

The Impact of Brokers on the Future of Content Delivery

Matthew K. Mukerjee^{*}, Ilker Nadi Bozkurt[†],
Bruce Maggs^{†‡}, Srinivasan Seshan^{*}, Hui Zhang[§]

^{*}Carnegie Mellon University [†]Duke University [‡]Akamai Technologies [§]Conviva Inc.
{mukerjee, srini}@cs.cmu.edu {ilker, bmm}@cs.duke.edu hzhang@conviva.com

Abstract

Various trends are reshaping content delivery on the Internet: the explosive growth of traffic due to video, users' increasing expectations for higher quality of experience (QoE), and the proliferation of server capacity from a variety of sources (e.g., cloud computing services, content provider-owned datacenters, CDNs, and ISP-owned CDNs). In order to meet the scale and quality demands imposed by users, content providers have started to spread demand across multiple CDNs using a broker. Brokers break many traditional CDN assumptions (e.g., unexpected traffic skew and significant variance in demand over short timescales). Through an analysis of data from a leading broker and a leading CDN, we show the potential challenges and opportunities that brokers impart on content delivery. We take the first steps towards improvement through a redesigned broker-CDN interface.

1 Introduction

Content delivery on the Internet has been constantly evolving to scale better, accommodate new workloads (e.g., streaming video), and incorporate new actors (e.g., CDNs), new protocols (e.g., HTTP chunk-based video), and new algorithms (e.g., video rate adaptation). The introduction of each new content delivery technique has had major impact on flows across all layers of the network stack. For example, the introduction of content delivery networks (CDNs) dramatically changed traffic patterns that ISPs handled, users' performance expectations, and the sheer volume of content that the Internet could deliver to end users.

Internet content delivery is in the middle of another such major change. Until recently, major content providers either contracted with a single CDN, such as Akamai [28], Level 3 [2], or CloudFront [4], or deployed their own CDN, such as Google [17] and Netflix [26]. Recently, the rise of CDN management services ("brokers"), such as Conviva [3], and CDN federation techniques [1, 27] are making it easier for content providers to enlist multiple CDNs to deliver content. Broker or federation techniques provide a way to stitch together multiple CDNs to provide the best performance, lowest cost, or any other strategy at specific locations.

At first glance, it may seem that the addition of brokers to content delivery is a minor change; however, federation is a surprisingly complicated process. For example, this stitching together of CDNs draws strong parallels to the stitching together of independently operated IP networks, which spurred more than a decade of research in routing protocols. Much like IP network federation, brokers enable new styles of small-scale deployment (regional CDNs, city-centric CDNs, etc.). However, the differences are overwhelming: IP network federation focused on reachability, whereas CDN federation is further constrained by specific performance and cost goals. Brokers additionally have the potential to make CDN workloads unpredictable, unintentionally exploit current CDN pricing schemes, and bias CDN traffic patterns in unusual ways. This paper serves as a call to arms; federation is already happening in the CDN space, yet there is little research being done to address it.

The primary goal of this paper is to take the first steps in this space by characterizing what kinds of problems brokers and CDNs face. We do this by examining data from both a broker and a CDN, illuminating a variety of problems stemming from an overly narrow interface between brokers and CDNs. We propose changes to this interface that potentially mitigate the negative impact of brokers and provide benefit to both brokers and CDNs. We believe that changes to the broker-CDN interface could: 1) improve broker use of CDN resources leading to a better performance and/or cost tradeoff for content providers, 2) increase workload stability for CDNs, and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotNets-XV, November 09-10, 2016, Atlanta, GA, USA

© 2016 ACM. ISBN 978-1-4503-4661-0/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/3005745.3005749>




Figure 1: Traditional content delivery.

3) increase stability between CDNs and brokers. While interface changes may address some of the immediate problems, addressing the issues of content delivery with multiple brokers and highly specialized CDNs is likely to require more sophisticated solutions, providing potential avenues for future work.

To summarize, we make the following contributions:

1. We identify a set of workload challenges that brokers create for CDNs by analyzing broker and CDN data.
2. We identify numerous ways that current CDN interfaces limit broker behavior.
3. We define a new candidate interface based on advertisement exchanges that allows CDNs to manage their workloads more explicitly and gives brokers finer-grained control on where the content will be served from.

