The Internet Ecosystem: a Traffic Measurement Analysis

Manuel Palacin

UPF PHD THESIS / YEAR 2015

SUPERVISOR
Miquel Oliver
Department of Information and Communication Technologies

Universitat Pompeu Fabra Barcelona

”The only predictable thing about the Internet is that it is unpredictable”
Leonard Kleinrock

Agradecimientos

En estas líneas quería dar las gracias a todas las personas que me han ayudado y me han dado su apoyo durante todos estos años. En primer lugar quería agradecer a mi supervisor de Tesis, Dr. Miquel Oliver, por todo su esfuerzo y dedicación. Sus consejos, conocimientos y sobretodo su paciencia han sido fundamentales en mi formación como investigador. De igual manera, quería hacer una mención especial tanto a Alex Bikfalvi como a Jorge Infante, a quienes profeso una gran admiración y de quienes he aprendido muchísimas cosas.

Tambien quería dar las gracias a los revisores, miembros del tribunal y a todos mis compañeros y excompañeros de departamento por aguantarme durante todos estos años: Albert Domingo, Luís Sanabria, Boris Bellalta, Javier González, Jaume Barceló, Simon Oechsner, Gabriel Martins, Geordie George, Ratheesh Mungara, Pascal Landry, Nikos Makriyannis, Matteo Signorini Albert Bel, Toni Adame, Ruizhi Liao (Super Ray para los amigos), Trang Cao, Sougata Pal, Davinia Hernández, Joan Melià, Núria Garcia, Aurelio Ruíz, Anna Sfairopoulou, Cristina Cano, Michail Batikas, Ayman Moghnieh, Miguel Ángel Carralero, Stefan Bott, Azadeh Faridi, Maddalena Nurchis, Iván Fernández, Pedro Vílchez, Iván Jiménez, Javier Ortega, a toda la gente del CAU, a Susana Domingo, Rosa Barragán, Sarah Collins, a todos mis compañeros de los cursos de idiomas, a toda la gente de secretaría DTIC, Lydia García, Jana Safrankova, Bea Abad, Judith Champion, Vanessa Jiménez, Joana Clotet y a todos mis compañeros de los diferentes proyectos europeos en los que he participado. Gracias tambien a la Universitat Pompeu Fabra por brindarme la oportunidad de poder contribuir con mi investigación y trabajo.

Por último, pero no por ello menos importantes, quería dar las gracias a mis padres Maribel y Antonio por apoyarme de manera incondicional en cada una de mis decisiones, a toda mi familia y a los que ya no estan, y a todos mis amigos a los cuales he dejado de ver en alguna que otra ocasión porque el deber no me lo permitía. Pero sobretodo quería dar las gracias a una persona muy especial que ha sido mi fuente de inspiración y sin la cual no habría superado este desafío: TQ Andrea ;-)!

Esta tesis doctoral ha sido posible gracias a:

- Ministerio de Ciencia e Innovación del Gobierno de España mediante el programa GEPETO (TEC2008-06055/TEC)
- Agència de Gestió d'Ajuts Universitaris i de Recerca de la Generalitat de Catalunya mediante el programa AGAUR (2009 SGR 617)
- Ministerio de Economía y Competitividad del Gobierno de España mediante el programa CISNETS (TEC2012-32354)
- Universitat Pompeu Fabra Department of Information and Communication Technologies Doctoral Scholarship
- OpenCities - Open Innovation mechanisms for Smart Cities, CIP-ICT-PSP-2010-4-270896.

Abstract

The analysis of the interconnection status quo between content providers and the different networks of the Internet Service Providers (ISPs) that transport Internet traffic to end users is essential in better understanding the evolution of the Internet ecosystem. In the recent years, we have witnessed a spectacular increase in the consumption of Internet traffic, especially multimedia content, which has driven both content providers and network operators, including transit, access ISPs and third party CDNs, to rethink their interconnection models.

Internet video services demand particular network characteristics for which the original Internet was not initially developed. During the last twenty years of commercial Internet, the research community has contributed multitudes of hardware and software improvements. From a hardware perspective, user devices have augmented their processing power to unimaginable limits, while prices have been placed within the reach of the vast majority. Meanwhile, Internet networks have increased their capacity and provided the gift of ubiquity thanks to wireless technologies. From a software perspective, Internet applications have motivated the appearance of new hardware and network protocols because of their requirement to provide innovative user experiences. Moreover, virtualisation and cloud technology have dramatically decreased the servers costs in facilitating the proliferation of new services. Surprisingly and despite the implementation of optimal and specific Internet protocols by the research community, the industry has opted for using existing robust and flexible protocols to transport their low-latency Internet content, HTTP over TCP. Certainly, current TCP and HTTP are not the same as the originals. However, they maintain their original essence and provide interoperability among millions of devices around the world. In addition, these protocols of the TCP/IP suite have a strategic advantage over others: they are an intrinsic part of the operating system and Internet browsers are present in almost all devices. Therefore, technology has evolved and matured when necessary, and we currently observe a scenario where new applications and new hardware solutions (networks and devices) fit perfectly into a stable

environment (HTTP over TCP). Now, therefore, we wonder if the evolution of Internet traffic has influenced the interconnection structure of the Internet.

Internet players have also evolved the ways in which they interact with the rest of the ecosystem because of the increase in Internet demand, and this fact has impacted the Internet topology. Originally, the Internet was configured by a set of operators organised in a multi-tier hierarchic structure, where users and content were located at the bottom. When they needed to reach each other, it was necessary to go up and down the different tiers of the hierarchy. The network expansion of access ISPs and large content providers motivated by the growth of Internet traffic caused an evolution of the Internet landscape to a more meshed model bringing end users closer to content. In this context, CDNs emerged as a technical solution to deal with the massive demands of popular content and to provide reliable services for low latency requirements. Many global content providers deploy their own CDNs, while others prefer to contract third parties with the objective of being closer to the end users. Moreover, the introduction of paid peering as an interconnection agreement for ISPs that want to terminate the traffic exchange in an access ISP has revolutionised the market, as peering was founded on a settlement-free basis. After this point, transit ISPs began to claim because they were required to pay to deliver traffic in an end network when traditionally the economic transaction was in the other direction. This new situation, in which employing CDNs or establishing peering agreements redefined the content delivery models, forced transit ISPs to reinvent themselves. Transit ISPs needed to adapt by introducing new services to their catalogue and addressing new market segments because global content providers began to bypass them to reach the access ISPs. All of these changes raise the question of what interconnection differences exist between different ISPs with different roles in the Internet ecosystem.

In this doctoral thesis, we perform an extensive traffic analysis from two perspectives to better understand the rationale behind the different Internet players. First, we analyse the Internet traffic from the perspective of the evolution of the Internet protocols. In analysing the protocols we

attempt to observe whether the traffic pattern has changed as new applications have emerged and the Internet demand have exploded. Second, we collect a data set of Internet traces that allows us to evaluate the connectivity between access ISPs and the most popular content providers. By analysing the Internet traces, we want to identify the differences and correlations in the interconnection models used by different Internet players.

The main contributions of this work are the empirical corroboration that Internet traffic is mostly formed by HTTP over TCP, two robust and flexible protocols that have found their best ally in the Web, and that from the perspective of web content, the Internet market provides sufficient interconnection alternatives for the different profiles of Internet players. Accordingly, after twenty years of commercial Internet, we confirm that the content delivery market is sustainable and offers plenty of opportunities.

Resumen

El análisis del statu-quo de las interconexiones entre los proveedores de contenidos y las diferentes redes de los operadores de Internet (en inglés ISPs) que transportan el tráfico de Internet hacia los usuarios es esencial para entender mejor la evolución del ecosistema de Internet. En los últimos años hemos sido testigos de un espectacular crecimiento en el consumo de Internet, especialmente contenidos multimedia, que ha llevado tanto a los proveedores de contenidos como a los operadores de red a replantearse sus modelos de interconexión.

Los servicios de video por Internet exigen unas características de red particulares para las cuales Internet no fue inicialmente diseñado. En los últimos 20 años de existencia del Internet comercial, la comunidad científica ha contribuido generando una gran cantidad de avances tanto en hardware como en software. Desde un punto de vista del hardware, los dispositivos que utilizan los usuarios han incrementado su potencia de procesado hasta límites inimaginables a la vez que sus precios se han puesta al alcance de la gran mayoría. De la misma forma, las redes de comunicaciones también han aumentado su capacidad de transmisión y nos han brindado el don de la ubicuidad gracias a los avances en tecnologías inalámbricas (redes GSM y WiFi). Desde una perspectiva del software, las aplicaciones de Internet han forzado la aparición de nuevo hardware y nuevos protocolos de red que permiten poder ofrecer innovadoras experiencias de usuario. Además, las tecnologías de virtualización y computación en la nube (cloud computing) han reducido espectacularmente los costes de servidores e infraestructura de red favoreciendo la aparición de nuevos servicios. Sin embargo y a pesar de la aparición de nuevos protocolos de Internet más eficientes, la industria ha optado por el uso de protocolos de Internet maduros como son TCP y HTTP que pese a que originalmente no estaban destinados para el transporte de tráfico multimedia, gracias a su robustez y flexibilidad dan un resultado óptimo en redes que ofrecen baja latencia. Si bien es cierto que las versiones actuales de TCP y HTTP no son las mismas que las originales, éstas continúan manteniendo la misma esencia y proporcionan interoperabilidad entre millones de dispositivos alrededor

del mundo. Por otro lado, estos protocolos de la suite TCP/IP disponen de una ventaja estratégica respecto a otros: forman parte de las funciones de red del sistema operativo y se ven beneficiados con la presencia de navegadores web en la mayoría de dispositivos. Por lo tanto, la tecnología ha evolucionado y madurado cuando ha sido necesario y esto nos da pie a que actualmente observemos un escenario donde las nuevas aplicaciones y soluciones hardware (tanto redes como dispositivos) encajan perfectamente en un entorno estable (HTTP sobre TCP). En este contexto nos preguntamos si la evolución del tráfico ha influenciado en la estructura de interconexión de Internet.

