req} = \frac{U_2 - U_4}{QoE^-} = \frac{b_{\text{distant}} \cdot p_{\text{caching}}}{QoE^-} \tag{35}$$

Assuming that both $QoE^- \leq QoE^+$ and $b_{\text{distant}} \cdot p_{\text{caching}} \leq QoE^+$, δ_{req} is below 1 when the ISP chooses the penalty to be as high as possible. How much below 1 it actually is,

¹Note that the game model could also be improved in other ways to make it more realistic. For example, the deployment of caching services takes some time, and thus the ISP first decides whether or not to deploy caching and after the caching system is up and running, the CP decides whether or not to pay for it. This could be modeled as a two-stage game but it would not change the fact that the CP has no incentive to pay for caching.

however, depends on how the costs of caching to the CP relate to its gain in service quality.

We note that there are also other mechanisms for supporting the (PAY FOR CACHING, CACHE) outcome of the business model game. For example, we could model the situation as a cooperative game, where the ISP promises to deploy the caching system if the CP pays for it and both negotiate a suitable contract. However, this means that the ISP would have to negotiate with multiple CPs for a payment scheme, which is more complicated than negotiating with only one party. It might actually be more convenient for the ISP to raise a fund and promise to deploy caching if it can collect enough money from CPs, e.g. onetime payments or some usage-based contracts. But as pointed out before, these options require a huge upfront commitment by the CPs. A long-term incentive to pay for caching may also exist when several CPs compete with service quality and the ISP's caching system has a capacity too limited to serve them all. The CP has no incentive to stop paying, if the caches of a CP will be replaced by data of another client when the payment stops.

7.3 Extensions to the model

The simple game model that we have presented can be extended in many ways. We could add more players to the game by introducing multiple ISPs and CPs with different preferences that could possibly compete with each other. ISPs may have different caching costs and CPs may value differently the better QoE under caching. In this advanced model, it may be that only some ISPs with low costs (or with enough CPs paying for it) deploy caching or only some CPs, who benefit enough from caching, pay for it. This raises the issue of violating net neutrality and not treating all content the same way. The system may work more efficiently if the ISPs were allowed to and could technically differentiate the CPs and their content. The ISPs may prefer a model where they could offer better QoE to only some CPs and only these CPs would pay more for the service. Another extension is to model vertically integrated players, who take both roles of ISPs and CPs. These issues are left for future research. One problem with the more general model is to get realistic estimates for the increased number of parameters and the players' payoffs.

The game can also be improved by modeling the players' payoffs in more detail. The data traffic is bursty, uncertain and changes in time. We could build different scenarios for the future traffic and evaluate the costs and benefits in each scenario. Combining these values and the probability estimates for the scenarios, we could calculate the expected utilities for each business model. We could also build a dynamic model, where the decisions are made in multiple periods, like in the repeated game, but the parameters could

change from one period to another and these changes could be stochastic. This way the players could make their decisions taking into account the possible scenarios of the future data traffic.

8 Conclusions

This paper provides a business perspective on ISPs in a caching environment. It analyzes and evaluates a new business model in a game-theoretic model. We showed that a market where CPs contribute to the cost of ISP caching (Business Model 3) is potentially unstable because the service without caching payment is always better for the CP, when the ISP caches anyways (Business Model 2). However, Business Model 3 can be Pareto-optimal, which gives ISPs and CPs the incentive to cooperate and establish ISP caching with long-term contracts and an appropriate pricing.

Further, we have shown how ISP caching can be stabilized in order to work without an upfront commitment by the CP. Our solution encourages the ISPs to terminate the positive effects of ISP caching for those clients, who do not pay for the system in order to force CPs to pay their share.

Additionally, ISP caching may use cloud services to obtain the required resources. However, according to our research, the relatively small expected amount of cached data makes it currently financially more attractive for providers of cloud storage to choose higher economies of scale instead of distributing capacity over several cache-friendly physical locations.

After our demonstration that ISP caching is a conclusive business model with mutual benefits to both ISP and CPs, further studies have to show whether content distribution via ISP caching is better from a technical perspective than the current practice with CDNs. For instance, a different number and location of caches could be better for different content (e.g. (Chen et al. 2002) describes an optimization regarding client latency and server capacity). Since ISPs have a higher degree of freedom in cache placement due to their ownership of the network infrastructure, ISPs do not have to compromise as much on the number and the placement of caches compared to CDNs. In addition, because ISPs are in charge of routing, they might cache more efficiently than pure-play CDNs. A combination of strategic cache locations and clever routing can disperse network capacity utilization and, hence, keep infrastructure costs down for the ISP. It would also be interesting to investigate the concept of horizontal provider federation in this regard, either in the sense of a federation of CSPs (similar to (Hassan et al. 2014), but for storage) that enable better or cheaper ISP caching, or in the sense of CDN federation as a reaction to the competition by the ISPs.