2 Content Delivery: The Past and the Present

The landscape of the Internet changed dramatically after the advent of CDNs. In this section, we first briefly discuss how content providers have traditionally utilized CDNs to help deliver content and then show how content delivery is changing with the addition of brokers.

2.1 Traditional Content Delivery

To better understand content delivery, we describe the players involved before describing how they fit together.

Content providers (CPs), such as ESPN, Netflix, and HBO, create or license content that users are interested in. Serving content to a global audience with good “quality of experience” (QoE)¹ would require massive infrastructure investment, so instead CPs rely on CDNs. CPs generate revenue through subscription-based premium services and/or by showing advertisements, and try to minimize delivery cost. CDN pricing models vary,

¹QoE is generally used in the context of video delivery (as a complex combination of metrics such as average bitrate, buffering ratio, and join time [8]), but in this paper we consider a more general notion of QoE compatible with content delivery at large.

but CPs often pay CDNs based on bandwidth usage using a 95/5 model [24].

CDNs specialize in reducing latency and improving throughput to clients by delivering content to them from geographically close clusters located in datacenters, peering points, universities, large businesses, or ISP networks. CDN deployment models vary: some CDNs opt to deploy many servers across a large number of geographic regions (e.g., Akamai [28] and Google [11]); others operate in a small number of strategic locations (e.g., Level 3 and CloudFront [31], and other “ISP CDNs” (e.g., Comcast [6]) are extremely localized, only serving their ISP customers within cities where they operate. CDNs determine which cluster to serve a user request from using, e.g., static assignments or network quality measurements. For example, Akamai uses latency and loss measurements from clusters to gateway routers in the network (not individual clients) [12, 22] to decide on a cluster. CDNs wish to provide reasonable performance to users while minimizing their bandwidth and (perhaps) energy costs. The prices CDNs offer CPs may reflect server location, but are generally flat across very coarse geographic regions (e.g., continents) [5, 25], regardless of the actual delivery cost. Contracts with CDNs are typically negotiated on a long time scale (e.g., months, years).

In the traditional setting (see Figure 1), CPs contract with a single CDN and express very broad policy goals (e.g., what content can be served by the CDN). The content is placed on geographically-diverse CDN front-end servers so that users get content from nearby locations. Clients learn of the content typically through an HTML page or a video manifest from the CP’s website that indicates which CDN to contact for content. Clients reach CDN servers using some resolution mechanism implemented by the CDN (e.g., DNS). The CDN chooses a server based on its measurement infrastructure, static assignments, etc. The user contacts this server and retrieves the content.

2.2 Brokers and Delivery Today

Today, many large CPs use multiple CDNs [9] as CDN price and performance shows spatial and temporal variations [23] requiring CDN selection to be dynamic to provide best possible QoE to users. **Brokers** (e.g., Conviva [3]) measure user QoE and build predictive models to determine which CDN can deliver the best performance to a particular user based on a variety of factors (e.g., user’s location, ISP, etc.) [15, 20]. Brokers have an important role in content delivery as they act on the CP’s behalf for CDN selection and QoE measurements, typically moving clients between different CDNs in real-time. The broker’s global view of clients aids in meeting CPs’ QoE and cost goals.




Figure 2: Brokered content delivery.

Figure 2 shows how content delivery works with a CP using multiple CDNs with a broker mediating clients. Brokers are a level of indirection for CDN selection. Instead of querying the appropriate CDN’s DNS server directly, clients first query the broker to learn which CDN is appropriate for them at the current time.

3 Potential Problems and Opportunities

Despite the seeming simplicity of an additional layer of indirection, brokers greatly complicate the content delivery ecosystem. Brokers and CDNs try to maximize their own objectives independently, but their decisions directly impact one another, potentially leading to sub-optimal decisions by both entities. Brokers have a global view of all user performance (QoE), but can only actuate their decision by selecting a different CDN. Conversely, CDNs have a wide range of possible servers to select, but have to make decisions based only on the users they currently see (a subset of users seen by the broker). This mismatch between data richness (brokers) and selection richness (CDNs) is the fundamental cause of many potential problems.