Los operadores de Internet también han evolucionado la manera de interactuar con el resto del ecosistema debido al incremento de la demanda de Internet. Este hecho ha tenido un impacto significativo en la topología de Internet. Originalmente, estaba configurado por un conjunto de operadores organizados siguiendo una estructura jerárquica en la cual tanto los usuarios como los contenidos estaban situados en la zona inferior de esta estructura y necesitaban escalar y descender la jerarquía para ponerse en contacto. La expansión de las infraestructuras de red de los operadores de acceso y grandes proveedores de contenidos fue motivada por el crecimiento significativo del tráfico de Internet y esto provocó la evolución de la estructura de Internet hacia un modelo más mallado que acercaba los contenidos a los usuarios. En este escenario aparecieron las CDNs (redes de distribución de contenidos) como solución técnica para tratar con la demanda masiva de contenidos de Internet ofreciendo un servicio de altas prestaciones en entornos que requieren baja latencia. Muchos proveedores de contenidos de ámbito global han decidido desplegar sus propias redes CDN, mientras que otros han preferido subcontratar estos servicios a operadores terceros con tal de estar más cerca de los usuarios finales. Además, la introducción del paid peering como método de interconexión para aquellos operadores que necesitan terminar la entrega de tráfico en un operador de acceso ha revolucionado el mercado ya que el este tipo de relaciones originalmente se acordaban sin contraprestación económica entre ninguna de las partes involucradas. A raíz de la aparición de esta nueva modalidad, los operadores de Internet encargados de ofrecer el servicio de tránsito empezaron a mostrar su preocupación ya que a partir de entonces fueron obligados en algunos casos a pagar un peaje por la entrega de tráfico cuando tradicionalmente las transacciones económicas se efectuaban en sentido inverso. Esta nueva situación donde proliferan el uso de redes CDN o el establecimiento de acuerdos de peering ha redefinido los modelos de distribución de tráfico de Internet y a la vez ha forzado a diferentes operadores a reinventarse. Los operadores de tránsito necesitaron adaptarse a las nuevas amenazas proporcionando un nuevo catalogo de servicios dirigido a un nuevo segmento de mercado debido a que los grandes proveedores de contenidos empezaron a omitirlos a la hora de conectarse con las redes de acceso. Todos estos cambios nos hacen preguntarnos cuales son los diferentes métodos de interconexión entre los diferentes operadores que forman el ecosistema de Internet.

En esta tesis doctoral se realiza un extenso análisis del tráfico de Internet des de dos puntos de vista con el objetivo de entender mejor la razón de ser de cada uno de los principales operadores de Internet. Primero, se analiza el tráfico de Internet desde un punto de vista de los protocolos. Gracias al análisis de los protocolos se pretende observar cómo ha cambiado el patrón de tráfico debido a la irrupción de nuevas aplicaciones y al incremento de la demanda de Internet. Segundo, se ha recogido un conjunto de muestras con trazas de Internet que nos permite evaluar la conectividad entre diferentes operadores de acceso y proveedores de contenidos. El análisis de estas trazas nos permite identificar las correlaciones entre los diferentes modelos de interconexión que utilizan los operadores. Las principales contribuciones de este trabajo son la corroboración empírica de que el tráfico que circula actualmente por Internet está principalmente formado por HTTP sobre TCP, dos protocolos robustos y flexibles que han encontrado en la Web su principal aliado, y la constatación de que el mercado de Internet proporciona suficientes alternativas de interconexión para los diferentes perfiles de operadores de Internet. Por lo tanto, podemos confirmar que después de veinte años de Internet comercial, el mercado de distribución de contenidos es sostenible y está lleno de oportunidades.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

1.1 Background

The Internet is very much alive as can be observed in the interconnection agreements that Internet Service Providers (ISPs) are continuously arranging. We define the Internet interconnection ecosystem as a set of business entities or organisations (called players or actors throughout the text) specialised in covering different technical needs and providing Internet services. We identify players specialised in generating Internet content: the content providers (CPs), which include over-the-top (OTT) providers, such as Google, Yahoo, Microsoft, Amazon, Facebook, Netflix, BBC or any other type of content-based service (e.g. web sites); players specialised in providing Internet connectivity to the end users, the access ISPs, which include both small and large operators, the latter including incumbent operators also known as SMP (significant market power) operators such as Telefonica, BT, Orange, Vodafone, Deutsche Telekom or AT&T; players specialised in transporting aggregated Internet traffic between networks, the tier-1s or transit carriers such as Level3 or Cogent; players specialised in optimising the distribution of Internet content using cache servers, the content delivery/distribution networks (CDNs) such as Akamai, Limelight or Edgecast; players specialised in providing connectivity and storage for hosting Internet content to those

CPs without network infrastructure such as Peer1, OVH or 1and1; and entities specialised in providing locations or neutral points to facilitate the interconnection between different players, the Internet eXchange Points (IXPs), also known as Network Access Points (NAPs). Although this enumeration aims to simplify the different main activities of the Internet players, it seems reasonable that most of these actors usually play more than one role and that their business activities evolve continuously according to their market opportunities and necessities. In addition, any actor can play different roles depending on the geographic area we analyse (e.g. in Spain, Telefonica offers Internet subscriptions to end-users, video content through their IP-TV platform, CDN and hosting services, and wholesale transit services through their backbone). Therefore, the Internet topology is geography dependent because not all actors are organised in the same way to cover national markets. Small markets tend to have fewer competitors for each role, with players even assuming more than one activity. Larger markets allow more competitors for each role and a greater dynamism in terms of agreements and acquisitions.

In Figure 1.1, we observe a snapshot of a set of Internet players interacting with each other and the users in the centre of the semi-circle. The different border,lines separate the different Internet players while the dotted lines separate different roles of the same Internet player. Figure 1.1 shows that the end users are surrounded by a set of access ISPs (ACCESS), and we then find the CDNs (SP CDNs are specialised CDNs such as Akamai, which mainly focuses on the CDN sector), hosting companies (HOST) and tier-1 ISPs. In some intersections, we find the IXPs that interconnect different ISPs. Finally, at the outer part of the semi-circle, we find the content providers. The diagram shows how different actors play different roles (dotted lines), e.g. at the bottom-left part of the figure, we can see an operator (e.g., Telefonica or AT&T) that performs all business activities, including access, tier-1 (international carrier), CDN, hosting, and content production (IP-TV). We also see another example in the bottom-right part of the figure of a global content provider (e.g., Google, Amazon or Yahoo) that expands its activities by having an international backbone network and its own CDN. Finally, we observe the

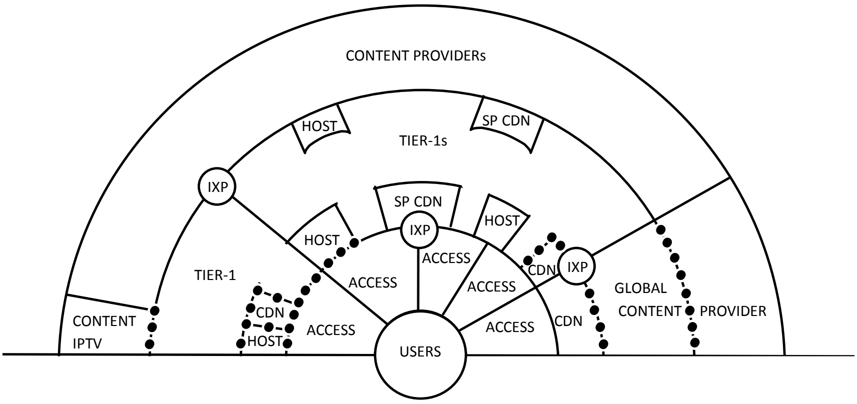

Figure 1.1: ISPs interaction.

specialised CDNs and hosting companies that can be connected to the access ISPs directly or using tier-1s as intermediaries.

1.1.1 Simplicity as a key to success

The simplicity and stability of the original Internet protocols have facilitated the expansion of the Internet, which in turn has increased business opportunities and triggered the appearance of new participants. The Internet is currently witnessing a transformation in which traditional Internet players are changing their roles and activities [43]. In recent years, the commercial Internet structure has evolved from a network-oriented topology to a more content-oriented topology, which means that new deployments consider how to reach content more than how to reach networks. This new approach means that ISPs no longer want to reach a specific network because content is there; instead, ISPs just want to reach content and do not care which network(s) serve(s) the content. Then, the increasing popularity of web content has enabled the positioning of content providers as key players in the Internet market [63], which has dramatically changed the Internet structure motivating the appearance of new forms of interconnection, such as paid peering, and the emergence of the CDNs.

The emergence of specialised CDNs and the network expansion of access operators have changed the traditional multi-tiered or hierarchical structure of the Internet. These Internet players are moving a significant amount of the Internet traffic avoiding, when possible, the use of intermediaries. These intermediaries are the tier-1 ISPs that used to play the role of offering transit services to content providers. However, the new models of interconnection based on direct interconnections such as (paid) peering are removing leadership from the tier-1s because they are being by-passed by the content providers and the access operators. These new forms of interconnection are threatening tier-1s causing them to rethink their current business models.

Content providers also evolved by diversifying their business models and increasing their service catalogue. In recent years many content providers have observed that they can be more than simple Internet content producers. They have seen that investing in deploying their own network infrastructures can provide strategic advantages: i) the better and faster the network, the more users they attract to their services, ii) the more developed a network they have, the better interconnection alternatives they have, and iii) in some cases they can transform their business and re-sell their own network resources to third content providers. However, there are many other content providers that, for different strategic reasons or because they lack the technical/economic resources, cannot invest in deploying their own network and require other interconnection models to offer their content.

Internet users are shifting to an Internet-oriented way of life in which media entertainment dominates their social activities. Several studies [30] revealed that people have increased their use of the Internet for both consumer and professional purposes. Web-based applications such as social networks, large-file downloads and video streaming are now part of our life, and the Internet has adapted its topology and network capacity to this increasing demand.

1.2 Motivation

This thesis is focused on addressing the following points:

1.2.1 Understanding how the Internet structure

The Internet topology has evolved to accommodate the massive demand of multimedia traffic. We have learned from academia that the Internet was a hierarchical model [63], where users and content were at the access ISPs located at the bottom of the hierarchy and national and international transit ISPs were at the upper layers handling the transportation of traffic between users and content. This theoretical model has evolved towards more interconnected structures (almost mesh model) [36] where links between ISPs are critical to their competitiveness and where new players like the CDNs have appeared to compete in the market optimising the traffic transportation. The evolution of the technology is another element affecting the Internet topology, mainly, the last upgrades made at the backbones, incrementing the network capacity beyond the *terabit per second* rate and the deployments made at the access networks, where optical fibre is reaching the end client (FTTH). Finally, we must mention the evolution of the Internet market as a key factor affecting how the Internet is organised. The size and maturity of the local and national markets deeply affect the topology in the sense that we find different types of network infrastructures and technologies, different levels of public administration involvement in promoting shared points of interconnection (IXPs) and different national laws. All these factors make today's picture of the Internet unclear, and the motivation and challenge are to obtain a brighter image of the Internet players and their relationships.

1.2.2 Mixed technical and commercial factors affect the topology

Most of the research efforts to understand how the Internet works are from a technical perspective. Researchers have actively analysed Internet topology by inspecting network protocols, interconnections, link usage and other information related to Internet traffic. There is also work done from the market perspective trying to understand the Internet organisation and the motivations behind each Internet player when it comes to interaction with other players. However, to our knowledge, there has been little focus on considering market and technology jointly. We believe both areas may explain more completely how the Internet is organised. Furthermore, we recognise *policy* as an important issue but avoid including it because it exceeds the scope of this thesis. Territorial regulations, mostly focused on the access part of the network, do not apply to the core of the Internet where interconnections take place. *Network neutrality* is introduced in some parts of the thesis by analysing the major interconnection decisions arising from the conflicts of interest among some ISPs. However, we do not enter into the debate to assess which is the most appropriate operation model from a *policy* perspective.

