UNIVERSIDAD CARLOS III DE MADRID

TESIS DOCTORAL

Alternative Revenue Sources for Internet Service Providers

Autor:
Pradeep Bangera

Director:
Dr. Sergey Gorinsky, IMDEA Networks Institute

DEPARTAMENTO DE INGENIERÍA TELEMÁTICA

Leganés (Madrid), 2016

UNIVERSIDAD CARLOS III DE MADRID

Ph.D. Thesis

Alternative Revenue Sources for Internet Service Providers

Author:
Pradeep Bangera

Director:
Dr. Sergey Gorinsky, IMDEA Networks Institute

DEPARTMENT OF TELEMATIC ENGINEERING

Leganés (Madrid), 2016

Alternative Revenue Sources for Internet Service Providers

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Prepared by
Pradeep Bangera

Under the advice of
Dr. Sergey Gorinsky, IMDEA Networks Institute

Departamento de Ingeniería Telemática, Universidad Carlos III de Madrid

---

Date: September, 2016

Web/contact: pradeep.bangera@imdea.org

This work has been supported by IMDEA Networks Institute.

TESIS DOCTORAL

Alternative Revenue Sources for Internet Service Providers

Autor: Pradeep Bangera
Director: Dr. Sergey Gorinsky, IMDEA Networks Institute

Firma del tribunal calificador:

Firma:

Presidente:

Vocal:

Secretario:

Calificación:

Leganés, de de

Dedication

I dedicate this work to The True and The Living God who revealed himself to me in the person of The Lord Jesus Christ by The Holy Spirit to the ultimate Truth which is able to save anyone who believe and practice, at one's own free-will.

I sought the Lord, and he heard me, and delivered me from all my fears [psalm 34:4]. The Lord redeemeth the soul of his servants: and none of them that trust in him shall be desolate [psalm 34:22]. Thy word is a lamp unto my feet, and a light unto my path [psalm 119:105]. Strengthened with all might, according to his glorious power, unto all patience and longsuffering with joyfulness; Giving thanks unto the Father, which hath made us meet to be partakers of the inheritance of the saints in light: Who hath delivered us from the power of darkness, and hath translated us into the kingdom of his dear Son: In whom we have redemption through his blood, even the forgiveness of sins: Who is the image of the invisible God, the firstborn of every creature: For by him were all things created, that are in heaven, and that are in earth, visible and invisible, whether they be thrones, or dominions, or principalities, or powers: all things were created by him, and for him: And he is before all things, and by him all things consist [Colossians 1:11-17].

Disclaimer: The Spiritual gratitude expressed in this page rests solely with the author. The dedication does not intend to represent the views of the other individuals and organizations affiliated with this thesis

Acknowledgements

From a formless to a near definite structure, the molding of this thesis to its current state has been possible only because of tremendous amount of encouragements from many helping hearts and valuable advices from specific intellectual minds. In the following, I would like to express my heartiest gratitude and sincere respect to each one of them.

First, all due credit goes to my doctoral advisor Dr. Sergey Gorinsky for any achievements (publications) arising from this thesis, because he introduced me to the Internet economics research and skillfully edited our published papers. Without him, I would never had a thought to study this domain nor would have published a single paper in a good competitive conference. I am also extremely grateful to him for enabling me in scaling-up my English writing skill, from a school-level to a more acceptable standard expected for writing quality articles. Still, I am learning and improving with every new piece of writing, including this thesis, to scale the summit of art of writing quality articles.

My heartfelt prayers and gratitudes to my beloved brothers in the Lord, J. S. Kumar, Chitharanjan and Antony, for providing immense emotional support, love and encouragement until this day and even for evermore. I particularly express my eternal gratitude to my brother J. S. Kumar for enduring and guiding me with great patience during my troublesome times. I am thankful to my parents for their care and concern for me. My heartiest gratitude for my uncle Sitharam Bangera for taking care of my financial needs during my first trip to Europe that led to my first step in the scientific research marathon. To my good friends, Vinodkumar, Srivatsa, Abhik Ghosh, Sudhirkrishna, Pradeep Tiwari and Andres Segura Aragones, I thank you deeply for remembering me amidst your busy work and family life.

My sincerest thanks to my fellow research colleagues at IMDEA Networks Institute: Kshitiz Verma for brainstorming on the concepts of competition and cooperation in the Internet, and Ignacio Castro for enlightening me on the basics of economics. I want to specially thank Syed Hasan– my fellow research colleague at IMDEA Networks, my good friend and my flat-mate for 3 years in Leganes; for besides research-related brainstorming, also for inspiring me on several aspects of life– sincerity, honesty, selflessness, patience, humility and determination.

Tons of thanks to the management of the IMDEA Networks Institute for driving,

encouraging and supporting quality research activities, and giving opportunity to people from across the world and culture, and providing a pleasant work atmosphere. My sincere and humble thanks to Arturo Azcorra for helping me and making it possible to get access to the RedIRIS traffic data, which otherwise was impossible, and also many thanks to Jose Felix and Joel Rosental for providing me with all the technical assistance during the traffic data collection process. Much appreciation and thanks to all the PhD students, Research Assistant Professors, Research Associate Professors, Research Professors, Visiting Professors, Admin and Research Support team whom I met and interacted during my stay at IMDEA Networks Institute between early November, 2009 until the end of September, 2013. I personally want to thank Rebecca de Miguel and Ana Gonzalez for exhorting me on my last working day at IMDEA Networks, to finish my Ph.D. thesis, at which point in time the status of this thesis was not only formless but also needed a very brutal uphill run towards completion. Your exhortations really played a vital role in diligently pursuing and completing this thesis. I also sincerely thank all the teaching and research staff at Telematic Engineering department at Carlos III University of Madrid for sharing your valuable knowledge during my master studies and during the weekly seminars on every Wednesday.

Lastly, my inexpressible apologies to Dr. Mounir Boussedjra, my former supervisor at ESIGELEC-IRSEEM, Rouen (France), for not wisely attending to your earnest requests, which I now deeply regret. But, I am very thankful to you for being a very wonderful and a cool boss with a very friendly attitude, that I ever worked for. Thanks a million for appreciating and inspiring me to pursue Ph.D. studies which motivated my journey to Spain.

Abstract

The Internet has evolved from a small research network towards a large globally interconnected network. The deregulation of the Internet attracted commercial entities to provide various network and application services for profit. While Internet Service Providers (ISPs) offer network connectivity services, Content Service Providers (CSPs) offer online contents and application services. Further, the ISPs that provide transit services to other ISPs and CSPs are known as transit ISPs. The ISPs that provide Internet connections to end users are known as access ISPs. Though without a central regulatory body for governing, the Internet is growing through complex economic cooperation between service providers that also compete with each other for revenues. Currently, CSPs derive high revenues from online advertising that increase with content popularity. On other hand, ISPs face low transit revenues, caused by persistent declines in per-unit traffic prices, and rising network costs fueled by increasing traffic volumes.

In this thesis, we analyze various approaches by ISPs for sustaining their network infrastructures by earning extra revenues. First, we study the economics of traffic attraction by ISPs to boost transit revenues. This study demonstrates that traffic attraction and reaction to it redistribute traffic on links between Autonomous Systems (ASes) and create camps of winning, losing and neutral ASes with respect to changes in transit payments. Despite various countermeasures by losing ASes, the traffic attraction remains effective unless ASes from the winning camp cooperate with the losing ASes. While our study shows that traffic attraction has a solid potential to increase revenues for transit ISPs, this source of revenues might have negative reputation and legal consequences for the ISPs. Next, we look at hosting as an alternative source of revenues and examine hosting of online contents by transit ISPs. Using real Internet-scale measurements, this work reports a pervasive trend of content hosting throughout the transit hierarchy, validating the hosting as a prominent source of revenues for transit ISPs. In our final work, we consider a model where access ISPs derive extra revenues from online advertisements (ads). Our analysis demonstrates that the ad-based revenue model opens a significant revenue potential for access ISPs, suggesting its economic viability.

Contents

Dedication **ii**

Acknowledgements **v**

Abstract **vi**

Contents **ix**

List of Figures **xii**

List of Tables **xiii**

1 Introduction **1**

1.1 Motivation 2

1.2 Contributions 3

1.3 Thesis overview 4

1.4 Research publications out of our contributions 5

2 Background and Basic Information **7**

2.1 Basic building blocks of the Internet 7

2.1.1 Internet Protocol (IP) 7

2.1.2 Border Gateway Protocol (BGP) 8

2.1.3 Inter-AS business relation types 8

2.1.4 Network and content hosting infrastructure 9

2.1.5 Content service providers and online advertisements 10

3 Economics of Traffic Attraction by Transit ISPs **13**

3.1 Introduction 13

3.2 Real case study of YouTube traffic attraction 14

3.2.1 Methodology 14

3.2.2 Impact of traffic attraction on Inter-AS paths and traffic 16

3.3 Traffic attraction by large transit ISPs 18

3.3.1 Model 20
3.3.2 Evaluation methodology 22
3.3.3 Evaluation results 23
3.4 Attraction viability 34
3.5 Related work 35
3.6 Summary 36

4 Dissecting the Online Content Hosting Ecosystem 39
4.1 Introduction 39
4.2 Methodology 41
4.2.1 Data collection 41
4.2.2 Identifying the ad URLs 42
4.2.3 Identifying the regular URLs 42
4.3 Discovering the content infrastructure 44
4.3.1 Eliminating the misleading DNS servers 44
4.3.2 Resolving the URLs using open recursive DNS 44
4.3.3 IP-to-hosting infrastructure mapping 45
4.4 Measurement results 45
4.4.1 Hosting-infrastructure characteristics 45
4.4.2 Byte volume and location of contents 52
4.4.3 Content delivery performance 55
4.5 Discussion: Economic impact of Internet advertising 56
4.6 Related work 57
4.7 Summary 58

5 An ad-based Revenue Model for Access ISPs 59
5.1 Introduction 59
5.2 Background 61
5.3 Model 62
5.3.1 Ad revenue potential 62
5.3.2 Utility of the users 63
5.4 Data 64
5.4.1 Financial data of access ISPs 64
5.4.2 Ad pricing data 65
5.5 Analysis 66
5.5.1 Ad revenue potential 66
5.5.2 Economic viability of the ad-based revenue model 68
5.5.3 Incentives for users 70
5.6 Related Work 72
5.7 Summary 73

6 Conclusions and Future Work **75**

References **86**

List of Figures

3.1 Number of BGP paths through transit ASes (a) before the YouTube's prefix hijacking by PTCL and (b) after PTCL hijacks all YouTube-bound traffic. 16
3.2 Percentage transit traffic changes for the top-5 winning and loosing transit ASes. 17
3.3 Example of traffic attraction via prefix deaggregation: (a) before attraction; (b) after attraction. 19
3.4 Distribution of origin traffic shares in our traffic matrices. 21
3.5 Payment change for the attracting AS. 24
3.6 Payment changes for the tier-1 attractors. 24
3.7 Payment changes for the winning and losing ASes when the attracting AS is from tier 1, 2, or 3. 25
3.8 Changes in the upstream, downstream, and peering traffic of the 10 largest losers when the attractor is a tier-1 AS: (a) without filtering; (b) with filtering. 25
3.9 Payment change for the attracting AS when all losers do the filtering. . . 26
3.10 Payment change for attractor T1b with multi-stage filtering. 27
3.11 Payment change for attractor T1b with multi-stage disconnection by losing customers of the attractor. 28
3.12 Payment changes for 10 ASes when the traffic attraction is done by T1b only vs. all 10 ASes. 29
3.13 Changes in the upstream, downstream, and peering traffic of 10 ASes when the attraction is done by: (a) T1b only and (b) all 10 ASes. 29
3.14 Sensitivity of payment changes to the AS-level Internet topology. 30
3.15 Sensitivity of the attractor's payment change to the traffic matrix. 31
3.16 Sensitivity of the attractor's payment change to transit pricing. 32
3.17 Sensitivity of the attractor's payment change to attraction intensity. . . . 33
3.18 Distributions of AS-level path lengths. 33

4.1 Data collection using a VPN 41

4.2 Distribution of (a) IP-address count and (b) AS count for the ad and regular contents of the websites. 46
4.3 Hosting ASes arranged according to the number of websites served in each country: (a) ad-hosting ASes and (b) regular-hosting ASes. 47
4.4 Cumulative distribution of IP addresses for different clusters of: (a) ad-hosting ASes and (b) regular-hosting ASes. 48
4.5 Cumulative fraction of IP addresses arranged according to the customer cone size of: (a) ad-hosting ASes and (b) regular-hosting ASes 48
4.6 Cumulative distribution of the customer-cone size of: (a) ad-hosting ASes and (b) regular-hosting ASes. 50
4.7 Number of websites served by: (a) ad-hosting clusters and (b) regular-hosting clusters, arranged according to the number of IP addresses per website. 51
4.8 Fraction of online content bytes hosted by the ASes according to (a) customer-cone size; and (b) peering coefficient. 53
4.9 Distribution for the number of hosting countries with respect to: (a) fraction of ASes and (b) fraction of hosted bytes. 54
4.10 Delivery performance for regular contents and ads: (a) response times (b) Download times. 55

5.1 Responses of the survey: (a) questions Q1 and Q2 and (b) question Q3. . 62
5.2 Ad revenue potential as a percentage of the average revenues, CapEx, and OpEx of the ISPs. 66
5.3 Ad revenues as a percentage of the CapEx for different CPMs and 3 distributions of online time. 66
5.4 CPM variation for Airtel for different average online times and different fractions of the CapEx. 67
5.5 CPM variation for Den Networks for different average online times and different fractions of the CapEx. 68
5.6 Ad revenues as a percentage of CapEx for different ad frequencies. 68
5.7 User incentives that Airtel and Den can offer: (a) speed and (b) data cap. 69
5.8 Utility changes for providing extra speed as incentives by (a) Airtel and (b) Den. 70
5.9 Utility changes for offering extra data as incentives by (a) Airtel and (b) Den. 71

List of Tables

4.1 Top 10 ASes arranged according to the number of IP addresses used for hosting ads and regular contents in clusters A_h and R_t 49
4.2 Top 10 ASes arranged according to the bytes served for ads and regular contents 52
4.3 Top 5 CDNs arranged according to the served bytes 53
4.4 Top-10 hosting countries ranked by the AS counts and fractions of the hosted bytes. 54

5.1 Notations used for deriving the users-utility function. 63
5.2 Monthly average financial data as of March 31, 2015 for 2 Indian access ISPs. 64
5.3 CPM prices received by a popular Indian entertainment website. 65
5.4 Ad revenue potential of the 2 access ISPs with $t = 60$ minutes and $f_a = 1$ impression/minute. 65
5.5 Speed and data prices offered by Airtel and Den Networks. 69

Chapter 1

Introduction

The Internet began with the interconnections of few computers as part of a research project by DARPA (Defense Advanced Research Project Agency). It rapidly evolved from being a small research network towards a large globally distributed network of millions of computers. With the deregulation of the Internet, many commercial entities entered the field to provide various services for profit. The entities that provide network infrastructure and connectivity services are called Internet Service Providers (ISPs). The entities that provide online contents and application services are referred to as Content Service Providers (CSPs). An ISP manages its network by segmenting the network into one or more Autonomous Systems (ASes) – collection of network routers and switches managed under a common administrative policy. From an economic perspective, Internet infrastructure services can be classified into 3 categories: transit, hosting, and access. While access ISPs supply Internet connections for residential and enterprise users, hosting ISPs serve CSPs by providing computing, storage and networking resources for hosting online contents on the Internet. Transit ISPs offer connectivity to the global Internet and operate geographically distributed communication infrastructures. During the early stage, the Internet topology evolved to form a hierarchical structure, wherein large transit ISPs, known as tier-1 providers, interconnected with each other to form the core of the Internet. Surrounding the core, the intermediate layer of the Internet topology consisted of medium and small transit ISPs, known as tier-2 and tier-3 providers respectively. On the periphery of the intermediate layer, CSPs, access, and hosting ISPs formed the outer edge of the Internet.

The creation and maintenance of the ISP network infrastructures involve substantial costs. It is common for the ISPs to recover the costs by charging customers for connectivity services. Access ISPs charge monthly subscription fees to end users. Hosting ISPs charge CSPs for the provided computing, storage and networking resources. CSPs rely primarily on online advertisements (ads), besides selling goods online and offering paid online services. For traffic delivery, access, and hosting providers pay transit providers

for traffic delivery across the Internet.

As alluded, the Internet is a network of thousands of interconnected ISPs operated by profit-driven business entities. Each ISP meticulously engineers its own physical network to enable efficient paths for its internal traffic flows. The interconnections between ISPs are mainly driven by economic interests and business negotiations. To establish inter-ISP business relations, an ISP negotiates bilateral contracts with its neighbors based on many factors, such as potential traffic demand. Therefore from an economic point of view, the Internet is a loosely connected network of inter-ISP business relations established on a basis of complex economic policies.

1.1 Motivation

Currently, the Internet does not have a central governing body to directly regulate all the service providers across the globe. Yet, the Internet is growing through complex economic cooperation between service providers that also compete with each other for revenues. In the current Internet economy, the CSPs' revenue growth is 2 times higher than the ISPs' revenue growth [1]. The CSPs' economic growth is mainly boosted by online ad revenues which grow with the online content consumption by end users. Therefore, the CSPs have incentives to create advanced innovative online contents, which often consume substantial network capacity, e.g., online 3D video streaming.

On the other hand, transit ISPs' revenues mainly depend on the traffic volume on their links with their transit customers. While access ISPs' revenues depend on fixed monthly subscription fees paid by end users, their network costs consisting of transit payments and capital expenses are traffic-dependent. Thus a larger traffic volume provides higher incentives to transit ISPs in the form of transit revenues and lower incentives to access ISPs due to rising network costs. Currently, ISPs generate a lower Return on Investment (RoI) due to falling per-unit traffic prices in the transit sector, and rising network costs in the access sector due to increasing traffic volume in the access ISPs' networks [1, 2]. Besides, the capital and operational expenditures (CapEx and OpEx respectively) of ISPs are higher compared to the expenditures of CSPs [3].

In the Internet core, persistent declines in the transit traffic prices [2] as well as extensive peering and content caching at the Internet edge [4–7] have recently put transit businesses under significant financial pressure [8]. To boost revenues, transit providers experiment with new economic models, such as paid peering [9], remote peering [10, 11], group purchase of transit [12, 13] and partial transit [14], to diversify the offered connectivity products. Further, there exists anecdotal evidence that few transit providers use unorthodox routing techniques to attract extra traffic of transit customers [15, 16]. Hence, this thesis begins by studying incentives for transit ISPs to boost their revenues via traffic attraction.

As a consequence of the financial challenges in Internet transit, transit providers have branched out into providing hosting services to CSPs. To achieve higher performance for the bandwidth-intensive online contents, CSPs approach transit ISPs to leverage the global transit infrastructure for caching their online contents geographically closer to end users. Thus, the content hosting by transit ISPs have transformed the content traffic sources from the earlier remote CSP-originated to distributed sources of content traffic across the Internet transit hierarchy. Our thesis explores this paradigm shift in the content hosting. This shift necessitates modeling of the content traffic matrix to also consider traffic sources in transit networks, which were earlier viewed as mere forwarders of end-to-end traffic.

Next, because of the increasing traffic originated by CSPs and their caches on transit networks, access ISPs face rising network costs due to frequent backbone capacity upgrades to accommodate the growing traffic. As access ISPs earn revenues in the form of fixed subscription fees, their profit margins decrease due to rising network costs. Recently, access ISPs demand financial payments from CSPs for allocating more capacity for the bandwidth-intensive content traffic or subtly degrade the content delivery performance at the access backbone [17]. Such tussles have led to many net-neutrality wars, where CSPs accuse access ISPs of traffic discrimination by deliberate performance degradation [18–20]. On the other hand, access ISPs are also in disputes with their transit providers regarding violations of prior-agreed traffic ratios [21]. Thus, in this thesis, we propose and analyze how access ISPs can use online ads as an alternative revenue source.

1.2 Contributions

In this thesis, we study alternative revenue sources that might help ISPs to economically sustain their network infrastructure. In our first work, we assess economic incentives for transit providers to boost transit revenues by attracting customer traffic. As mentioned earlier, a transit provider's revenue depends on the volume of customer traffic, i.e., the transit revenue increases as the customer-traffic volume increases. Therefore, in order to generate more revenues, a transit provider may have a financial interest in attracting extra customer traffic [22, 23]. First, we analyze a real incident of YouTube's traffic attraction that occurred in early 2008. Next, motivated by the insights from this real case study, we conduct in-depth studies for hypothetical scenarios of traffic attraction by ISPs across the transit hierarchy. Based on real empirical data for the Internet's AS-level topology and traffic, our extensive simulations and sensitivity analyses reveal that traffic attraction and reaction to it redistribute traffic on the inter-AS links and create camps of winning, losing and neutral ASes with respect to transit payments. Despite various countermeasures by losing ASes, the traffic attraction remains effective unless ASes from the winning camp cooperate with the losing ASes. Although the traffic attraction has

a potential to significantly increase revenues for transit providers, this technique might have negative reputation and legal consequences for the ISPs.

Next, we explore content hosting as a source of extra revenue. In our second work, we dissect the global hosting ecosystem to discover a wide presence of transit ISPs offering hosting and content delivery services. To explore the global content hosting ecosystem, we conduct real Internet-scale measurements using a novel technique that leverages a Virtual Private Network (VPN) to collect real online contents from top 2,165 websites across 52 countries, and use a vast network of around 22,000 open recursive Domain Naming System (DNS) vantage points to discover the infrastructures hosting the online contents. Our results show that ISPs host online contents across the transit hierarchy, significantly on tier-2 and tier-3 transit ASes, followed by access ASes, CSPs, and tier-1 transit ASes. The latter results are valuable for realistically modeling of the content traffic matrix. Besides, our analyses also provide realistic understanding on the network characteristics of the hosting infrastructures and content delivery performance as perceived by end users.

While transit providers seek extra revenues in hosting of online contents, access providers face rising network costs due to the increasing volume of bandwidth-intensive content traffic. In our third and final work, we evaluate an economic model for access ISPs to derive extra revenues from online advertisements. First, we evaluate an ad-based revenue model for access ISPs. Then, we analyze its revenue potential and economic viability for different access ISPs. We validate our model using financial data collected from two, one large and one medium-sized, access ISPs operating in India. Our analyses demonstrate a significant revenue potential up to 50% of the capital expenses) for the large access ISP and non-trivial gain up to 5% of the capital expenses for the medium-sized access ISP. Next, we establish conditions for economic viability of the model. Lastly, we demonstrate various incentives (6–9 Mbps extra speed or 12–20 GB extra data) that an access ISP can offer to subscribers of an ad-subsidized Internet plans. Unlike the content-sponsorship models where CSPs directly compensate access ISPs for content delivery costs, our ad-based revenue model relies on direct collaboration of access ISPs with advertisers, lessening the concerns about network neutrality.

1.3 Thesis overview

The thesis is organized as follows. Chapter 2 presents background information related to our research area. In Chapter 3, we present our studies on economics of customer-traffic attraction by transit providers to boost transit revenues. In chapter 4, we conduct a real Internet-scale measurement study of the global content hosting ecosystem to explore the presence of transit providers in the hosting of online contents. In chapter 5, we derive an ad-based revenue model for access ISPs. Chapter 6 concludes the thesis with a discussion of our research findings and future work on economic, technical, and security aspects of

the Internet infrastructure spectrum.

1.4 Research publications out of our contributions

Accepted papers:

- Pradeep Bangera and Sergey Gorinsky, "Traffic Attraction by Internet Transit Providers: An Economic Perspective", in the Proceedings of IFIP Networking 2014 [24].
- Pradeep Bangera and Sergey Gorinsky, "An Economic Perspective on Traffic Attraction by Internet Transit Providers", in the Proceedings of IEEE ICNP 2013 (Poster) [25].
- Pradeep Bangera and Sergey Gorinsky, "Impact of Prefix Hijacking on Payments of Providers", in the Proceedings of COMSNETS 2011 [26].

Working papers:

- Pradeep Bangera, and Sergey Gorinsky, "Content versus Ads: Dissecting the Hosting Ecosystem" *Part of this work was accepted by IFIP Networking 2016 and later withdrawn from publication.*
- Pradeep Bangera, Syed Hasan and Sergey Gorinsky, "AdBroadband: Analysis of Economic Viability of Ad-based Revenue Model for Access ISPs"

Chapter 2

Background and Basic Information

In this chapter, we present and discuss basic information pertaining to our research problems.

2.1 Basic building blocks of the Internet

While the complex interconnections of routers and switches have spun the Internet, the AS-level graph is commonly employed for visualizing the Internet topology. For universal end-to-end communication, ISPs utilize the Internet Protocol (IP) [27] for addressing network interfaces in the Internet, and Border Gateway Protocol (BGP) [28] for establishing AS-level paths to destinations.

2.1.1 Internet Protocol (IP)

An IP prefix is a succinct representation for contiguous IP addresses owned by an AS. An IP router contains a routing and forwarding table. The IP prefixes associated with various destination networks are stored in the routing table. For traffic engineering, an AS might deaggregate an IP prefix into multiple longer prefixes and announce these longer deaggregated prefixes to other ASes via BGP [29]. The routing table contains separate entries for multiple paths to the same IP prefixes learned by a router from different neighboring routers. The forwarding table in a router compactly maps each IP prefix to an output interface of the router. During packet transmission, an IP router quickly forwards an incoming IP packet to the output link determined through the longest-prefix match rule, which selects the longest among the matching prefixes.

2.1.2 Border Gateway Protocol (BGP)

While IP addresses identify network-device interfaces in the Internet, BGP establishes inter-domain connectivity between neighboring ASes to enable traffic between them. BGP is a path-vector protocol that serves as a de facto standard protocol for routing between ASes and exchanges AS-level best-path announcements between neighboring ASes to support global reachability of IP prefixes. The announcing AS either owns the advertised prefixes or learns about the best paths to the prefixes from other neighboring ASes. Generally, BGP selects the best path to a destination IP prefix by comparing the path attributes of different AS-level paths. Usually, network operators configure routing policies in the router to set operator-preferred path attributes instead of relying on default BGP values.