In this section, we explore these potential problems separately, using data from a leading CDN and a leading broker. We examine the broker data to understand when it uses different CDNs (e.g., over geographic regions, time, etc.). We examine the CDN data to understand its use of different server clusters. The key conclusion is that there are not only potential problems but also new opportunities to make more informed pricing and provisioning decisions for CDNs, and to help brokers improve QoE and reduce cost for CPs.

3.1 Traces

Broker: We collect trace data from a video delivery broker. The trace includes an entry for each user session containing the request arrival time, which video was requested, the average bitrate, session duration, the user city and AS, the initial CDN contacted, and the current




Figure 3: Cost per byte serving users geolocated in various (anonymized) countries (A-T) relative to the average.




Figure 4: Brokers can greatly skew CDN traffic patterns making pricing difficult.

CDN delivering the video. The data covers roughly an hour of off-peak requests (33.4K total) for one content provider (a music video streaming website).

The data exhibits similar trends to those seen in other works [7]: video popularity is a power-law distribution, as is the distribution of requesting cities. Most users abandon almost immediately (around 78%). The distribution of bitrates is bimodal with peaks at the lowest and highest bitrate. The trace identifies three large CDNs (here “A,” “B,” and “C”) directly and lists the rest as “other.” CDN A is a CDN with clusters in many locations. CDN B and C deploy large amounts of capacity in a small number of locations.

CDN: We collect Internet mapping data from a major CDN to compare performance estimates across US-based clusters. The data provides a score estimating the performance between blocks of IP addresses and candidate CDN clusters. This score is a function of latency and packet loss. Measurement happens periodically and frequently (several times per minute) through pings from clusters to routers with large networks of clients behind them.

From the same CDN we collect data on the cost per byte for the 20 countries with the highest volume of traffic, using client geolocation to bin requests into countries. We then compare them to the average delivery cost. We anonymize this data and present it in Figure 3.

3.2 Potential Problems for a CDN

Brokers create problems for CDNs for two main reasons: 1) brokers divide CDN workloads in potentially biased ways (e.g., sending most traffic from small towns to one CDN) that shift over time and 2) brokers make




Figure 5: Broker’s usage of CDNs for a sampling of countries based on request count.




Figure 6: Sessions moved between CDNs by the broker in our trace in 10s intervals.




Figure 7: Broker’s usage of CDNs, sorted by requests per city in the US. Dotted lines are best-fit linear regressions.

smaller CDNs viable (e.g., ones that do not provide global coverage), increasing competition. We highlight the implications of these problems.

Unintentional exploitation of pricing: Figure 4 illustrates a scenario that can cause CDN pricing issues. CDN Y provides adequate performance (at low cost) for all clients except for the left-most group, who must be served by CDN X. Unfortunately for CDN X, this group is served by a very expensive cluster. CDN X could serve that group from a cheaper cluster, but reserves those clusters for the other groups, which unfortunately are sent to CDN Y by the broker.

We see potential pricing problems occur at the country level. Figure 5 shows how utilization of CDNs in our broker trace differ in different countries. The remaining percentage of users are serviced by other smaller CDNs. We show all countries that originated 100 or more requests during our trace, in random order. Note that utilization varies significantly: e.g., CDN B barely serves 7, yet almost entirely serves 8; CDN A is rarely used in 8, 11, and 15, etc. Different countries around the globe can have markedly different pricing for transit bandwidth. CDN CloudFlare claims that compared to the US (and Europe), Asia, Latin America, and Australia, cost 7x, 8x, and 20x more respectively [13]. Within a CDN the data paints a similar picture, showing up to a ~30x disparity in pricing between countries (Figure 3).

Most CDNs allow CPs to pay a flat rate per traffic delivered (based on a 95/5 model [10, 24]) with price changes (e.g., 2-7x) depending on very coarse geographic regions [5, 25]. With this flat rate, maximizing CDN profit means minimizing expenses such as transit costs. If (for example) a broker mainly uses a CDN in an expensive region (e.g., Australia) and nowhere else (see Figure 4), it might have a significant impact on the profitability of the CDN. This may drive the CDN to raise its prices in response to this traffic demand, which may cause the broker to shift demand away from them.

