Computer Communications 67 (2015) 23–33

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Routing-policy aware peering for large content providers

Qixin Gao [a], Feng Wang [b], Lixin Gao [c,*]

[a] Department of Computer Science, Northeastern University, China
[b] School of Engineering and Computational Sciences, Liberty University, USA
[c] Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, USA

ARTICLE INFO

Article history:
Received 30 April 2014
Revised 11 February 2015
Accepted 26 June 2015
Available online 3 July 2015

Keywords:
Content providers
Peering connection
Peering relationships
Routing policy

ABSTRACT

Large content providers, such as Google, Yahoo, and Microsoft, aim to directly connect with consumer networks and place the content closer to end users. Exchanging traffic directly between end users and content providers can reduce the cost of transit services. However, direct connection to all end users is simply not feasible. Content providers by-and-large still rely on transit services to reach the majority of end users. We argue that routing policies are an important factor in considering the selection of ISPs for content providers. Therefore, determining which ISP to peer or use as a transit becomes a key question for content providers. In this paper, we formulate the policy-aware peering problem, in which we determine not only which ISP to connect with, but also the kind of peering agreement to establish. We prove that such a policy-aware peering problem is NP-complete, and propose a heuristic algorithm to solve the problem. Further, we perform a large-scale measurement study of the peering characteristics of five large content providers, and evaluate the existing peering connections deployed by the content providers. Our measurement results show that changing the existing peering agreements or adding as little as 3–5 new peering connections can enhance the connection between content providers and end users significantly.

1. Introduction

The rapid growth of large content providers, such as Google, Yahoo, and Microsoft, is changing not only the inter-domain traffic patterns but also the hierarchical topology of the Internet. Content providers are using the Internet to deliver videos, TV shows, movies, sports and live programs to mobile phones, televisions, tablets and computers. A recent measurement study [39] has shown that the majority of the inter-domain traffic is between large content providers and consumer networks. At the same time, the Internet topological hierarchy becomes flatter and denser due to the inter-connection strategies employed by large content providers [17,28,39]. Furthermore, large content providers (i) have built their own global backbones, (ii) directly peer with consumer networks, (iii) move the content closer to the end-users (through content delivery networks for example), and (iv) adopt various peering policies in order to ensure small latency in delivering content. For example, Google moves the majority of its video and search traffic away from transit providers to its own backbone infrastructure and directly connects with consumer networks [39].

However, direct connection to all customer networks is simply not feasible. Large content providers, while still relying on transit services of large Internet Service Providers (ISPs) to have global reachability, may select some access networks to connect directly with end users in order to provide better content delivery. For example, Netflix, a television and movie streaming content provider, agreed to pay Comcast millions of dollars annually to deliver its content directly [54]. Netflix previously connected to Comcast via its service provider, Cogent Communications. The new connection to Comcast's network will speed up video streaming to Netflix customers. Therefore, determining which providers to connect to and the corresponding contractual agreement of the connections becomes a key question for large content providers.

In this paper, we formulate the problem of peer selection for large content providers. A content provider can choose to connect to an ISP using its transit service for global reachability. Alternatively, the content provider can connect to an ISP to reach its customers only. Our goal is to determine the ISPs that a content provider should connect with and the kind of peering agreement that the content provider should establish with the selected ISPs.

On the surface, our problem might be solved with content placement algorithms [9,21,34,40,41,47,51,58]. However, the fact that routing policies and peering agreements are taken into account in the selection of peers makes the problem much more challenging. Different from the content placement problem, we determine not only

* Corresponding author. Tel.: +1 413 545 4548.
E-mail addresses: gaoqx@mail.neuq.edu.cn (Q. Gao), fwang@liberty.edu (F. Wang), lgao@engin.umass.edu, lgao@ecs.umass.edu (L. Gao).

http://dx.doi.org/10.1016/j.comcom.2015.06.020

who to peer with but also the peering agreement to establish. We further consider routing policy in quantifying the distance between the content providers and an end user. It is well known that an AS may take a longer path rather than the shortest path, possibly as a result of the routing policies [46,56,59]. We prove that such a policy-aware peering problem is NP-complete, and propose a heuristic algorithm to solve the problem.

To evaluate our heuristic algorithm, we first perform a large-scale measurement study of peering characteristics of five large content providers. We then evaluate our heuristic algorithm in the context of current Internet topology. We derive the minimum number of peering ISPs with different peering agreements to cover today's Internet. We investigate the existing peering agreements employed by the five large content providers. We find that it is still possible for those content providers to cover their clients within the same distance if they maintain less transit service agreements. In addition, our measurement shows that adding as little as 3–5 new peering connections can enhance the connection between content providers and end users significantly.