BGP uses various attributes that can influence the best path of the incoming IP traffic, such as *AS-PATH length, origin-type, and MED*. The *AS-PATH length* denotes the number of AS hops to reach a destination IP prefix. The *origin-type* attribute can assume three values: *IGP* (BGP-originated routes), *EGP* (Exterior Gateway Protocol routes) and *INCOMPLETE* (for other protocols). The Multi-Exit Discriminator (*MED*) is used to assign priorities among multiple links of the AS with its immediate neighbors. During the selection of the best path to an IP prefix, BGP first prefers the shortest AS path to an IP prefix. To break ties between two paths with equal lengths, BGP prefers the path having the IGP attribute over the path with *EGP* and *INCOMPLETE* attributes. To receive traffic from a neighboring AS on the most-prioritized link, operators assign the smallest *MED* value for the preferred link.

For storing and processing the path attributes, BGP maintains 3 logical tables called RIBs (Routing Information Bases) [30, 28, 31]: (1) Adj-RIB-In (Adjacent RIB Input) stores all learned AS level paths, including multiple paths per prefix; (2) Loc-RIB (Local RIB) remembers the best path for every prefix; (3) Adj-RIB-Out (Adjacent RIB Output) stores the paths announced to other routers.

2.1.3 Inter-AS business relation types

Before employing IP and BGP for traffic delivery, each ISP first negotiates and forms economic agreements for interconnections with neighboring ASes. In the transit agreement, customer ASes pay their transit providers for traffic delivery. Usually, the business-relation type between a customer AS and its transit provider is called customer-provider (or provider-customer) [32] relation. Unlike transit agreements, peering agreements do not result in inter-AS traffic delivery payments, but ASes equally share peering infrastructure costs [33]. The inter-AS relation type in the latter is called a peering relationship. Further, a link between two ASes that belong to the same ISP is called a sibling link. Traffic on sibling links does not result in any inter-provider payment.

2.1.4 Network and content hosting infrastructure

While ISPs negotiate and interconnect their networks, each ISP offers specialized services to cater to different sections of customers, such as transit customers, CSPs, and residential users.

Transit ASes are interconnected in a hierarchical pattern, with large tier-1 ASes at the top of the hierarchy (i.e., the Internet core) and smaller tier-3 ASes at the bottom of the hierarchy. The tier-1 transit providers operate large network infrastructures, consisting of one or more ASes spanning the globe and having Points of Presence (PoP) in multiple countries. The tier-1 providers do not pay any other ISPs for traffic exchange. They usually maintain settlement-free peering relationships with other tier-1 providers and receive payments from their transit customers for sending and receiving traffic. There are about a dozen of ISPs that are tier-1 [34]. Tier-2 transit providers are smaller than the tier-1 providers, have continental coverage, and pay tier-1 providers for traffic delivery. Tier-3 providers are usually regional and national transit providers that buy transit from tier-2 or tier-1 ISPs. The transit providers commonly employ the 95th-percentile billing model for charging for transit of customer traffic [35, 36]. This method measures the average traffic rate over each 5-minute interval during a 1-month billing period and then uses the 95th percentile of all the individual traffic-rate samples as a basis for calculating the monthly bill.

Hosting ISPs and Content Delivery Networks (CDNs) [37, 38] specialize in hosting and delivering the online contents of CSPs. The CDNs distribute the online contents by deploying caches in different ASes across the Internet hierarchy and across geography. As a result, end users enjoy low latencies and high throughput because of the geographical proximity of the content caches. To fetch the contents from the nearest caches, CDNs usually employ IP anycast forwarding [39] or their proprietary protocols by using a Canonical Name (CNAME) record to select the best cache. The selection considers various metrics, such as the user location, cache load, inter-AS path length, and link latencies. While CDNs carefully select in which ASes to place caches to improve content delivery performance, access ISPs also have economic incentives to offer their networks for cache placements in order to reduce their transit costs by delivering most traffic locally rather than through a transit-provider link.

Demand for Internet access is rising as online contents prove themselves greatly valuable for satisfying information, entertainment and business needs. The technological advancement in the Internet access infrastructure has evolved from a dial-up Internet access with the maximum data speed of 56 Kbps to high-speed broadband access with data speeds over 100 Mbps, suitable for the current bandwidth-intensive online contents. With the growing demand for higher Internet speeds as well as increasing competition in the access market, access ISPs have been continuously upgrading their network capacity and geographical coverage by making huge capital investments year-on-year. Driven by in-

tense competition, most of the access ISPs have adopted the fixed monthly payment model that is largely traffic-agnostic. Therefore, as end users consume more online contents, the access ISPs experience higher traffic delivery costs. Besides cache placement collaborations with CDNs to reduce traffic costs, access ISPs also started negotiating with CSPs to sponsor the delivery of online contents to end users, i.e., CSPs are asked to pay some fraction of the cost incurred by access ISPs for delivering the content traffic. The latter content sponsorship model recently attracted much academic research interests [40–42] and also intense debates on protecting network neutrality.

2.1.5 Content service providers and online advertisements

While transit and access ISPs face low returns on their network infrastructure investments due to the unsustainable traffic costs, CSPs earn high revenues through online advertisements. Presently, almost all of the online ads are published by CSPs via their websites and mobile applications.

In the Internet economic ecosystem as a whole, a major share of the revenues come from online ads. Online ads enable CSPs to offer end users free access to a vast number of online contents and applications. Unlike the advertising in printed and broadcast electronic media, advertising on the Internet provides better control over the delivery of ads to the right set of audience, thus maximizing the Return-on-Investment (ROI) for the advertisers. The entities that advertise their products are known as advertisers. The CSPs who offer space on their websites for displaying ads are known as publishers. The entity that links the advertisers and publishers for displaying ads is known as an ad network. The ad network aggregates ads from several advertisers and selects a relevant website for displaying an ad. The process of displaying a single ad on a website is referred to as an ad impression. CSPs attract billions of end users to their websites, providing a highly appealing platform for the advertisers to publicize their products.

Thus, CSPs maintain contracts with one or more ad networks or may operate their own ad network to receive relevant ads. Besides aggregating ads from several hundreds of advertisers, an ad network also partners with behavioral targeting agencies to gain insights into user interests. Upon receiving the user visits, a website sends a request to the ad network to fetch relevant ads. The ad network then selects a relevant ad in real time, after processing a number of information, such as the advertisers' bid prices, ad quality score, website ranking, website category, and user profile. This information is gathered by the behavioral targeting agencies using web cookies. Finally, a relevant ad from the highest bidder is displayed on the target website.

For pricing, the advertisers are charged by the ad networks using one of the three popular ad pricing models, such as cost per mille (CPM), cost per click (CPC) and cost per user action (CPA). Each pricing model serves different business objectives. The CPM model is typically employed in brand awareness campaigns. The CPC and CPA

models are generally used for attracting traffic and sales respectively. In the CPM model, the advertisers pay for every 1000 ad impressions on the websites. In CPC and CPA respectively, the advertisers pay for every ad click and purchase made by end users on the websites. The ad revenues are then shared between the CSP and ad network based on their contract.

Chapter 3

Economics of Traffic Attraction by Transit ISPs

3.1 Introduction

Internet transit is a massive decentralized economy where thousands of ISPs sell and resell traffic-delivery services. Transit revenues depend on traffic rates: the more traffic the provider transits, the larger the revenue is. As more Internet traffic bypasses the transit services due to peering arrangements, the transit ISPs miss on earning transit revenues on this traffic. Data from TeleGeography [8] reveals that the share of traffic passing through the transit networks reduced from 47% in 2010 to 41% in 2014. Besides, the per-unit transit traffic prices also decline.

Traffic attraction refers to a family of BGP [28] techniques enabling an AS to receive traffic that would otherwise flow elsewhere. Because a transit provider receives payments for transit of customer traffic, the provider has a direct financial interest in attracting extra traffic of customers, e.g., when competing with another provider for transit traffic of a multihomed customer [22, 23]. While some of prior studies on traffic attraction [43–46] indicate that attraction of extra traffic provides economic benefits to transit providers, the prior work mostly focuses on security rather than traffic economics. The previous security-themed work examines various traffic-attraction techniques and argues that secure versions of BGP or configuring prefix-filtering policies do not neutralize them. There is also a body of related game-theoretic and simulation-based studies [23, 47, 48] that include economic considerations but analyze traffic attraction in small-scale artificial settings.

In this chapter, we study the economics of traffic attraction by transit providers to increase their transit revenues. We begin by analyzing a real traffic-attraction incident in the Internet. Motivated by the insights obtained from the analysis of this real incident, we conduct further studies, from economic and technical perspectives, of hypothetical

scenarios of traffic attraction by providers across the transit hierarchy.

3.2 Real case study of YouTube traffic attraction

In early 2008, YouTube owned AS 36561[1] with five assigned prefix spaces according to the RIPE RIS Dashboard [49]. 208.65.152.0/22 represented one such space and is the prefix that attracted a majority of YouTube-addressed traffic. On the 24th of February in 2008, AS 17557 belonging to PTCL (Pakistan Telecommunication Company Limited) hijacked YouTube traffic for approximately two hours and fourteen minutes by announcing the more specific prefix 208.65.153.0/24. The intention of the hijacking was to block access to YouTube within the state of Pakistan but the impact was significantly more far-reaching because PTCL announced 208.65.153.0/24 also to its provider PCCW Global (AS 3491), and the latter advertised globally the bogus PTCL paths for the longer prefix. Consequently, PTCL became a black hole that attracted and discarded packets sent to YouTube from all over the global Internet. YouTube detected the sharp decrease in its incoming traffic and reacted by announcing the even more specific prefix 208.65.153.0/25. The countermeasure restored some traffic flow to YouTube, yet PTCL remained able to attract a nontrivial fraction of YouTube-addressed traffic due to the path length and other factors that affect the routing policies of various ASes [50].

3.2.1 Methodology

The specific basis for our traffic attraction investigation is the above real incident of prefix hijacking. The incident attracted significant attention and was actively discussed by operators and researchers on NANOG (North American Network Operator Group) mailing list [51]. Whereas, the Internet incorporates BGP announcement monitoring systems such as PHAS (Prefix Hijack Alert System) [52], RIPE RIS (Réseaux IP Européens Routing Information Service) [49] and BGPmon (BGP monitoring and analyzer tool) [53], our analysis relies on actual announcement data collected by the monitoring systems during the hijacking incident.

The adopted research method supplements the real-data analysis with simulations for two reasons. First, the available real data do not paint the full picture of the traffic attraction incidents. Second, simulations enable us to examine suppositional scenarios that are more suitable for revenue-boosting traffic attraction. Our choice for the simulation platform is C-BGP [54], a widely used simulator for BGP routing problems. C-BGP determines AS-level paths for all traffic and bidirectional traffic rates for every inter-domain link in the simulated topology. The BGP routing policies [32] in C-BGP satisfy the valley-free routing conditions [55]. Although real Internet routing occasionally deviates

[1] Before Google restructured YouTube.

from such policies, valley-free routing constitutes a reasonable approximation because we are mostly interested in qualitative insights.

3.2.1.1 Topology

To initialize C-BGP with a realistic contemporary AS-level topology of the Internet, we use a data set collected by CAIDA (Cooperative Association for Internet Data Analysis) [56] at the time of the real hijacking incident. The AS-level Internet topology is as per the CAIDA data set dated 21 February 2008 and captures relationships between 27184 ASes. The CAIDA data set classifies relationships between a pair of ASes as customer-provider (encoded as -1 in the data set), provider-customer (encoded as 1), and peering (encoded as 0). We remove from each data set all customer-provider pairs because of their redundancy: for every provider-customer relationship provided to C-BGP, the simulator automatically configures a transit link associated with both provider-customer and corresponding customer-provider relationships. Since C-BGP does not recognize sibling relationships between ASes, we also substitute sibling relationships (encoded as 2) with peering relationships. The number of sibling relationships in the CAIDA data set is small, and the substitution has a negligible impact on the fidelity of the simulations.

3.2.1.2 Traffic

While YouTube-bound traffic contains video clips uploaded by YouTube users as well as requests for clip downloads, the uploads are likely to dominate the requests in terms of the traffic volume, and we focus only on this former type of traffic. According to [57], video was uploaded to YouTube in 2008 at the rate of 12 hours per minute, meaning that the volume of video clips uploaded every minute was such that playing them one after another would take 12 hours. After analyzing a collection of video clips in the FLV (FLash Video) format with playing times in the range from 1 minute to 1.3 hours, we estimate that 1 hour of playing time corresponds to 100 MB of data. Hence, our estimate for the average year-2008 rate of YouTube-addressed traffic is 160 Mbps.

Determining the origins of the YouTube-addressed traffic is a more challenging task. ISPs commonly perceive exchanged traffic volumes as sensitive information. While we are aware of anonymized data sets that quantify the relative potency of various ASes to generate traffic, the goal of our study necessitates associating a generated traffic volume with each specific AS. Without having access to real data sets of the latter type, we allocate the generated traffic to all BGP-connected ASes uniformly, i.e., each AS in our C-BGP simulations generates YouTube-addressed traffic at the same rate of 6 Kbps.

Our simulations rely on the aforementioned synthetic traffic to evaluate the impact of prefix hijacking on inter-provider links. With each AS in the Internet-scale topologies, we associate the traffic demand for the advertised prefix. Then, we utilize C-BGP to

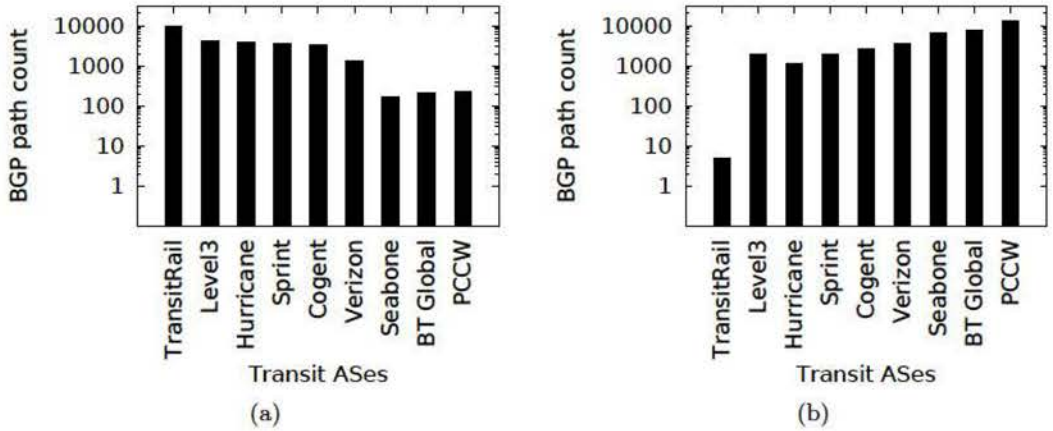

Figure 3.1: Number of BGP paths through transit ASes (a) before the YouTube's prefix hijacking by PTCL and (b) after PTCL hijacks all YouTube-bound traffic.

determine the rate of traffic flowing in both directions of each inter-ISP link. We refer to this traffic rate as a *link load* of the inter-provider link.

3.2.2 Impact of traffic attraction on Inter-AS paths and traffic

In accordance with our general simulation methodology from Section 3.2.1, we perform two simulation runs. In the first run, YouTube (AS 36561) announces its prefix 208.65.152.0/22. In the subsequent run, PTCL (AS 17557) additionally advertises its sham ownership of the more specific prefix 208.65.153.0/24 to hijack YouTube-addressed traffic. The C-BGP simulations reveal the complete success of the PTCL hijacking attempt: as a result of announcing the longer prefix, PTCL starts receiving all YouTube-addressed traffic with no continued delivery to YouTube itself. This simulation outcome is consistent with the historical accounts of the actual hijacking incident [50]. After each of the runs, C-BGP identifies exactly 100 ASes as being unable to reach any announced prefix. Hence, the number of BGP-connected ASes in the reported simulations stands at 27,084.

3.2.2.1 Connectivity of transit ASes to the advertised prefix

In the following, we examine the impact of the hijacking on the BGP connectivity of all transit providers to the announced prefix, i.e., we focus on the transit ASes which forward traffic from other ISPs to the advertising entity. Using C-BGP, we determine all converged BGP paths in the simulated scenario. Then, for each transit AS, we count the number of paths from other ASes through this transit AS to the advertising entity. Below, we interchangeably refer to this value as the number of served BGP paths or *BGP path count* of the transit AS. Figure 3.1 presents the number of BGP paths through the transit

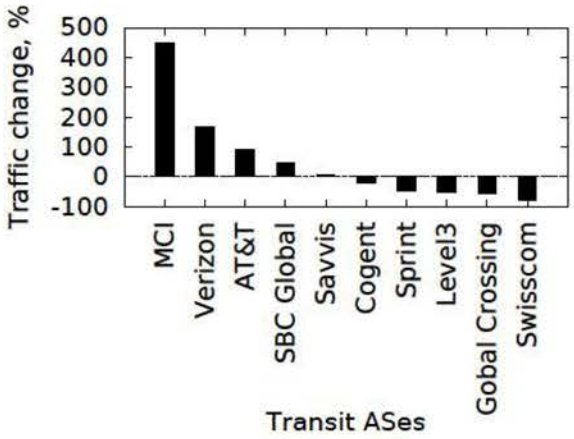

Figure 3.2: Percentage transit traffic changes for the top-5 winning and loosing transit ASes.

ASes, before and after the PTCL hijacks YouTube's prefix. The name of the organization that owned the transit AS number during the time of the prefix-hijacking incident in year 2008 is reported.

Before the prefix hijacking by PTCL, the converged routing involves 2,878 transit ASes. TransitRail (AS 11164) serves 9,733 paths and is the largest last-hop aggregator of YouTube-bound traffic. Level3 (AS 3356), Hurricane (AS 6939), Sprint (AS 1239), and Cogent (AS 174) are also connected directly to YouTube and constitute the next four biggest carriers of its incoming traffic with 3,983, 3,863, 3,487, and 3,395 served BGP paths respectively, as shown in figure 3.1(a). After PTCL hijacks all YouTube-bound traffic, the number of transit ASes decreases to 2,760. Figure 3.1(b) shows the number of BGP paths served through these large transit ASes to PTCL after the hijacking. The BGP path counts for the five giants TransitRail, Level3, Hurricane, Sprint, and Cogent shrink to 5, 1,915, 1,124, 1,854, and 2648 respectively. The preserved paths do not lead to YouTube anymore but instead contribute to the hijacking success of PTCL. The top-3 providers with the largest BGP path counts after the hijacking are PCCW Global (AS 3491), BT Global (AS 5400), and Telecom Italia Seabone (AS 6762): the hijacking boosts their BGP path counts from 219, 206, and 164 to 12,942, 7,793, and 6,319 respectively.

3.2.2.2 Traffic changes of the transit ASes

Whereas the BGP path counts highlight the transit potential of the ASes, the actual quantity of traffic passing through the transit ASes depends on the traffic engineering goals set by the operators based on their economic policies. Hence, this section utilizes the computed BGP paths to derive the link loads of all inter-AS links of each transit AS to compute the percentage of transit traffic changes arising because of the traffic attraction by PTCL. As shown in figure 3.2, the largest transit traffic increase of 450% is observed

for MCI (AS 702)[2], followed by Verizon (AS 701), AT&T (AS 7018), SBC (AS 7132)[3], and Savvis (AS 3561) at 170%, 91%, 48% and 7%, respectively. Among the top traffic losers, Swisscom (AS 3303) faces the largest 80% reduction, followed by Global Crossing (AS 3549)[4], Level 3 (AS 3356), Sprint (AS 1239) and Cogent (AS 174) losing 58%, 52%, 47% and 22% respectively.

3.3 Traffic attraction by large transit ISPs

While the above work – by using the real data to drive the simulations – offers interesting insights, it also identifies promising directions for further research of the topic. Our results suggest that large transit ISPs face higher traffic changes. Hence, they have the strongest financial incentives to attract traffic to increase transit revenues. This indicates that traffic attraction can create a fertile ground for tussles between ISPs. Next, to accurately translate the BGP path counts and inter-AS link load into transit revenues, it is important to account for Internet cross-traffic. As our experience shows, scalability of C-BGP simulations are serious obstacles to incorporate the cross-traffic in C-BGP.

In this section, we conduct in-depth studies from an economic perspective, where transit ASes try to boost their revenues by attracting transit traffic and then delivering it to the proper destinations. We strive to model the traffic matrices realistically to account for the cross-traffic and overcome C-BGP scalability challenges by enhancing the memory management in C-BGP. We also enable C-BGP to provide the attract-then-deliver feature for our economic analysis of traffic attraction by the transit ASes. We focus on the economics of customer-traffic attraction by providers across the transit hierarchy and report extensive C-BGP [54] simulations in an Internet-scale model configured with realistic data on inter-domain traffic, topology, and pricing. We consider attractors from the top 3 tiers of the transit hierarchy as well as 3 types of reactions by other ASes to the attraction: (1) filtering, i.e., discarding the BGP announcements that trigger the attraction, (2) disconnection by discontented customers, i.e., severance of their business relationships with the attractor altogether, and (3) attempts of discontented ASes to attract extra traffic to themselves. The broader scope and higher realism – combined with sensitivity studies – enable our work to offer deeper quantitative insights into traffic-attraction economics and reach reliable qualitative conclusions.

The specific BGP technique for traffic attraction in our study is prefix deaggregation. Although BGP is a sophisticated protocol, it does not provide an AS with a reliable mechanism to validate the path information announced by neighboring ASes. Due to the longest-prefix match rule of IP forwarding, the BGP announcement of a longer deaggregated prefix steers traffic to the announced path. Multihomed ASes routinely employ

[2] AS 702 is now operated by Verizon Business
[3] AS 7132 is now operated by AT&T Internet Services
[4] AS 3549 is now operated by Level 3 Communications

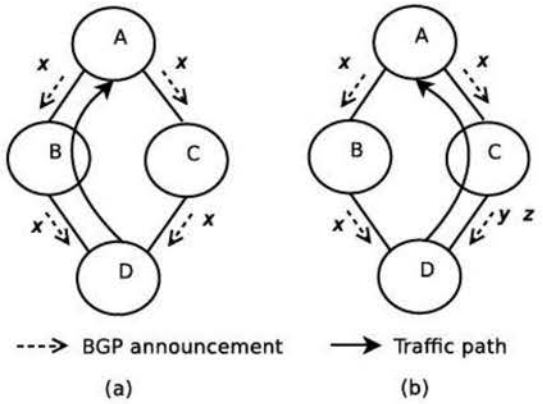

Figure 3.3: Example of traffic attraction via prefix deaggregation: (a) before attraction; (b) after attraction.

prefix deaggregation to balance their incoming traffic among their multiple connections to the Internet. Instead of the load balancing by multihomed ASes, our paper studies attraction of additional customer traffic. We consider the kind of prefix deaggregation where an intermediary AS learns a prefix from a customer, deaggregates the prefix, and announces all longer prefixes to each of its other customers. In particular, the traffic-attracting AS splits a learned prefix equally into 2 longer prefixes and announces both longer prefixes to the customers (Note that the AS deliberately does not announce the deaggregated prefixes to any of its peers so that none of its current traffic shifts from its customer links to its peering links).

To illustrate how prefix deaggregation enables traffic attraction, we consider a simple scenario in figure 3.3, where networks B and C directly learn prefix x from their mutual customer A. In figure 3.3(a), providers B and C propagate prefix x to their another mutual customer D, network D sends traffic to network A through provider B, and network B collects transit payments from both customers A and D. In figure 3.3(b), when provider C deaggregates prefix x into longer prefixes y and z, and announces these prefixes y and z to customer D, the traffic from network D to network A is attracted to flow through provider C rather than provider B; consequently, the traffic-delivery payments from both customers A and D go to provider C instead of provider B.

While the considered prefix-deaggregation method is easily implementable in practice, a transit AS can also attract extra traffic by employing a different BGP technique. For example, the intermediate AS can attract additional traffic by reducing the path length in a propagated BGP announcement. There is anecdotal evidence that prominent ASes attract extra customer traffic by rewriting the origin attribute in propagated BGP announcements [15,16]. These and other alternative techniques for traffic attraction represent an interesting topic for future studies on traffic economics.

For the role of traffic attractors, we select ASes throughout the Internet transit hierar-

chy. Whereas tier-1 networks are in the strongest position to attract significant amounts of extra traffic, our interactions with industry experts suggest that traffic attraction by tier-2 ASes is more common. In addition to traffic attraction by tier-1 and tier-2 transit providers, we also examine traffic attraction by tier-3 networks.

3.3.1 Model

Our modeling strives to abstract the highly complicated problem into a manageable representation for realistic study and discussion. Instead of focusing on a single setting, we parameterize our AS-level Internet model to experiment with realistic ranges of parameter settings.

3.3.1.1 Topology

Despite a decade of intensive research, the AS-level Internet topology is not known accurately, e.g., due to missed links [58]. To deal with this uncertainty, we consider 3 alternatives including topologies reported by CAIDA on 2012/6/1 and UCLA [59] on 2012/5/1. The third topology, to which we refer as UCLA+, is synthetic. We derive it from the UCLA topology by adding links that UCLA reported for at least 15 out of the 30 subsequent days. The enhancement contributes 1,656 peering and 75 transit links to the UCLA+ topology.

3.3.1.2 Traffic

Characterizing the inter-domain traffic is another notoriously hard problem. Our study considers 9 traffic matrices guided by empirical Internet data. The matrices generalize the measurement results suggesting that (a) the fraction of traffic originated by the largest content source was about 5% in 2009 and growing [60], and (b) the distribution of traffic from origin ASes is Zipf-like with the shape parameter between 0.9 and 1.1 [60,61]. We obtain the $9 = 3*3$ traffic matrices by combining 3 settings along each of the following 2 dimensions: (1) overall distribution function of origin traffic shares and (2) assignment of the traffic shares to specific ASes.