Unpredictable workloads: Figure 6 shows a time-series graph of the percentage of user sessions within a 10 second interval that have been shifted from one CDN to another midstream. We see this value is surprisingly

high throughout (averaging ~40%). We note, however, that at some points this dips to ~20% and at other times rises above ~60%. This indicates that not only do brokers often move traffic around, but the rate at which they do is highly variable. This potentially makes load-balancing and long-term provisioning difficult for CDNs.

Biased traffic patterns: Figure 7 shows the utilization for the CDNs in our trace, plotted as a function of number of requests per city. The dotted lines are best-fit lines. We see that regardless of city size, CDN B and CDN C’s usage does not change, whereas CDN A is strongly favored in smaller cities. This is perhaps unsurprising due to CDN A’s broader worldwide deployment. CDN A is also generally more expensive compared to CDN B and CDN C, suggesting that a broker will try to avoid CDN A where other options are available. Succinctly, brokers do not merely split traffic evenly among CDNs; traffic may be arbitrarily divided in geographic regions and change over time.

This leads us to believe that brokers and other CDNs can cause a CDN difficulty in planning and long-term provisioning (e.g., cluster location and size). For example, if the broker decides to stop using CDN A in big cities (e.g., CDN B deploys more servers) this will impact CDN A’s provisioning. If CDN B then raises its prices, the broker may move more traffic back to CDN A, again impacting provisioning.

Pricing pressure from specialization: Traditionally, CDNs required global coverage to provide adequate performance to all users. Brokers allow CPs to use many smaller CDNs, as long as in combination they provide global coverage. Brokers make smaller specialized CDNs feasible; for example, one could focus on delivery in remote areas, or on being the best within a given city. Although pricing pressure from the increased diversity of CDNs could lower delivery costs for CPs, it might lead to lower profit margins for CDNs.

3.3 Potential Problems for a Broker

CDNs create problems for brokers for two main reasons: 1) CDNs make decisions that impact end-to-end delivery without information about all clients seen by the broker,




Figure 8: Problems with coarse CDN-level selection granularity.

1 Alternate Choice	2 Alts.	3 Alts.	4 Alts.
77.8%	64.5%	53.7%	43.8%

Table 1: How often alternate CDN clusters with similar performance scores exist.

and 2) CDNs share little information with the broker. We highlight the implications of these problems.

Coarse resource selection: CDNs and brokers only directly interact with each other through a narrow interface: clients (on behalf of the broker) send DNS queries to the CDN to locate a server cluster. The CDN’s DNS response maps the user to a small number of clusters (e.g., one). Unfortunately, this means that when performance is inadequate based on the CP’s objectives, the broker’s only recourse is to switch CDNs (even if other better choices exist within the current CDN; see Figure 8). Effectively, the granularity of change a broker can make is very coarse.

Table 1 shows how often there are alternate clusters with similar estimated performance (based on latency and loss measurements) in the CDN data. We find that on average there are four server clusters (i.e., 3 alternate choices) that have similar scores (within 25% of the best), yet typically only one choice is returned. This data indicates a potential opportunity; for example, during a transient failure or performance issue, brokers could contact the CDN for one of these alternate clusters. Currently, a broker would move the traffic to a different CDN, which is unnecessarily extreme. Asking the CDN for an alternate cluster (rather than suddenly removing users) could make load-balancing easier for the CDN.

Inaccuracies in measurement: Both brokers and CDNs spend much effort building maps of the Internet to predict performance to and from clients and servers. This is by no means a small task; in recent work a broker claimed that they regularly handle 100M user sessions per day, 3M users concurrently during peak hours, and 10s–100s of thousands of users entering and exiting per minute [15]. They also imply that this leads to 50–100 GB of new sample data to process per minute. Sharing mapping information could greatly improve the accuracy of the data as both CDNs and brokers have limited vantage points into the network. Namely, CDNs such as




Figure 9: Sample design to fix the broker-CDN interface based on online ad exchange networks.

Akamai typically measure (in advance of connections) from clusters to gateway routers [22], whereas brokers generally only measure (during a connection) from users to the servers made available by the CDNs.