1.2.3 The Internet, a highly evolving dynamic structure

The Internet has become a dynamic structure, very unpredictable and challenging to analyse. Therefore, we need to understand the dynamics of the Internet ecosystem where there are constant movements between actors and roles generated by mergers and acquisitions that directly impact the Internet topology. Even during the development of the thesis work, we have observed how some of these operations affected the Internet traffic and the topology in relatively short periods of time. We witnessed the acquisition of the transit ISP Global Crossing by Level3, the mergers of Vodafone with Ono and Orange with Jazztel in the Spanish access market and the global expansion of Google's network.

Finally, once the Internet players, their roles, their relationships and their interconnection models are identified, our motivation is to understand the rationale behind each player, why certain strategic decisions are made and how technology affects (or not) those decisions. In other words, we aim to identify why each actor interacts in a particular way by focusing on how Internet content is delivered to the end users.

1.3 Research questions

According to these motivations we present the following research questions that we aim to address in this document:

- What is the evolution of the Internet protocols behind Internet traffic?
- What are the different interconnection models among Internet players?
  - Is there any interconnection pattern for the different types of content providers?
  - What are the interconnection alternatives for regional and local content providers?
  - Is direct interconnection a common trend among access ISPs and content providers?
  - How are tier-1s reacting to attract and retain content providers?
  - What is the role of the Internet eXchange Points?

1.4 Goal and objectives

This thesis aims to further examine how the Internet is organised by analysing its topology from the perspective of the market and technology. The thesis intends to provide an overview of how content providers interconnect with the rest of the Internet players and proposes an innovative methodology for discovering these interconnections. Below, we present the main objectives of the thesis that are connected to the above research questions:

- **To perform a study on the evolution of the Internet traffic:** this study helps to better understand the evolution of the Internet traffic in recent years in terms of traffic consumption, network technologies, types of Internet protocols and types of applications. With these data, we can explain why there have been so many changes in the Internet topology and why interconnection solutions such as CDNs appear. In this part of the work, we have observed the establishment of the TCP/IP suite as the de facto standard of communications, where TCP is the most used transport protocol and the evolution of the HTTP protocol as the predominant application. In this evolution, we observe that legacy protocols such as SMTP or POP3 for email and FTP for file sharing have been migrated to the web using the HTTP protocol.

- **To analyse the Content Distribution Networks:** this study presents the emergence of content distribution networks (CDNs) as one of the consequences of the re-organisation of roles in the Internet ecosystem. Here, we describe the different types of CDNs and how they work. We also propose a taxonomy of CDNs based on the type of Internet player that implements them. Finally, we analyse how the CDNs deliver Internet traffic compared with other mechanisms such as traffic engineering techniques.

- **To offer a measurement-based methodology for discovering the paths that Internet traffic follow and the relationships between players:** due to the lack of public information to evaluate how Internet players are interconnected and how traffic flows along the different networks, we propose a new tool to discover how content travels from the servers of the content providers to the end users at the access ISPs. This tool consists of a platform for performing and analysing Internet measurements. This platform is used to measure the Internet flows, to store large amounts of bulk data, and to visualise and aggregate the data based on different criteria. Once the data is organised in the platform, we can then analyse the different interconnection trends.

- **To discover the different models of interconnection of the Internet players:** this study applies the acquired knowledge to define different models of interconnections and we survey a large number of popular content providers to observe their interconnections.

We focus on discovering how different content providers with different business activity, size and geographic coverage use different approaches to reach their target audience, i.e., we observe a set of common patterns among the same type of data such as the use of direct (paid)peering interconnections between access ISPs and large content providers, the delegation of content to specialised CDNs, the interconnection through shared facilities such as IXPs and the use of solutions like transit services offered by tier-1 ISPs or hosting served by specialised data centres.

1.5 Methodology

To achieve these objectives, we follow a methodology based on the analysis of the problem, design of a platform for collecting Internet measurements and a data processing to provide an interconnection model. This methodology is based on the best practices used by previous well-known projects described in the literature.

First, we performed an in-depth study of the literature to identify the current state of the art in the Internet ecosystem, the different interconnection models, and the existence of different Internet measurement platforms. Throughout the literature, we have identified a number of difficulties in identifying the different interconnection models of the different Internet players. For example, we use the *CAIDA AS relationships* [18] dataset, which aggregates a huge amount of interconnections between ISPs but sometimes misses some critical links. Similarly, the *PeeringDB database* [81] provides a large dataset of direct links but fails to offer a general view of the Internet topology because it only contains links of ISPs present in IXPs that have decided to make their interconnection agreements public. Both approaches individually fulfill their purposes but cannot offer a general view of the end-to-end paths that IP traffic follows based on their stored interconnections. Therefore, we need to combine them with other techniques (such as *traceroute*) to discover the complete route that content follows to reach end users. Once the shortcomings are identified, we re-adapt some of these experiences to unearth the hidden interconnections of these Internet players.

Second, we analysed the evolution of the Internet traffic to understand the reasons that traffic consumption is continuously increasing. In this analysis we look for evidence about traffic patterns that will help us to describe why new Internet players such as CDNs have appeared and why traditional players have re-organised their businesses. For this purpose, we have used a 10-years public data set from the Internet 2 project [55]. By leveraging this dataset, we conducted a study on the emergence of the Content Delivery Networks and classified them based on their structure and the Internet players that operate them. This study was performed because of the noticeable impact of these traffic delivery solutions on the Internet market.

Third, because of the lack of public information about the Internet topology and the interconnection models of the different Internet players, we designed a platform for collecting Internet measurements from different end-points located in commercial networks. We use this tool to obtain broad information focused on revealing the routes that traffic follows from the content providers to reach the end users.

Fourth, we leveraged the measurement platform with the objective of observing the multiple interconnection models of different Internet players. Through the analysis of the collected dataset of measurements we obtained the sufficient information to extract solid conclusions and answer the stated research questions.

Thesis contribution

In Figure 1.2 we observe the different steps explained previously along with the different work done during the course of this PhD Thesis (publications, released projects and workshops).

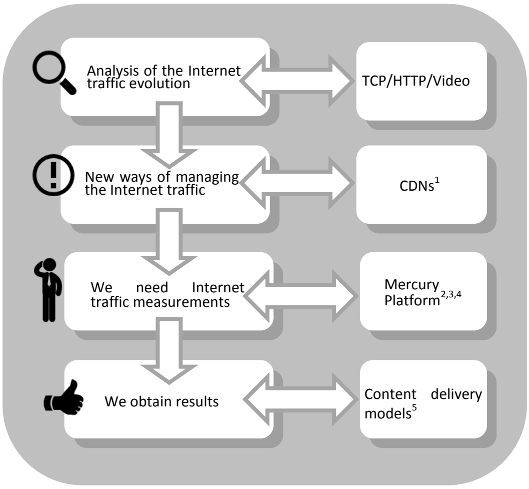
Analysis of the Internet traffic evolution
TCP/HTTP/Video
New ways of managing the Internet traffic
CDNs[1]
We need Internet traffic measurements
Mercury Platform[2,3,4]
We obtain results
Content delivery models[5]

Figure 1.2: Thesis outline.

1. Palacin, M., Oliver, M., Infante, J., Oechsner, S., Bikfalvi, A.. *The Impact of Content Delivery Networks on the Internet Ecosystem.* Journal of Information Policy, North America, 3, jul. 2013.

2. Manuel Palacin and Alex Bikfalvi and Miquel Oliver . *The Mercury Platform.* `https://mercury.upf.edu/mercury`, 2013.

3. Manuel Palacin and Alex Bikfalvi and Miquel Oliver. *Rethinking PlanetLab for Internet Measurements.* Internet Measurement Conference, Workshop Barcelona, Spain, 2013.

4. M. Palacin, A. Bikfalvi, and M. Oliver. *Mercury: Revealing Hidden Interconnections Between Access ISPs and Content Providers.* In Proceedings of the 20th EUNICE Conference, number LNCS number 8846 in 20th EUNICE Conference. Springer's Lecture Notes in Computer Science, 2014.

5. M. Palacin, A. Bikfalvi, and M. Oliver. *Analyzing the Interconnection Models of Content Providers.* PhD School Workshop at the ACM Traffic Monitoring and Analysis (TMA), 2015.

1.6 Thesis outline

This thesis is a multidisciplinary effort. Readers from different backgrounds and levels of expertise will find some chapters more interesting than others. The thesis is organised as follows. Chapter 2 surveys the most important contributions in the literature about Internet interconnection. This chapter will be useful for a novice researcher who wants to study the interconnection market and the different existing measurement tools. Chapter 3 presents the evolution of the Internet traffic in the last decade. This chapter will be useful for under graduate students who want to discover the most important events in the life of the Internet and the characteristics of the traffic. Chapter 4 introduces the CDNs as key players in the Internet ecosystem. This chapter is interesting for those readers who want to know how CDNs work and which Internet players use these technical solutions. Chapter 5 proposes a new platform for conducting Internet measurements. This chapter will be interesting to both novice and expert researchers who want to use a measurement platform to examine Internet paths from the end users to the content providers. Chapter 6 evaluates the different Internet players based on the results obtained from the measurement platform and models the existing interconnection strategies. This chapter could be useful for network administrators and specialised media, as it shows how content is distributed throughout the Internet. Finally, chapter 7 summarises the work done in the thesis and proposes future lines of research.

Chapter 2

STATE OF THE ART

This chapter is an overview of the state of the art, presenting the most significant research related to the topic of this thesis. The first section presents the related work regarding how Internet content has impacted the whole Internet structure. The research done by authors such as Labovitz, Faratin and Clark provides interesting insights about the substantial changes occurred in the last years. The second section explains a literature review about the Internet interconnection. This section includes information about the different types of interconnection agreements, about how to infer interconnection relationships, and more specifically about how to detect interconnection structures and strategies that exchange or deliver huge fractions of Internet traffic (IXPs and CDNs). Finally, the third section shows different measurement tools and approaches that other researchers are using to gather information from the Internet. All these sections present different approaches connected to the Internet ecosystem. From the point of view of this thesis, the study of the Internet ecosystem from a global perspective can contribute to better understand why Internet is organised in such way and why Internet players act as they do.

2.1 The Internet Ecosystem

In the last years we have witnessed several changes in the Internet ecosystem. Labovitz et al. [63] analysed 200 exabytes of Internet traffic (which represents 20% of all Internet inter-domain traffic) during a two-year period (2007-2009). They deployed a platform that monitored the traffic of 110 cable operators, international carriers, regional networks and content providers. The methodology consisted in collecting traffic flow samples and BGP routing information from routers of the participant providers. By analysing these traffic flows they observed interesting findings: i) the Internet traffic grew annually at an average of 44.5% during the analysed period, ii) the majority of the inter-domain traffic has migrated to a reduced number of procotols such as TCP or UDP, which includes video over HTTP, iii) content has migrated to a relatively small number of hosting infrastructures (by July 2009, 30 ASes contribute to the 30% of the traffic), iv) Google represents the fastest and largest growing traffic contributor, and v) the majority of today's traffic flows directly between content providers, CDNs and access ISPs by-passing the hierarchical tier structure. They highlight the emergence of large content providers (Google, Microsoft) and CDNs (Akamai, Limelight) as key players of the Internet ecosystem and they stand out the infrastructure expansion of the access ISPs (Comcast).