For the overall distribution of the origin traffic shares, we consider 3 instances of the Zipf-Mandelbrot function defined as $K/(r+v)^D$, where r refers to the rank of the AS, D denotes the shape parameter, and K and v are constants. The instance settings are as follows:

[Z1] $D = 1.1$, $K = 0.181$, and $v = 0.5$;

[Z2] $D = 0.9$, $K = 0.072$, and $v = 0$;

[Z3] $D = 0.7$, $K = 0.108$, and $v = 0$.

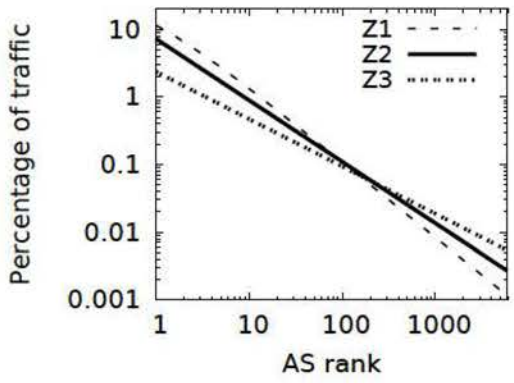

Figure 3.4: Distribution of origin traffic shares in our traffic matrices.

Figure 3.4 depicts the 3 overall traffic-share functions that cover the realistic ranges for the shape parameter (from 0.7 to 1.1) and largest traffic share (from 3% to 12%).

To assign the traffic shares to specific ASes, i.e., to set rank r of each AS, we consider the following 3 diverse options for the ranking metric:

[R1] peering coefficient, which is the number of peers divided by the total number of providers and customers of the AS; the peering coefficient tends to be the highest among small/intermediate ASes that provide access to content hosts (these ASes have many peers but few transit links);

[R2] IP address count, i.e., the number of IP addresses originated by the AS;

[R3] website visit frequency, i.e., the number of users who visit the websites hosted by the AS.

We use notation ZγRδ to denote the traffic matrix with overall traffic-share function Zγ and AS ranking metric Rδ.

In each of the 9 traffic matrices, we distribute the originated traffic of an AS among destinations (i.e., content-consuming ASes) in proportion to the IP address counts of the destination ASes. In addition to the above content traffic, our traffic matrices also incorporate peer-to-peer traffic, with the peer-to-peer traffic between a pair of ASes set proportionally to the IP address count of each AS in the pair.

3.3.1.3 Pricing

In comparison to traffic, pricing of inter-AS relationships is even more difficult to infer due to confidential bilateral agreements guarding them. Empirical evidence indicates that IP transit is subject to subadditive pricing where prices per Mbps are lower for higher traffic rates. In particular, the empirical data suggest as reasonable the following pricing function [62, 5, 4]: the monthly payment from the customer to the provider of transit

link l is $t_l = b \cdot L^m$ where L denotes the traffic rate in Mbps, $b = 45$, and $m = 0.75$ (the computed payment is in U.S. dollars). Our study generalizes this transit pricing function by considering the range of m values from 0.4 to 1, i.e., from heavy economies of scale to linear pricing which does not offer any price discount for higher traffic rates.

While the 95th-percentile billing [35,36] is common in IP transit charging, this method requires a large number of individual traffic-rate samples. Instead of modeling the traffic at a fine-grained time scale, we directly represent the traffic of each transit link as the 95th-percentile rate because: (a) fine-grained traffic statistics are not publicly available for most ASes, and (b) simulations for a detailed traffic model do not scale to a large population of ASes.

Peering relationships are usually free of financial settlements between the peers. On the other hand, maintaining a peering link does involve costs, e.g., payments to an IXP (Internet eXchange Point) [58] that provides the physical infrastructure for peering. To represent the peer's cost of maintaining its peering link l, we adopt peering-cost function $p_l = 20 \cdot L^{0.4}$ suggested in prior work [62, 5, 4]. In comparison to the transit-link payment for the same traffic rate, the peering cost is always lower but not negligible.

To determine the overall traffic-delivery payment for each AS, we partition all external links of the AS into 3 sets: set $\mathbb{V}$ contains the transit links where the AS is a provider, set $\mathbb{U}$ is for the transit links where the AS acts as a customer, and set $\mathbb{G}$ includes all its peering links. Then, the monthly traffic-delivery payment of the AS is $P = \sum_{l \in \mathbb{V}} t_l - \sum_{l \in \mathbb{U}} t_l - \sum_{l \in \mathbb{G}} p_l$ with positive values denoting traffic-delivery revenues, and negative values representing traffic-delivery expenses.

3.3.2 Evaluation methodology

To evaluate traffic attraction and countermeasures by other ASes, we conduct large-scale simulations in C-BGP [54]. We optimize C-BGP to overcome its scalability limitations. For example, even if each AS announces a single prefix only, the standard C-BGP can exhaust a relatively large physical memory before computing stable AS-level paths. We improve the memory management of C-BGP to scale up the simulations to at least 6,000 prefix announcements. We integrated the memory-scalability feature into C-BGP and supplied the modified version to the C-BGP code repository [63].

While our optimizations alleviate – but not eliminate – the C-BGP scalability limitations, we conduct the simulations by focusing on the core of each topology and representative prefixes of the ASes. We extract the topological core by excluding all the stub ASes and all their links. For the 3 examined topologies, their cores contain around 6,000 transit ASes. We determine traffic-delivery payments for all these transit ASes. Although the simulations do not directly consider the stub ASes outside the topological core, we account for the traffic of these ASes when computing the traffic-delivery payments for the

transit ASes. While real ASes generally own and announce several prefixes, the C-BGP scalability limitations prevent announcing multiple prefixes from each AS. On the other hand, a single prefix per AS is sufficient for C-BGP to simulate all inter-domain communications in the topological core. Therefore, we associate the total inter-domain traffic of each AS with a single representative prefix of this AS.

To configure our traffic matrices, we rely on real data. Based on the Cisco-VNI statistics [64], we set the total rates of the inter-domain content and peer-to-peer traffic to 45 Tbps and 15 Tbps respectively, with 4 times more content traffic flowing from servers to clients than from clients to servers. We specify the IP address counts of the ASes according to the CIDR report [65]. The peering-coefficient ranks of ASes are in accordance with the CAIDA topology. We calculate the website visit frequencies based on the data provided by the Alexa Web Information Service from Amazon for top 100,000 websites [66].

Among the 6,000+ transit ASes in the topological cores, we select 30 ASes to act as traffic attractors. Each of tiers 1, 2, and 3 in the transit hierarchy contributes exactly 10 ASes (those with the largest numbers of customers) to the attractor set. We refer to the 10-AS groups as T1, T2, and T3 respectively. To denote AS β from tier α, we use notation T$\alpha\beta$.

Our main metric is payment change. It measures the relative change in the traffic-delivery payment of the AS in comparison to the baseline scenario where no AS tries to attract additional traffic.

Unless explicitly stated otherwise, we report the results for the following default settings. The topology is from CAIDA. The results are averaged over the 9 traffic matrices. Transit-pricing exponent m is set to 0.75. When deaggregating prefixes to attract traffic, the attracting AS deaggregates prefixes announced by its 100 largest customers. We also study sensitivity of the results to the topology, traffic, and pricing.

3.3.3 Evaluation results

3.3.3.1 Attraction by a single AS

We start by examining what happens when a single AS attempts to attract traffic. We repeat this experiment for the 30 attractors with each of the 9 traffic matrices and record the payment change for the attracting AS. Using box plots, figure 3.5 presents the results arranged according to the tier of the attracting AS. The plots demonstrate that transit ASes have significant financial incentives to attract traffic: the median payment change is 148%, 38%, and 21% for T1, T2, and T3 respectively. The tier-1 networks are in the strongest position to benefit from traffic attraction because they sell transit to numerous customers but do not buy transit themselves. Being less central in the transit hierarchy, all considered tier-2 ASes are still able to raise their revenues by attracting

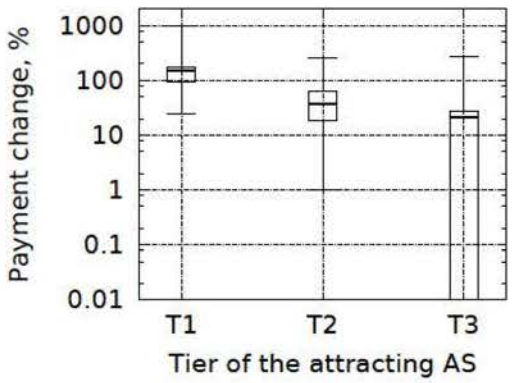

Figure 3.5: Payment change for the attracting AS.

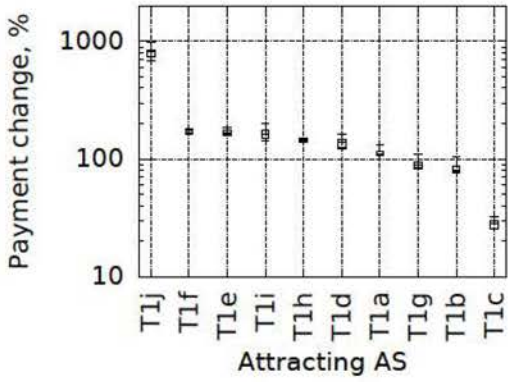

Figure 3.6: Payment changes for the tier-1 attractors.

traffic. While some tier-3 networks also benefit substantially from the traffic attraction, it is not beneficial for other attracting ASes from T3.

Focusing on the experiments with the tier-1 attractors, figure 3.6 shows that while T1j and T1c are 2 opposite extremes with the median payment changes of 830% and 28% respectively, the payment changes for T1b and the 7 other ASes from T1 are rather similar to each other. The large gap between the payment changes for T1j and T1c is mostly due to the different sizes of these ASes. In comparison to T1j, T1c serves more transit traffic and attracts a larger amount of extra traffic in absolute terms. Nevertheless in relative terms, the payment gain is much higher for the smaller T1j.

3.3.3.2 Winners, losers, and neutrals

By redistributing traffic in the AS-level topology, traffic attraction by an AS affects traffic-delivery payments for other ASes. We refer to the ASes with increased traffic-delivery payments as winners, and to the ASes with negative payment changes as losers.

Figure 3.7: Payment changes for the winning and losing ASes when the attracting AS is from tier 1, 2, or 3.

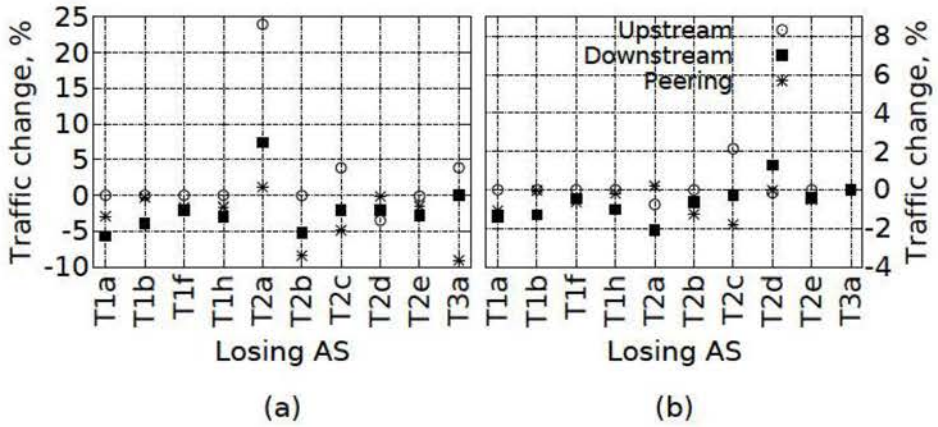

Figure 3.8: Changes in the upstream, downstream, and peering traffic of the 10 largest losers when the attractor is a tier-1 AS: (a) without filtering; (b) with filtering.

Neutrals are the ASes with unchanged traffic-delivery payments. Figure 3.7 plots the payment changes for the winning and losing ASes in the experiments of section 3.3.3.1. The traffic attraction by tier-1 networks makes the most divisive impact on the payments: the fractions of winners, losers, and neutrals are 17%, 26%, and 57% respectively. When the attracting AS is from T2, the fractions of winners and losers decrease to 11% and 8% respectively. After the traffic attraction by tier-3 networks, the impact is highly local: either winners or losers comprise only 1% of the AS population.

3.3.3.3 Impact on traffic

To understand where the attracted traffic comes from, we classify the inter-domain traffic of an AS into 3 types: (1) upstream, i.e., traffic from the AS to its providers,

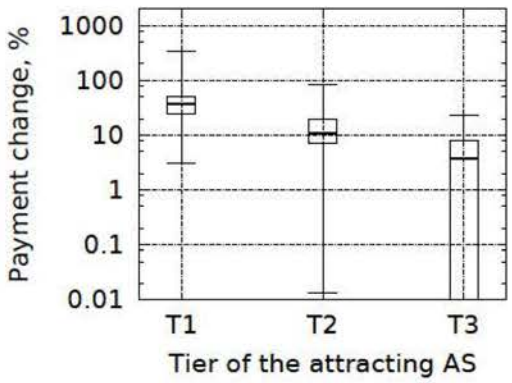

Figure 3.9: Payment change for the attracting AS when all losers do the filtering.

(2) downstream, i.e., traffic from the AS to its customers, and (3) peering, i.e., traffic on the peering links of the AS. Then, we consider the 10 ASes that suffer the largest declines (in absolute terms) of their traffic-delivery payments after the traffic attraction by the T1 networks. These largest losers are 4 tier-1, 5 tier-2, and 1 tier-3 ASes. Figure 3.8a depicts changes in the upstream, downstream, and peering traffic of these 10 losers after the attraction. While the losing tier-1 ASes have no upstream traffic either with or without the attraction, their downstream traffic decreases, suggesting that the lost downstream traffic is acquired by the attractor. The attraction makes a qualitatively similar impact on the traffic of T2b and T2e. For the 3 other largest losers from the lower tiers, the upstream traffic increases due to sending more traffic to the tier-1 attractor, and the downstream traffic increases to a smaller extent (e.g., for T2a) or even decreases (e.g., for T2c). From an overall perspective, the traffic attraction reduces peering traffic and pulls extra traffic up the transit hierarchy.

3.3.3.4 Filtering by losing ASes

To analyze responses of other ASes to the traffic attraction, we first consider filtering, i.e., discarding the deaggregated prefixes announced by the attractor. Figure 3.9 presents the payment change for the attracting AS when all losing ASes from the experiments in section 3.3.3.1 do the filtering. Comparing the results in figures 3.5 and 3.9, we see that the filtering reduces but does not remove the financial benefits for the traffic attractor. With the filtering, the median payment change for the attracting AS is 37%, 11%, and 4% for T1, T2, and T3 respectively.

Figure 3.8b offers insights into the inability of the filtering to negate the attraction. Although the filtering can help a losing AS – e.g., T2a – to reduce its upstream traffic to the attractor, the filtering by the loser does not prevent its customers from their switching to alternative paths via the attractor.

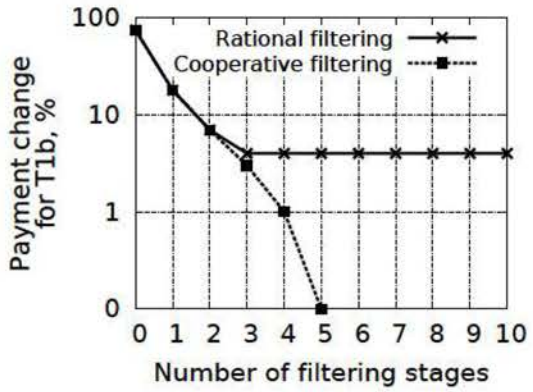

Figure 3.10: Payment change for attractor T1b with multi-stage filtering.

3.3.3.5 Multi-stage rational filtering

The filtering examined in section 3.3.3.4 redistributes traffic in the AS-level topology and creates additional losers. In this section, we extend the above filtering scenario into a multi-stage reaction where the group of filtering ASes on each stage expands by incorporating the additional losers from the previous stage. We refer to the filtering by losers as rational because it is done only by the ASes that financially suffer from the traffic redistribution. To assess the effectiveness of multi-stage rational filtering, we consider the setting where stage 0 corresponds to the traffic attraction by T1b without filtering. Figure 3.10 shows that while the payment change for attractor T1b decreases on stages 1 through 3, the payment gain for the attractor stabilizes at 4% after stage 3, which yields no additional losers. Hence, even if the filtering is done by all losing ASes, the multi-stage rational filtering does not eliminate financial incentives for the traffic attraction. The use and propagation of the deaggregated prefixes by the winners and neutrals allow the attractor to increase its traffic-delivery revenues.

3.3.3.6 Cooperative filtering

While section 3.3.3.5 demonstrates the inability of rational filtering to negate the financial benefits of the attractor, we now explore what happens if the other winning ASes go against their own financial interests and also react by filtering the deaggregated prefixes. Referring to such filtering as cooperative, we consider the multi-stage cooperative filtering where stage 0 corresponds to the traffic attraction by T1b without filtering, stage 1 involves filtering by all losers and neutrals, and each subsequent stage expands the group of the filtering ASes with additional winners (selected in the increasing order of their payment gain). Figure 3.10 shows that the cooperative filtering negates the payment gain of attractor T1b on stage 5 where the filtering is done by all customers of the attractor.

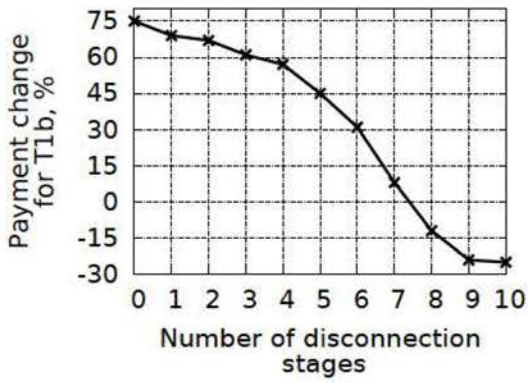

Figure 3.11: Payment change for attractor T1b with multi-stage disconnection by losing customers of the attractor.

3.3.3.7 Disconnection

The evaluation in sections 3.3.3.4 through 3.3.3.6 shows that filtering is not an effective countermeasure unless the winners resist the traffic attraction against their own financial interests. Now, we examine a more severe reaction by losers where losing customers sever their business relationships with the attracting AS altogether. Again, we consider a multi-stage version of the response where stage 0 corresponds to the traffic attraction by T1b without filtering (and without disconnection). On stage 1 of the disconnection response, the attractor is disconnected from the 1% of its losing stage-0 customers that are selected in the decreasing order of their absolute losses. On each of stages 2 through 7, the cumulative number of the disconnected customers of T1b doubles. On stage 7 where the attractor is disconnected from 45% of all its original customers (i.e., 64% of its losing stage-0 customers), the attractor still has the payment gain of 8%. On each of the subsequent stages, we disconnect all remaining T1b losing customers from the previous stage. Both stages 8 and 9 create additional losers among the connected T1b customers. On stage 10 where the attractor is disconnected from 85% of all its original customers, no new losers emerge, and the payment change of T1b stabilizes. The remaining 15% of all original T1b customers are either winners or neutrals on stage 10 and hence do not disconnect from the attractor. Figure 3.11 depicts the dynamics of the multi-stage disconnection response by the losing customers. The results demonstrate that the disconnection by losing customers is ineffective unless a large portion of them terminate their business relationships with the attractor.

3.3.3.8 Attraction by multiple ASes

The previous sections show that neither filtering nor disconnection eliminates the financial incentives for traffic attraction unless participation in the response is broad. Now,

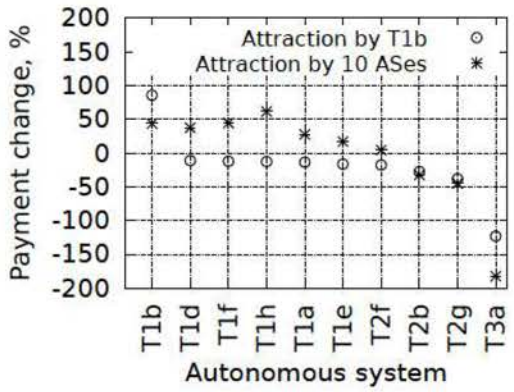

Figure 3.12: Payment changes for 10 ASes when the traffic attraction is done by T1b only vs. all 10 ASes.

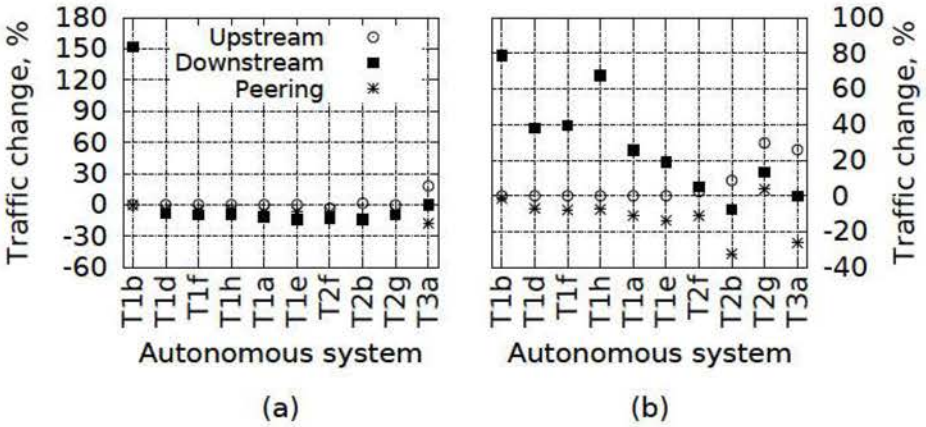

Figure 3.13: Changes in the upstream, downstream, and peering traffic of 10 ASes when the attraction is done by: (a) T1b only and (b) all 10 ASes.

we consider a different reaction where a losing AS defends its traffic-delivery payment by attracting extra traffic to itself. Specifically, we consider the scenario where in response to the traffic attraction by T1b the 9 largest losers (in absolute terms) try to attract traffic as well. The expanded set of 10 attractors includes 6 tier-1, 3 tier-2, and 1 tier-3 ASes. Figure 3.12 shows the payment changes for these 10 ASes when the traffic attraction is done by T1b only vs. all 10 ASes. When all 10 ASes try to attract traffic, all 6 tier-1 ASes and T2f gain from the traffic attraction but the payment changes for T2b, T2g, and T3a are negative. The results confirm our earlier observation that tier-1 networks are in the strongest position to benefit from traffic attraction. Furthermore, figure 3.12 demonstrates that ASes from lower tiers are not assured to gain from traffic attraction when multiple networks attempt to attract extra traffic.

To understand why the transit ASes from different tiers fare differently when multiple

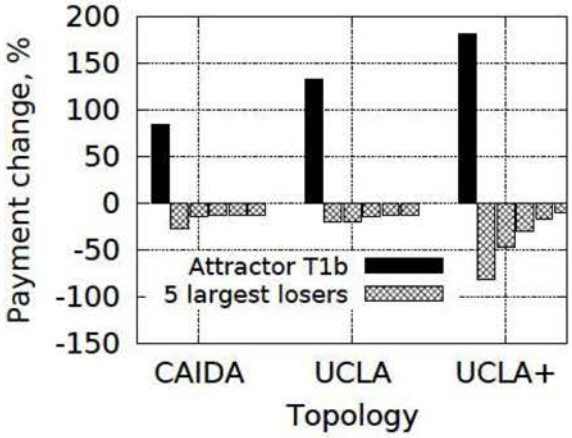

Figure 3.14: Sensitivity of payment changes to the AS-level Internet topology.

networks try to attract traffic, figure 3.13 plots the changes in the upstream, downstream, and peering traffic of the 10 ASes. When T1b acts as the only attractor, it greatly increases its own downstream traffic and decreases the downstream traffic for the other ASes except T3a that suffers the loss due to an increase in the upstream traffic. When all 10 ASes try to attract traffic, all 6 tier-1 networks (which never have any upstream traffic) win by increasing their downstream traffic. For T2f, the downstream traffic grows as well, and this growth outweighs the increase in its upstream traffic. On the other hand, T2b, T2g and T3a remain losers because their upstream traffic grows significantly while their downstream traffic increases less (if at all). Some attractors are more powerful than others.

3.3.3.9 Sensitivity to the topology

To study how sensitive our model is to its assumptions, we first examine its sensitivity to the AS-level Internet topology in the scenario where only T1b attracts traffic. In addition to the CAIDA topology, we also consider the UCLA and UCLA+ topologies. Figure 3.14 presents the payment changes for the attractor (which is the largest winner) and 5 largest losers. In quantitative terms, the topology has a substantial impact, e.g., the payment change for the attracting AS varies from 85% to 180% with the CAIDA and UCLA+ topologies respectively. On the losing side, the maximum loss by an AS varies from -9% to -81% with UCLA and UCLA+ topologies respectively. This latter result highlights the importance of the topology in general and peering links in particular for economic outcomes of traffic attraction. By adding the relatively small numbers of 75 transit and 1,656 peering links, the enhancement of the topology from UCLA to UCLA+ makes the traffic attraction significantly more powerful. The distributions of payment changes for all ASes are qualitatively the same in the 3 examined topologies.

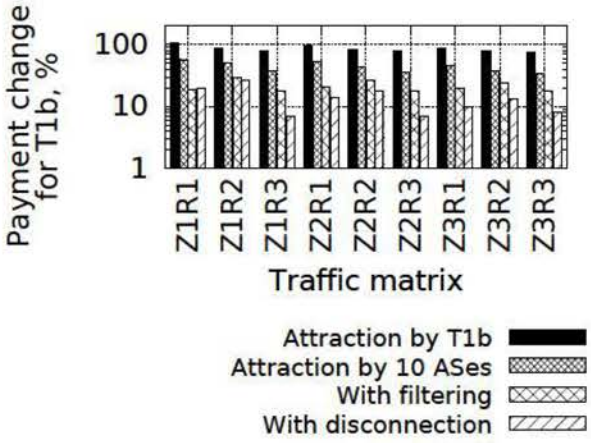

Figure 3.15: Sensitivity of the attractor's payment change to the traffic matrix.

