Content-aware Traffic Engineering*

Benjamin Frank
T-Labs/TU Berlin
bfrank@net.t-labs.tu-berlin.de

Ingmar Poese
T-Labs/TU Berlin
ingmar@net.t-labs.tu-berlin.de

Georgios Smaragdakis
T-Labs/TU Berlin
georgios@net.t-labs.tu-berlin.de

Steve Uhlig
Queen Mary, University of London
steve@eecs.qmul.ac.uk

Anja Feldmann
T-Labs/TU Berlin
anja@net.t-labs.tu-berlin.de

ABSTRACT

Recent studies show that a large fraction of Internet traffic is originated by Content Providers (CPs) such as content distribution networks and hyper-giants. To cope with the increasing demand for content, CPs deploy massively distributed server infrastructures. Thus, content is available in many network locations and can be downloaded by traversing different paths in a network. Despite the prominent server location and path diversity, the decisions on how to map users to servers by CPs and how to perform traffic engineering by ISPs, are independent. This leads to a lose-lose situation as CPs are not aware about the network bottlenecks nor the location of end-users, and the ISPs struggle to cope with rapid traffic shifts caused by the dynamic CP server selection process.

In this paper we propose and evaluate Content-aware Traffic Engineering (CaTE), which dynamically adapts the traffic demand for content hosted on CPs by utilizing ISP network information and end-user location during the server selection process. This leads to a win-win situation because CPs are able to enhance their end-user to server mapping and ISPs gain the ability to partially influence the traffic demands in their networks. Indeed, our results using traces from a Tier-1 ISP show that a number of network metrics can be improved when utilizing CaTE.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design; C.2.3 [**Network Operations**]: Network Management

Keywords

Traffic Engineering, Content Delivery, Load Balancing, Network Optimization, ISP-CDN Collaboration.

1. INTRODUCTION

Today, a large fraction of Internet traffic is originated by a small number of Content Providers (CPs) [5]. Major CPs are highly popular rich media sites like YouTube and Netflix, One-Click Hosters (OCHs), e. g., RapidShare and MegaUpload, as well as Content Delivery Networks (CDN) such as Akamai and Limelight, and hyper-giants, e. g., Google, Yahoo!, and Microsoft. To cope with the increasing demand for content, CPs deploy massively distributed server infrastructures to replicate content and make it accessible from different locations in the Internet [2].

*A full version of this paper is available in [4].

SIGMETRICS'12, June 11–15, 2012, London, England, UK.
ACM 978-1-4503-1097-0/12/06.

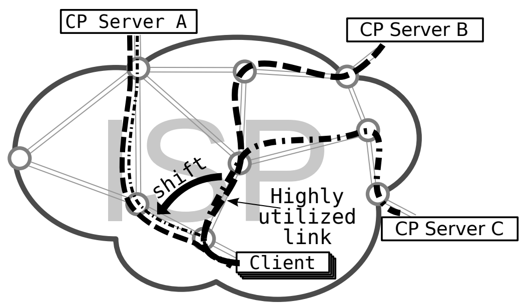

Figure 1: Content-aware traffic engineering process.

The unprecedented growth of demand for content and the resulting massive deployment of content delivery infrastructures pose new challenges both to CPs and to ISPs. For CPs, the cost of deploying and maintaining such a massive infrastructure has significantly increased during the last few years [8] and the price charged for delivering traffic to end-users has decreased due to the intense competition. Furthermore, CPs struggle to engineer and manage their infrastructures, replicate content based on end-user demand, and assign users to appropriate servers. The latter is challenging due to the mis-location of end-users [1]. Furthermore, inferring the network conditions within an ISP without direct information from the network is difficult [6].

Moreover, due to highly distributed server deployment and adaptive server assignment, the traffic injected by CPs is volatile. For example, if one of its locations is overloaded, a CP will re-assign end-users to other locations, resulting in large traffic shifts in the ISP network within minutes. Current ISP traffic engineering adapts the routing and operates on time scales of several hours [3], thus, is too slow to react to rapid traffic changes caused by CPs.

2. THE CaTE APPROACH

The pressure for cost reduction and customer satisfaction that both CPs and ISPs are confronted with, coupled with the opportunity that massively distributed server infrastructures offer, motivate us to propose a new tool in the traffic engineering landscape. We introduce *Content-aware Traffic Engineering* (CaTE). CaTE leverages the location diversity offered by CPs and, through this, enables adaptation to traffic demand shifts. In fact, CaTE relies on the observation that by selecting an appropriate server among those available to deliver the content, the path of the traffic in the network can be influenced in a desired way. Figure 1 illustrates the basic concept of CaTE. The content requested by the client is in principle available from three servers (A, B, and C) in the network.

However, the client only connects to one of the network locations. Today, the decision of where the client will connect to is solely done by the CP and is partially based on measurements and/or inference of network information and end-user location. With CaTE the decision on end-user to server assignment can be enhanced by recommendations offered by the ISP. In the illustration, a highly utilized path can be avoided by mapping users to server A and as a consequence shift traffic to less congested links.

CaTE complements the existing traffic engineering solutions by focusing on traffic demands rather than routing, by combining (i) the knowledge of CPs about their location diversity and server load, with (ii) the ISPs detailed knowledge of the network conditions and end-user location. This can be achieved without either party sharing sensitive operational information. To this end the ISP ranks candidate servers that are communicated by the CP based on the network conditions and the source of demand that can be either the resolver or the end-user [7]. CaTE offers additional traffic engineering capabilities to ISPs to better manage the volatility of CP traffic. Also, thanks to the recommendations offered by ISP networks, CPs gain the ability to better assign end-users to their servers and better amortize the cost of deploying and maintaining their infrastructure. Furthermore, the burden of measuring and inferring network topology, and the state of the network, both challenging problems, is removed from the CPs. In short, all involved parties, including the end-users, benefit from CaTE, creating a win-win situation for everyone.