This increase of Internet traffic observed by the authors matches temporarily with the increase of popularity of video providers like YouTube (acquired by Google in 2006) and with the massive upgrade of access networks (3.5G and xDSL2+). Another important point about this research is the concentration of Internet content held by few players, which have become major contributors of inter-domain traffic (see Table 2.1).

In 2007 Faratin et al. [43] and in 2011 Clark et al. [26] also identified the emergence of CDNs and the network expansion of access ISPs but they concern more about economic impact and policy issues. They determined that CDNs and access ISPs would prefer a direct interconnection between them instead of using an intermediary. However, the significative market power held by the access ISPs (they directly connect the end-

Rank	Provider	Percentage
1	Google	5.03
2	ISP A	1.78
3	LimeLight	1.52
4	Akamai	1.16
5	Microsoft	0.94
6	Carpathia	0.82
7	ISP G	0.77
8	LeaseWeb	0.74
9	ISP C	0.73
10	ISP B	0.70

Table 2.1: Top ten origin ASNs as an average weighted percentage of all inter-domain traffic in July 2009 (adapted from [63])

users) and the asymmetric pattern of the Internet traffic, allows them to demand an economic compensation for terminating the traffic within their networks. Then, revenue neutral peering cannot be applied for these direct interconnections and paid-peering appears as a solution. According to the authors paid peering will increase the direct connections, reduce transit costs, reduce operation costs, and increase transaction costs around peering agreements. All these consequences have an impact on the Internet topology as all of them facilitate a more meshed Internet.

In 2008, Gill et al. [46] identified that large content providers have begun to deploy their own wide-area networks (WANs) bringing them closer to end-users, and bypassing tier-1 ISPs in many paths. They initially discussed about the consequences of video services for the Internet topology and how networks are evolving to accommodate such amount of traffic. They noticed that the Internet topology is becoming flatter as large content providers are relying less on tier-1 ISPs, and peering directly with more access ISPs. They examined large content providers like Google, Yahoo and Microsoft and they evaluate their connectivity degree. To address this, they probed the 20 most popular web sites from 50 different traceroute servers and then they created a metric for counting the presence of tier-1 ISPs in the AS hops and a metric for counting the different AS interconnections of the content providers.

In 2012, Shavitt and Weinsberg [94] extracted similar conclusions as previous authors and also identified that content providers are trying to by-pass transit networks to get closer to the final users. With this strategy, content providers aim to save transit costs and obtain a better content delivery performance. They performed a 5-year study that collected interconnection measurements from different platforms launched in mid-2000s (DIMES [92] and iPlane [65]) and they seek for topological trends. They concluded that content providers are becoming key players of the Internet: their results denote that the content providers are increasing and diversifying their interconnections while transit carriers are losing dominance. They also observed that content providers are increasing their presence in IXPs (40% of the links), obtaining more interconnection agreements with small transit and access providers, and spreading their network infrastructures.

Gill and Shavitt concluded that the Internet is turning into a mesh. These authors used different methodologies and observed that large content providers are becoming more influential in the Internet ecosystem. In the same line as Labovitz and Clark, the authors detected empirically the impact of the interconnection strategies of content providers on the Internet topology: large CPs are upgrading their own network infrastructures and increasing their direct interconnections with access ISP instead of using intermediaries (transit ISPs). This new situation bring us to a new Internet scenario where traditional players need to re-adapt their roles in order to continue competing.

The different research studies presented in this section has had a great influence on the thesis. The works done by authors like Labovitz and Clark identified movements in the Internet structure and this thesis aims to bear out these observations and to contribute modeling the interactions between Internet players.

2.2 The Internet Interconnection

Traffic exchange between two different Autonomous Systems requires interconnection agreements where the conditions and policies are explicitly declared. Internet interconnection agreements are based on the Public switched telephone network (PSTN) interconnection agreements. Basically, there are two types of agreements between PSTNs , the Bill And Keep (BAK) and the Calling Party Network Pays (CPNP). BAK consists in bilateral agreements between two operators or ISPs for exchanging voice traffic from their customers in pure reciprocity. In a CPNP agreement an operator agrees to carry the voice traffic from another operator to a third operator or interconnection point. In the latter case the calling operator (caller) pays the total cost of the point-to-point connection.

Internet agreements have some parallelism with the PSTN agreements and we can make an analogy between them. The Internet Peering agreement could be considered as an evolution of the BAK and the Internet Transit could be seen as an evolution of the CPNP. Next, we describe the standard Internet interconnection agreements are described:

- Full Transit Interconnection: a transit interconnection is an agreement in which one autonomous system (AS) agrees to carry the traffic that flows between a customer AS and all other networks and receives a fee for this service. Technically, it consists of advertising the customer AS routes to other ASes and advertising other AS routes to the customer AS using a single default route or a set of routes. Typically, different size operators exchange traffic using transit agreements where small operator pays to the big operator for this service.

- Settlement Free Peering Interconnection: a peering interconnection is an agreement in which two or more autonomous systems interconnect physically with each other to exchange traffic without charging each other. An AS guarantees access to its network to another AS in a reciprocal way. Therefore, with this practice, each AS assumes the cost of using its network by other AS in exchange for the benefits of being able to use the networks of other providers. Typically, similar size operators exchange traffic using peering agreements.

These agreements perform correctly in the presented situations where the size and range of the operators are known and rigid. In summary, operators of similar size and traffic volume would implement peering agreements whereas operators of different size and traffic volume would implement transit agreements.

However, the emergence of new parties such as CDNs and the asymmetry of Internet traffic between ISPs require new agreements adapted to the new conditions. The ongoing evolution of the Internet market was noticed by Besen et al. [13] in the early 2000s. Besen et al. evaluated how these new conditions affect the higher-tier ISPs and concluded that small ISPs can exchange traffic among them without passing through a transit operator and creating peering arrangements. This fact causes that the Internet becomes more flexible and it would incentive larger operators not to refuse peering with the smaller ones.

Furthermore, there are several studies from the mid-2000s that propose alternatives to face the asymmetric loads between ASes and the most common are the paid peering and the partial transit regimes. Jahn and Prufer [57] and Shrimali and Kumar [96] evaluated these new interconnection agreements and concluded that paid peering benefits both parties of the agreement and increases the demand because it encourages larger operators to peer with smaller ones. Faratin et al. [43] introduced new concepts about paid peering and partial transit and they suggested situations in which peering should not be done, but also revealed incentives in favor of peering.

In the following, we describe new forms of Internet interconnection agreements:

- **Partial Transit Interconnection:** a partial transit agreement is conceptually similar to the full version but if differs in that a seller AS provides only access to a subset of Internet routes to the customer AS. Therefore the seller provides only connectivity to a part of the Internet. This agreement is a response to two competing commercial pressures. On one hand, in asymmetric environments where peering or transit agreements are set, providers with significant amounts of in-bound traffic could resell the extra out-bound capacity to a third party. On the other hand, partial transit could be used to balance peering ratios.

- **Paid Peering Interconnection:** a paid peering agreement is a business relationship whereby companies such as Internet Service Providers (ISPs), Content Distribution Networks (CDNs), Large Scale Network Content Providers, reciprocally provide access to each others' customers, with or without some form of compensation. Paid peering is identical to free peering in terms of how routes are announced and how traffic is processed. Paid peering permits ASes, who otherwise would fail to negotiate a peering agreement, to better accomplish their interconnection needs. This type of interconnection is widely offered by access ISPs (tier-3) for accessing to the end subscribers (eye-balls).

The evolution of agreement types may be viewed as a rational expansion in order to accommodate the wide diversity of needs. This is consistent with market competition forcing participants to innovate towards more efficient cost-saving contracts.

Inferring the AS Relationships

Determining the interconnection relationships between two adjacent ASes is not an easy task because of the lack of public information from the involved ASes and because of the complexity of inferring the AS relationships based on the analysis of the BGP paths. ASes commonly use private agreements and the information about them is confidential. This drawback only leaves the possibility of inferring the AS relationships by monitoring the routing messages that the edge-routers exchange. This has the limitation that multiple monitors deployed at various vantage points are needed for obtaining a reliable view of the interconnection market.

Otherwise, one can only obtain partial or inconsistent results. Next we present the related work:

In 2001, Gao [45] was a pioneer and proposed new heuristics for revealing AS relationships. Gao focused on detecting customer-to-provider (c2p), sibling-to-sibling (s2s) and peering (p2p) relationships. She presented an heuristics algorithm that infers the Autonomous Systems relationships from BGP routing tables in Route Views servers. AT&T confirmed 99.1% of the inferred results.

The following year Subramanian et al. [100] considered the AS relationships as a graph theory optimisation problem, the Type of Relationship (ToR) problem. They proposed a new inference algorithm which aggregates AS paths from multiple vantage points but they did not consider the s2s relationships.

Battista et al. [37] confirmed in 2003 that the ToR problem is NP-complete. They also demonstrated that ToR formulation cannot be applied for inferring peering links. However, they proposed a new inference algorithm that improved the results of previous works for c2p and p2c links.

Later in 2004, Xia and Gao [115] improved the algorithm and evaluated it against previous Gao's work and Subramanian et al. heuristics. Their results revealed that this new algorithm performs better for few cases (inferring p2p links). They concluded that the rapid growth of the Internet is hampering the inference of AS relationships, e.g., two linked ASes could have different AS relationships at different interconnection points.

Between 2005 and 2007, Dimitropoulos et al. [39, 38] detected other issues of the ToR formulation that led to obvious incorrect inferences such as well-known large providers are inferred as customers of small ASes. They proposed multi objective optimisation techniques that minimised the number of invalid paths. Using these new techniques they identified not only more c2p adjacencies but also more p2p links. They compared their results with existing BGP routing tables and they observed that these tables miss a large amount of AS adjacencies, mostly p2p links. The results from Dimitropoulos et al. are publicly available for the research community at the AS Relationships project [18] from CAIDA. They archive the inferences in this database on a weekly basis.

All these works contributed to infer the nature of the interconnection relationships between ASes. The maximum completion of these contributions is the CAIDA AS Relationships database. This database has helped greatly to feed the measurement platform developed during this thesis and to provide a reliable source of information that validates the methodology of our work.

The Internet eXchange Points

The Internet eXchange Points (IXPs) are physical premises where ISPs exchange Internet traffic between their administrative domains (ASes). The main purpose of the IXPs is to facilitate the traffic exchange between two different AS without an intermediary. Some of the advantages of using these locations are the operation cost reduction, the higher bandwidth and the lower latency. There are different models of IXPs [112]: the US model which is administrated by a private organisation, and the European model which is administrated by a public organisation (public involvement and funding) with the collaboration of the participants present at the location.