Future studies also have to address caching between networks of several ISPs. We assume in our business model game that the ISP of a CP takes care of this, for instance, by peering with other ISPs (Section 3). While peering agreements are a possibility for cooperation between ISPs of the same size and cache utilization, compensation policies in case of unbalanced use of caches have to be discussed. Agreements also have to be made in case the individual ISPs do not participate in caching. Service level agreements and violation penalties can help to ensure a certain quality of the service between different ISPs and between ISP and CP as well.

References

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2), 383–396.
- Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutscher, D., & Ohlman, B. (2012). A survey of information-centric networking. *IEEE Communications Magazine*, 50(7), 26–36.
- Akamai (2015). Aura licensed cdn (lcdn). <https://www.akamai.com/us/en/solutions/products/network-operator/>. Accessed on Dec 1st 2015.
- Akamai (2015). Customer stories: Mtv networks. <https://www.akamai.com/us/en/our-customers/customer-stories-mtv-networks.jsp>. Accessed Dec 1st 2015.
- Akamai (2015). Peering networks detailed view. <http://www.peeringdb.com/view.php?asn=20940>, Accessed Dec 1st 2015.
- Armbrust, M., Fox, A., Griffith, R., Joseph, A., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58.
- AT&T (2015). Content delivery and digital media services. <http://www.business.att.com/enterprise/Portfolio/content-delivery/>. Accessed Feb 16th 2015.
- AWS. (2015). Amazon s3 pricing. <http://aws.amazon.com/s3/pricing/>. Accessed on June 29th 2015.
- Barish, G., & Obraczka, K. (2000). World wide web caching: Trends and techniques. *IEEE Communications Magazine*, 38, 178–184.
- Berg, K., & Kitti, M. (2013). Computing equilibria in discounted 2×2 supergames. *Computational Economics*, 41, 71–78.
- Berg, K., & Kitti, M. (2014). Equilibrium paths in discounted supergames. <http://econpapers.repec.org/paper/tkkdpaper/dp96.htm>. Work in progress.
- Carofoglio, G., Morabito, G., Muscariello, L., Solis, I., & Varvello, M. (2013). From content delivery today to information centric networking. *Computer Networks*, 57(16), 3116–3127.
- Chen, Y., Katz, R.H., & Kubiawicz, J. (2002). Dynamic replica placement for scalable content delivery. In Druschel, P., Kaashoek, F., & Rowstron, A. (Eds.), (Vol. 2429 pp. 306–318). Berlin: Springer. Peer-to-Peer Systems. <http://link.springer.com/chapter/10.1007>.
- Cho, K., Jung, H., Lee, M., Ko, D., Kwon, T.T., & Choi, Y. (2011). How can an isp merge with a cdn *IEEE Communications Magazine*, 49(10), 156–162.
- Cisco (2014). Cisco global cloud index: Forecast and methodology, 2013–2018. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.pdf. Accessed on June 29th 2015.
- Cisco (2015). Cisco visual networking index: Forecast and methodology, 2014–2019. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.pdf. Accessed on June 29th 2015.
- Comcast: Plans & pricing (2015). <http://business.comcast.com/internet/business-internet/plans-pricing>.
- Costa-Requena, J., Kimmerlin, M., Manner, J., & Kantola, R. (2014). Sdn optimized caching in lte mobile networks. In *2014 International Conference on Information and Communication Technology Convergence* (pp. 128–132). South Korea: Busan.
- Dán, G. (2011). Cache-to-cache: Could isps cooperate to decrease peer-to-peer content distribution costs *IEEE Transactions on Parallel and Distributed Systems*, 22(9), 1469–1482.
- Dilley, J., Maggs, B., Parikh, J., Prokop, H., Sitaraman, R., & Wehl, B. (2002). Globally distributed content delivery. *IEEE Internet Computing*, 6(5), 50–58.
- DiPalantino, D., & Johari, R. (2012). Traffic engineering with semiautonomous users: A game-theoretic perspective. *IEEE/ACM Transactions on Networking*, 20(6), 1938–1949.
- Faratin, P. (2007). Economics of overlay networks: An industrial organization perspective on network economics. In *Proceedings of the Joint Workshop on The Economics of Networked Systems and Incentive Based Computing (NetEcon+IBC), in conjunction with ACM Conference on Electronic Commerce (EC'07)*.
- Faratin, P., Clark, D., Bauer, S., & Lehr, W. (2008). The growing complexity of internet interconnection. *Communications & Strategies*, 72, 51–71.
- Fitchard, K. (2012). Forget the cdn players, netflix is caching its own video. <http://gigaom.com/video/forget-the-cdn-players-netflix-is-caching-its-own-video/>.
- Frank, B., Poesse, I., Lin, Y., Smaragdakis, G., Feldmann, A., Maggs, B., Rake, J., Uhlig, S., & Weber, R. (2013). Pushing cdn-isp collaboration to the limit. *SIGCOMM Comput. Commun. Rev.*, 43(3), 34–44. doi:10.1145/2500098.2500103.
- Fudenberg, D., & Tirole, J. (1991). *Game Theory*: MIT Press.
- Google (2015). Google caching overview. <https://peering.google.com/about/ggc.html>. Accessed on Feb 16th 2015.
- Grove, N., Agic, D., & Sedlmeir, J. (2013). Reporting policies of isps: Do general terms and conditions (gtcs) match with the reality? In *TPRC 41: The 41st Research Conference on Communication, Information and Internet Policy*.
- Hamilton, J. (2010). Cloud computing economies of scale. <http://channel9.msdn.com/events/MIX/MIX10/EX01/>. MIX'10 talk.
- Hassan, M.M., Hossain, M., Sarkar, A., & Huh, E.N. (2014). Cooperative game-based distributed resource allocation in horizontal dynamic cloud federation platform. *Information Systems Frontiers*, 16(4), 523–542. doi:10.1007/s10796-012-9357-x.
- Jiang, W., Zhang-Shen, R., Rexford, J., & Chiang, M. (2009). Cooperative content distribution and traffic engineering in an isp network. *ACM SIGMETRICS Performance Evaluation Review - SIGMETRICS '09* pp. 239–250. doi:10.1145/1555349.1555377.
- Katsaros, K., Xylomenos, G., & Polyzos, G.C. (2011). Multicache: An overlay architecture for information-centric networking. *Computer Networks*, 55(4), 936–947.
- Khare, V., & Zhang, B. (2011). Making cdn and isp routings symbiotic. In *Proceedings of the 2011 31st International Conference on Distributed Computing Systems, ICDCS '11*. 869–878. *IEEE Computer Society, Washington, DC, USA*.
- Kimmerlin, M., Costa-requena, J., & Manner, J. (2014). Caching using software-defined networking in lte networks. In *Proceedings of IEEE International Conference on Advanced Networks and Telecommunications Systems*.
- Künsemöller, J., Zhang, N., & Soares, J. (2013). Isp business models in caching. In *Proceedings of the 29th International Conference on Data Engineering, Workshops (ICDEW 2013)* (pp. 279–285).

- Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., & Jahanian, F. (2010). Internet inter-domain traffic. *ACM SIGCOMM Computer Communication Review*, 40(4), 75–86.
- Lee, S., Jiang, J., Chiu, D.M., & Lui, J. (2008). Interaction of isps: Distributed resource allocation and revenue maximization. *IEEE Transactions on Parallel and Distributed Systems*, 19(2), 204–218. doi:10.1109/TPDS.2007.70714.
- Level3 (2010). News archive: Level3 releases statement to clarify issues in comcast/level3 interconnection dispute. <http://news.level3.com/index.php?s=23600&item=65047>. Accessed Dec 1st 2015.
- Netflix (2015). Netflix open connect content delivery for isps. <https://openconnect.itp.netflix.com/deliveryOptions/index.html>. Accessed Dec 1st 2015.
- Norton, W.B. (2010). Internet transit prices - historical and projected <http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php>.
- Norton, W.B. (2014). The 2014 Internet Peering Playbook: Connecting to the Core of the Internet. DrPeering Press.
- Pardalos, P., Migdalas, A., & Pitsoulis, L. (2008). *Pareto Optimality, Game Theory and Equilibria*. Springer Optimization and Its Applications: Springer.
- Peterson, L., & Davie, B. (2014). Rfc 7336: Framework for content distribution network interconnection (cdni). <http://tools.ietf.org/html/rfc7336>.
- Pham, T.M. (2015). Analysis of isp caching in information-centric networks. In *Proceedings of the IEEE RIVF International Conference on Computing Communication Technologies - Research, Innovation, and Vision for the Future (RIVF)* (pp. 151–156). doi:10.1109/RIVF.2015.7049891.
- Psaras, I., Chai, W., & Pavlou, G. (2012). Probabilistic in-network caching for information-centric networks. In *Proceedings of the second ACM SIGCOMM workshop on Information-Centric networking*.
- Rayburn, D. (2015). The state of the cdn market: Video pricing, contract, volume and market sizing trends. www.cdnpricing.com.
- Shakkottai, S., & Srikant, R. (2006). Economics of network pricing with multiple isps. *IEEE/ACM Transactions on Networking*, 14(6), 1233–1245. doi:10.1109/TNET.2006.886393.
- Shrimali, G., Akella, A., & Mutapcic, A. (2010). Cooperative inter-domain traffic engineering using nash bargaining and decomposition. *IEEE/ACM Transactions on Networking*, 18(2), 341–352. doi:10.1109/TNET.2009.2026748.
- Soares, J., Carapinha, J., Melo, M., Monteiro, R., & Sargento, S. (2012). Resource allocation in the network operator's cloud: A virtualization approach Computers and Communications (ISCC), 2012 IEEE Symposium on, pp. 000,800–000, 805.
- Telefonica (2014). Akamai and telefonica enter into global content delivery alliance. <https://www.globalsolutions.telefonica.com/en/news/2014/03/25/telefonica-and-akamai/>. Accessed Dec 1st 2015.
- Vakali, A., & Pallis, G. (2003). Content delivery networks: Status and trends. *IEEE Internet Computing*, 7(6), 68–74.
- Vaquero, L., Rodero-Merino, L., Caceres, J., & Lindner, M. (2009). A break in the clouds: towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1), 50–55.
- der Veen, V. (2011). Orange and cogent fight over traffic management before french competition watchdog. <http://www.futureofcopyright.com/home/blog-post/2011/09/08/orange-and-cogent-fight-over-traffic-management-before-french-competition-watchdog.html>.
- Wang, X., Chen, M., Taleb, T., Ksentini, A., & Leung, V.C.M. (2014). Cache in the air: Exploiting content caching and delivery techniques for 5g systems. *IEEE Communications Magazine*, 52(2), 131–139.
- Wong, R.K. (2013). Feasibility and a case study on content optimization services on cloud. *Information Systems Frontiers*, 15(4), 525–532. doi:10.1007/s10796-012-9379-4.
- Wulf, J., Zarnekow, R., Hau, T., & Brenner, W. (2010). Carrier activities in the cdn market - an exploratory analysis and strategic implications. In *2010 14th International Conference on Intelligence in Next Generation Networks* (pp. 1–6). doi:10.1109/ICIN.2010.5640943.