Formally, we model the network as a directed graph $G(V, E)$ where V is the set of nodes and E is the set of links. An origin-destination (OD) flow f_{od} consists of all traffic entering the network at a given point $o \in V$ (origin) and exiting the network at some point $d \in V$ (destination). The traffic on a link is the superposition of all OD flows that traverse the link. The relationship between link and OD flow traffic is expressed by the routing matrix A. If $A_{ml} = 1$ the OD flow m traverses link l. The routing matrix A can be derived from routing protocols, e.g., OSPF, ISIS, BGP. Let **y** be a vector of traffic counts on links and **x** the vector of traffic counts in OD flows, then $\mathbf{y}$=$A\mathbf{x}$. Traditional traffic engineering reduces to controlling and optimizing the routing function and to steering traffic through the network in the most effective way. Translated into the above matrix form, traffic engineering is the process of adjusting A, given the OD flows **x**, so as to influence the link traffic **y** in a desirable way. In CaTE, we revisit traffic engineering by focusing on traffic demands rather routing changes:

Definition 1: Content-aware Traffic Engineering (CaTE) is the process of adjusting the traffic demand vector **x**, given a routing matrix A, so as to change the link traffic **y**.

Not all traffic can be adjusted arbitrarily. Only traffic for which location diversity is available can be adjusted by CaTE. Therefore, $\mathbf{x}$=$\mathbf{x}_r$+$\mathbf{x}_s$ where $\mathbf{x}_r$ and $\mathbf{x}_s$ denote the content demands that can be adjusted and can not be adjusted (as there is only a single location) respectively. The degree of freedom in adjusting traffic highly depends on the diversity of locations from which the content can be obtained. We can rewrite the relation between traffic counts on links and traffic counts in flows as follows: $\mathbf{y}$=$A(\mathbf{x}_s + \mathbf{x}_r)$. CaTE adjusts the traffic on each link of the network by adjusting the content demands $\mathbf{x}_r$: $\mathbf{y}_r$=$A\mathbf{x}_r$ to satisfy a traffic engineering goal. In the full version of this paper [4] we show how to assign demands nearly optimal to available network locations for a number of CPs in regards to different network metrics.

3. EVALUATION

CaTE allows ISPs and CPs to optimize for a number of network metrics, such as link utilization, path length, and path delay. We quantify the effects of CaTE when using each one of these metrics individually in a large Tier-1 ISP. Note that our system design also allows for a combination of metrics. We focus on the top 10 CPs, as these have a significant share of the overall traffic [4]. To evaluate the effectiveness of CaTE, we observe network traffic and measure network quantities such as maximum link utilization as the CPs apply their server selection algorithms. We assume that the CPs follow the recommendations provided by the ISP, per request, and estimate the effects on these network quantities. Similar observations are made when applying CaTE to the top 1 and 100.

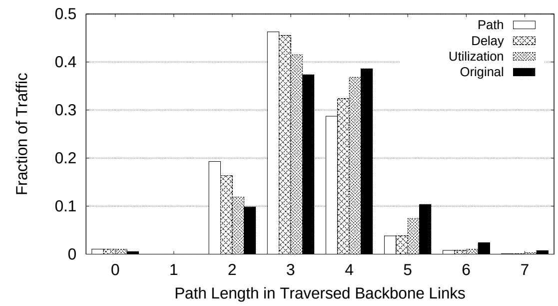

Figure 2: Traffic distribution per path length in a Tier-1 ISP.

In Figure 2, we plot the traffic distribution per path length before we apply CaTE (Original) and when we apply CaTE using one of the network metrics mentioned above. In all cases, CaTE redirects the traffic towards paths with the same or even shorter path, thus the overall traffic is reduced. Our results also show that the delay as well as the maximum link utilization are reduced as the traffic is shifted from highly utilized links to less utilized ones. In the full version of this paper [4] we present a much larger set of results and evaluate CaTE in other networks and under different demands.

4. FUTURE WORK

To capitalize on the substantial performance benefits that CaTE offers to both CPs and ISPs as well as end-users, we are currently building a system that allows all the involved parties to utilize CaTE. Moreover, we are investigating how to install CaTE in a number of ISPs and CPs and report the performance benefits for both ISPs and CPs in the wild. We are also interested in evaluating CaTE for ISP-ISP collaboration.

5. REFERENCES

[1] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Comparing DNS Resolvers in the Wild. In *Proc. of ACM IMC*, 2010.

[2] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Web Content Cartography. In *Proc. of ACM IMC*, 2011.

[3] B. Fortz and M. Thorup. Internet Traffic Engineering by Optimizing OSPF Weights. In *Proc. of IEEE INFOCOM*, 2000.

[4] B. Frank, I. Poese, G. Smaragdakis, S. Uhlig, and A. Feldmann. Content-aware Traffic Engineering. *CoRR*, abs/1202.1464, 2012.

[5] A. Gerber and R. Doverspike. Traffic Types and Growth in Backbone Networks. In *Prof. of OFC/NFOEC*, 2011.

[6] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS Oper. Syst. Rev.*, 44(3), 2010.

[7] I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. Improving Content Delivery using Provider-Aided Distance Information. In *Proc. of ACM IMC*, 2010.

[8] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the Electric Bill for Internet-scale Systems. In *Proc. of ACM SIGCOMM*, 2009.