Figure 2.1 shows the typical structure of an IXP. The physical interconnection between two participants (ASes) is performed through a switch usually via optical fibre Gigabit ports. There are two types of interconnection between IXP participants: i) *private peering* in which two ISPs interconnect directly using a link between their routers (see ISP B and ISP D in Figure 2.1), and ii)*public peering* in which two or more ISPs connect their routers to a shared switch (see ISPs A, C and D in Figure 2.1). *Private peering* offers more reliability and is easier to monitor and debug whereas *public peering* is easier to administrate and the most cost effective way when it comes to interconnect multiple ISPs. In addition, *public peering* is a selection criteria for enterprise customers as this assures that the ISPs provide interconnection diversity.

The typical operation of the IXPs is that participants pay a member-

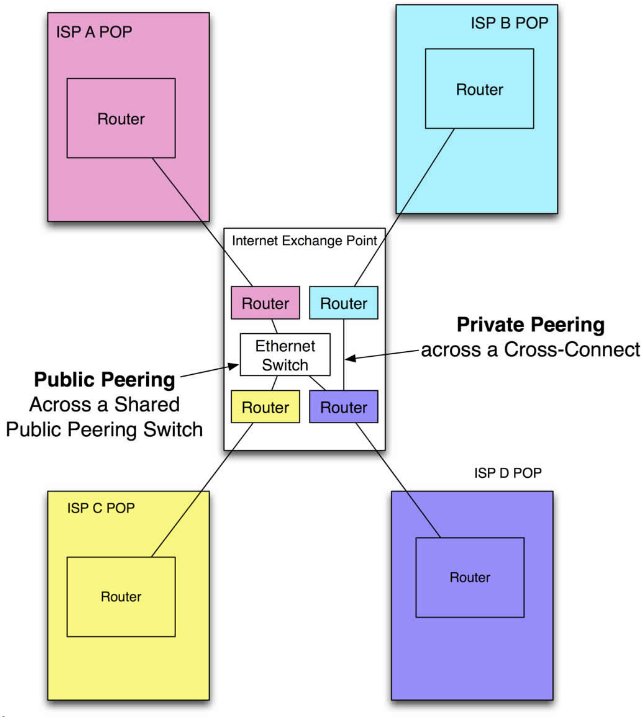

Figure 2.1: Internet eXchange Point structure (source: [112])

ship fee and for the use of the switch ports and/or collocation of network equipments. The traffic exchange is metered depending on the interconnection agreement. Typically, ASes perform settlement-free peering agreements which encourage the use of IXPs due to cost-saving reasons. IXPs reduce the latency because the two ISPs can interconnect at the same city. IXPs can also offer higher bandwidths at lower costs to areas with poorly developed long-distance links. In these areas the price per Megabit of the long-distance link costs more than connecting through an IXP.

As we have previously mentioned, most of the participants at the IXPs perform peering agreements. Most of these agreements are publicly available due to the participants follow an open peering policy. This fact has facilitated the identification of a high number of peering relationships.

Many projects such as PeeringDB [81], PCH [78] and EuroIX [42] have been developed to facilitate the interconnection between ASes. From this initiatives, we highlight PeeringDB which is commonly used by network administrators to publish their peering preferences and their presence in the different IXPs. They put a special emphasis on peering interconnections, such that participant ASes can register and advertise their IP prefixes. Based on this data, in the recent years there have been many new contributions augmenting the AS topology with relationship information [76].

Furthermore, there are also authors putting efforts into mapping all the IXPs and identifying the derived peering relationships. On one hand, Augustin et al. provided a new approach for detecting IXPs and inspecting their AS participants based on these databases and traceroute measurements [11]. They detected 223 out of 278 IXPs and demonstrated that most of the remaining IXPs are invisible to tracerouting. On the other hand, Ager et al. [3] focused on how the IXPs operate. They analysed 9 months of traffic in a large European IXP. They clarified some myths about IXPs, e.g., tier-1 ISPs do not peer at IXPs (they do), IXPs are not used for transit (they are), the number of p2p links in the Internet is larger than what has been assumed or IXP peerings are mostly used for back-up (they are not). They also revealed the existence of a very diverse ecosystem in terms of the participant ASes' business types, peering strategies, traffic exchanges, and geographic coverage.

IXPs are key elements of the Internet ecosystem as they promote the use of shared infrastructures to facilitate and reduce the price of interconnection between ISPs. The task done by projects like PeeringDB, in which anyybody can visualise the public interconnections between ISPs, is really important for other ISPs because they can see who is present at that IXP and to whom they can peer. In addition, the released information from PeeringDB is really useful for the research community in order to study the relationships between actors and, to then generate Internet graphs. In this thesis we use PeeringDB as a source of information to complement the measurements with peering information. Thanks to this data, we have been able to identify that some type of ISPs tend to use

IXPs as the main way of interconnection to others.

2.2.1 Detecting CDNs

In the previous sections we have presented many works that point to discover the Internet topology in a global way, not focusing on any specific player. In this section, we focus on the Content Delivery Networks under hypothesis that they have substantially modified the Internet interconnection structure. Although there are many efforts dedicated to describe and propose new structures of Content Delivery Networks, here we will only focus on the methodologies to detect these networks.

Mao et al. [66] evaluated three of the major commercial CDNs to determine the distance between web clients and their local DNS servers. They use four measurement metrics: i)*AS clustering* consists of determining if a client is in the same AS as its local DNS server, ii)*Network clustering* refers to observing whether a client is in the same network (IP prefix) as its local DNS server, iii)*Traceroute divergence* consists on counting the number of IP hops between a client and its local DNS server, and iv)*Round-trip time (RTT) correlation* refers to examining the correlation between the message round-trip times from a probe point to the client and its local DNS server. Their results show that clients and DNS servers use to be in the same AS but not in the same network domain, that they are close in terms of number of IP hops and that real-time RTT measurements is a good indicator for determining latency but that they cannot evaluate the accuracy of this metric due to their lack of measurement locations.

Huang et al. [53] conducted extensive measurements that accurately characterise the performance of two large-scale commercial CDNs: Akamai and Limelight. They located all the content and DNS servers of these two CDNs and they assessed the server availability and the server delay. They also identified that these two CDNs have different deployment philosophies, Akamai deploys its servers inside the client ISPs whereas Limelight deploys its servers at vantage points close to the client ISPs. To locate the servers, they used multiple distributed nodes of the PlanetLab platform (see subsection 2.3) for executing multiple DNS queries.

They observed that CDNs collocate DNS and content servers and determined the IP addresses of 27,000 Akamai content servers (6,000 DNS servers) and all 4,100 Limelight servers (3,900 DNS servers). Based on delay measurements, they also identified that Limelight uses anycast IP addressing.

Su et al. [99] also used PlanetLab nodes to probe Akamai CDN. They aimed to understand how the dynamic of Akamai's redirection works. They observed that proximity between clients and Akamai servers is a key factor e.g. nodes far away from the Akamai network see considerably more Akamai servers than those which are closer when they observe long-time periods. They also confirmed that Akamai DNS servers redirect clients to the optimum content server based on latency measurements and they isolated the corresponding latency to network-side effects by using ping probes.

Adhikari et al. analysed the infrastructures of the two major video content providers, Youtube and Netflix. Regarding YouTube [2] they used PlanetLab nodes and conclude that 80% of the analysed IP addresses belong to YouTube/Google whereas the other 20% belong to other ISPs like Comcast or Bell Canada. This means that Google has cache servers within the client ISP following the Akamai strategy as well as in its own vantage points. They also showed an extensive geographic diversity with 47 distinct cache locations around the world. Regarding Netflix [1], they conducted an experiment from both residential homes and PlanetLab nodes across the United States. They used DNS resolutions to obtain the IP addresses of the cache servers, and then they used WHOIS lookups to determine the owners of the addresses. They observed that this video service uses three different CDNs at the same time for video playback, Akamai, Limelight and Level3. However, there was only one active CDN whereas the other two act as backup when the primary cannot provide the video minimum quality. They also proposed that combining the three CDNs simultaneously, there will be an improvement of more than 50% in the average bandwidth of the video playback.

Calder et al. [19] focused on Google infrastructure and propose a methodology based on a novel geolocation technique to determine the IP addresses and the geographic location of the cache servers. They detected an expansion of the Google infrastructure between November 2012 and August 2013. They used EDNS [33] from multiple PlanetLab and MLab nodes to enumerate the number of IP addresses and they observed a sevenfold increase during this period. Then they presented a new geolocation approach called client-centric-geolocation CCG that combines MaxMind [69] and latency measurements to determine the location of Google cache servers. Its accuracy outstands over previous approaches but its effectiveness decreases in low-population regions because of the large distances to data centres. They also confirmed that Google is deploying cache servers inside ISPs. Furthermore, they noticed that Google is expanding to remote regions of countries that already hosted servers and to countries that previously did not appear to serve Google's homepage.

Ager et al. [4] proposed a new methodology for detecting web content infrastructure based on the evaluation of BGP snapshots and DNS queries to the most popular and long-tail web sites. Their methodology performed successfully even using few vantage points if they are strategically distributed. According to their results, they reveal that 46% of the popular content is served from North America, 20% from Europe and 18% from Asia. They detect that few hosting infrastructures are serving a large number of hostnames (Top20 is serving 20%): they highlighted that Akamai and Google delivered a significant fraction of the content from their infrastructures. Moreover, they also observed that these organisations split their infrastructures in multiple clusters where they allocate their content depending on their popularity.

This thesis focuses on identifying the CDNs as they have been key elements in the transformation of the Internet structure. Such content delivery solutions have been the response to the large demand of multimedia traffic and somehow, it has led to the emergence of new specialised Internet players or to the evolution of the business activities of legacy ISPs. The impact of the CDNs can be observed in the relationships between ISPs as the asymmetry of the exchanged traffic is a matter of dispute as we will see in future chapters. Table 4.1 summarises this section.

Year	Author	Objective	Methodology	Conclusion
2002	Mao et al. [66]	To determine the proximity between clients and local DNS	4.2M client to LDNS associations in 3 months	DNS is good for obtaining a LDNS in the same AS but not for obtaining one of the same network
2008	Huang et al. [52, 53]	To characterise the performance of AKamai and Limelight CDNs	*DNS lookups* from multiple PlanetLab nodes	They discovered thousands of content and DNS servers and the CDN structure of Akamai and Limelight
2009	Su et al. [99]	To undestand the dynamics of Akamai CDN redirection	*DNS lookups* from multiple PlanetLab nodes	Akamai DNS servers redirect to the optimum content server based on latency measurements
2011	Ager et al. [4]	To detect web content infrastructures	BGP snapshops and *DNS lookups* from few but strategically positioned vantage points	66% of total internet content is served from North America and Europre and Top 20 Content Providers are serving 20% of content
2012	Adhikari et al [2, 1]	To analyse the infrastructure of YouTube and Netflix	*DNS and WHOIS lookups* from multiple PlanetLab nodes	80% of YouTube servers are in Google AS whereas the other 20% is served using servers in other networks (Akamai strategy). Netflix uses threee CDNs (Akamai, Limelight and Level3) at the same time but only one is active and the other two are backup
2013	Calder et al. [19]	To determine the location of Google cache servers	*EDNS* from PlanetLab and MLAB nodes to enumerate the IP addresses	Significant expansion of Google network between November 2013 and August 2013

Table 2.2: The most relevant CDN research analysis

2.3 Measurement Tools

In this section we will describe the major initiatives, projects and derived measurement tools devoted to the Internet Topology. Discovering the Internet structure at the AS-level is a challenge for the research community due to the lack of public information and due to the inaccuracy of the current approaches. Historically, the Route-Views project provided one of the first datasets expressing the AS topology of the Internet, based on BGP measurements [71]. Route-Views deploys a handful of BGP monitors in different vantage locations (most of them at IXPs) around the globe and collects BGP sessions. The aim of this project is to analyse how IP prefixes are propagated and to help ISPs to debug and optimise their network. However, the academic community uses the released datasets to conjecture how autonomous systems are interconnected.