3.3.3.10 Sensitivity to the traffic matrix

Expanding the above sensitivity study, we now examine the role of the traffic matrix. For each of our 9 traffic matrices, figure 3.15 plots the payment change for attractor T1b in 4 scenarios: (1) only T1b attracts traffic, (2) attraction is by 10 networks as in section 3.3.3.8, (3) filtering is done by all losing networks as in section 3.3.3.4, and (4) 50% of the T1b losing customers disconnect from the attractor. The 9 considered traffic matrices are diverse with respect to both overall distribution of origin traffic shares and assignment of the traffic shares to specific ASes. Despite this diversity, the results are qualitatively the same and relatively stable: the payment change for T1b varies from 75% to 105% when only T1b attracts traffic, from 34% to 57% with the 10 attractors, from 18% to 27% with the filtering, and from 7% to 26% with the disconnection. The quantitative outcomes of the traffic attraction are less sensitive to the traffic matrix than to the topology.

3.3.3.11 Sensitivity to pricing

To assess the sensitivity of our model to pricing, we consider the same 4 scenarios as in section 3.3.3.10 and reduce transit-pricing exponent m from 1 to 0.9, 0.8, 0.75, 07, 06, 05, and finally 0.4. Figure 3.16 tracks the payment change for attractor T1b and exhibits substantial quantitative variations in the outcomes. In the first 3 scenarios, the payment change for T1b remains positive but decreases greatly: from 249% to 25% when T1b is the only attractor, from 126% to 14% when the 10 ASes attract traffic, and from 42% to 4% with the filtering. With the disconnection by 50% of the losing customers, the payment change not only decreases as the transit-pricing exponent is reduced but also becomes negative: the payment change for T1b is 169% for $m = 1$ and -27% for $m = 0.4$. The

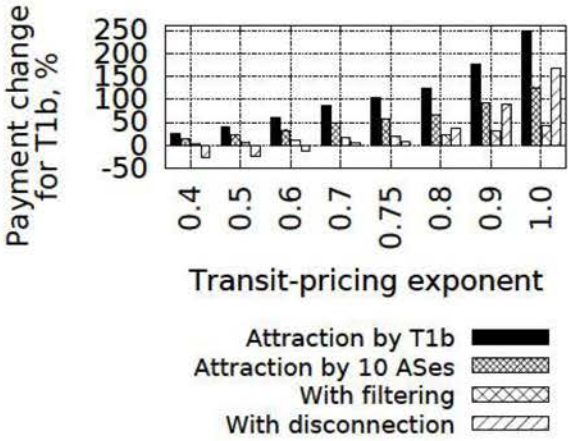

Figure 3.16: Sensitivity of the attractor's payment change to transit pricing.

payment change becomes 0 when the transit-pricing exponent is around 0.65. According to Telegeography data, the transit-pricing exponent is currently 0.8, 0.73, 0.69, and 0.65 for Oceania, Asia and South America, Europe, and North America respectively. Thus, our results suggest that at least a half of the losing customers needs to disconnect from the attractor to eliminate the benefits of the attraction under the current pricing.

3.3.3.12 Sensitivity to attraction intensity

While sections 3.3.3.5 through 3.3.3.7 explore how intensive the filtering and disconnection should be to negate the benefits of traffic attraction, we now examine the sensitivity of the attractor's gain to the intensity of attraction. In these experiments where only T1b attracts traffic (as in section 3.3.3.1), T1b changes its attraction intensity by deaggregating a different number of prefixes announced by its largest customers. Specifically, the attractor deaggregates 1, 10, 100, or 1,000 prefixes. Figure 3.17 presents the payment change for T1b. The median payment change is 7%, 30%, 78%, and 108% for 1, 10, 100, and 1000 deaggregated prefixes respectively. The marginal utility of the attraction intensity diminishes quickly. Even by deaggregating a relatively small number of prefixes (e.g., 100 as in the default setting of our studies) the attractor obtains most of its maximum possible gain.

3.3.3.13 Impact on path lengths

Finally, we evaluate how the traffic attraction and countermeasures affect the lengths of AS-level paths. Figure 3.18 plots the distributions of the AS-level path lengths for 5 scenarios: (1) baseline without any attraction; (2) attraction by T1b only, (3) attraction by 10 networks as in section 3.3.3.8, (4) with stabilized multi-stage rational filtering by all losing networks as in section 3.3.3.5, and (5) with stage-7 disconnection by losing

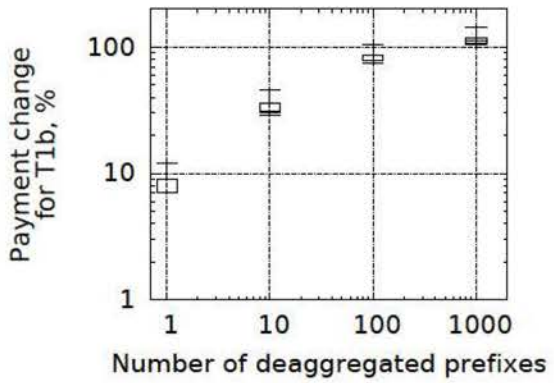

Figure 3.17: Sensitivity of the attractor's payment change to attraction intensity.

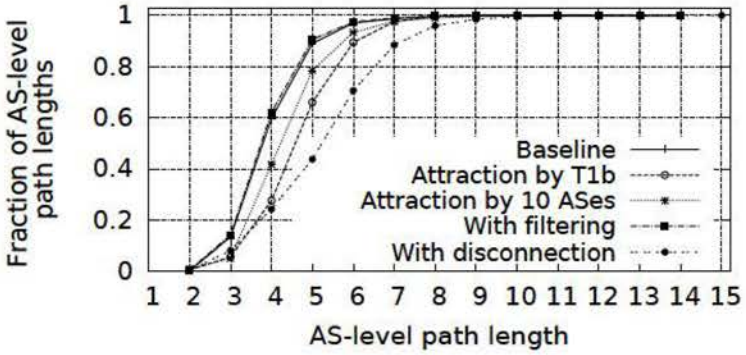

Figure 3.18: Distributions of AS-level path lengths.

customers of T1b as in section 3.3.3.7. The path-length distributions are mostly similar to each other. The average path length is 4.4 hops in the baseline scenario. When only T1b attracts traffic, paths elongate, with their average length increasing to 5.1 hops. The attraction by the 10 ASes reduces this elongation, with the average path length becoming 4.8 hops. The filtering almost restores the path-length distribution of the baseline scenario and yields the average path length of 4.5 hops. The disconnection has an opposite effect, with the average path length increasing to 5.7 hops. Overall, the elongation of paths under traffic attraction is not significant to deter this revenue-increasing behavior.

3.3.3.14 Impact on router complexity

While the traffic attraction via prefix deaggregation can boost the transit revenue of the transit ASes, the deaggregation also increases their router memory requirements. This section analyzes the extra router-memory costs imposed by the traffic attraction.

To assess the memory costs, one needs to understand how IP routers store prefixes.

Typically, a BGP-speaking router has BGP RIBs (Adj-RIB-In, Loc-RIB and Adj-RIB-Out) and Global RIB (all-protocol routing table) [67, 68]. The router inserts the best paths indicated by Loc-RIB into Global RIB. Global RIB also includes static paths as well as paths learned from intra-domain routing protocols. For actual forwarding of IP datagrams, the router uses FIB (Forwarding Information Base) derived from Global RIB. The router stores BGP RIB in RP (Route Processor card) and FIB in LC (Line Card).

The examined traffic attraction strategy requires storing extra paths per each deaggregated prefix in RP and LC. Let q be the number of the more-specific announcements, ϕ and φ denote the total number of BGP paths that RP and LC respectively can accommodate, ρ and ϱ represent the price of RP and LP respectively, and Y refer to the lifetime of the cards. Then, amortized monthly cost R of storing the extra paths in the router is

$$R = \frac{q}{Y}\left(\frac{\rho}{\phi} + \frac{\varrho}{\varphi}\right). \tag{3.1}$$

To quantify the extra router-memory cost R imposed by the prefix deaggregation, we use $\phi = 1.5 \cdot 10^7$ and $\varphi = 7 \cdot 10^5$ according to [69], $\rho =$ \$32K and $\varrho =$ \$480K (price of RP and LC for Juniper T1600 routers [70]), and $Y = 60$ months (average router lifetime [71]). While the largest number of more-specific announcements in our experiments is $q = 2{,}000$ (i.e., 2 more-specific announcements per each of the 1,000 deaggregated prefixes), the corresponding extra memory cost for each BGP-speaking router amounts to only \$23 per month. This cost is tiny in comparison to the additional traffic-delivery revenues offered by the traffic attraction.

While the above analysis shows that the increased memory needs are insignificant as an economic impediment for the traffic attraction, the prefix deaggregation also increases the processing load on the routers. Because a transit AS can attract substantial additional traffic by deaggregating a relatively small number of prefixes, the increased processing is unlikely to be an effective deterrent against the traffic attraction either. We will analyze the router processing costs in more detail in our future work.

3.4 Attraction viability

The viability of traffic attraction via prefix deaggregation is a multidimensional issue. Without pretending to be comprehensive, this section examines the issue from technical, legal, and business angles.

Technically, it is easy for a transit ASes to deaggregate prefixes. Other ASes can detect the deaggregation with tools that monitor BGP prefix announcements and IP datagram flows. While even the detection is not straightforward, the owner of the deaggregated prefix does not have effective technical means for stopping the traffic attraction, as our simulations show in the scenarios with prefix filtering and counter-announcing the deag-

gregated prefixes. While the owners of deaggregated prefixes are in the most legitimate position to complain, the origin ASes are not the parties most afflicted by the traffic attraction. Besides, because the origin ASes conduct prefix deaggregation themselves for traffic engineering [72], the argument that intermediary ASes may not do the same for the same reason becomes weaker.

The afflicted parties can try to neutralize the traffic attraction via prefix deaggregation through litigation. The global nature of the Internet complicates the judicial process for prefix-deaggregation cases. There seem to be no solid legal grounds for objecting against the studied deaggregation-based form of traffic attraction. In particular, the deaggregator and prefix owner can operate under different national legal systems. While the body of laws governing the Internet is generally slim but growing, we are not aware of any precedents of prefix-deaggregation litigation. Prior laws and guidelines issued by various governments for other Internet-related disputes demonstrate that the outcomes of the legal battles are highly unpredictable. It is not clear whether the contested actions deviating from an expected Internet behavior will be ruled illegal or legitimate, or even deserving a special protection by the law.

The business side of the Internet is likely to serve as an important sphere for settling the traffic attraction tussles. In order to offer the universal Internet connectivity to own customers, any AS anywhere in the routing hierarchy has to maintain business relationships with other ASes. While different ASes have clearly different negotiation power, losing a customer or peer is rarely a desirable outcome even for a huge transit ISPs. In the business world, reputations are tangible assets: a bad reputation can severely diminish the ability of the ISP to negotiate transit and peering contracts. Hence, if the ISP community as a whole starts to deem deaggregation-based traffic attraction unacceptable, the risk of a bad reputation can serve as a strong disincentive for an ISP to boost revenues through deaggregation-based traffic attraction.

3.5 Related work

Unlike our economic investigation of traffic attraction, prior studies of the subject approach it mostly from security perspectives [43–45]. Ballani et al. [43] explore the ability of an AS to attract traffic to itself for either discarding the attracted traffic or delivering the traffic to the destination. The paper considers 2 different attraction techniques where the attracting AS announces an invalid path to a prefix: by spuriously claiming to be either the prefix owner or an intermediary on the invalid path. [43] estimates the feasibility of such traffic attractions but does not study their economic impact. Nordstrom and Dovrolis [45] explore various attraction techniques and 2 countermeasures: filtering and adoption of S-BGP, a secure version of BGP. Again with an exclusive focus on security rather than traffic economics, [45] concludes that filtering is ineffective and that S-BGP

is too heavy to get deployed. Goldberg, Schapira, et al. [44] study robustness of S-BGP and other secure routing protocols to traffic attraction. [44] argues that the secure routing protocols fail to neutralize traffic attraction and need to be supplemented with defensive filtering.

Few other prior studies driven by security considerations aim at detecting the prefix hijacking. Lad, Massey, et al. [52] develop PHAS (Prefix Hijack Alert System), an online system that notifies the prefix owner when the BGP path to the prefix changes. McArthur and Guirguis [73] explore stealthy forms of prefix hijacking that attract small amounts of traffic and thereby avoid detection. Zhang, Zhao, and Wu [74] investigate an attack where an AS selectively drops BGP announcements to severely disrupt the routing.

There is also related work that includes economic considerations [23, 48, 46, 47, 75]. Gill et al. [46] analyze a game where the financial benefits of traffic attraction serve as incentives to deploy S-BGP. [46] does not evaluate the economic incentives but simply use them as a basis for the analyzed S-BGP deployment scenario. Lutu et al. [75] study whether traffic engineering via prefix deaggregation can reduce the transit expenses of prefix owners. In contrast, our work evaluates traffic attraction by intermediary ASes that seek to increase their transit revenues. We consider multiple attractors from different transit hierarchy and evaluate various countermeasures, such as multi-stage prefix filtering and link disconnection. Goldberg, Halevi, et al. [23], Levin et al. [47], and Kalogiros et al. [48] present game-theoretic investigations of economic incentives in inter-domain routing. While these papers provide thorough analyses for small-scale artificial settings, we conduct Internet-wide simulations driven by realistic data on inter-domain traffic, topology, and pricing.

The broader scope and higher realism, combined with the sensitivity studies, enable our work to yield new insights. For example, while [43] refuses to consider deaggregation-based attraction because the attractor is presumably unable to deliver the attracted traffic to the destination, our study demonstrates that the deaggregation-based attraction and delivery are not only feasible but also highly beneficial for the transit revenues of the attractor.

3.6 Summary

Relying on the extensive modeling and C-BGP simulations, this chapter presents an economic perspective on traffic attraction and countermeasures. Motivated by the insights obtained from the simulation and analysis of a real traffic-attraction incident of YouTube's prefix hijacking by Pakistan Telecom, we conducted an in-depth study of various hypothetical scenarios to understand economic and technical aspects of customer-traffic attraction by transit providers.

Our work shows that attraction and reaction to it redistribute traffic in the AS-level

topology and create numerous winners and losers in the AS population. The results demonstrate that tier-1, tier-2, and tier-3 ASes have significant financial incentives to attract traffic. In comparison to ASes from the lower tiers, the tier-1 ASes are in a stronger position to benefit from traffic attraction with respect to: (a) the degree of the attainable gain, (b) impact on other networks, and (c) preserving their own gain when multiple ASes attract traffic. The traffic attraction provides the financial gains by pulling extra traffic from peering links up the transit hierarchy.

The traffic attraction remains effective despite countermeasures unless the participation by ASes is very broad. Rational filtering does not remove the attraction incentives even when all losing ASes do the filtering. Only if winning ASes go against their own financial interests and join the filtering, such cooperative filtering eliminates the financial benefits of the attractor. The disconnection by losing customers is ineffective unless a large portion of them terminate their business relationships with the attractor. The increased router complexity and elongation of paths are too insignificant to be strong deterrents against the traffic attraction. Our studies of the model sensitivity to the topology, traffic, and pricing show that the quantitative outcomes are less sensitive to the traffic matrix than to the topology and pricing; qualitatively, the results remain consistent.

Our work strives to foster a discussion on customer-traffic attraction for revenue gain. While the financial benefits for the attractors are substantial, the cooperative filtering or disconnections can negate these benefits only if a very large number of ASes, including winners participate in the counter-responses. While we do not advocate (or oppose) traffic attraction, our results raise the possibility that the increasing financial pressure on IP transit business might prompt transit providers to attract traffic. Also, while the wide scope and sensitivity analysis make our results fairly generic, no simulation study can cover the full set of potential behaviors.

Chapter 4

Dissecting the Online Content Hosting Ecosystem

4.1 Introduction

While the previous chapter examined attraction of customer traffic as a means to boost revenues of transit providers, this chapter looks at hosting as an alternate source of revenues for transit providers. Usually, transit ISPs specialize in transit services. But due to the declining profit margins in the transit sector, transit ISPs start to offer hosting and content delivery services to earn extra revenues needed to sustain the ISPs' network infrastructures. The entrance of transit ISPs into the hosting sector changes the way the Internet delivers online contents, i.e., from a remote content origination to a localized distribution from ISPs throughout the transit hierarchy [76]. Hence, this new evolving trend in hosting is important for understanding the content traffic distribution across the Internet and geography. The realistic insights from such understanding are valuable for modeling the Internet-scale content traffic matrix more realistically by including transit ASes, which were earlier ignored as content traffic sources.

The online content sector is richly financed by online ad revenues. In 2015, the online ad revenues in the USA totaled up to $59.6 billion (US dollars) [77]. The online contents rapidly evolved from light-weight text files to dense rich-media and videos. Hence, the performance levels needed to deliver contents also significantly increased. To minimize latencies, ad networks and CSPs also widely employ CDNs to cache contents closer to end users. To cache contents across the Internet, CDNs use third-party network infrastructures, e.g., transit ASes. Since transit providers operate global high-speed network infrastructures managed under one or more ASes with geographically distributed points of presence (PoPs), the CDNs optimally cache the bandwidth-intensive contents across geographically distributed PoPs of the transit ASes [76]. In return, the transit providers earn revenues by using their network infrastructures for hosting online contents, and the

CDNs and also CSPs and ad networks benefit by saving on the large capital expenditures for setting up their own global network infrastructures to deliver contents. Besides, the CDNs also cache online contents on the access ISPs, located at the edge of the Internet. Such caching, also benefits the access ISPs by reducing their transit costs and CDNs by supporting higher content throughput.

To explore global content hosting, we use a novel measurement approach that leverages a VPN to collect real online contents from top 2,165 websites across 52 countries. Next, using a network of around 22,000 open recursive DNS servers spread across 172 countries and 8,500 ASes, we discover a vast ecosystem of hosting ASes and CDNs across the Internet hierarchy and geography. Since online advertisements have been the driving force behind the immense growth in the online content sector, our study pays special attention to online ads and classify online contents into ads and regular contents. We also study qualitative and quantitative differences in hosting characteristics, such as IP address deployment, content distribution across Internet, and delivery performance between the two types of contents.

Our analyses of the real measurement data reveal that online contents are widely distributed across the Internet transit hierarchy and geography. We observe several clusters of ASes with similar hosting characteristics. While ads are distributed across a large number of smaller AS clusters, regular contents are distributed across a small number of bigger AS clusters. These results show trends of replicating ads locally and regular contents globally. We also observe that the AS clusters have different mixes of ASes from different tiers, e.g., with more core ASes or edge ASes. AS clusters in their position of ASes in the transit hierarchy also reflect on their IP address usage, i.e., the IP usage is the highest for the AS cluster containing the highest number of core ASes and the lowest in the AS cluster with the highest number of edge ASes. Further, our results reveal most significant hosting at the intermediate layer of the Internet hierarchy, followed by the Internet edge and core. Though the aggregate fraction of contents hosted at the core is smaller, the density of content bytes per AS is the highest at the core. The distributed hosting across the transit hierarchy reflects in content delivery time; on average 22% of ads and 62% of regular contents are downloaded within a second from requesting the web page. Next, we observe that ads use a higher number of IP addresses and more ASes per website compared to regular contents, suggesting that ads use more servers for load distribution. As a result, the average latency of 1.6 seconds for ads is lower than 3.5 seconds for regular contents. From an economic perspective, around 60% of ads and 55% of regular contents originate from ASes that have more peering links than transit links, indicating that more than a half of ad and regular content traffic might incur less transport costs due to likely routing through peering links.

The main contributions of this work are as follows:

1. Hosting is pervasive throughout the transit hierarchy, including the tier-1 networks.

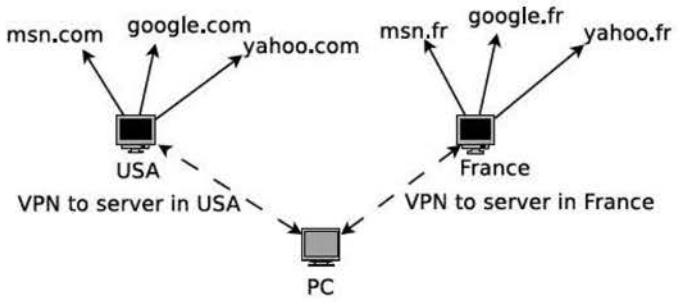

Figure 4.1: Data collection using a VPN

Our results confirm the trend towards increasing the number of roles an AS plays in the Internet ecosystem.

2. Ad and regular contents are hosted by significantly different populations of ASes. Replication is global for regular content and local for ads.

3. Reflecting the differences in the hosting AS populations, delivery performance for ads and regular contents also differ substantially. While the initial response to web requests is quicker for ads, the download time is lower for regular contents.

4.2 Methodology

In this section, we describe our methodology for measuring the hosting infrastructures and IP addresses used by online contents. First, we collect content data from a set of top websites for different countries. Since online ads have been the major source of revenues for CSPs, our work also focuses on studying the hosting characteristics of the ads. Therefore, we classify online contents into ads and regular contents and analyze qualitative and quantitative differences in hosting characteristics between them. By performing Internet-scale measurements, we discover the underlying hosting infrastructures and IP address usage of online contents.

4.2.1 Data collection

A general aspiration in online advertising is that users receive ads that are most relevant to the users' interest and locations. The users' interests are mostly determined by tracking which websites the users visit. The locations are determined from the source IP addresses of the users. Based on these data, ad networks choose to send ads from servers near to users' locations. Also, popular CSPs (e.g., Yahoo) host country-specific versions of their websites providing unique regular contents specific to the region and hosting the contents on nearby infrastructures.

In order to collect location-specific ads and regular contents, we use a novel VPN-based approach to collect online contents from top websites in many countries. Specifically, we use the VPN service from HMA.com [78] that operates servers in 52 different countries around the world. Next, we select the top 50 Alexa [79] websites for each of these 52 countries and automate the browsing of all the 50 websites from our data-collecting PC via a VPN connection to the server in each of the 52 countries, as shown in figure 4.1. Our novel VPN-based methodology enable us to collect the country-specific online contents delivered to users in the locations of the VPN presence.

After connecting our PC to a remote VPN server, we retrieve all the hyperlinks on the landing page of a website using Lynx (Linux-based text browser). Since CSPs use URLs embedded inside scripts on their websites for fetching ads, videos, and other contents, and because Lynx does not execute XML, javascript, and flash objects, we employ the methodology proposed in [80] to collect the embedded URLs of each website. These hyperlinks are then opened in the Firefox browser configured with Firebug, FireStarter, and NetExport add-ons. Using the above methodology, we browsed 2,600 websites in 52 countries and collect around 300 GB of HyperText Transfer Protocol (HTTP) and HyperText Transfer Protocol Secure (HTTPS) header data from 2,165 websites during mid-January 2013.

4.2.2 Identifying the ad URLs

A website is a complex collection of contents including third-party contents. The third parties usually serve ads, Application Program Interfaces (APIs), widgets, etc. to the websites. To identify the ad URLs in our dataset, we use the filtering rules from the widely used filtering plugins, such as AdBlock Plus and Ghostery. Using these filtering rules, we identify around 7,380 ad URLs in our dataset.

4.2.3 Identifying the regular URLs

Identifying regular URLs among third-party URLs is a complicated task. In addition to ad URLs, websites also contain third-party URLs for widgets and external websites which the browser plugins do not filter. Intuitively, we can identify the regular URLs of a website by matching their Second Top-Level Domain (STLD) names with the STLD of the website's primary URL, e.g., mail.yahoo.com and www.yahoo.com have common STLD yahoo.com. On other hand, URLs native to the website also can have STLDs, e.g., us.yimg.com, that cannot be straightaway identified by matching the primary STLD. We refer to such URLs as alternate URLs.

The methodology in [80] relies on the authoritative nameservers to identify the alternate URLs of websites. If any URL on a website shares its authoritative nameserver with the website's primary URL, then such URL is flagged as a native regular

URL. This method works for websites that manage their own authoritative nameservers but fails to distinguish the URLs of websites hosted by third-party hosting providers. For example, third-party hosting provider Amazon hosts websites www.9gag.com and www.rockmelt.com and also supplies an authoritative name service for them. The URL-identification method based on the authoritative nameservers flags all the URLs of www.9gag.com and www.rockmelt.com as native to both the websites. Besides, www.9gag.com uses alternate URLs with the d*.cloudfront.net pattern to host static contents from Amazon's Cloudfront CDN service. This makes the task of distinguishing between alternate and third-party URLs on www.9gag.com difficult because the cloudfront.net STLD is registered and administered by Amazon. Therefore, even a *whois* query for this domain name does not reveal the host website of such URLs.

We use a combination of techniques to identify regular and alternate URLs of each website and separate the third-party URLs of the website. In the following, we describe our procedure for the regular and alternate URL identification.

Step 1: Matching the STLD. First, we flag a URL as regular if the STLD of the URL matches the STLD of the website's primary URL.

Step 2: Verifying the referrer. The URLs left after the previous step are now subjected to a referrer verification where the URLs that are not referred by the website's primary URL are filtered out. This is achieved by checking the Referrer field in the HTTP request headers. The check filters out most of the third-party URLs referred by third-party domains. For example, website www.dropbox.com has external links to several third-party websites such as nytimes.com. Third-party URLs such as css.nyt.com belonging to nytimes.com can be present in the HTTP data collected for the dropbox.com. However, the Referrer field of such third-party URLs is nytimes.com. Therefore, third-party URLs can be filtered out by selecting only those URLs that are referred by the website's primary URL.