The rest of the paper is organized as follow: In Section 2, we introduce routing policies and peering relationships guided by commercial relationships. In Section 3 we formally define the policy-aware content placement problems. In Section 4, we propose a heuristic algorithm to address the problems. In Section 5, we perform a large-scale measurement study of peering characteristics of five large content providers, calculate the minimum number of peering ISPs for today's Internet, and compare the peering selection return by our algorithm with that of the content providers. Section 6 presents the related work. We conclude the paper in Section 7 with a summary.

2. Background

In this section, we introduce the concepts and terms used throughout the paper: content providers, routing policies and peering connections. Finally, we use two examples to demonstrate the impact of routing policies and peering relationships on the selection of peering locations.

2.1. Content providers

An Internet content provider is an organization, such as Google or Facebook, which supplies content on the Web, such as web pages, video, and movies. It is likely that a content provider uses one or several content delivery networks (CDNs), for example Akamai, to deliver its content. A CDN is a shared distributed system deployed across the Internet for delivering content to end-users with high availability and high performance. CDNs can help content providers reduce the cost because CDNs offload the traffic from the content provider's infrastructure. Content providers typically subscribe to one or several CDNs services. Then, content is replicated over several mirrored servers, which are closer to end users. In order to meet its Service Level Agreement (SLA) for its customers, CDN providers usually cooperate with each other so that content providers can utilize services of multiple CDNs to deliver contents.

While CDNs help to improve the content delivery, there is a recent trend [39] for large content providers to bring content closer to end-users. For example, in 2011, Google peered with Puerto Rico Bridge Initiative (AS23114) to improve speeds and costs in Puerto Rico [36]. Before peering with PRBI, the traffic from Puerto Rico had to be routed to some transit ISPs at the US mainland, which can cause congestion on these transport links. By peering with PRBI in Miami, Google was able to reduce the amount of Internet traffic traversing these ISPs so that millions of residents of the island have a better user experience at a lower cost [36]. Furthermore, a content provider can also directly connect with access networks to deliver the content to end users. For example, Netflix established an agreement with Comcast for a direct connection [54].

Content providers typically adopt an open peering strategy to assess potential peering relationships [3,43,48]. Under the open peering strategy, content providers are willing to peer with any AS that requests a peering relationship. For example, Google's peering strategy requires at least 100Mbps peak traffic to establish a bilateral BGP peering over an IXP. If an AS with less than 100Mbps traffic can peer via the route servers at any participating IXP [2]. In addition, Google recommends that peers advertise all their prefixes over all peering sessions with Google.

Content providers usually establish a connection with an ISP at Internet Exchange Points (IXPs) or a private peering facilities. Note that while the traffic exchange might be free of charge, it is also possible to establish paid peering agreements between content providers and ISPs. For example, Netflix pays Comcast for its direct connection [54]. However, even under paid peering, the connection is only used to exchange traffic between Netflix and Comcast's customers. That is, Netflix does not use Comcast to reach the global Internet.

2.2. Routing policies and peering policies

Routing policy typically conforms to the commercial relationships between Autonomous Systems (ASes). A customer pays its provider for connectivity to the rest of the Internet. A pair of peers agree to exchange traffic between their respective customers free of charge. A mutual-transit agreement allows a pair of administrative systems to provide connectivity to the rest of the Internet for each other. The commercial contractual relationships between ASes translate into the export rule that an AS does not transit traffic between two of its providers and peers, which is called *no-valley* routing policy [24].

In addition to the no-valley routing policy, an AS typically chooses a customer route over a route via a provider or peer since the AS does not have to pay its customer to carry traffic or maintain a traffic volume ratio between the traffic from and to a peer. This policy is referred to as *prefer-customer* routing policy. For example, one of the peering policies adopted by content providers is that no transit or third party routes are allowed to be announced to the content providers. In other words, all routes exchanged between a content provider and the peering AS must be the peering AS's and the peering AS's customers' routes.

In this paper, the following definitions are used to represent content providers' peering strategy. A *peering connection* is a logical connection between a content provider and an ISP. A content provider can establish two kinds of peering connection with an ISP. A *peering-to-peering connection* means that the content provider and the ISP agree to exchange their own and their customers' routes. We refer to the connection as *local connection*. A *transit connection* allows the content provider to access the rest of the Internet and advertises the content provider's prefixes to the rest of the Internet, through the ISP. We refer to the connection as *global connection*.

2.3. Impact of routing policy and peering relationships

We use two examples to demonstrate the impact of routing policy and peering relationships on the content placement decision. To simplify our examples, we determine a peering location and the peering relationship such that the AS level distance between end users and a content provider is minimized, which resembles the well-known facility location problem. In the next section, we will formulate a cost function to take in account other factors. Note that both examples are real scenarios derived from our dataset, which will be presented in Section 5.

The first example demonstrates the effect of routing policies on the placement of peering connections. In Fig. 1(a), Google (AS15169)