Despite their pioneering work, researchers realised that measuring the Internet topology from few vantage points leads to partial results [20]: in general, peering relationships are difficult to identify. Similar projects to Route-Views are CIDR Report [108] and BGP Routing Table Analysis [82] supported by APNIC the Asia-Pacific Internet Registry. CIDR Report analyses the Internet routing table from a global perspective while BGP Routing Table Analysis focuses on the regional perspective. To perform this task both initiatives analyse the BGP messages from various APNIC location a Route-Views BGP monitors.

Analysing BGP sessions is really useful to identify how ASes are interconnected but it does not show you how IP traffic is routed throughout the different ISPs from an end-point perspective. To address this, Mao et al. developed a traceroute tool at the AS-level. They created a database of IP-to-AS mappings based on the observation of BGP announcements in combination with IP traceroute measurements [67, 68]. However, the authors noticed the difficulty of detecting IXPs and sibling relationships, as well as mapping mismatches due to measurements in a limited geographical region.

Recognising the benefit of measuring the Internet from the edge, Shavitt et al. proposed DIMES, a measurement infrastructure using a large number of software agents [92]. The main objective of the project is to obtain the Internet graph at the AS and IP level. Their methodology is based on collecting traceroute and ping measurements from their agents and then they translate the IP addresses to the corresponding AS number. Although the project is open and its data is freely available, the information does not include the relationship between ASes. Later, the same author [93] combined DIMES measurement with BGP information from Route-Views to infer the AS relationships with successful results.

In parallel, Dimitropoulos et al. focused on this issue. Their work started initially as an effort to model and generate synthetic but realistic AS topologies [41]. Subsequently, they attempted a classification of the Internet ASes using data collected from the Internet Routing Registries [56] and Route-Views. Their data in combination with the active measurements of the Archipielago (Ark) project sponsored by CAIDA [16] has contributed to the improvement in the knowledge of the AS Relationships (see CAIDA AS Rank [18, 40]). The Ark project is an evolution of the Skitter project [17] and it aims to provide a large-scale measurement infrastructure that allows researchers to collaborate in a distributed way. Ark uses distributed monitors to collect measurements about IP topology, DNS resolution and AS links.

He et al. [50] merged both data from these databases with their own traceroute tool, called RETRO. They aimed to reveal missing peering relationships and they identified that many of these occur at IXPs. They initially tested Skitter to collect the traceroute measurements but they realised that this tool was not suitable because of the small number of vantage points. For this reason, they implemented RETRO which uses public traceroute servers to collect measurements from many diverse locations. This diversity of vantage points (more than 1200) facilitated that IP traffic traverses throughout IXPs.

Chen et al. [21] increased the number of traceroute sources by developing a measurement add-on, called ONO, for a popular BitTorrent client. This methodology has a two-fold objective. Firstly, they obtain traceroute measurements from end-points. Secondly, they facilitate the adoption of the tool as it is packaged with the BitTorrent client and it operates in background. The add-on performs traceroute measurements to a set of random destinations from the established P2P connections and then it translates the IP addresses of the path to the corresponding AS numbers using BGP data from Route-Views. Once they have the AS path, they apply an inference algorithm which focuses on detecting hidden p2p links combining data from public IXP and s2s links using information from CAIDA datasets. They concern about detecting loops in the AS-paths (they discard these traces) and they also identify multiple situations where missing and extra hops appear in the AS-level traceroute when we compare with BGP AS-path.

There are also other multi-purpose platforms that can be used by the research community for executing network measurements. PlanetLab [83] is an open platform with hundreds of nodes around the world for implementing and testing new Internet services. Research institutions that want to execute any type of experiment, contribute with at least a couple of nodes (PlanetLab servers). To set up an experiment, researches select the servers that they require from a list of available nodes and they obtained an slice of resources (CPU, RAM, HD, network, user account...) from these nodes. This means that researchers can remotely connect to the selected nodes and test/execute any kind of application.

There are many projects that use PlanetLab. iPlane [65] is a service that provides accurate predictions about Internet paths and it uses PlanetLab for collecting traceroute measurements from different nodes. CoDeeN [86] is an academic testbed Content Distribution Network (CDN) built on top of PlanetLab. CoDeeN consists of a network of high performance proxy servers deployed on many PlanetLab nodes. The Measurement Lab (MLab) [106] is another open distributed platform but unlike PlanetLab, MLab is only focused on network measurements. It provides a set of network tools to perform many tests between the MLab nodes and the client machines. The tests provide real-time performance information from speed, throttling, blocking, and rich diagnostic metrics. Finally all the collected data from the tests is publicly available and can be accessed via a WS-REST api. One of the tools that is available by default on MLab is Paris *traceroute* [12]. Paris traceroute is a new implementation of the original Van Jacobson *traceroute* that performs better to the per-packet load balancing of IP routers. Paris traceroute modifies the structure of the IP packets to create traffic flows that routers tend to forward following the same IP-path. The drawbacks of Paris traceroute are that it does not come with the OS network suite, that it require root privileges to be executed (it uses raw sockets) and it does not performs perfectly for all situations (but better than the classic *traceroute*) because some routers along the packet path follow a per-packet load balancing policy. The Reseaux IP Europeens Network Coordination Centre (RIPE NCC), the Regional Internet Registry (RIR) in Europe, provides a collaborative tool called ATLAS [89] that consists in a global network of probes that measure Internet connectivity and reachability. By the end of 2013, there are more than 5000 participants around the world in many commercial networks. Participants can conduct distributed measurements using the different Atlas nodes and execute *ping* and *traceroute* measurements using a simple REST API. Thanks to ATLAS, researchers have generated many topology maps that identify the performance, in terms of round-trip-time (RTT), for different DNS root servers. In addition, ATLAS provides a large and public data set with measurements of previous participants that can be used via the API. Table 2.3 summarises the most relevant Internet measurement platforms and projects.

All these platforms and tools have influenced on the measurement platform designed in the framework of this thesis. The measurement platform uses many conceptual ideas from these projects like the use of distributed clients and a central database, the use of trusted databases like CAIDA AS Relationships and PeeringDB to detect the AS relationships, the use of RouteViews to translate IP addresses to AS numbers and the use of Paris traceroute to perform a more accurate probe. All these previous works use one or many of these ideas, but none of them use all these concepts with the specific purpose of identifying the interconnection models of different type of Internet players.

Year	Project	Objective	Methodology
1999	CIDR Report [108]	To encourage ISPs to drop classful prefixes in favour of classless aggregates and to examine the Internet Routing table on a global scale	Takes BGP feeds from multiple APNIC locations and from Route-Views
1999	BGP Routing Table [82]	To analyse the Internet Routing table from a regional perspective	Takes BGP feeds from multiple APNIC locations
2001	RouteViews [74]	To allow Internet users to view global BGP routing information from the perspective of other locations around the internet	Around 20 routers, mostly in US. Routers establish BGP sessions but do not announce prefixes. Researchers can see the public BGP sessions archives (2001-up to now)
2003	PlanetLab [83]	To develop new technologies for distributed storage, network mapping, peer-to-peer systems, distributed hash tables, and query processing	1350 nodes at 677 sites, mostly research centres and universities. Each node is a Unix virtual machine that can be accessed by any registered researcher using API to perform any type of experiment
2009	DIMES [104]	To obtain the Internet graph at the AS and IP level	Distributed agents (over 200) collecting *ping* and *traceroute* measurements and then translate the IP addresses to the corresponding AS number
2005	RETRO [51]	To infer potential peering relationships in IXPs	Uses public Traceroute servers and then translates IP-level paths to AS to apply inferring algorithms
2006	iPlane [75]	To provide accurate predictions of Internet path performance for emerging overlay services. to accurately and efficiently predict latency, bandwidth, capacity and loss rates between arbitrary Internet hosts	Takes *traceroute*) measurements from different PlanetLab nodes and Traceroute servers and publishes results on a daily basis
2007	ONO [110]	To improve the download speed in BitTorrent by discovering optimal paths	Plugin in BitTorrent client that performs traceroute measurements between P2P nodes. There are over 100k acties clients. To generate paths it uses IP-to-AS translation using RouteViews datasets
2007	CAIDA Ark [16]	To reduce the effort needed to develop and deploy sophisticated large-scale measurements, and provide a community-oriented measurement infrastructure on a security-hardened distributed platform. Is the evolution of the Skitter platform	Uses distributed monitors (Raspberry Pi) (61 in 32 countries, August 2012) to collect measurements about IP topology, DNS resolution and AS links. Registered users can access the monitors through Ark services like Vela
2008	MLAB [106]	To advance network research and empower the public with useful information about their broadband and mobile connections	Measurement tools hosted in MLAB servers that measure the speed and evaluate the performance of certain applications. Users can access to the public output of the different tools
2010	ATLAS [89]	To create the world's largest Internet measurement network	More than 5000 participants around the world in many commercial networks. Participants can conduct distributed measurements (*ping* and *traceroute*) using a simple REST API. Any user can browse and visualise the public results

Table 2.3: Most relevant Internet measurement platforms

Chapter 3

EVOLUTION OF THE INTERNET TRAFFIC

With the aim of observing the increasing Internet usage, paying special attention to the web content, we consider that it is essential to perform a study of the Internet traffic evolution focusing on the analysis of the Internet protocols and network technologies (see Figure 3.1). To this end, we have analysed an eight-year data set of Internet traffic collected from the Internet 2 project. This study confirms the predominance of the TCP and HTTP (web) protocols witnessed by previous publications like Labovitz et al. [63] or CISCO [23]. Our work concludes that new (over-the-top) Internet services rely on the mature and stable TCP and HTTP protocols of the TCP/IP stack and these protocols are held on emerging technologies (optical fibre, wi-fi, 4G, etc.). Hence, the stability and flexibility of TCP and HTTP has helped in their success, leaving the room-for-technical-improvement to the upper (application) and lower (physical and link) layers.