Difficult debugging: When users are getting bad QoE, CPs contact the broker. Figuring out what went wrong is difficult; the problem may lie with the broker software, the CDNs, the users’ ISPs, etc. This is further complicated as brokers will typically move these users to other CDNs, thus a poorly performing CDN would no longer see the problematic clients. Matching logs between brokers and CDNs to find the root cause is difficult as brokers and CDNs generally don’t share this kind of information.

4 The Path Forward

The problems seen are exacerbated by the narrow broker-CDN interface (DNS mapping users to clusters). Since minimal information is shared, brokers can’t make decisions explicit to CDNs and CDNs have little input in these decisions. We take the first steps towards fixing these issues by presenting a possible design of a much wider interface, based on ad exchanges.

4.1 Ad exchange-style Interface

We propose the following interface, which allows CDNs and brokers to share information more freely, allowing both parties to cooperatively decide on the assignment of users to clusters. We leverage techniques similar to those used in online advertisement exchanges. In our exchange, the broker “advertises” groups of users to CDNs, which make (multiple) cluster-level “bids” to serve these customers. CDNs express policy, performance, and cost considerations directly through their bids. The broker leverages many CDNs to meet CPs’ QoE and cost goals.

Our broker-CDN interface is an eight-step protocol (see Figure 9):

1. **Estimate:** CDNs estimate performance from candidate clusters to clients.
2. **Gather:** Broker counts users in each region.
3. **Announcement:** Broker advertises user counts to CDNs.
4. **Marriage:** CDNs match users to clusters.

5. **Bid**: CDNs send “bids” for users to broker.
6. **Optimize**: Broker decides which bids to accept.
7. **Accept**: Broker tells CDNs which bids are accepted.
8. **Setup**: Broker indicates to clients which CDN clusters to connect to.

Although our design is still preliminary, we try to directly address the bulk of the problems shown in §3. Namely, we focus on fine-grained CDN cluster selection and making resource allocation explicit (which solves pricing, unpredictability, biasing, and makes the effects of specialization more apparent). An out-of-band long-term control protocol could be used to address sharing measurement data and logs.

4.2 Examples

To provide an example of the power of this design, consider the scenario in Figure 4. Recall that today CDN X serves the client in the left group with a cluster that is expensive to delivery from, as it reserves the other clusters for clients in the other two groups. CDN X never sees these clients, as the broker decides that CDN Y can provide adequate service for less cost. With our design CDN X is given a more proactive role; it is asked by the broker to explicitly allocate clients to clusters. Through bidding CDN X can make its concerns clear to broker by explicitly telling it to have clients in the left group use one of the cheaper CDN X clusters if the broker sends clients in the other groups to another CDN.

Many other problems in §3 can be solved similarly. The impact of a shift of traffic to a CDN (Figure 6) is greatly reduced as CDNs are asked explicitly before the broker moves any traffic. Traffic pattern biasing (Figure 7) is solved similar to pricing concerns problem explained above; CDNs provide different client-server mapping strategies to the broker based on how the broker divides up incoming traffic. CDNs provide the broker cluster-level bids, providing a simple way to switch clusters as needed (Figure 8).

4.3 Discussion

This “auction”-style interface might appear complex compared to a more intuitive solution where CDNs simply provide dynamic pricing to CPs. However, dynamic pricing solutions provide their own set of difficulties (e.g., convergence, fairness, and stability [16, 18, 21]) that would require similar amounts of thought. Additionally, “auction”-style systems allow CDNs to directly adjust how their resources are being utilized over time and provide a simple solution for stitching together multiple CDN clusters within a CDN to serve one group of requests.

While this interface intuitively allows brokers to perform better, the incentives for CDNs to join are more subtle. In particular, it’s worth questioning if there are benefits that are worth giving up sensitive internal CDN

information to a broker. Some benefits are straightforward: CDNs like CDN B in Figure 7 can compete with other CDNs in smaller cities, as they did before brokers were introduced. Also, CDNs can now appropriately charge for difficult or hard-to-reach clients.