Jörn Künsemöller finished his PhD on the Infrastructure-as-a-Service market within the research training group Automatisms and the collaborative research centre On-the-fly Computing at the University of Paderborn, Germany, in 2014. Next to cloud infrastructure, his research interests focus on game theory and human-computer interaction.

Nan Zhang received her M.Sc. (Tech.) in Communications Engineering from Aalto University School of Science and Technology, Finland, in 2010 with major in Data Networks and minor in Strategy and International Business. She is currently finalizing her postgraduate studies at the Department of Communications and Networking, Aalto University, Finland. Her research interests include techno-economics of Internet content delivery solutions, such as content delivery networks, information-centric based solutions and software defined networks, and next generation mobile networks.

Kimmo Berg received the PhD degree in operations research from Aalto University School of Science, Finland, in 2011. He has coauthored papers in international journals and conferences like IEEE Communications Magazine, IEEE DySPAN and ICEC. His research interests include game theory, nonlinear pricing and applications in telecommunications: cognitive radio, dynamic spectrum access, mobile offloading and caching.

João Soares PhD, worked as a Research Engineer at Portugal Telecom Inovação e Sistemas and was awarded his PhD at the University of Aveiro, Portugal. He now works for Ericsson. His research interests are in cloud computing, software defined infrastructure (SDI), software defined networking (SDN), and network function virtualization (NFV).