Step 3: Leveraging the request frequency. While the previous step does not handle the case where a website refers URLs to multiple third-party websites, the native STLDs of a website get high request counts compared to third-party STLDs. The following condition separates all the alternate URLs from the website-referred extraneous URLs. Each STLD i is selected if,

$$y_i \cdot n_i > \sigma \tag{4.1}$$

where y_i is the total request count of STLD i, n_i is the total number of URLs with STLD i, and σ is the standard deviation for the request counts of all the STLDs on the website. The URLs of the selected STLD i are considered to be regular URLs of a website.

In total, we identified 19,140 regular URLs using the above 3 steps.

4.3 Discovering the content infrastructure

This section describes how we discover the network resources (i.e., IP addresses, CNAMEs, and ASes) and hosting-infrastructure geographic footprints of the online contents. The discovery utilizes a network of open recursive DNS vantage points to resolve each URL in [81]. We start with around 130,000 open recursive DNS servers as globally distributed vantage points across the Internet. Then, we eliminate around 36,500 DNS servers for reasons such as recursion unavailable (23,000), unreachable (12,500) and invalid DNS response (1,080).

4.3.1 Eliminating the misleading DNS servers

While small in number, the servers with invalid responses are particularly important to filter out because of their potentially large negative impact on the measurement accuracy. We detect DNS servers that inject fake IP addresses during the URL resolution. Two types of invalid responses are observed. The first type consists of invalid replies for only particular URLs, such as www.facebook.com and www.youtube.com, and is also observed in [82, 83]. Invalid responses of the second type come from DNS servers that fail to resolve a URL or are unwilling to perform DNS recursions. Among the latter, the invalid IP addresses point to ad-displaying search pages managed by the entities that operate the misleading DNS servers.

To detect misleading servers of the first type, we send a single DNS query to each of the 130,000 DNS servers to resolve the primary www.facebook.com URL of Facebook. The IP address returned by a DNS server is then used to launch a reverse DNS query. The reverse DNS request to a valid IP address of the www.facebook.com URL should return a name record containing facebook.com as the STLD. On the other hand, an invalid IP address yields a name record with an irrelevant STLD. Employing the above method, we detect around 630 misleading DNS servers. Surprisingly, the largest number of such DNS servers are hosted in the United States with 417 servers in 15 different ASes. China stood second with a total of 180 such misleading DNS servers hosted by 32 ASes.

To detect misleading DNS servers of the second type, we send a single DNS query to resolve a non-existent URL, e.g., pppqqppqqp.com. An open recursive DNS server is expected to respond to such query with an NXDOMAIN error code. A misleading DNS server replies with an IP address pointing to an ad-displaying search page. Using this method, we detect 450 misleading DNS servers hosted across 48 countries and 210 ASes, bringing the total of eliminated misleading DNS servers to 1,080.

4.3.2 Resolving the URLs using open recursive DNS

After the above filtering steps, we have about 94,000 open recursive DNS servers spread across approximately 8,500 ASes, 22,040 prefixes, and 172 countries. By selecting

a single IP address from each prefix, we use around 22,040 DNS servers as vantage points. Each ad and regular URL is resolved from these vantage points to obtain the CNAMEs and IP addresses associated with the URL. The DNS resolutions of 19,140 regular and 7,380 ad URLs yield 102,600 and 90,000 IP addresses respectively, with the average of 12 IP addresses per ad URL and 5 IP addresses per regular URL.

4.3.3 IP-to-hosting infrastructure mapping

Finally, the IP addresses obtained in the previous section are mapped to their respective ASes, hosting organizations, and geographyic locations. We utiliz the IP-to-AS mapping service from [84] to map each IP address to its hosting AS, prefix, registry, and hosting organization. To map an IP address to its hosting country, we use the GeoIP tool from MaxMind [85] which is largely accurate on the country level compared to its city-level and finer resolutions. When possible, we also map each URL of a website to its CNAME, hosting AS, IP address, prefix, hosting country, registry, and hosting organization to form a network-level record of the URL. An URL with multiple IP addresses is mapped to multiple records. Finally, the total number of ASes discovered for the 19,140 regular and 7,380 ad URLs are 2,177 and 2,272 ASes respectively, and the total number of host countries are 115 and 134 respectively.

4.4 Measurement results

In this section, we analyze the measurement data to study the hosting infrastructures with respect to their ASes, CDNs, geography and inter-AS business relations. We also study top hosting networks and delivery performance for online contents. Our analytic objective is to understand the online hosting in its relationship with the ISPs across the transit hierarchy.

4.4.1 Hosting-infrastructure characteristics

Mechanisms for hosting of contents range from centralized to geographically distributed. To achieve higher performance by minimizing latencies, CSPs and ad networks distribute their contents across the Internet and geography by collaborating with different CDNs and ISPs. The consumed network resources, such as IP addresses and ASes, depend on the hosting mechanism. The CDNs mainly leverage the infrastructures of many ISPs across the Internet to cache contents closer to the end users.

IP address and AS counts. Figure 4.2 presents the cumulative distribution for the number of IP addresses and AS counts for the ads and regular contents of the examined websites. Figure 4.2(a) reveals that the ad contents served on the median fraction of websites use up to 3,000 IP addresses, while the regular contents use only up to 100 IP

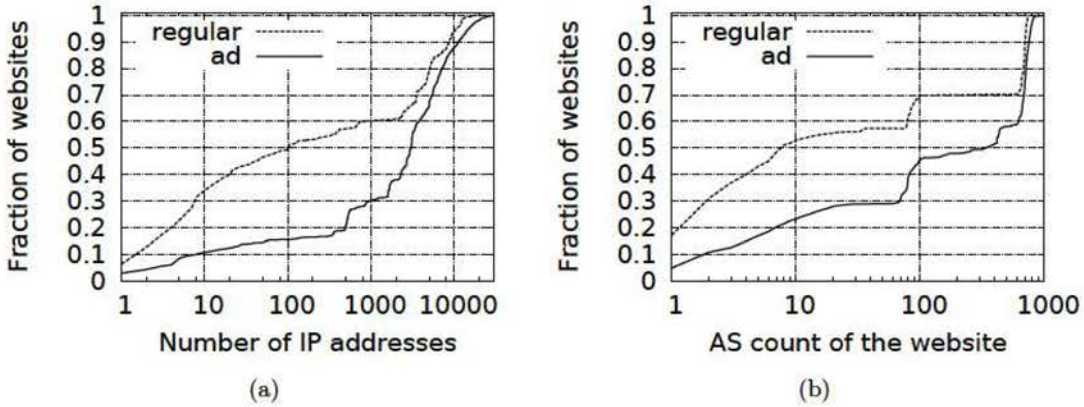

Figure 4.2: Distribution of (a) IP-address count and (b) AS count for the ad and regular contents of the websites.

addresses. The results suggest that ad contents are distributed across a larger number of servers than regular contents. In figure 4.2(b), up to 70% of the websites receive ad contents distributed across around 700 ASes and regular contents from about 100 ASes. An interesting characteristic is the step increases in the AS counts of the websites. This behavior arises because of the presence of different CDNs that host contents according to their geographical popularity. The steps indicate clusters of websites hosted by CDNs that replicate online contents across multiple ASes as per the websites' geographical popularity. Online contents of very popular websites are globally replicated across a large number of ASes. The contents of less popular websites are mostly confined within a continent or country, and hosted by a smaller number of ASes. Figure 4.2(b) shows respectively, 4 and 3 clusters of websites for the ad and regular contents according to the AS count of the websites: clusters with 1 to 30 ASes (local hosting), 70 to 100 ASes (regional hosting), 600 to 1000 ASes (global hosting) for both the ad and regular categories, and the 4th cluster with 400 to 500 ASes (continental) hosting for ads.

Clustering of hosting ASes. Next, we analyze features shared by the hosting ASes. While our data are for the top-50 websites in 52 countries, ASes might share hosting characteristics such as the number of websites, geographical reach, and number of IP addresses. To explore the shared features, we record how many of the popular websites each hosting AS serves in each of the countries. Figures 4.3(a) and (b) show these data for all the 2,272 ad-hosting ASes and 2,177 regular-hosting ASes on a scatter plot with a colored density. Each tile in the above plot represents a hosting AS according to the number of websites served in a country. The tile color ranges from white through blue to black and signifies the density of ASes that share the respective website count per country. The higher the AS density is, the darker the tile color becomes.

Figure 4.3(a), shows two distinct clusters A_h and A_l and four faint clusters A_t, A_i,

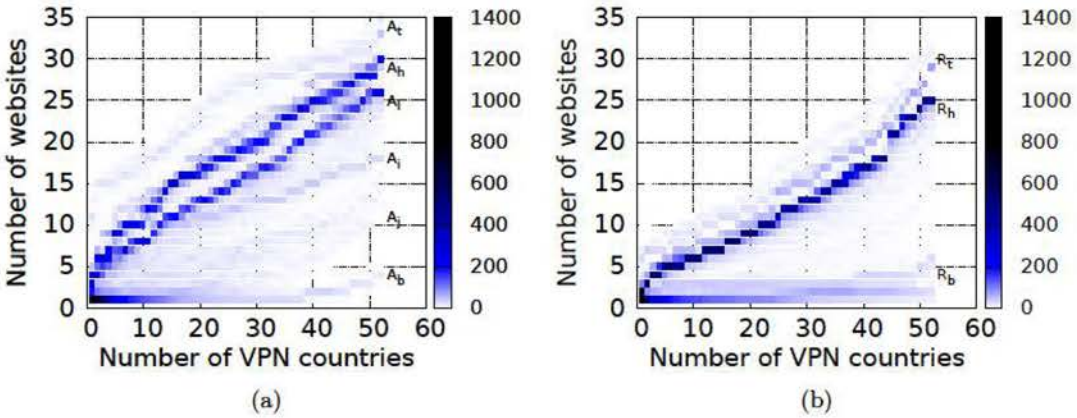

Figure 4.3: Hosting ASes arranged according to the number of websites served in each country: (a) ad-hosting ASes and (b) regular-hosting ASes.

A_j, A_b of ASes serving ad contents to websites across the countries. On the other hand, figure 4.3(b) shows three clusters R_t, R_h and R_b of ASes serving regular contents. Hence, while the ad-hosting ASes form a larger number of smaller clusters, and the regular-hosting form a smaller number of bigger clusters, the results suggest that ad contents are locally replicated to serve regional users and that regular contents are globally replicated to cater to users across the Internet.

Next, we investigate differences between the AS clusters. We focus on four significant clusters A_t, A_h, A_l, and A_b of ad-hosting ASes and three clusters R_t, R_h, and R_b of regular-hosting ASes. The four ad-hosting clusters comprise of 47, 362, 278, and 1,456 ASes with the average number of websites per AS per country equal to 24, 18, 14, and 4 respectively. The three regular-hosting clusters consist of 103, 759, and 1,252 ASes with the average number of websites per AS per country being 15, 11, and 2 respectively.

IP-address count of AS clusters. Figures 4.4(a) and (b) present the cumulative distribution of IP addresses for the ad-hosting and regular-hosting AS clusters respectively. Clusters A_b and R_b contain the highest fraction, around 40%, of the ASes that use a single IP address for hosting ad and regular contents. On the other hand, clusters A_t and R_t employ a large number of IP addresses compared to other clusters: 30% of the ASes in A_t and 60% of the ASes in R_t use more than 100 IP addresses. Interestingly, regular-hosting cluster R_h and ad-hosting cluster A_l have a similar profile: around 96% of the ASes in cluster A_l are present in cluster R_h, which in turn contains about 86% of the ASes in cluster A_h.

Position in the transit hierarchy. Next, we analyze the IP-address distribution for each AS cluster with respect to its position in the transit hierarchy. We use the AS-level dataset from Center for Applied Internet Data Analysis (CAIDA) [56] to categorize the ASes into edge, intermediate, and core ASes. Using the CAIDA dataset, we compute

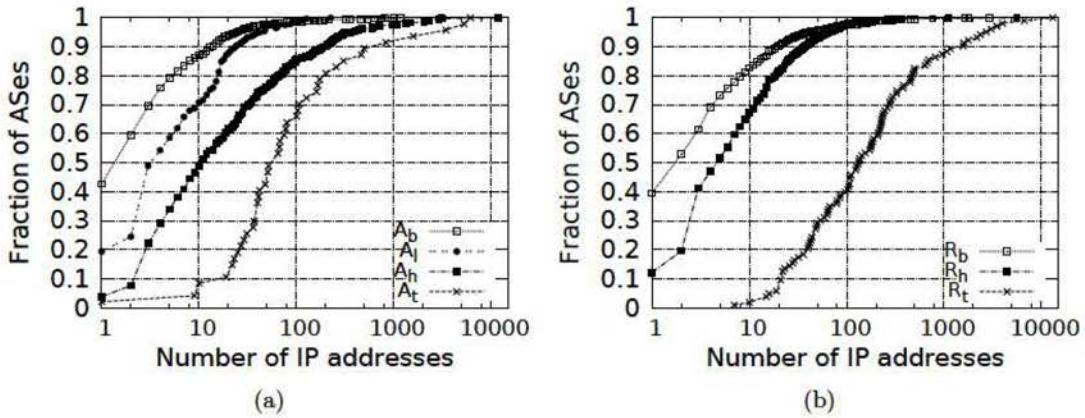

Figure 4.4: Cumulative distribution of IP addresses for different clusters of: (a) ad-hosting ASes and (b) regular-hosting ASes.

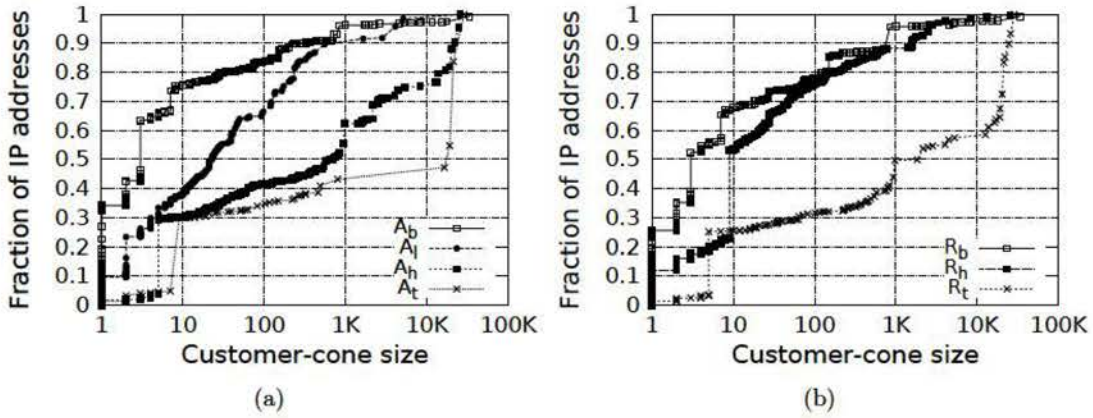

Figure 4.5: Cumulative fraction of IP addresses arranged according to the customer cone size of: (a) ad-hosting ASes and (b) regular-hosting ASes

for each AS, the customer-cone metric i.e., the number of direct and indirect transit customers of the AS [86]. Based on the studies in [34], we label all the tier-1 ASes as core ASes. Those core ASes usually have very large customer-cone values exceeding 20,000. Using the method in [87], any AS with the customer-cone value smaller than 5 is labeled as an edge AS. The edge ASes are normally regional ISPs, CSPs, enterprise, or campus networks. All the ASes, not labeled as edge and core ASes, are classified as intermediate ASes. These ASes are typically operated by tier-2 or tier-3 ISPs that connect edge ASes with core ASes.

Figures 4.5(a) and (b) present the cumulative fraction of IP addresses arranged according to the customer-cone size of the the examined ad-hosting and regular-hosting ASes respectively. The fractions of IP addresses used by edge ASes is the highest in

Table 4.1: Top 10 ASes arranged according to the number of IP addresses used for hosting ads and regular contents in clusters A_h and R_t

AS number	Organization	IPs (regular)	IPs (ad)	Customer-cone	Service (Hierarchy)
20940	Akamai	13,182	11,832	5	CDN (Edge)
4436	GTT (nLayer)	3,513	3,196	975	Transit (Intermediate)
209	CenturyLink	3,054	2,922	20,294	Transit (Core)
3257	GTT (Tinet)	2,725	2,340	24,729	Transit (Core)
7922	Comcast	2,368	2,202	2,164	Access (Intermediate)
1299	TeliaSonera	2,196	2,054	25,753	Transit (Core)
7843	Time Warner	1,603	1,500	932	Access (Intermediate)
1273	Vodafone	1,577	1,346	13,780	Transit (Intermediate)
1239	Sprint	1,167	1,070	22,042	Transit (Core)
5511	Orange	1,073	989	4,431	Transit (Intermediate)

clusters A_b and R_b, followed by clusters A_l, R_h, and A_h, and the lowest in clusters A_t and R_t. In contrast, the fractions of IP addresses used by intermediate and core ASes are the highest in clusters A_t and R_t and lowest in clusters A_b and R_b. This indicates that the IP-address usage in AS clusters depends on the position of the ASes in the transit hierarchy.

The IP-address usage in ad cluster A_h and regular-hosting cluster R_t is qualitatively similar. The total number of IP addresses used by all ASes in clusters A_h and R_t is around 47K and 61K respectively and is the largest among all the clusters. To understand the close similarity of the IP-address usage between clusters A_h and R_t, we examine which ASes form these two clusters. The clusters share 50 prominent ASes that collectively account for around 40K IP addresses. 80% of the ASes are intermediate ASes, followed by 10% core ASes and 10% edge ASes.

For clusters A_h and R_t, table 4.1 presents the top-10 ASes among these 50 prominent ASes according to the number of IP addresses, along with the customer-cone size, service, and hierarchy types of the ASes. The top-10 set consists of 7 transit networks, including 4 core and 3 intermediate ones, 2 access ASes, which are both intermediate, and 1 CDN, which is an edge hosting AS. The prominent 50 ASes form 14% of the ASes in ad-hosting cluster A_h. In contrast, regular-hosting cluster R_h contains the other 86% of the ASes in cluster A_h. Similarly, regular-hosting cluster R_t contains about 85% of the ASes in ad-hosting cluster A_t . This reveals a big difference from cluster R_t: while R_t contains the prominent 50 ASes, A_t does not. In ad-hosting cluster A_t which collectively accounts for 21K IP addresses, 5 prominent ASes consisting of Google (AS 15169) and 4 large transit ASes: NTT (AS 2914), Level3-GBLX (AS 3549), Deutsche Telekom (AS 3320), and Telecom Italia (AS 6762) contribute up to 81% of the total IP addresses that serve the ad contents.

Clusters A_b and R_b, where edge ASes predominantly contributed IP addresses, also

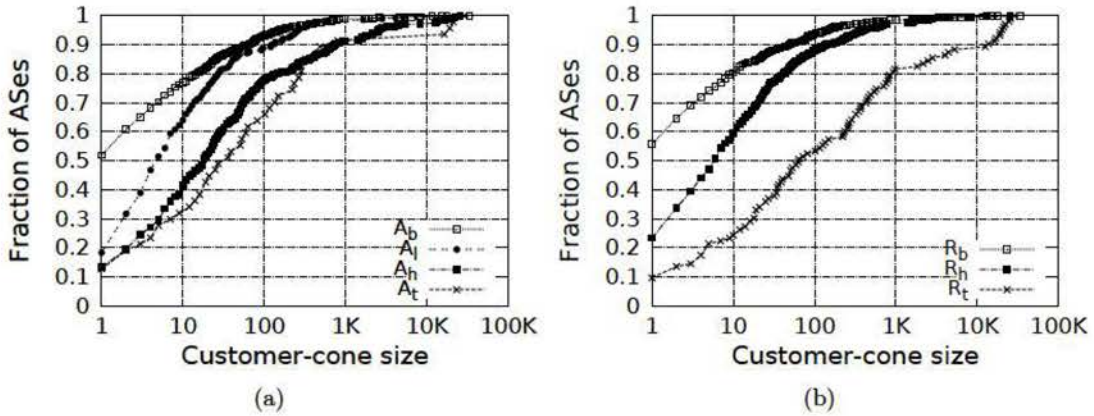

Figure 4.6: Cumulative distribution of the customer-cone size of: (a) ad-hosting ASes and (b) regular-hosting ASes.

display qualitatively similar IP usage across the Internet hierarchy and collectively account up to 18K and 21K IP addresses respectively. In both clusters, Amazon (AS 16509) and Microsoft (AS 8075) are the 2 ASes with the largest IP count. While Akamai (AS 16625), CDNetworks (AS 36408), and Amazon (AS 14618) are the next top ASes in ad cluster A_b, China Telecom (AS 4134), China Unicom (AS 4837), CDNetworks (AS 36408) are the next top ASes in regular-hosting cluster R_b with respect to the IP-address count.

Transit hierarchy in the clusters. Figures 4.6(a) and (b) plot the cumulative distribution of the customer-cone size of the ASes in all ad-hosting and regular-hosting clusters respectively. Both ad and regular contents are hosted by ASes from across the Internet hierarchy. The fraction of edge ASes varies from 18% in cluster R_t to 44% in cluster R_h, and up to 72% in cluster R_b, as shown in figure 4.6(b). Figure 4.6(a) depicts a qualitatively similar trend for the ad-hosting clusters. The fraction of edge ASes is the smallest for cluster A_t and increases in clusters A_h, A_l, and A_b. Complementing the fraction of edge ASes, the fractions of intermediate and core ASes follow the opposite trend. While R_t contains 75% intermediate and 7% core ASes, R_b has 27.7% intermediate and 0.3% core ASes. The above results show that the AS clusters preserve the transit-hierarchy positions of their ASes.

Websites hosted on the AS clusters. Figures 4.7(a) and (b) plot the number of websites served by the ad-hosting and regular-hosting clusters respectively, arranged according to the number of IP addresses on the website. The IP-address counts of the websites correlate highly with the AS position in the transit hierarchy. The clusters consisting mostly of core and intermediate ASes use a larger number of IP addresses per website. For example, A_h and R_t use more IP addresses per website than A_b and R_b. In contrast to the IP count of the websites, the website count of the clusters has a reverse dependence on the transit-hierarchy position: while clusters A_h, A_l and A_b serve 757, 780,

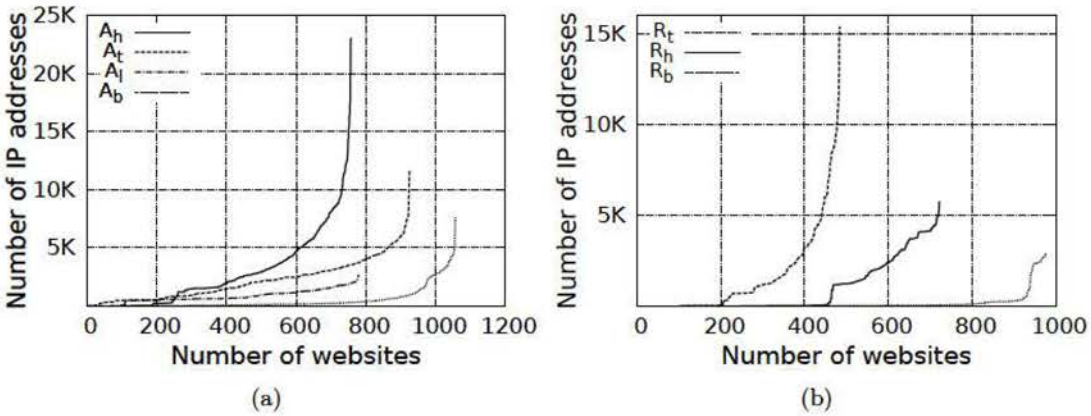

Figure 4.7: Number of websites served by: (a) ad-hosting clusters and (b) regular-hosting clusters, arranged according to the number of IP addresses per website.

and 1,057 websites respectively, clusters R_t, R_h, and R_b serve 484, 722, and 977 websites respectively. Cluster A_t is an exception from this trend and serves 925 websites. In all 4 ad-hosting clusters, www.yahoo.com and www.msn.com are the 2 websites that receive ads from the largest number of IP addresses. In clusters R_t and R_h, www.weather.com and www.msn.com have the highest IP address counts. In cluster R_b, www.bing.com and www.lemonde.fr are the websites with the largest number of IP addresses.

The above results can be summed up as follows:

1. Ads use higher numbers of IP addresses and AS counts per website than regular contents, suggesting that ads use a higher number of servers to distribute the load more broadly.

2. While ad-hosting ASes form a larger number of clusters, the clusters of regular-hosting ASes are larger in size, indicating that replication is global for regular contents and local for ads.

3. Clusters with a majority of ASes from a higher tier of the transit hierarchy serve more websites per AS per country and employ a larger number of IP addresses per website.

4. The fractions of core, intermediate, and edge ASes in the clusters vary for different clusters, demonstrating that AS clusters reflect the hierarchical positions of their ASes.

5. With a single exception, clusters from a higher tier of the transit hierarchy serve a smaller number of websites.

Table 4.2: Top 10 ASes arranged according to the bytes served for ads and regular contents

Regular content			ad content		
Organization	AS number	Bytes (%)	Organization	AS number	Bytes (%)
EdgeCast	15133	10.5	Google	15169	24.0
Akamai	20940	5.2	EdgeCast	15133	4.2
Level 3	3356	3.3	Akamai	20940	2.8
Wikimedia	43821	3.3	Amazon	16509	2.7
Microsoft	8075	1.7	Level 3	3356	1.3
NTT	2914	1.6	NTT	2914	1.2
China Telecom	4134	1.5	NetVision	1680	1.2
CDNetworks	36408	1.3	UAB Hostex	47205	1.2
Amazon	16509	1.3	Telia	1299	1.2
China Unicom	4837	1.2	fibre one	24961	1.0

4.4.2 Byte volume and location of contents

Now, we study the hosted contents in regard to their byte volume and location.