3.1 Introduction

Almost 40 years after the birth of the Internet Protocol, this veteran architecture keeps evolving. There are and have been dozens of research

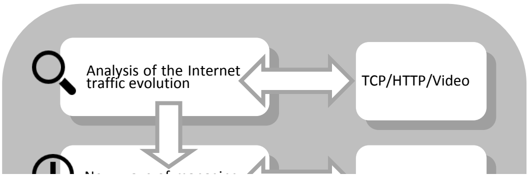

Figure 3.1: Thesis outline: evolution of the Internet traffic.

programs that test new architectures and protocols (e.g. the GENI [109] [44] and FEDERICA projects [102] [101]), but despite these attempts to improve the Internet, the TCP/IP stack continues livelier than ever. The Internet protocol suite was initially designed as an US Department of Defense project [25], but it evolved to the commercial worldwide network that we currently know. When the TCP/IP suite was designed, it was created with the capacity to evolve, and this principle is still valid.

The maturity that the Internet has reached offers a solid basis for drafting an accurate analysis of Internet traffic, focusing on the period from 2002 to 2010. This study offers a perspective to understand what has happened to the Internet in terms of its protocols and applications. The spectacular growth of the number of Internet users and traffic shows a strong correlation. However, the leading applications have changed since the beginning of the analysed period. New services and applications have proliferated, while the protocols remain almost the same as those based on the TCP/IP stack.

This chapter illustrates why the TCP protocol remains the predominant protocol in the Internet, although more application-oriented transport protocols have been developed. It then analyses the application layer to identify what applications are being served on top of TCP. The chapter focuses on why applications are migrating to web protocol HTTP and conjectures the location of the hidden peer-to-peer (P2P) traffic during the last few years.

The chapter includes an analysis of the statistical data from the IP traffic evolution, Internet protocol distribution and HTTP predominance. To perform this analysis, we have used data and empirical evidence gathered from different publications, papers and network statistics. Finally, the study extracts some conclusions and proposes future research topics related to Internet architecture.

3.2 (R)evolution of the Internet

3.2.1 A brief history

The TCP/IP suite was developed as a result of a military research project of the Defense Advanced Research Projects Agency (DARPA) of the United States in the early 1970s. The research project created the ARPANET [64] network, the first packet switching network, which is considered the predecessor of the Internet. At the same time, Ethernet technology, the dominant standard for local area networks (LANs), was being developed. The successes of both the TCP/IP suite and the Ethernet are due to the support of major vendors in the telecom industry and acceptance from end users.

When the TCP/IP suite was being implemented, a set of protocols and services were developed that used it. Services include email (SMTP), remote command execution (Telnet, RSH and recently SSH), file transfer (FTP), dynamically obtaining IP address (RARP and DHCP) and the domain name system (DNS). These services helped to increase the popularity of the Internet and LANs based on Ethernet technology. See the timeline of the Internet in Figure 3.2.

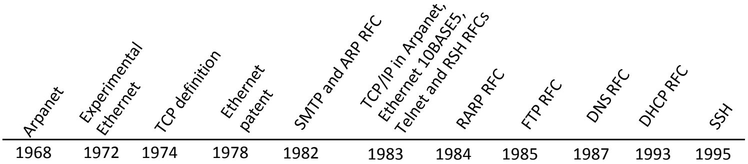

Figure 3.2: Most significant developments in the Internet from 1968 to 1995

3.2.2 The birth of the World Wide Web (WWW)

In the early 1990s, the commercial Internet, best known as the World Wide Web (WWW), was born at the CERN institute in Geneve. It appeared as a result of implementing a hypertext protocol, a hypertext language for representing resources and a definition of a unique identifier for these resources (URI). In 1994, the World Wide Web Consortium (W3C) was created. The W3C was responsible (in coordination with the IETF) for managing the standardisation of the hypertext language HTML and web protocol HTTP. In the following years, several HTML language and HTTP protocol versions appeared. This fact facilitated expanding the Internet towards companies and end users. The success of HTML and HTTP in the commercial sector (what we commonly know as the Web) is comparable to the obtained by the telephone or the television. See the timeline of the evolution of the World Wide Web in Figure 3.3.

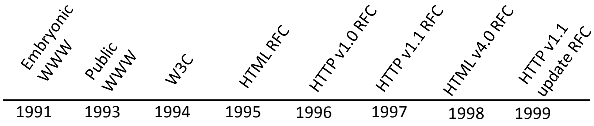

Figure 3.3: Most significant developments in the Internet from 1991 to 1999

3.2.3 The Web 2.0

In 2004, Tim O'Reilly [107] first mentioned the term Web 2.0. This term did not define a protocol or a technology. It defined an attitude and a way of creating the Web. Traditionally, the Web followed the client-server paradigm, whereas another term appeared in Web 2.0: the prosumer. This new actor simultaneously played the roles of a producer (server) and consumer (client). In this scenario, the user acts as a content generator. Examples of this kind of content include wikis, social networks or blogs

[72]. Moreover, the early 2000s were a stage of proliferating new Web technologies and sharing protocols. Examples include technologies such as XML, RSS, AJAX, FLASH, FLEX, DOM, JSP, PHP, Perl, .NET or different peer-to-peer (P2P) technologies. The emergence of new technologies, in combination with a new user role, made the Web more humanlike, dynamic and participative; in other words, it became more social.

3.2.4 The Multimedia Age

Since the late 2000s, the boom in multimedia content has been the most important occurrence. The current Web environment and Internet technologies (mobile and fixed) provide sufficient capability for delivering multimedia content with a high-quality user perception. The expansion of HFC, DSL and 3.5G lines has provided broadband access that allows users to download thousands of terabytes of data and play millions of Internet videos [23]. Video on top of the Internet (Video over HTTP) acts as a killer application that consumes thousands of Mbps and generates large revenues for the content providers and all companies responsible for its delivery.

Furthermore, new Web standards such as HTML5 are handling interoperability issues and providing default multimedia support for the HTML standard. HTTP over TCP/IP is the predominant protocol that allows encapsulating any type of Internet content. In this context, the Content Delivery Networks (CDNs) and large Content Providers (e.g., Google, Yahoo, MSN, Amazon, Facebook) arose, and they positioned themselves in both the applications market and network ecosystem [63]. See the timeline of the evolution of the Internet content in Figure 3.4.

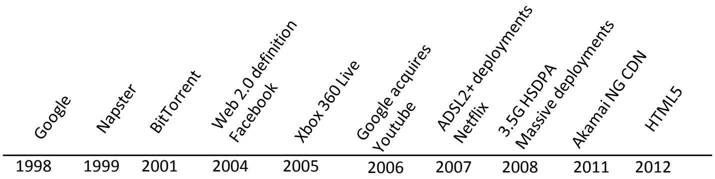

Figure 3.4: Most significant developments in the Internet from 1998 to 2012

3.3 Internet Protocols

This section describes the main reasons for TCP/IP predominance and the flexibility of the HTTP protocol in transporting any kind of application.

3.3.1 TCP/IP is still the predominant protocol

The TCP/IP suite is a set of communication protocols developed along the 1970s and 1980s. Why is an ”old” set of protocols still alive in such a changing world as the communications sector? We do not use the same primitive TCP/IP protocols, but the basis and the concept remain the same. We can remark on the adaptability [70] of the suite to the different situations and the requirements of each period, presumably having left behind other obsolete optimal technologies. Below, the most outstanding TCP/IP features are listed:

- **A complete suite:** The TCP/IP suite is a technological solution to an engineering problem in the real world formed by four layers while the OSI model (ISO/IEC 7498-1) is the standard reference model of interconnection formed by seven layers and describes the functionalities that each layer must cover to provide interaction in a network communication. TCP/IP covers almost all layers of the OSI networking model (see Figure 3.5). The *link layer* defines the protocols to describe the local network topology and the interfaces to transmit the datagrams and corresponds to the *physical* and *data link* layers of the OSI model. The *Internet layer* performs the task of exchanging datagrams between networks and defines the addressing and routing structures and corresponds to the *network* layer. The *transport layer* provides the host-to-host communication creating a channel for the applications and it has its homonym layer in the OSI model with the same functions. The transport *UDP* protocol provides basic and unreliable datagram service while *TCP* protocol provides flow-control, connection establishment, and reliable transmission. The *application layer* is responsible to represent the user data and to communicate with other applications using the underlying services of lower layers. This layer is represented by the *session*, *presentation* and *application* layers of the OSI model which separates more specifically the communication functionalities.

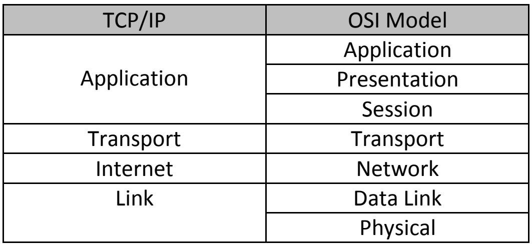

TCP/IP	OSI Model
Application	Application
	Presentation
	Session
Transport	Transport
Internet	Network
Link	Data Link
	Physical

Figure 3.5: TCP/IP stack and OSI model

- **Adaptability and Flexibility:** TCP/IP evolves when any networking drawback appears. We observe this in the migration from IPv4 to IPv6 because of the IP address depletion of the version 4 of the IP protocol and to enhance the security. We also see the adaptability of the TCP/IP stack in the different TCP versions. TCP Tahoe, Reno, New Reno, Vegas, SACK and the more recent versions BIC, CUBIC, Westwoodm, Fast TCP or High Speed TCP have continuously introduced new improvements to their congestion avoidance mechanisms with the objective of reacting better to the packet loss and providing more fairness. Another point is the widely use of the TCP/IP for wireless LANs. The Wi-Fi standard (IEEE 802.11) typically uses the TCP/IP protocols on top of the Ethernet wireless technology despite it is well known that is not the most optimal approach. In addition we see the TCP/IP flexibility in the capacity to adapt to many traffic patterns, e.g., video streaming was thought to be transmitted using optimal real-time protocols (RTP over UDP) but finally, industry relies on the *mature but robust* TCP that meets almost any application requirement.

- **Ease of implementation:** It is really simple to implement new services as almost all operating systems in any type of hardware device support the TCP/IP stack. Programming new applications that communicate with other ones is a feasible task because operating systems and software development kits provide natively APIs to code. This is one the keys of success of the TCP/IP, its simplicity, that has helped to the development of thousands of services supporting this model.

- **Present and future:** Hundreds of RFCs (Request For Comments) and vendors support TCP/IP as a standard. Software companies and hardware vendors are developing many applications and devices for this architecture, and thousands of research projects are based on it. There is no short-term plan for moving the Internet industry to a new clean-slate Internet design although there are some recognised researchers claiming for this change [88]. Therefore the future of this stack will be the same as always, reinventing itself and providing new improvements without changing the basis.

As we have previously mentioned, the stability and high adoption of the TCP/IP model have facilitated the proliferation of both software and hardware, building a rich ecosystem around the Internet industry despite not being the most optimal solution. In this context, content providers have a great opportunity to spread their services as the Internet acts as a proven and reliable platform from an economic prospective.