We argue, however, that while large CDNs might not like the idea, their customers (CPs) are already using multiple CDNs to achieve their specific QoE and cost goals. As CPs enlist more small-scale CDNs, larger CDNs may have difficulty competing against increased specialization. Our proposed interface allows large CDNs to remain competitive by making the specific needs of the customer (CP) more transparent between the broker and CDNs (e.g., which bids from a CDN are accepted reveal which regions need higher QoE or lower cost than the CDN currently provides). Conversely, general adoption of such an interface would also make it easier for smaller CDNs to compete, assuming some CPs allow them to serve a fraction of their traffic.

Parallels to failed CDN-ISP or ISP-P2P collaboration proposals are superficially similar, yet broker-CDN collaboration differs as both run at the application layer and both have the same customer (the CP).

5 Related work

The most relevant related work looks at widening interfaces in content delivery through collaboration. This includes alternative CDN designs, such as federated Telco-CDNs [7] and P2P-CDN hybrids [7, 32], and the potential benefits of CDN/ISP collaboration [14, 29, 30] These works show that ISPs can aid CDNs with assigning users to CDN clusters. Experience Oriented Network Architecture (EONA) [19] argues that content owners and infrastructure owners should collaborate to improve the end user’s QoE. Though similar, our work focuses on analyzing data to understand the specific problems between CDNs and brokers.

6 Conclusion

The limited interface between CDNs and brokers is causing a variety of problems for both parties. We conclude that a better broker-CDN interface could help mitigate problems. However, this paper only scratches the surface, as we only focused on single broker systems. As the content-delivery ecosystems evolves to include multiple brokers and specialized CDNs, new complex interactions requiring more analysis and new interfaces are likely to arise.

Acknowledgments

The authors would like to thank Raja Sambasivan, Peter Steenkiste, Bala Chandrasekaran, KC Ng, Junchen Jiang, Yuchen Wu, David Naylor, and the anonymous reviewers for their help and feedback. This work is supported in part by NSF awards #CNS-1345305 and #CNS-1345284.