Top content-hosting ASes. In this section, we consider ASes individually without clustering. We estimate the volume of bytes hosted by the ASes as follows. First, we compute the total bytes served from each URL. Then, we uniformly distribute the URL's total bytes among all the IP addresses of the URL. Finally, we aggregate the bytes associated with the IP addresses to determine the content volumes of the ASes. Table 4.2 presents the top-10 ASes arranged according to the fraction of bytes delivered for ads and regular contents. These top-10 ASes contribute up to 31% of all regular-content bytes and 41% of all ad bytes. AS 15133 operated by EdgeCast serves 10.5% of the total regular-content bytes, which is the largest in our dataset, and predominantly hosts contents for websites such as Pinterest, Twitter, and WordPress. Google serves up to 24% of the total ad bytes in our dataset.

Top CDNs. Table 4.3 presents the top 5 ad-hosting and regular-hosting CDNs in our measurements. Akamai is the biggest CDN for regular contents, serving 34% of the websites containing 27% of the total regular-content bytes. It serves 5% of the regular-content bytes from its flagship AS 20940 and other 22% from more than 700 third-party ASes. Unlike Akamai, EdgeCast serves all content bytes from its servers located across the Internet and managed under its single flagship AS 15133. For ads, Google serves 61% of the websites and 29% of the total ad bytes in our dataset. Google's ads are split between its primary AS 15169, which serves 24% of the total ad bytes, and more than 400 third-party ASes, which serve the remaining 5% of the total ad bytes.

Content in the transit hierarchy. Next, we assess the fractions of byte volumes

Table 4.3: Top 5 CDNs arranged according to the served bytes

Regular content				ad content			
CDN	Byte(%)	AScount	Website(%)	CDN	Byte(%)	AScount	Website(%)
Akamai	27.5	730	34.2	Google	29.0	448	61.5
EdgeCast	10.5	1	8.1	Akamai	21.5	730	49.3
Level 3	3.4	15	10.8	EdgeCast	4.2	1	16.3
Microsoft	2.4	11	4.2	Amazon	3.2	26	26.1
Amazon	1.5	23	13.0	Level 3	1.4	11	15.6

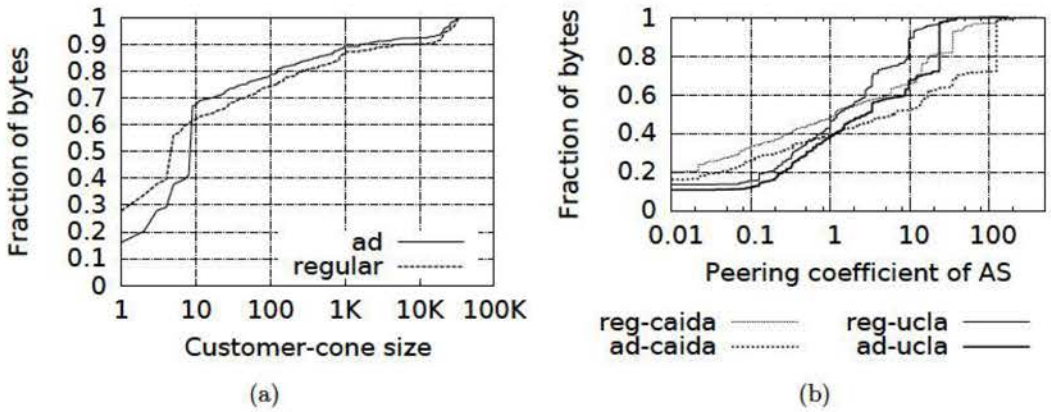

Figure 4.8: Fraction of online content bytes hosted by the ASes according to (a) customer-cone size; and (b) peering coefficient.

hosted by core, intermediate, and edge ASes. Figure 4.8(a) shows that nearly 33% of the ad bytes and 40% of the regular-content bytes are hosted by edge ASes which constitute 55% and 59% of the total 2,272 ad and 2,177 regular ASes respectively. The core ASes host roughly 4% of ad and 6% regular-content bytes, which is significant because there are only 12 such ASes. Intermediate ASes host the remaining 63% of ad and 54% of regular-content bytes. Though the content volume is the lowest in the core, the volume density is the highest for the core ASes and equals 0.5% for ads and 0.3% for regular contents. Both ad and regular-content density per AS approximately equals 0.06% for intermediate ASes and 0.03% for edge ASes.

Byte volume vs. business relationships. To understand the relation between the hosted content volumes and business profiles of the ASes, we compute the peering ratio for each AS, defined as the number of peering links divided by the number of transit links. To compute peering ratios, we use the inter-AS relationship datasets from CAIDA and UCLA [59], which are similar in the number of transit links and largely different in the number of peering links. A peering ratio less than 1 implies that an AS has more transit links and less peering links. A large transit AS has a large number of transit

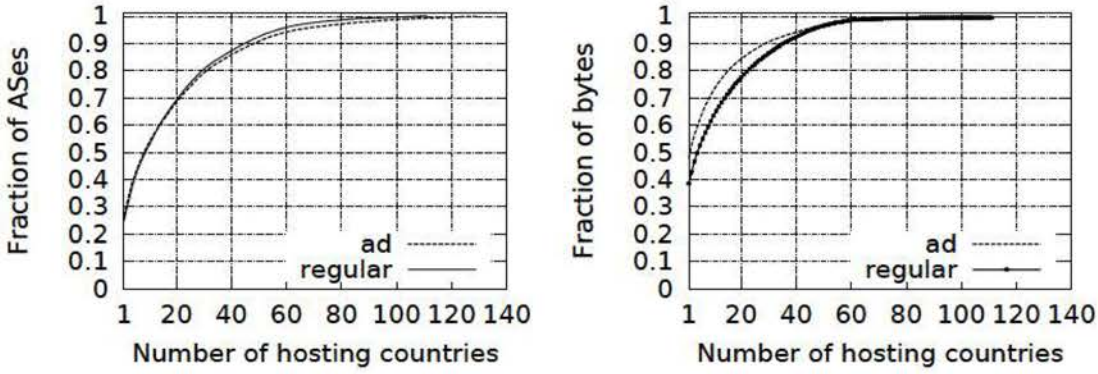

(a) (b)

Figure 4.9: Distribution for the number of hosting countries with respect to: (a) fraction of ASes and (b) fraction of hosted bytes.

Table 4.4: Top-10 hosting countries ranked by the AS counts and fractions of the hosted bytes.

Regular content				ad content			
Ranking by AS count		Ranking by bytes		Ranking by AS count		Ranking by bytes	
Country	AS count	Country	Bytes(%)	Country	AS count	Country	Bytes(%)
USA	511	USA	38.6	USA	558	USA	47.8
UK	109	Netherlands	4.1	Germany	110	Ireland	4.8
Germany	102	China	3.9	UK	108	Germany	3.6
Russia	96	Russia	3.0	Russia	93	UK	2.9
Japan	84	France	2.8	Japan	76	Denmark	2.8
France	71	Germany	2.2	France	66	Russia	2.7
Canada	59	Estonia	2.1	Netherlands	61	Israel	1.9
Australia	54	UK	2.0	Canada	61	Lithuania	1.9
Singapore	49	Sweden	1.9	Poland	52	Japan	1.7
Poland	49	Poland	1.7	Australia	51	Austria	1.7

customers and relatively small number of peers. Therefore, the peering ratio of large transit ASes is below 1. Small ASes have a peering ratio above 1 because they peer extensively to reduce their transit costs [5, 6]. Figure 4.8(b) presents the distribution of ad and regular-content bytes hosted by ASes in relation to peering ratio of the ASes. In both CAIDA and UCLA datasets, around 60% of ads and 55% of regular contents originate from ASes with a peering ratio greater than 1. Thus, using these fractions together with the previous results for fractions of byte volumes hosted by intermediate ASes, we can derive the specific fraction of byte volumes hosted by intermediate ASes with a peering ratio below and above 1 to be 36% and 27% for ads respectively, and 39% and 15% for regular contents respectively. These specific results are useful for modeling a content traffic matrix which considers inter-AS business relationships.

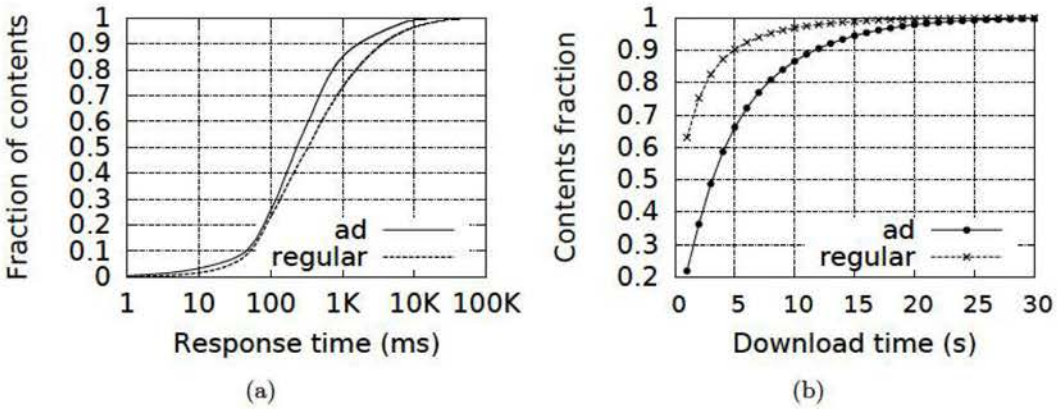

Figure 4.10: Delivery performance for regular contents and ads: (a) response times (b) Download times.

Geographic locations of contents. Figure 4.9(a) plots the cumulative fraction of hosting ASes in regard to the hosting countries. For both ad and regular contents, the top-10 countries ranked by the AS counts account for about 54% of the hosting ASes. Figure 4.9(b) demonstrates that the top-10 countries ranked by the fraction of hosted bytes account for around 62% for regular contents and 72% for ads. Table 4.4 shows that the sets of top-10 hosting countries are substantially different for the two metrics: the countries with a smaller number of ASes and higher byte volumes host popular CSPs, such as Google and Yahoo. The above results might help an ISP to decide in which country to deploy its infrastructure.

4.4.3 Content delivery performance

While the previous studies characterize content hosting across the transit hierarchy and geography, we now examine delivery performance for the hosted contents.

The end users, publishers, and advertisers are interested in the contents to be delivered as promptly as possible. We measure response and download times for regular contents and ads.

Response times. Figure 4.10(a) plots the distribution of response times for ads and regular contents. The profiles for both types of contents are qualitatively similar up to response times around 100 ms and diverge beyond that point. The 90th percentile of response times for ads and regular contents are 1.6 seconds and 3.5 seconds respectively. The smaller content size and more local hosting are likely reasons for the faster responses to ad requests.

Download times. The download time of a content is the time between requesting its webpage and arrival of the content. This time depends on the priority given to the content in regard to the other contents of the webpage. Figure 4.10(b) presents the distribution of

the download times for all ads and regular contents. The 90th percentile of the download times for regular and ad contents are 4 and 12 seconds respectively, suggesting that a majority of ads arrive long after regular contents. Also, nearly 22% of ads and 62% of regular contents are downloaded within one second of requesting the web page. Hence, ads have smaller response times and larger download times. One plausible reason is a lower priority given to ad requests on a majority of the websites.

4.5 Discussion: Economic impact of Internet advertising

Over the last 2 decades, entrepreneurs leveraged the growing Internet to provide various web-based services, such as online banking, shopping, and search. The new Internet-enabled electronic marketplace attracted numerous advertisers, thereby bringing billions of revenues. In return, the financing of the online contents by the online advertisers drove the Internet infrastructures growth and content-hosting innovation, indirectly contributed to transforming the Internet from a transit hierarchy towards a flat mesh of peering links, and also fueled inter-ISP relationship disputes and debates over net-neutrality [88–90].

Internet infrastructure growth: Millions of individuals and enterprises produce websites to earn ad revenues. The demand for website hosting brings thousands of entities into the hosting business. [91] reveals that the hosting-provider population grows with a factor 1.97 since January 2002. While the web traffic constitutes 52% of all Internet traffic [60], the transit ISPs invest heavily into upgrading their network capacity to accommodate the swelling traffic. The international Internet capacity nearly quadrupled from 24 Tbps in 2009 to 104 Tbps in 2013 [92].

Impact on transit prices: The increasing demand for higher network capacities and advancements in network technologies fueled intense competition in transit markets and persistently reduced transit prices [92]. As transit and access ISPs experienced a decreasing Return on Investment (ROI), many of them overhauled their business models to cope with the new market realities. Supported by ad revenues, content hosting proved more profitable than relying only on transit services. Many incumbent transit providers, such as Level 3 and Cogent, started their own CDN businesses and collaborated with many ad networks and popular content providers, e.g., Netflix.

Impact on inter-AS business relationships: Traditionally, the inter-AS business relationships were mostly limited to transit and peering. The augmented content hosting and caching created a need for several new types of inter-AS economic relations, such as paid peering, partial transit, and remote peering, to deal with increasing heterogeneity in business profiles and geographic footprints of ASes.

Business relationships and net-neutrality disputes: The entrance of transit ISPs into content hosting payed way for several inter-ISP business relation disputes in the Internet. Generally, it is very uncommon for access ISPs to be transit customers of

CSPs. On the other hand, CSPs usually pay access ISPs either through paid-peering or partial-transit relationships. Due to content hosting by transit ISPs, the access ISPs that buy transit connections from such transit ISPs dispute their relationships because of content traffic originating from the transit ISPs, e.g., Level 3 vs. Comcast dispute [93]. Also, large access ISPs demand payments from remote CSPs in return for providing premium quality of service (QoS) to the content traffic, e.g., Google vs. Free dispute [94], or otherwise degrade the content traffic performance by not allocating sufficient network capacities [17]. These developments fuel debates on net neutrality, an informal code for transport networks to treat all the Internet traffic equally without discriminating on the basis of content, application, site, and protocol.

Potential revenue source for access ISPs: Recently, few UK-based ISPs have trialed an ad-sponsored 3G Internet access to heavily subsidize the monthly subscription fee to as low as $0 per month [95, 96]. In this model, the advertisers pay the ISP for displaying ads to end users. The ad revenues might help the ISP to partly recover the costs and also enable it to provide affordable Internet access, thereby attracting more end users.

4.6 Related work

Online contents were the focus of many research works. [97, 98, 80, 99] examined web-based contents and services along with their features such as content types, number of requests, bytes, and servers. [81, 100, 101] carried out measurement-based studies of real operational CDNs. A common focus is on few selected CDNs. [102, 103] explored the footprints of hosting infrastructures across the Internet and geography via real measurements relying on either volunteers [102] or traffic traces [103]. Our work uses a simple, and yet novel, VPN-based approach to collect online contents and then discover their hosting infrastructures from a large number of geographically distributed DNS servers. We detected a large ecosystem of hosting infrastructures, including those invisible from a large portion of the Internet. We also studied differences in hosting of ads versus regular contents. Unlike the previous studies that derive conclusions from meta data such as the number of served URLs, we measure byte volumes to characterize hosting more accurately.

Recently, few works explored content delivery performance for ad traffic. [104] analyzed delivery of ads with a focus on effectiveness of ad blocking. [105] charted 3 prominent ad networks and evaluated their latency and effectiveness of user targeting. [106] characterized mobile ad traffic using data collected in an operational ISP to study the traffic frequency, content types and energy implications for mobile devices. That work also briefly glimpsed into the hosting infrastructures of ad contents. Our work distinguishes itself from the above by collecting contents from the top 2,165 websites in 52 countries

to characterize their hosting infrastructures. Our work also explicitly compares hosting properties for ads versus regular contents.

4.7 Summary

In this chapter, we explore online content hosting and the role of transit ISPs in it. We analyze the hosting infrastructures with respect to their IP resources, positions in the transit hierarchy, geographic locations, and user-perceived delivery performance. We classify online contents into ads and regular contents and study differences in their hosting. Using a novel VPN-based measurement approach, we collect ads and regular contents from the top 2,165 websites in 52 countries. Our characterization of hosting infrastructures utilizes a vast network of open recursive DNS vantage points spread across 8,500 ASes and 172 countries.

Our analyses reveal that content hosting is pervasive throughout the transit hierarchy, including the tier-1 networks. Our results confirm the trend towards an increasing number of roles an AS plays in the Internet ecosystem. We observe that ads and regular contents are hosted by significantly different populations of ASes. While replication is global for regular contents, ads are locally replicated. Reflecting the differences in the hosting AS populations, performance for ads and regular contents also differs substantially. While responses to individual content requests are quicker for ads, the download time is lower for regular contents of a webpage. Our measurement data and analytic results are valuable for deriving a realistic matrix for content traffic, where transit ISPs are included as sources of content traffic.

Chapter 5

An ad-based Revenue Model for Access ISPs

5.1 Introduction

While chapter 4 shows that transit ISPs broadly rely on content hosting as a source of extra revenues, this chapter analyzes a model where access ISPs derive revenues from online ads.

The demand for Internet access by residential users is increasing and driven mainly by unlimited availability of online contents. A majority of the CSPs that offer free online contents derive their revenues from online ads. By financing the CSPs, the ads indirectly incentivize technological advancements in content development and hosting. Because online ads provide the CSPs with higher revenues when the consumption of the online contents is higher, the CSPs have incentives to create more advanced and innovative online contents, delivery of which might require significant network capacities. As chapter 4 reveals, the CSPs widely collaborate with different CDNs and transit ISPs to cache bandwidth-intensive contents across the transit hierarchy. On other hand, end users demand faster Internet access to experience a good quality of browsing the bandwidth-intensive contents. To satisfy the growing demand for faster Internet access and to sustain increasing traffic volumes, access ISPs have to periodically upgrade their network capacities.

The network upgrades are usually expensive, and the access ISPs are vocal about their rising network costs. To recover the costs of upgrading and operating their networks, the access ISPs monthly charge subscription fees to end users. The two most common billing models for access ISPs are usage-based and flat-rate billing [107, 108]. In usage-based billing, the users are charged according to the consumed traffic volume. With flat-rate billing, the access ISPs charge fixed monthly fees regardless of the consumed traffic. Due to intense competition in the access market, a majority of access ISPs have gravitated from the usage-based to flat-rate billing [109]. Unlike the flat rate for revenues, the costs

of the access ISPs depend on the peak traffic rates.

With the flat-rate revenues and usage-based costs, access ISPs find themselves in a difficult economic situation [1]. Responding to the challenge, many access ISPs adopt a combination of flat-rate and usage-based billing, i.e., by introducing data caps [110]. In the data-cap model, a user pays a fixed monthly fee but if the consumed traffic exceeds a monthly quota, the Internet access is either throttled to a lower speed or terminated altogether until the end of the billing cycle. Alternatively, large access ISPs demand financial compensation from popular CSPs for the costs of delivering the CSPs' bandwidth-intensive contents [111]. Meanwhile, some CSPs collaborate with access ISPs to subsidize end users subscriptions [112]. These initiatives trigger concerns about net neutrality [19, 20]. On the other hand, because transit ISPs increasingly host online contents, access ISPs also demand financial compensation from such transit networks [21].

While online ads proved to be a rich source of revenues, this chapter analyzes a model where an access ISP exploits online ads to earn extra revenues. We assess the revenue potential and economic viability of the ad-based model for access ISPs with different customer bases. While end users are sensitive to their subscription costs, this chapter also studies the utility of the ad-based revenue model for the end users. In the considered model, an access ISP directly collaborates with advertisers to display ads to end users without engaging CSPs. The ads are displayed in a dedicated space in the browser. Technical details of displaying the ads are discussed later in the chapter. We investigate the following questions about the ad-based revenue model:

1. What is the revenue potential of this model for differently sized access ISPs under current market prices of online ads?

2. What is the target ad revenue and per-unit price needed to offset the access ISP's network costs?

3. What is the maximum number of ads that an access ISP can display to the end users without degrading the quality of their browsing experience, and what is the ad revenue potential on this advertising level?

4. What are the incentives that an access ISP can offer to its ad-subsidized Internet subscribers?

We evaluate our model using financial data from one large and one medium-sized access ISPs operating in India. Our analysis shows a significant revenue potential of around 50% of the capital expense for large access ISP which serves millions of Internet subscribers. For the medium-sized access ISP which has tens of thousands of end users, the ad revenue potential is about 5% of its capital expense. Under the current Cost-Per-Mille (CPM) prices for ads and typical time spent by end users on the Internet, our

analyses reveal that the ad-based revenue model is economically viable for access ISPs. The ad revenues enable access ISPs to offer up to 6–9 MBps of additional access speed or 12–20 GB additionally consumed data as incentives for users to subscribe for an ad-subsidized plan. Also, we conduct a market survey of user interests in subscribing to ad-subsidized Internet plans. Our survey reveals that a significant majority of users are interested in trying this option.

This chapter is organized as follows. In section 5.2, we briefly discuss practicality of ad-subsidized Internet plans. The model implementation and user-utility formulation are presented in section 5.3. Section 5.4 reports the data collected for evaluating the model. Section 5.5 analyzes the revenue potential and economic viability of the model and also assesses the user incentives. Section 5.6 briefly discusses prior works, and section 5.7 summarizes the chapter.

5.2 Background

Generally, Internet users detest online ads because many poorly designed ads hinder the browsing experiences by floating and rolling across the screen or auto-playing audio-video files without an explicit consent from the users. However, online ads enable CSPs to serve online contents freely, making the Internet popular among residential users. Recently, few access ISPs have started to insert pop-up ads into browsers by tampering with content traffic without a user consent, triggering security and privacy concerns [113].

The considered ad-based revenue model gives users an explicit choice by offering transparent access plans subsidized by ad revenues. The users who choose to receive ads benefit from extra speed or data cap incentives. This model allows ISPs to earn ad revenues with an explicit consent of the user. To display ads, the ISPs provide mandatory plugins to be installed on browsers by the ad-subsidized Internet subscribers. The browser plugins create a dedicated space within the browser window, e.g., a horizontal or vertical panel at the window edge. The plugins periodically fetch ads from ad servers of the ISP or an ad network, and display the ads in this dedicated space. The model allows the users to control the number and frequency of received ads. To make the model robust to misuses, disabling or tampering with the browser plugins automatically disconnects the ad-subsidized subscriber from the Internet. While security and privacy aspects of the model are clearly important, their detailed study is a topic for future work.

Another important question is whether end users are willing to adopt ad-subsidized plans. Hence, we conduct a survey (https://www.surveymonkey.com/r/8BQ8TNN) by posing the following 3 questions to more than 100 residential users:

Q1: If your ISP proposes you to sign up for a new Internet plan that is cheaper or has faster speed or provides more data than your existing Internet plan, but comes with non-intrusive online ads sent by the ISP, would you sign up?

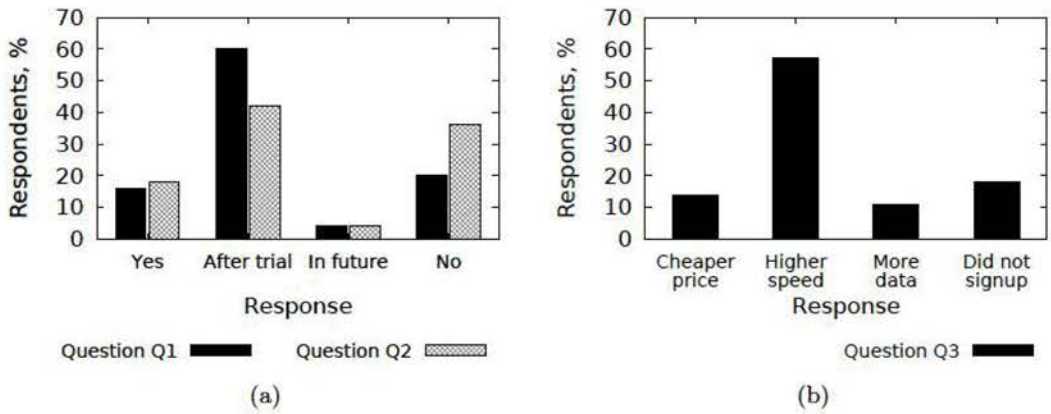

Figure 5.1: Responses of the survey: (a) questions Q1 and Q2 and (b) question Q3.

Q2: If receiving more ads brings you more discounts in the form of a smaller price or speed increase or data-limit increase, would you be interested in receiving more ads?

Q3: If you decide to sign up for a new ad-sponsored Internet plan, which one of the following will be your primary criterion to migrate from your current plan to the new plan: cheaper price, higher speed, or more data?

Figure 5.1 presents the survey results. For question **Q1**, 76% of the respondents express willingness to try an ad-subsidized Internet plan, either unconditionally or after a trial. Around 20% of the respondents are not interested in an ad-subsidized plan. Though the fraction of disinterested users increases by 16% in question **Q2**, 62% of the responses are positive. For question **Q3**, 57% of the respondents prefer a higher speed, another 14% opt for price discounts, and 11% are interested in data-cap upgrades. The above results clearly substantiate an interest of residential users in ad-subsidized Internet plans.

Last but not least, the ad-based revenue model for access ISPs lessens concerns about net neutrality in comparison to direct sponsorship by CSPs. Direct payments by CSPs to ISPs are viewed problematic because prominent CSPs might use them to obtain preferential delivery of their online contents. In the ad-based revenue model, ISPs make no agreements with CSPs and deal directly with advertisers or ad networks only.

5.3 Model

5.3.1 Ad revenue potential

Let n be the number of Internet users served by an access ISP, and let t be the average number of minutes spent by the users daily on the Internet. Let CPM denote the Cost-per-Mille market price of ads. Whereas the access ISP periodically sends ads to the users, we use f_a measured in impressions per minute to denote the frequency of

Table 5.1: Notations used for deriving the users-utility function.

Notations	Description
U	Utility of a user
h	Higher traffic speed
d_h	Data cap for the higher speed
l	Throttled speed after the data cap is exhausted
d_l	Data cap for the throttled speed
P	Monthly fee of the Internet plan
p	Per-unit price

the ad impressions per user. Then the average ad revenue per user (AARPU) per month (assumed to last 30 days) is:

$$r_a = 30 \cdot t \cdot f_a \cdot c_a \tag{5.1}$$

where $c_a = CPM/1000$ is the cost of a single ad impression.

Then, the total ad revenue potential for an access ISP serving the n users is expressed as:

$$R_a = n \cdot r_a. \tag{5.2}$$

Instead of using the market-driven CPM prices, the access ISP can set its own ad price to reflect its network cost. We model such ad price CPM_{ISP} as a function of its capital expenditure E_c:

$$CPM_{ISP} = \frac{1000 \cdot E_c}{30 \cdot t \cdot f_a \cdot n}. \tag{5.3}$$

5.3.2 Utility of the users

A majority of the access ISPs around the world employ the data-cap billing model [110]. As discussed in section 5.1, the data-cap billing model involves higher traffic speed h provided until monthly data cap quota d_h is consumed, after which the speed is throttled to lower rate l until the end of the billing cycle. Data cap d_l for the throttled speed depends on the number of days between the throttling moment and next billing cycle. Usually, access ISPs offer multiple tariff plans with different combinations of traffic speeds and data caps for different monthly fees P. Table 5.1 sumps up the model notations.

We model the utility of a user as a standard alpha-fair function [108]. The user derives a utility equivalent to the monthly fee of the selected Internet tariff. The utility with the higher and throttled speed respectively is captured by:

$$U_h = a_h \cdot (1-\alpha)^{-1} (h \cdot d_h)^{(1-\alpha)} - (p \cdot h \cdot d_h) \tag{5.4}$$

$$U_l = a_l \cdot (1-\beta)^{-1} (l \cdot d_l)^{(1-\beta)} - (p \cdot l \cdot d_l) \tag{5.5}$$

where $\alpha = \beta = 0.5$ denote the price sensitivity of the utility functions, constants a_h and

Table 5.2: Monthly average financial data as of March 31, 2015 for 2 Indian access ISPs.

Parameters	Notation	Units	Access ISP	
			Airtel	DEN
Service type	–	–	DSL	Cable
Users	n	Thousand	1508	23
Revenue	R	Million Rs.	1560	17.25
ARPU	r_u	Rs.	1034	750
OpEx	E_o	Million Rs.	2195	45.96
CapEx	E_c	Million Rs.	422	57.16

a_l represent the users utility levels for per-unit price p, speed limits h and l, and data caps d_h and d_l. After solving ${U_h}' = 0$ and ${U_l}' = 0$, we express a_h and a_l as follows:

$$a_h = p \cdot (h \cdot d_h)^{\alpha}, \quad a_l = p \cdot (l \cdot d_l)^{\beta}. \tag{5.6}$$

Per-unit price p is derived by computing the weighted average of monthly fee P proportionally to traffic speeds h and l and data caps d_h and d_l:

$$p = \frac{P}{h \cdot d_h + l \cdot d_l}. \tag{5.7}$$

Finally, the aggregate utility of a user for any billing month is:

$$U = U_h + U_l. \tag{5.8}$$

5.4 Data

We instantiate the model with real financial data of 2 prominent access ISPs in India. Financial data of access ISPs, such as average monthly revenues, operational expenses (OpEx), and capital expenses (CapEx) are often treated as confidential records. Nevertheless, few ISPs publish these data in their annual financial reports. We use annual reports of the two considered ISPs to obtain the financial information. Unlike the revenues separated into mobile, DSL, and other service types, the expenditures are published without breaking them into service segments. Therefore, our work utilizes the consolidated CapEx of each access ISP as a baseline for the ISP's network costs.

5.4.1 Financial data of access ISPs

Airtel [114] and Den Networks [115] are the 2 access ISPs in our study. Airtel is a large access ISP serving millions of users. Den is a medium-sized access ISP with thousands of users. Table 5.2 presents their financial data for March 2015. The revenues and costs in this chapter are quoted in Indian Rupees (Rs.). During the time of validation of this model, Rs. 66 was equal to 1 US dollar.

For each ISP, we collect the number of retail Internet users, average revenue per user (ARPU), and corresponding service type. The monthly average revenue of the ISP for

Table 5.3: CPM prices received by a popular Indian entertainment website.

Year-Month	CPM (Rs.)
2015-06	68.25
2015-05	52.65
2015-04	85.8
2015-03	86.45
2015-02	85.15
2015-01	86.45

Table 5.4: Ad revenue potential of the 2 access ISPs with $t = 60$ minutes and $f_a = 1$ impression/minute.

CPM	Units	Ad revenue potential	
		Airtel	DEN
CPM_{max}	Million Rs.	235	3.5
CPM_{avg}	Million Rs.	210	3.2
CPM_{min}	Million Rs.	143	2.2

a service type is computed by multiplying the ARPU and total number of users of the corresponding service, i.e., $R = n \cdot r_u$. Due to unavailability of the expenditure data for individual service types, we use the monthly consolidated Capex and OpEx which might include expenses of enterprise and other non-retail Internet services of the ISP. The OpEx values are computed by subtracting revenue R from EBITDA (Earning Before Interest, Tax, Depreciation and Amortization) published in the financial report.

5.4.2 Ad pricing data

The prices of online ads depend on the advertisers' budget and their willingness to pay. Also, the ad agencies can set different minimum starting prices of ads in different pricing models, such as CPM and CPC, and different formats, such as banner and video. This work uses real CPM prices of ads received by TellyReviews, a popular weekly entertainment website that hosts weekly reviews and updates of popular Indian TV shows [116]. During the time of this study, the website attracted on average 591,000 visitors and 1.6 million page views per month and received display-ads from 4 ad agencies, such as WordAds, AdSense, Gravity, and OnClickAds. Table 5.3 shows the CPM prices received by the website during the 6 months from January to June 2015. The maximum, average, and minimum CPM prices over this period are CPM_{max} = Rs. 86.45, CPM_{avg} = Rs. 77.35, and CPM_{min} = Rs. 52.65 respectively.

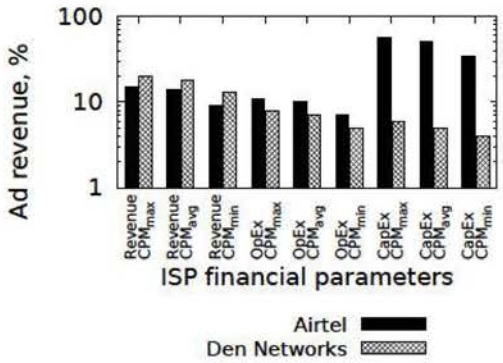

Figure 5.2: Ad revenue potential as a percentage of the average revenues, CapEx, and OpEx of the ISPs.

5.5 Analysis

5.5.1 Ad revenue potential

We estimate the ad revenue potential for Airtel and Den using equality 5.2. According to [117], the global average time spent per day by a user on the Internet is $t = 60$ minutes. With the impression frequency of 1 ad per minute, the revenue potential with CPM_{max} is estimated to be around Rs. 235 million and Rs. 3.5 million for Airtel and Den respectively. Table 5.4 presents the ad revenue potential of the 2 ISPs with CPM_{avg} and CPM_{min}.

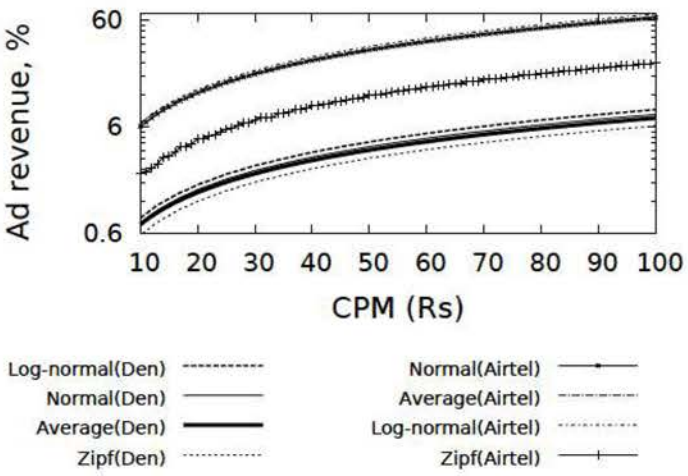

Figure 5.3: Ad revenues as a percentage of the CapEx for different CPMs and 3 distributions of online time.

To understand the significance of the ad revenue potential for the access ISPs, we compare it with the monthly average revenue, CapEx, and OpEx of the ISPs. Figure 5.2

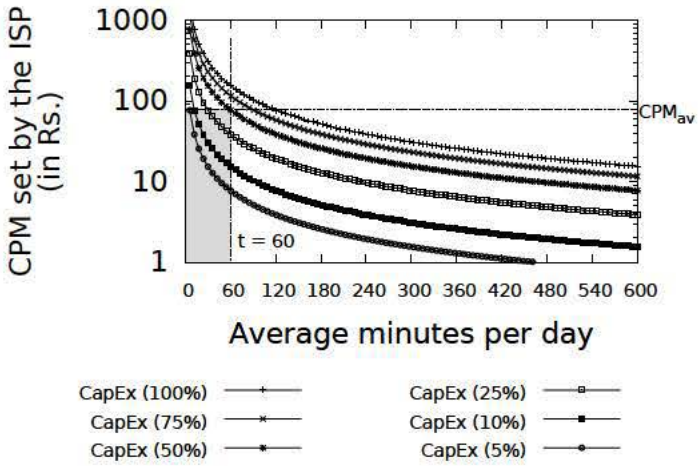

Figure 5.4: CPM variation for Airtel for different average online times and different fractions of the CapEx.

shows that Airtel's ad revenue potential ranges within 9–15% of its average revenue, 7–11% of its OpEx, and 34–56% of its CapEx. For Den, its revenue potential is 4% of its CapEx and 20% of its average revenue with CPM_{min} and CPM_{max} respectively.

While the above analysis assumes that all users spend online the same average time, this time actually varies across the users. Due to the lack of real data on the online time of the users, we employ several well-known synthetic distributions, such as normal, log-normal, and Zipf-Mandelbrot distributions, to represent the online times. Various reports published by telco operators and market-research firms estimate the average online time of Indian users to be between 20 and 360 minutes per day [118, 119, 117, 120]. We conservatively choose the online time to range between 20 and 120 minutes per day, so that the average online time is 60 minutes in all the 3 distributions.

Figure 5.3 presents the ad revenue potential of each ISP relative to its CapEx with 3 CPM prices and 3 distributions of online times. The plot also depicts the ad revenue potential when the online time kept at 60 minutes per day for all users. Except for Zipf-Mandelbrot distribution for Airtel, the ad revenue potential are qualitatively similar when the online times are constant or vary. The ad revenue potential as a percentage of the CapEx of Airtel varies from 7% at the lowest CPM price of Rs. 10 to 56% at the highest CPM price of Rs. 100.

The above analysis indicates a tangible revenue potential for access ISPs, particularly for Airtel-like large access ISPs. While a large user base is advantageous, small access ISPs can earn non-trivial ad revenues.

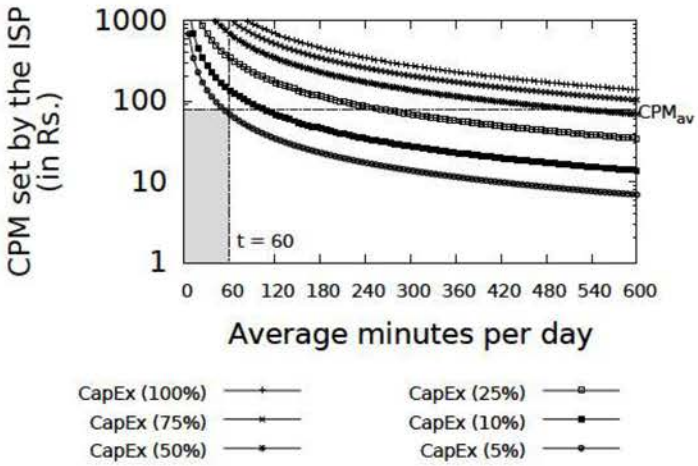

Figure 5.5: CPM variation for Den Networks for different average online times and different fractions of the CapEx.

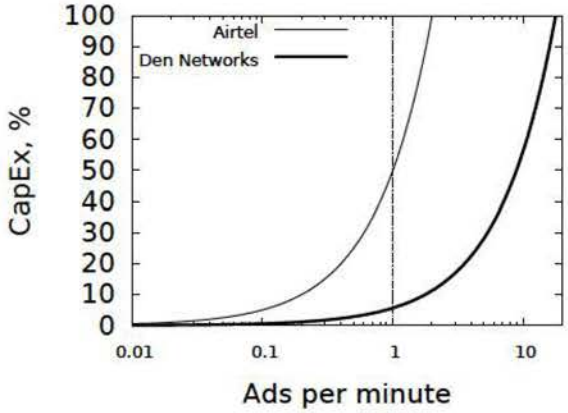

Figure 5.6: Ad revenues as a percentage of CapEx for different ad frequencies.

5.5.2 Economic viability of the ad-based revenue model

As ads are periodically displayed independently of the number of page views, the average online time of the users is a key parameter of our model. Figures 5.4 and 5.5 plot the ISP-selected CPM prices for Airtel and Den respectively as the average online time varies to cover a certain percentage of the CapEx.

In figure 5.4, we set CPM_{avg} = Rs. 77.35, $t = 60$ minutes, ad frequency $f_a = 1$ ad per minute as baselines. In order to earn ad revenues equivalent to 100% of Airtel's CapEx, the ISP-selected CPM price is Rs. 155. On the other hand, if CPM_{ISP} is equal to market average price CPM_{avg}, then the average online time should be around 121 minutes. At present, these settings are not economically feasible because the ISP-selected

Table 5.5: Speed and data prices offered by Airtel and Den Networks.

Airtel				Den Networks			
Tariff plan	Price (Rs.)	Speed h, l (Mbps)	Data-cap (GB)	Tariff plan	Price (Rs.)	Speed h, l (Mbps)	Data-cap (GB)
A1	1899	16, 0.512	80	D1	700	5, 1	30
A2	2199	16, 0.512	120	D2	1000	5, 2	50
A3	2499	16, 0.512	200	D3	1250	5, 2	100
A4	2099	24, 1	80	D4	900	20, 1	30
A5	2399	24, 1	120	D5	1300	20, 2	50
A6	2999	24, 1	200	D6	1550	20, 2	100
A7	2399	40, 1	80	D7	1200	50, 1	30
A8	2699	40, 1	120	D8	1600	50, 2	50
A9	3299	40, 1	200	D9	2100	50, 2	100

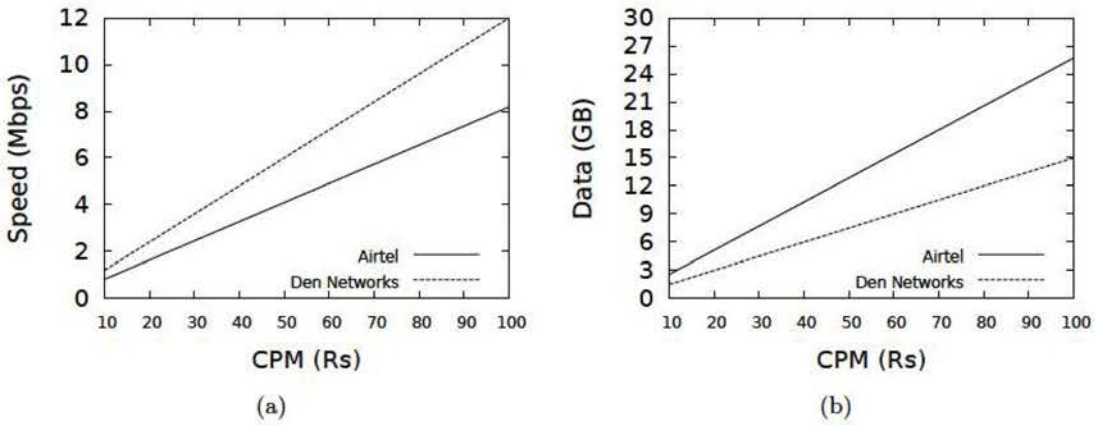

(a) (b)

Figure 5.7: User incentives that Airtel and Den can offer: (a) speed and (b) data cap.

CPM price is over 2 times higher than the average market price, and the average online time is much higher than the present global average time. Similarly, it is economically infeasible to reach ad revenues equivalent to 75% of the CapEx. The shaded area in the graph represents economically viable settings where ad revenues cover 50% of Airtel's CapEx. Similarly, figure 5.5 shows that economically viable settings allow Den to earn ad revenues covering 5.6% of its CapEx. Figure 5.6 plots the percentage of the CapEx that ad revenues can cover with different ad frequencies. For Airtel, covering 100% of the CapEx requires displaying 2 ads per minute. For Den, the respective frequency is 18 ads per minute, which is clearly not viable.

The above analyses show that the ad-based revenue model can produce higher ad revenues and can be more economically viable for larger ISPs due to lower average costs per user.

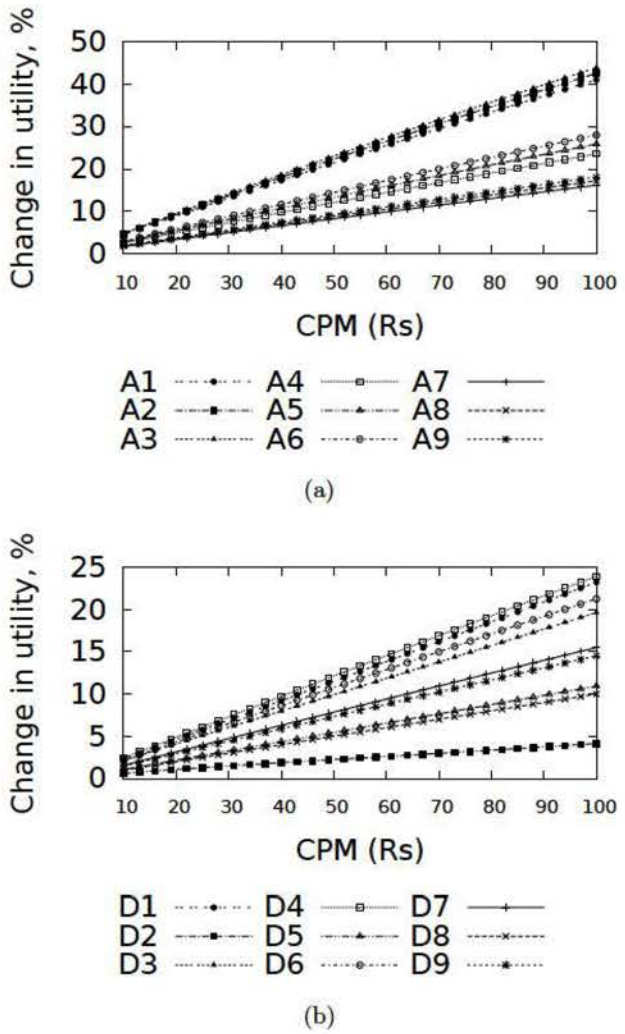

Figure 5.8: Utility changes for providing extra speed as incentives by (a) Airtel and (b) Den.

5.5.3 Incentives for users

By offsetting the ISP's expenses, the ad revenues enable the ISPs to offer attractive Internet access plans to the users and thereby increase its customer base.

The incentives can be in the form of higher access speeds or larger data caps. To quantify these potential expenses, we obtain real ISP tariff prices. Table 5.5 presents the tariff prices of Airtel and Den for different speeds, throttled speeds, and data caps. Then, we compute the average incremental price of the speed by adding all the price differences between the tariffs with the same data caps and dividing this sum by the sum of all the corresponding speed differences. Similarly, we compute the average incremental price of the data cap by adding all the price differences between the tariffs with the same speeds and dividing this sum by the sum of data-cap differences. Table 5.5 shows that average

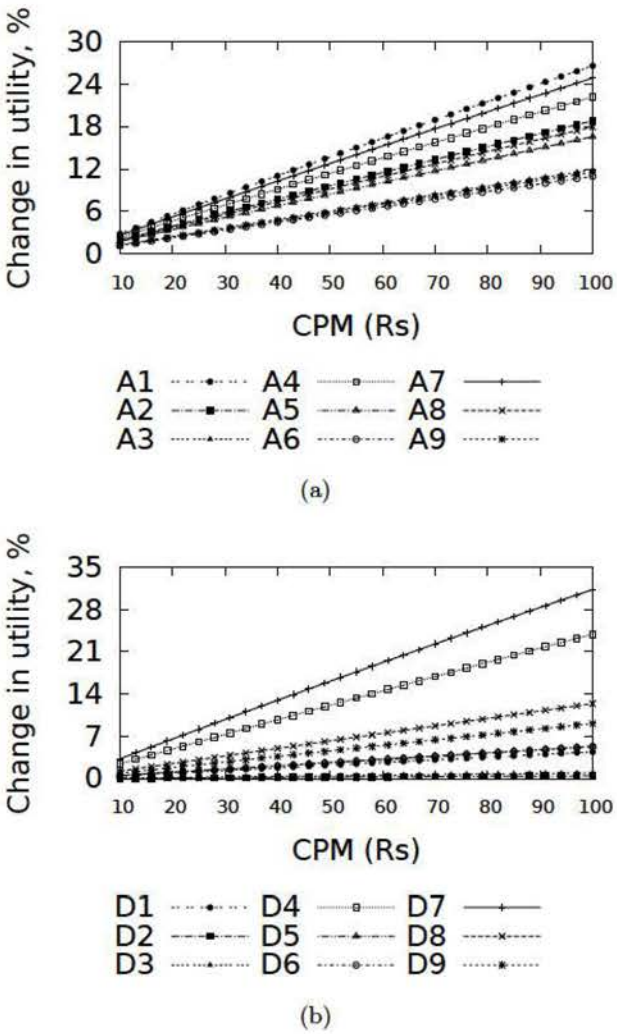

Figure 5.9: Utility changes for offering extra data as incentives by (a) Airtel and (b) Den.

incremental price of the speed and the data cap are Rs. 22 per Mbps and Rs. 7 per GB respectively for Airtel, and Rs. 15 per Mbps and Rs. 12 per GB respectively for Den.

Using the above incremental prices, we compute the speed and the data-cap incentives proportionally to the average ad revenue per user with r_a defined in equality 5.1. Airtel can offer speed incentive $\triangle h = r_a/22$ and data incentive $\triangle d_h = r_a/7$. Den can offer speed incentive $\triangle h = r_a/15$ and data incentive $\triangle d_h = r_a/12$. Figures 5.7(a) and 5.7(b) present the potential speed and data-cap incentives respectively, as CPM changes for Airtel and Den. The economically viable speed with CPM_{avg} is 6.3 Mbps and 9.3 Mbps for Airtel and Den respectively. The data-cap incentives for the same CPM price are 20 GB and 12 GB for Airtel and Den respectively. Thus, while Den can offer stronger incentives for speeds than Airtel, Airtel can offer stronger incentives for data caps.

Using equality 5.8, we analyze the impact of the incentives on the users' utility. The

utility changes are differences in the utility without and with the incentive. We compute them for all the 9 tariff plans of Airtel and Den in table 5.5. Figure 5.8(a) presents the relative utility changes for the speed as incentive of Airtel. The utility changes are highest for low-speed tariffs A1, A2, and A3 gaining up to 32-34% at CPM_{avg}. The utility changes are the smallest for high-speed tariffs A7, A8, and A9 achieving 12-14% gains at CPM_{avg}. For the Airtel tariffs with same speed, the utility gain increases with the increasing data cap. Figure 5.8(b) shows that Den can offer with CPM_{avg} the highest utility gains of 15–19% with tariffs D3, D6, D1, and D4 and the lowest gain of 3% with tariff D2. Unlike Airtel, the utility changes for the Den's tariff plans are not clustered on the basis of speed. Because Den uses smaller data caps and larger throttled speeds relative to the corresponding high-speed limits. As a result, the tariff plans with larger throttled speeds or smaller data caps offer lower utility changes.

Figure 5.9(a) depicts data cap incentives that Airtel can offer. Here, the utility gains are the highest for the tariffs with the smallest data caps across all the speed limits. Though the utility gain increases with the increasing speed among the small data-cap tariffs, the utility gain is sensitive to the throttled speeds of the tariffs. Therefore, the highest gain of 20% at CPM_{avg} is achieved by tariff A1 because of its smallest throttled speed. Tariffs A4 and A7 with similar data caps and higher throttles speeds offer smaller gains. The smallest gain at CPM_{avg} is 8% with tariff A6. Figure 5.9(b) presents the data-cap incentives for Den. At CPM_{avg}, tariffs D1, D2, and D3 offer negligible utility gains less than 1%. Tariff D7 offeres the highest gain of 24%.

The above analyses substantiate that the ad-based revenue model can enable large access ISPs to offer significant incentives to the users. Though the ad revenue potential is lower for small ISPs, the incentives can be tangible in their case as well.

5.6 Related Work

Many prior works on sponsoring ISPs costs pivot around partnerships between ISPs and CSPs. [40] proposes a sponsorship agreement between ISPs and CSPs that discounts the bandwidth costs incurred due to ads. Similarly, [41,42] analyzes the incentives derived from partnerships between ISPs and CSPs as well as implications for the users, ISPs, and CSPs. Charging of CSPs by ISPs to offset the costs of delivering content traffic is also the focus of [111,20,121,122]. Our work differs from the previous works by considering a direct business agreement between an ISP and advertisers without engaging the CSPs. Few ISPs have employed online ads to subsidize the monthly subscription fee of the users [95,96]. In these models, users are offered extra data caps in return for watching video ads. Our work, though similar in spirit, distinguishes from such models by adopting a different technique to display ads to users and we conduct extensive evaluations of ad-based model to analyze the revenue potential and economic viability conditions for large and medium-sized access

ISPs. Our evaluation reports significant benefits to the larger access ISPs.

5.7 Summary

While airports, shopping centers, and supermarkets offer ad-sponsored Internet access to promote local businesses, this work evaluates an ad-based revenue model where an access ISP earns ad revenues to offset its network expenses or to provide extra subscription incentives to the users. Seeded by real financial data of 2 prominent access ISPs in India and using real market-driven CPM prices of online ads, our analysis assesses the revenue potential and economic viability of the ad-based revenue model. We analyzed potential incentives in the form of speed, data-cap upgrades, and utility gains for the ad-subsidized Internet subscribers.