7 References

- [1] Cisco Report: The CDN Federation - Solutions for SPs and Content Providers To Scale a Great Customer Experience. http://www.cisco.com/c/dam/en_us/about/ac79/docs/sp/CDN-Federation_Phase-2-Pilot.pdf.
- [2] Level3 Communications. <http://www.level3.com/en/products/content-delivery-network/>.
- [3] Conviva. <http://www.conviva.com>, 2015.
- [4] AMAZON. Amazon cloudfront. <https://aws.amazon.com/cloudfront/>, 2016.
- [5] AMAZON. Amazon cloudfront pricing. <https://aws.amazon.com/cloudfront/pricing/>, 2016.
- [6] ARSTECHNICA. It's not a fast lane but comcast built a cdn to charge for video delivery. <http://arstechnica.com/information-technology/2014/05/its-not-a-fast-lane-but-comcast-built-a-cdn-to-charge-for-video-delivery/>, 2014.
- [7] BALACHANDRAN, A., SEKAR, V., AKELLA, A., AND SESHAN, S. Analyzing the potential benefits of cdn augmentation strategies for internet video workloads. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (New York, NY, USA, 2013), IMC '13, ACM, pp. 43–56.
- [8] BALACHANDRAN, A., SEKAR, V., AKELLA, A., SESHAN, S., STOICA, I., AND ZHANG, H. Developing a predictive model of quality of experience for internet video. In *SIGCOMM '13: Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM* (Aug. 2013), ACM Request Permissions.
- [9] BIZETY. The multi-cdn strategy. <https://www.bizety.com/2014/05/09/multi-cdn-strategy/>, 2014.
- [10] BUYYA, R., PATHAN, M., AND VAKALI, A. *Content delivery networks*, vol. 9. Springer Science & Business Media, 2008.
- [11] CALDER, M., FAN, X., HU, Z., KATZ-BASSETT, E., HEIDEMANN, J., AND GOVINDAN, R. Mapping the expansion of google's serving infrastructure. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (New York, NY, USA, 2013), IMC '13, ACM, pp. 313–326.
- [12] CHEN, F., SITARAMAN, R. K., AND TORRES, M. End-user mapping: Next generation request routing for content delivery. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (New York, NY, USA, 2015), SIGCOMM '15, ACM, pp. 167–181.
- [13] CLOUDFLARE. The relative cost of bandwidth around the world. <https://blog.cloudflare.com/the-relative-cost-of-bandwidth-around-the-world/>, 2014.
- [14] FRANK, B., POESE, I., SMARAGDAKIS, G., FELDMANN, A., MAGGS, B. M., UHLIG, S., AGGARWAL, V., AND SCHNEIDER, F. Collaboration opportunities for content delivery and network infrastructures. *Recent Advances in Networking 1* (2013), 305–377.
- [15] GANJAM, A., SIDDIQUI, F., ZHAN, J., LIU, X., STOICA, I., JIANG, J., SEKAR, V., AND ZHANG, H. C3: Internet-scale control plane for video quality optimization. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)* (Oakland, CA, May 2015), USENIX Association, pp. 131–144.
- [16] GIBBENS, R. J., AND KELLY, F. P. Resource pricing and the evolution of congestion control. *Automatica* 35, 12 (1999), 1969–1985.
- [17] GOOGLE. Google cloud cdn. <https://cloud.google.com/cdn/>, 2016.
- [18] HAN, D., GRANDL, R., AKELLA, A., AND SESHAN, S. Fcp: a flexible transport framework for accommodating diversity. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 135–146.
- [19] JIANG, J., LIU, X., SEKAR, V., STOICA, I., AND ZHANG, H. EONA: Experience-Oriented Network Architecture. In *HotNets-XIII: Proceedings of the 13th ACM Workshop on Hot Topics in Networks* (Oct. 2014), ACM Request Permissions.
- [20] JIANG, J., SEKAR, V., MILNER, H., SHEPHERD, D., STOICA, I., AND ZHANG, H. Cfa: A practical prediction system for video qoe optimization. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation* (Berkeley, CA, USA, 2016), NSDI'16, USENIX Association, pp. 137–150.
- [21] KELLY, F. P., MAULLOO, A. K., AND TAN, D. K. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society* 49, 3 (1998), 237–252.
- [22] LISIECKI, P. A., NICOLAOU, C., AND ROSE, K. R. Scalable, high performance and highly available distributed storage system for internet content, Nov. 24 2009. US Patent 7,624,169.
- [23] LIU, X., DOBRIAN, F., MILNER, H., JIANG, J., SEKAR, V., STOICA, I., AND ZHANG, H. A case for a coordinated internet video control plane. In *Proc. ACM SIGCOMM* (2012), pp. 359–370.
- [24] MAXCDN. Which cdn pricing model costs less. <https://www.maxcdn.com/blog/cdn-pricing-models/>, 2016.
- [25] MICROSOFT. Content delivery network (cdn) pricing. <https://azure.microsoft.com/en-us/pricing/details/cdn/>, 2016.
- [26] NETFLIX. Netflix open connect. <https://openconnect.netflix.com/en/>, 2016.
- [27] NIVEN-JENKINS, B., LE FAUCHUER, F., AND BITAR, N. RFC 6707: Content Distribution Network Interconnection (CDNI) Problem Statement. <https://tools.ietf.org/html/rfc6707>.
- [28] NYGREN, E., SITARAMAN, R. K., AND SUN, J. The akamai network: a platform for high-performance internet applications. *ACM SIGOPS Operating Systems Review* 44, 3 (2010), 2–19.
- [29] POESE, I., FRANK, B., KNIGHT, S., SEMMLER, N., AND SMARAGDAKIS, G. PaDIS emulator: an emulator to evaluate CDN-ISP collaboration. In *SIGCOMM '12: Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication* (Aug. 2012), ACM.
- [30] POESE, I., FRANK, B., SMARAGDAKIS, G., UHLIG, S., FELDMANN, A., AND MAGGS, B. Enabling content-aware traffic engineering. *SIGCOMM Computer Communication Review* 42, 5 (Sept. 2012).
- [31] USC-NETWORK SYSTEMS LAB. CDN Geographic Coverage. <http://usc-nsl.github.io/cdn-coverage/>.
- [32] ZHAO, M., ADITYA, P., CHEN, A., LIN, Y., HAEBERLEN, A., DRUSCHEL, P., MAGGS, B., WISHON, B., AND PONEC, M. Peer-assisted content distribution in akamai netsession. In *Proceedings of the 2013 conference on Internet measurement conference* (2013), ACM, pp. 31–42.