Our work demonstrates that the large ISP with millions of subscribers has a significant revenue potential up to 50% of its CapEx. The small ISP with tens of thousands of subscribers can use ads to earn up to 5% of its CapEx. Our results indicate that the ad-based revenue model might be economically viable for both large and small access ISPs. The ad revenues can enable access ISPs to offer users incentives in the form of extra 6–9 Mbps speed or 12–20 GB data cap. Our market survey shows that a significant majority of residential users are willing to consider ad-subsidized Internet plans. Besides, the ad-based model reduces concerns about net neutrality and also allows an access ISP to earn ad revenues by displaying ads with an explicit consent of the user.

Chapter 6

Conclusions and Future Work

In this thesis, we analyze various approaches that ISPs can use for earning extra revenues to economically sustain their network infrastructures. We begin our research by exploring the economics of customer-traffic attraction by transit providers to boost their transit revenues. This work reveals technical feasibility and significant revenue potential of traffic attraction despite broad countermeasures by other ISPs. Then, we consider hosting as an alternative way for transit ISPs to raise extra revenues. This measurement work reveals content hosting to be pervasive across the Internet transit hierarchy. Finally, we turn our attention to access ISPs that struggle to deal with their rising network costs in the traditional flat-rate subscription model. We evaluate an alternative revenue model where access ISPs earn ad revenues by displaying ads to the end users with their explicit consent and also offer attractive incentives for subscribing to the ad-subsidized Internet plan.

In our first work, presented in chapter 3, we examine how transit ISPs can attract transit traffic to boost their transit revenues and whether other ISPs have effective countermeasures against the traffic attraction. To deal with the vast size of the Internet ecosystem and lack of comprehensive real data, we combine available real data and traffic modeling to conduct realistic Internet-scale simulations of traffic attraction in our optimized version of C-BGP. We start our study of customer-traffic attraction by simulating the actual incident of YouTube's prefix hijacking by Pakistan Telecom in 2008. Motivated by the insights from this simulation, we conduct in-depth studies on various hypothetical scenarios of traffic attraction by means of prefix deaggregation by transit ASes from different tiers of the transit hierarchy. The prefix deaggregation causes traffic to divert to paths specified in announcements of deaggregated prefixes. We consider ASes from the top 3 tiers of the transit hierarchy to act as traffic attractors. We also enable other ASes to act against the traffic attraction through various mechanisms, such as filtering of the deaggregated prefixes, disconnections from the traffic-attraction AS, and attempts of losing ASes to counter-attract extra customer traffic to themselves.

Our simulations demonstrate that traffic attraction and reactions to it redistribute traffic on the inter-AS links and create camps of winning, losing and neutral ASes with respect to their transit revenues. The evaluation reveals that tier-1 ASes have significant financial incentives to attract traffic, with smaller benefits derived by tier-2 and tier-3 ASes. Despite various countermeasures by losing ASes, the traffic attraction via prefix deaggregation remains effective unless ASes from the winning camp cooperate with the losing ASes. Also, disconnections from the traffic-attracting AS does not completely eliminate the gains from the traffic attraction, unless a large portion of the ASes, including winners, terminate their business relationships with the attractor. Our sensitivity analysis for the topology, traffic, and pricing shows qualitatively similar results, even though the sensitivity to the traffic matrix is quantitatively smaller than to the topology and pricing. Our analysis also demonstrates that the extra router-memory costs are not an effective deterrent against traffic attraction.

This work strives to understand feasibility of earning extra transit revenues by customer-traffic attraction in the presence of countermeasures by other ISPs. While our work does not advocate, or oppose, prefix deaggregation for the traffic attraction, the results demonstrate significant financial benefits for traffic-attraction by transit providers.

Our second work presented in chapter 4 explores hosting of online contents as an alternative source of revenues for transit providers. To explore the global hosting ecosystem, we employ a novel VPN-based measurement approach to collect real online contents from top 2,165 websites in 52 countries. Then, we use a network of 22,000 open recursive DNS servers in 172 countries and 8,500 ASes to discover ASes and CDNs hosting the online contents of the 2,165 websites. Since online ads are the primary source of revenues in the online content sector, we classify online contents as ads and regular contents and analyze qualitative and quantitative differences in hosting of the two types of contents.

Our Internet-scale measurements discover a vast ecosystem of CDNs and ASes hosting online contents across the transit hierarchy and geography. The pervasive and significant hosting of online contents by transit networks raises a need for new models where the content traffic matrix includes transit ASes as sources of content traffic. Our work reveals that online contents are most commonly hosted at intermediate and edge layers, with smaller presence at the core of the transit hierarchy. Though the aggregate fraction of contents hosted at the core is smaller than at the intermediate and edge layers, the content densities in bytes per AS are the highest for the core ASes. We identify several AS clusters with different numbers of ASes and similar content-hosting characteristics. While ad contents are distributed across many smaller AS clusters, the regular contents are distributed across few bigger AS clusters. Hence, the replication is global for regular contents and local for ads. Compared to the regular contents, ad contents employ a higher number of IP addresses and ASes per website, suggesting that ads are hosted by a higher number of servers to distribute load broadly. Ads and regular contents also

differ substantially in their delivery performance. While responses to individual content requests are quicker for ads, the download time is lower for regular contents of a webpage.

While ad revenues help CSPs to produce more advanced bandwidth-intensive contents, access ISPs face rising costs to deliver the bandwidth-intensive contents. Besides, access ISPs struggle to earn sufficient revenues in the traditional flat-rate subscription model. While many access ISPs experiment with alternative pricing of end-user subscriptions, few large access ISPs demand content-based payments from CSPs or transit providers, triggering concerns about net neutrality and inter-ISP relation disputes.

In our third work presented in chapter 5, we evaluate an alternative economic model where access ISPs earn ad revenues by directly collaborating with advertisers and ad networks to display ads to end users. We assess the revenue potential and economic viability of the model for access ISPs. We also analyze incentives that access ISPs can offer ad-subsidized Internet subscribers. We evaluate our model using real financial data from 2 prominent access ISPs in India and real CPM ad prices from a popular Indian entertainment website. Our analysis demonstrates a tangible revenue potential for access ISPs up to 50% of CapEx for a large access ISP and up to 5% of CapEx for a medium-sized access ISP. If the access ISPs decide to use ad revenues for providing incentives to users, the ad revenues can support incentive up to 6–9 Mbps extra speed or 12–20 GB extra data cap. Lastly, our market survey of Indian retail customers shows their interest in trying an ad-subsidized access plan. While the model relies on direct collaboration of ISPs with advertisers without engaging CSPs, this model faces lesser concerns about net neutrality.

To reflect on the thesis as a whole, we can make the following generalizing observations. The ISPs' quest for finding alternative revenue sources increases diversity and inter-dependence of service providers in the transit, content, hosting, and access sectors. While service providers compete for higher profits, they also cooperate with each other to realize their individual objectives. The evaluation of traffic attraction demonstrates competition and cooperation among transit ISPs to attract traffic and defend against the attraction respectively. Besides, our findings shed light on understanding how vulnerable the inter-domain routing is to unorthodox announcements. Employing measurements, we demonstrate pervasive economic cooperation between CSPs and transit ISPs to host online contents. While such collaboration improves content delivery performance, it also fuels tussles of access ISPs with transit ISPs and CSPs. Several business models proposed to address such tussles, the proposals do not resolve all concerns, e.g., about net neutrality. We evaluate an ad-based revenue model with a proven track record in the content sector and demonstrate its tangible potential to be successfully leveraged by access ISPs.

By exploring different revenue sources for ISPs, this thesis paves way for future works on economic, technical, and security aspects in the Internet economic ecosystem. In the following, we briefly present our envisioned future goals stemming from this thesis.

While our traffic-attraction work extensively evaluates reactions of losing ISPs via technical means, this work does not evaluate legal repercussions. Whereas prefix deaggregation does not violate any law and is routinely used for traffic engineering, the traffic attraction via prefix deaggregation might face future legal challenges. This can be an interesting direction for future work on traffic attraction by transit providers. While several security proposals to detect deaggregation-based traffic attraction exist, none of them currently prevents traffic attraction that relies on a different technique, such as path shortening or origin AS spoofing. We envision a future security work to prevent traffic attraction regardless of the underlying attraction technique. Next, our work on exploring the content hosting provides only a single snapshot of the content-hosting ecosystem. To measure the evolution of the hosting ecosystem, we envision future work developing a fully automated tool to characterize hosting of online contents at weekly, monthly, and yearly granularities. Leveraging machine-learning methodologies, we also envision to use the real temporal data from the latter future work to develop security solutions that anticipate and prevent traffic attraction. Next, the ad-based model discussed in this thesis relies on a browser plugin to periodically fetch ads. Therefore, we plan to develop such browser plugin for ad-subsidized Internet access that preserves security and privacy of the users.

References

[1] A.T. Kearney, "Internet Value Chain Economics: Gaining a Deeper Understanding of the Internet Economy," https://www.atkearney.com/documents/10192/a70da6a8-aa98-4e43-999b-3a83a58d1c80, June 2009.

[2] DrPeering International, "Internet Transit Prices - Historical and Projected," http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php, August 2010.

[3] Cogent Communication Inc., "Contribution by Cogent Communications On DGCONNECTs Public Consultation on Specific Aspects of Transparency, Traffic Management and Switching in an Open Internet," *Digital Agenda for Europe, 2012.*

[4] H. Chang, S. Jamin, and W. Willinger, "To Peer or not to Peer: Modeling the Evolution of the Internet's AS-level Topology," *INFOCOM 2006.*

[5] A. Dhamdhere and C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh," *CoNext 2010.*

[6] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?" *PAM 2008.*

[7] L. Ramaswamy, L. Liu, and A. Iyengar, "Cache Clouds: Cooperative Caching of Dynamic Documents in Edge Networks," *ICDCS 2005.*

[8] TeleGeography, "IP Transit Revenues, Volumes Dependent on Peering Trends," https://www.telegeography.com/press/press-releases/2014/07/08/ip-transit-revenues-volumes-dependent-on-peering-trends/index.html, July 2014.

[9] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy," *TPRC 2007.*

[10] DrPeering, "The Great Remote Peering Debate," http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2012/9/18_The_Great_Remote_Peering_Debate.html, September 2012.

[11] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois, "Remote Peering: More Peering without Internet Flattening," *CoNext 2014.*

[12] I. Castro and S. Gorinsky, "T4P: Hybrid Interconnection for Cost Reduction," *NetEcon 2012.*

[13] I. Castro, R. Stanojevic, and S. Gorinsky, "Using Tuangou to Reduce IP Transit Costs," *IEEE/ACM Transactions on Networking*, October 2014.

[14] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, "How Many Tiers? Pricing in the Internet Transit Market," *SIGCOMM 2011*.

[15] The NANOG Archives, http://mailman.nanog.org/pipermail/nanog/2012-May/048564.html, May 2012.

[16] The NANOG Archives, http://mailman.nanog.org/pipermail/nanog/2011-December/043036.html, December 2011.

[17] MLab, "ISP Interconnection and its Impact on Consumer Internet Performance," https://www.measurementlab.net/publications/M-Lab_Interconnection_Study_US.pdf, October 2014.

[18] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate ~ A White Paper," *ACM SIGCOMM Computer Communication Review*, January 2007.

[19] K. Zhu, "Bringing Neutrality to Network Neutrality," *Berkeley Technology Law Journal*, January 2007.

[20] N. Economides and J. Tg, "Network neutrality on the Internet: A Two-sided Market Analysis," *Information Economics and Policy*, August 2012.

[21] arstechnica, "Peering problems: digging into the Comcast/Level 3 grudgematch," http://arstechnica.com/tech-policy/2010/12/comcastlevel3/, December 2010.

[22] L. Li and C. Chen, "Exploring Possible Strategies for Competitions between Autonomous Systems," *ICC 2008*.

[23] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright, "Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP," *SIGCOMM 2008*.

[24] P. Bangera and S. Gorinsky, "Traffic Attraction by Internet Transit Providers: An Economic Perspective," *Networking 2014*.

[25] P. Bangera and S. Gorinsky, "An Economic Perspective on Traffic Attraction by Internet Transit Providers," *ICNP 2013*.

[26] P. Bangera and S. Gorinsky, "Impact of Prefix Hijacking on Payments of Providers," *COMSNETS 2011*.

[27] Hinden, R. and Deering, S., "Internet Protocol Version 6 (IPv6) Addressing Architecture," *RFC 3513*, April 2003.

[28] Y. Rekhter and T. Li, "A Border Gateway Protocol (BGP-4)," *RFC 1771*, March 1995.

[29] L. Cittadini, W. Muhlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel, "Evolution of Internet Address Space Deaggregation: Myths and Reality," *IEEE Journal on Selected Areas in Communications*, October 2010.

[30] H. Berkowitz, E. Davies, S. Hares, P. Krishnaswamy, and M. Lepp, "Terminology for Benchmarking BGP Device Convergence in the Control Plane," *RFC 4098*, June 2005.

[31] T. M. Thomas, D. Pavlichek, L. H. Dwyer, R. Chowbay, W. W. Downing, and J. Sonderegger, "Juniper Networks Reference Guide: JUNOS Routing, Configuration and Architecture," August 2002.

[32] M. Caesar and J. Rexford, "BGP Routing Policies in ISP Networks," *IEEE Network*, November 2005.

[33] "The Folly of Peering Ratios (as a Peering Candidate Discriminator)," http://drpeering.net/white-papers/The-Folly-Of-Peering-Ratios.html.

[34] S. Hasan and S. Gorinsky, "Obscure Giants: Detecting the Provider-Free ASes," *Networking 2012.*

[35] X. Dimitropoulos, P. Hurley, A. Kind, and M. Stoecklin, "On the 95-Percentile Billing Method," *PAM 2009.*

[36] R. Stanojevic, N. Laoutaris, and P. Rodriguez, "On Economic Heavy Hitters: Shapley value analysis of the 95th-percentile pricing," *IMC 2010.*

[37] B. Krishnamurthy, C. Wills, and Y. Zhang, "On the Use and Performance of Content Distribution Networks," *IMW 2001.*

[38] G. Pallis and A. Vakali, "Insight and Perspectives for Content Delivery Networks," *Communications of the ACM*, January 2006.

[39] D. Katabi and J. Wroclawski, "A Framework for Scalable Global IP-anycast (GIA)," *ACM SIGCOMM Computer Communication Review*, October 2000.

[40] M. Andrews, U. Ozen, M. I. Reiman, and Q. Wang, "Economic Models of Sponsored Content in Wireless Networks with Uncertain Demand," *SDP 2013.*

[41] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring Mobile Data: An Economic Analysis of the Impact on Users and Content Providers," *INFOCOM 2015.*

[42] L. Zhang and D. Wang, "Sponsoring Content: Motivation and Pitfalls for Content Service Providers," *SDP 2014.*

[43] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *SIGCOMM 2007.*

[44] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How Secure are Secure Interdomain Routing Protocols," *SIGCOMM 2010.*

[45] O. Nordstrom and C. Dovrolis, "Beware of BGP Attacks," *ACM SIGCOMM Computer Communication Review*, April 2004.

[46] P. Gill, M. Schapira, and S. Goldberg, "Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security," *SIGCOMM 2011.*

[47] H. Levin, M. Schapira, and A. Zohar, "Interdomain Routing and Games," *STOC 2008.*

[48] C. Kalogiros, M. Bagnulo, and A. Kostopoulos, "Understanding Incentives for Prefix Aggregation in BGP," *ReArch 2009.*

[49] "RIPE RIS." [Online]. Available: http://www.ripe.net/ris/

[50] "YouTube Hijacking: A RIPE NCC RIS Case Study," http://www.ripe.net/news/study-youtube-hijacking.html, February 2008.

[51] "NANOG Mailing List," http://www.nanog.org/mailinglist/, 2008.

[52] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," *USENIX-SS 2006*.

[53] "BGPmon." [Online]. Available: http://bgpmon.net/

[54] B. Quoitin and S. Uhlig, "Modeling the Routing of an Autonomous System with C-BGP," *IEEE Network*, November 2005.

[55] L. Gao and J. Rexford, "Stable Internet Routing Without Global Coordination," *IEEE/ACM Transactions on Networking*, December 2001.

[56] CAIDA: Cooperative Association for Internet Data Analysis, http://as-rank.caida.org/.

[57] "Oops, Pow, Surprise...24 Hours of Video All Up in Your Eyes!" http://youtube-global.blogspot.com/2010/03/oops-pow-surprise24-hours-of-video-all.html, March 2010.

[58] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," *SIGCOMM 2012*.

[59] Internet Research Lab, UCLA, http://irl.cs.ucla.edu/.

[60] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-domain Traffic," *SIGCOMM 2010*.

[61] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," *IMC 2005*.

[62] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An Agent-based Model of Interdomain Network Formation, Traffic Flow and Economics," *INFOCOM 2012*.

[63] SCM Repositories – C-BGP, http://c-bgp.svn.sourceforge.net/viewvc/c-bgp/trunk/src/bgp/as.c?revision=1327&view=markup.

[64] Cisco Visual Networking Index: Forecast and Methodology, 2010-2015., http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf, June 2011.

[65] CIDR Report, http://www.cidr-report.org/as6447/bgp-originas.html, April 2013.

[66] Alexa Web Information Service, http://aws.amazon.com/awis.

[67] K. Fall, P. B. Godfrey, G. Iannaccone, and S. Ratnasamy, "Routing Tables: Is Smaller Really Much Better?" *HotNets 2009*.

[68] RIBs and FIBs (aka IP routing table and CEF table), http://blog.ioshints.info/2010/09/ribs-and-fibs.html.

[69] Hurricane Electric Internet Services, http://lg.he.net/.

[70] New York State Office of General Services, http://www.ogs.state.ny.us/purchase/prices/7701821350PL_Juniper.pdf.

[71] Juniper Networks, T-series Hardware EOS Dates and Milestones, http://www.juniper.net/support/eol/tseries_hw.html.

[72] S. Kandula, D. Katabi, B. Davie, and A. Charny, "Walking the Tightrope: Responsive Yet Stable Traffic Engineering," *SIGCOMM 2005*.

[73] C. McArthur and M. Guirguis, "Stealthy IP Prefix Hijacking: Don't Bite Off More Than You Can Chew," *GLOBECOM 2009*.

[74] K. Zhang, X. Zhao, and S. Wu, "An Analysis on Selective Dropping Attack in BGP," *IPCCC 2004*.

[75] A. Lutu, M. Bagnulo, and R. Stanojevic, "An Economic Side-Effect for Prefix Deaggregation," *NetEcon 2012*.

[76] S. Hasan, S. Gorinsky, C. Dovrolis, and R. K. Sitaraman, "Trade-offs in Optimizing the Cache Deployments of CDNs," *INFOCOM 2014*.

[77] IAB, "IAB Internet Advertising Revenue Report, 2015 Full Year Results," http://www.iab.com/wp-content/uploads/2016/04/IAB-Internet-Advertising-Revenue-Report-FY-2015.pdf, April 2016.

[78] HMA, "Pro VPN, 2013." http://www.hidemyass.com/vpn/.

[79] ALEXA, "Top Sites By Country, 2013." http://www.alexa.com/topsites/countries.

[80] M. Butkiewicz, H. V. Madhyastha, and V. Sekar, "Understanding Website Complexity: Measurements, Metrics, and Implications," *IMC 2011.*

[81] C. Huang, A. Wang, J. Li, and K. W. Ross, "Measuring and Evaluating Large-scale CDNs," *IMC 2008.*

[82] L. Yuan, C.-C. Chen, P. Mohapatra, C.-N. Chuah, and K. Kant, "A Proxy View of Quality of Domain Name Service, Poisoning Attacks and Survival Strategies," *ACM Transaction on Internet Technology*, May 2013.

[83] Anonymous, "The Collateral Damage of Internet Censorship by DNS Injection," *ACM SIGCOMM Computer Communication Review*, July 2012.

[84] Team Cymru Community Services, http://www.team-cymru.org/Services/ip-to-asn.html.

[85] MaxMind: GeoIP Databases & Services, https://www.maxmind.com/en/geoip2-precision-country-service.

[86] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM SIGCOMM Computer Communication Review*, January 2007.

[87] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Transactions on Networking*, February 2010.

[88] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate: a White Paper," *ACM SIGCOMM Computer Communication Review*, January 2007.

[89] J. Musacchio, G. Schwartz, and J. Walrand, "A Two-Sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue," *Review of Network Economics*, March 2009.

[90] S. Jordan, "Implications of Internet Architecture on Net Neutrality," *ACM Transaction on Internet Technology*, May 2009.

[91] A. Dhamdhere and C. Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem," *IEEE/ACM Transactions on Networking*, October 2011.

[92] TeleGeography, "Global Internet Geography Executive Summary," http://www.telegeography.com/page_attachments/products/website/research-services/global-internet-geography/0004/1851/GIG_Executive_Summary.pdf, 2013.

[93] arstechnica, "Peering problems: digging into the Comcast/Level 3 grudgematch," http://arstechnica.com/tech-policy/2010/12/comcastlevel3/, December 2010.

[94] The Economist, "An ad-block shock: France v Google," http://www.economist.com/news/business/21569414-xavier-niel-playing-rough-internet-giant-france-v-google, January 2013.

[95] Ovivo, http://www.ovivomobile.com, 2013.

[96] Samba, http://www.sambamobile.com/Home/Samba, 2013.

[97] P. Gill, M. Arlitt, N. Carlsson, A. Mahanti, and C. Williamson, "Characterizing Organizational Use of Web-Based Services: Methodology, Challenges, Observations, and Insights," *ACM Transaction on the Web*, October 2011.

[98] H. Khandelwal, F. Hao, S. Mukherjee, R. Kompella, and T. Lakshman, "CobWeb: In-network Cobbling of Web Traffic," *Networking 2012.*

[99] S. Ihm and V. S. Pai, "Towards Understanding Modern Web Traffic," *IMC 2011.*

[100] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. Bustamante, "Drafting Behind Akamai: Inferring Network Conditions Based on CDN Redirections," *IEEE/ACM Transactions on Networking*, December 2009.

[101] S. Triukose, Z. Wen, and M. Rabinovich, "Measuring a Commercial Content Delivery Network," *WWW 2011.*

[102] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Web Content Cartography," *IMC 2011.*

[103] I. N. Bermudez, M. Mellia, M. M. Munafo, R. Keralapura, and A. Nucci, "DNS to the Rescue: Discerning Content and Services in a Tangled Web," *IMC 2012.*

[104] B. Krishnamurthy and C. E. Wills, "Cat and Mouse: Content Delivery Tradeoffs in Web Access," *WWW 2006.*

[105] Y. Wang, D. Burgener, A. Kuzmanovic, and G. Macia-Fernandez, "Understanding the Network and User-Targeting Properties of Web Advertising Networks," *ICDCS 2011.*

[106] N. Vallina-Rodriguez, J. Shah, A. Finamore, Y. Grunenberger, K. Papagiannaki, H. Haddadi, and J. Crowcroft, "Breaking for Commercials: Characterizing Mobile Advertising," *IMC 2012.*

[107] G. Kesidis, A. Das, and G. de Veciana, "On Flat-rate and Usage-based Pricing for Tiered Commodity Internet Services," *CISS 2008.*

[108] P. Hande, M. Chiang, R. Calderbank, and J. Zhang, "Pricing under Constraints in Access Networks: Revenue Maximization and Congestion Management," *INFOCOM 2010.*

[109] A. M. Odlyzko, "Internet Pricing and The History of Communications," *Computer Networks*, August 2001.

[110] K. Poularakis, I. Pefkianakis, J. Chandrashekar, and L. Tassiulas, "Pricing The Last Mile: Data Capping For Residential Broadband," *CoNext 2014.*

[111] P. Hande, M. Chiang, R. Calderbank, and S. Rangan, "Network Pricing and Rate Allocation with Content Provider Participation," *INFOCOM 2009.*

[112] NDTV, "Flipkart Pulls Out of Airtel Deal Amid Backlash Over Net Neutrality," http://www.ndtv.com/india-news/flipkart-pulls-out-of-airtel-deal-amid-backlash-over-net-neutrality-754829, April 2015.

[113] D. DAS, "Airtel & MTNL injecting ads into websites you visit," http://www.techworm.net/2015/07/airtel-mtnl-injecting-ads-into-websites.html, July 2015.

[114] Bharti Airtel Limited, "Quarterly report on the results for the fourth quarter and year ended March 31, 2015," http://www.airtel.in/wps/wcm/connect/59d8b798-3a91-4818-b82e-dff388585829/Bharti+Airtel+Limited_Quarterly+Report_March+31-+2015.pdf?MOD=AJPERES&ContentCache=NONE, April 2015.

[115] "Den Networks Limited, Annual Report 2014-15," http://www.dennetworks.com/common/file/Annual-Report/Annual%20Report%202014-15.pdf, May 2015.

[116] "TellyReviews: The weekly reviews of your favorite Telly shows," http://tellyreviews.com/.

[117] McKinsey&Company, "Online and Upcoming: The Internet's Impact on India," http://www.mckinsey.com/-/media/mckinsey%20offices/india/pdfs/online_and_upcoming_the_internets_impact_on_india.ashx, December 2012.

[118] TATA Communications, "Connected World II: Where does the Internet come from?" http://www.tatacommunications.com/sites/default/files/ConnectedWorldII.pdf, October 2014.

[119] Ericsson, "Ericsson Mobility Report India," http://www.ericsson.com/res/docs/2015/ericss-on-mobility-report-june-2015-rina-appendices.pdf, June 2015.

[120] comScore, Inc., "India Digital Future In Focus 2013," http://www.comscore.com/content/download/21739/1126111/file/India-Digital-Future-in-Focus-2013.pdf, August 2013.

[121] L. Waverman, "Two-Sided Telecom Markets and the Unintended Consequences of Business Strategy," *Competition Policy International*, February 2007.

[122] L. He and J. Walrand, "Pricing and Revenue Sharing Strategies for Internet Service Providers," *INFOCOM 2005.*