3.3.2 HTTP as a multipurpose application protocol

HTTP protocol usage has not stopped growing since its definition in the early 1990s. The HTTP protocol was initially designed to perform on top of the TCP transport layer, and its main purpose was to encapsulate HTML pages. Over time, it has started to encapsulate other types of content. Files, images, videos and real-time applications such as audio or video streaming are some good examples of HTTP versatility [84] [85].

HTTP flexibility resulted in replacing protocols designed to fill the gap of the requirements of any application by HTTP. As examples, we can mention HTTP for transferring large files, instead of FTP or P2P networks [61], and streaming audio or video, instead of real-time protocols like RTP/RTSP.

Both examples are good illustrations of the present state of the Internet. Many applications that used to implement specific protocols are migrating to the HTTP Web protocol over TCP [49] [98]. This happens because of its simplicity and flexibility. Internet research groups are changing some of their premises, and they are trying to adapt their services to the HTTP protocol instead of creating new protocols for each type of service. This concept is called *webification*.

The late 1990s and the early 2000s were a period of emerging new protocols. Those protocols were created to better fulfill requirements of new services. They were more optimal, faster and more robust but sometimes more difficult to implement. The vast majority of users, however, did not adopt most of these protocols.

The idea was simple: all operating systems (OS) had the TCP/IP stack installed, and most had a Web browser installed. Why not try to design applications on top of the HTTP protocol? The answer is as simple as the question: there was not enough bandwidth.

Before the mid-2000s, the backbone was not sufficiently dimensioned, and the access connections did not offer enough capacity to support applications that consumed large amounts of bandwidth. A large list of protocols that better satisfied the requirements of each traffic pattern thus emerged. P2P protocols for file sharing or RTP over UDP for audio/video streaming are examples where HTTP over TCP did not work well.

With the upgrade to an over-provisioned backbone network and the speed increase of the access connections, using these specific protocols was no longer necessary. With sufficient bandwidth, it is not necessary to tune quality of service (QoS) parameters, traffic shaping or waste efforts in defining new protocols. Putting services on top of the TCP/IP using the HTTP protocol is thus viable.

In this context, where bandwidth and optimisation are not constraints, creating new services on top of HTTP was relatively easy. Video and audio streaming or massive online gaming are good examples of proliferating services that squeeze the most the HTTP protocol. However, there are other concerns that must be considered. Most of the shortcomings of the HTTP protocol are being replaced by adding patches like the HTTPS specification for security and the server tunning in the service development. Therefore, the HTTP requires a deep update to meet with the requirements of the new generation services: HTTP/2 (also named HTTP 2.0) that will renew the capabilities of the successful HTTP 1.1 version.

Before introducing HTTP/2, we have to present a Google research project called SPDY started in 2009 [47]. SPDY is an application-layer protocol aimed to improve the Web performance by reducing the web page load latency and enhancing the web security. It achieves reduced latency by compressing, multiplexing, and prioritising the transfer of web page sub-resources so that only one connection per client is required. The intention of SPDY is not to replace the HTTP protocol; the SPDY protocol encapsulates the HTTP protocol and modifies the way the transmission is done between both client and server sides. SPDY has obtained a large involvement from the Web players from the very beginning. Many browsers including Chrome, Firefox, Internet Explorer 11, Opera or Safari and the major server vendors (Apache, Nginx, or Jetty) included extensions to support it. In addition, large content providers like Google (obviously!), Twitter, Facebook, Wordpress or Yahoo! enabled SPDY in their servers.

The improvements introduced along with the wide acceptance of the SPDY protocol triggered to the creation of a working group in 2012 with the intention of defining the HTTP/2 protocol specification. In February

2015 Google announced the abandonment of support for SPDY by the beginning of 2016 to focus solely on HTTP/2. This new protocol uses the basis of SPDY and changes the header compression algorithm (Huffman code instead of dynamic streams) to increase the robustness against cyber-attacks.

The HTTP/2 protocol [48] aims to fit better with the delivery of modern web sites. Current web sites have few in common with a web site of 15 years ago. Most of the modern web pages are compounded by multiple resources (videos, images, CSS and JavaScript files or even vectorial animations) that require powerful browsers to manage multiple connections. Furthermore, browsers are required to process all these elements in parallel providing an adequate quality of experience to the end user. The HTTP 1.1 protocol provided lots of upgrades over the original HTTP 1.0 like the use of persistent connections to allow connection reuse, the request pipelining to allow parallel request processing, and the introduction of cache functionalities. However, these improvements began to be insufficient for serving modern services and although networks (optical fibre, CDNs or 4G) and devices (desktops, laptops, tablets and mobile phone) upgraded their capacity, there exists a practical limit that forces a deep change in the way the web content is served. Rich multimedia content such as Web (video)streaming, JavaScript animations or multi-resource web sites have nothing to do with the original static and plain-text web pages. Modern sites have a bunch of requirements that the continuous improvements in the software side (browsers and servers) and in the hardware side (networks and devices) are not capable to fulfill. Therefore, we can only focus on the HTTP protocol again adapting it to the TCP nature.

HTTP/2 deals with all of these issues and redefines completely the HTTP protocol providing substantial improvements that accommodate better the next generation of web content. HTTP/2 is backwards compatible and re-uses the simple, useful and well-known HTTP codes and messages. Then, the version 2 of the HTTP protocol introduces a new concept for serving the web traffic based on using streams. This new architecture enables the possibility of multiplexing parallel requests and responses in a single communication avoiding the use of multiple TCP

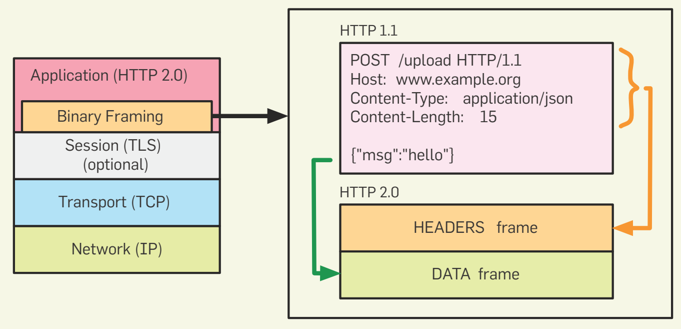

Figure 3.6: HTTP/2 binary framing (Source: [48])

connections and the limitation of the FIFO queue system of the HTTP 1.1 pipelining.

The new serving system defines a new sub-layer in the TCP Application lever called binary framing layer that controls how HTTP messages are sent between the client and the server (see Figure 3.6 an 3.7). The HTTP/2 protocol defines three elements in the communication:

- **Stream**: bi-directional virtual channel within a connection. It contains an identifier number and a relative priority value.
- **Message**: a complete sequence of frames that represent a logical message
- **Frame**: the smallest unit of communication in HTTP 2.0, each containing a frame header with at least the stream identifier to which the frame belongs, and carries a specific type of data (payload data, headers, priority, reset, setting, push, ping, goaway, window update and continuation).

Before sending any data, a new stream must be created and headers must be exchanged using frames. Then, payload data is sent using different frames within the same stream using a more efficient transmission. The increase of efficiency is provided by using a binary encoding of the data that allows to compress the information and by enabling the server

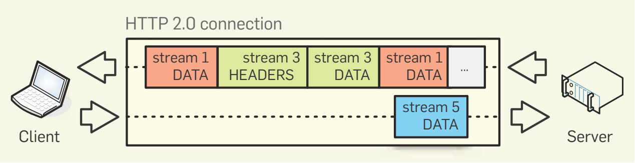

Figure 3.7: HTTP/2 client-server communication (Source: [48])

push capability which allows to serve multiple resources to a single resource. In addition, HTTP/2 includes new features such as a security layer or a method to prioritise frames (usually, the text of a web site is more important than the images).

In conclusion, HTTP/2 will provide more possibilities for the implementation of new web applications and services based on the web protocol and it will have a huge impact on the Internet industry: server software vendors will have to implement the new protocol, CDNs and hosting companies will need to re-adapt their infrastructures, developers will be required to update their skills to exploit the new capabilities and companies will have more opportunities to release low-latency services like online gaming, high definition video streaming or virtual reality. It is unknown how the transition to HTTP/2 will unfold, however what is certain is that it will be faster than the IPv4 to IPv6 migration because a large number of Internet players and stakeholders have interests on it (e.g. large content providers need to release their next-generation services).

To summary, following we list some of the main characteristics of the HTTP protocol:

- HTTP is a simple and robust protocol. HTTP is a state-less protocol that provides a full set of request methods and response codes to perform many tasks. HTTP/2 will introduce the concept of virtual channel to increase the efficiency of the transmission.

- HTTP adapts to any kind of traffic pattern (with enough bandwidth) and it will be ready for low-latency services when HTTP/2 exploits message compression and data prioritisation.

- HTTP is a de facto standard. Almost all operating systems of any device have a web browser installed and this facilitates the adoption and the future updates of new services. Furthermore, HTTP is not typically blocked by firewalls (well-known port 80) which frees the developer about networking concerns.

- HTTP can offer security in the data transmission using the HTTPS extension. With HTTP/2, security will be mandatory and all of messages will be encrypted.

3.4 Internet Traffic Statistics and Network Technologies

This section explains the methodology and describes the data used to perform the analysis of the Internet traffic. It then presents the obtained results, putting extra emphasis on the most significant points. Finally, the section includes a section about Internet application trends and one that analyses current network technologies.

3.4.1 Background and measurement methodologies

Traffic measurement is a complex task that must follow a detailed methodology to obtain precise results. The methodology used to gather a measurement is a fundamental factor to determine the quality of an experiment. For example, decisions on the number of samples and the techniques to discard erroneous measurements are crucial to avoid biased results. Another important factor is to use trusted sources when you combine multiple data sets. In this section, we present many relevant authors that have worked on defining methodologies related to the measurement of Internet traffic and then we explain the methodology that we follow to perform this study on the evolution of the Internet traffic.

Wamser et al. [111] define a complete methodology to classify network traffic and provide a full literature review of different works related to traffic measurement. As Karagiannis et al. [61], Parish et al. [80] and Archibald et al. [9] mentioned, it is sometimes difficult to classify IP traffic due to its characteristics. For example, P2P traffic is difficult to identify because it uses multiple protocols and ports and is hidden within the rest of the traffic. The usage variation of a certain protocol must therefore be inspected in detail before drawing a final conclusion.

This work classifies IP traffic using the Netflow network tool. Netflow was initially implemented by CISCO [22], but it has been standardised by the IETF as the Internet Protocol Flow Information eXport (IPFIX) [24]. This study uses weekly snapshots provided by the Netflow tool of the Internet 2 project [55] along the period from 2002 to 2010 [1]. The raw data is processed to create graphics and extract conclusions. This study also includes some data provided by the Internet World Stats [54] to correlate traffic measurements with the evolution of Internet users as well as data from other sources, including the CISCO [23] and Sandvine [90] annual reports.

Source	Internet2 Netflow weekly reports
Url	http://netflow.internet2.edu/weekly/
Period	From 2002 to 2010
Availability	Weekly
Scope	US Internet 2 network
Description	Netflow data from all core routers of the Internet2 network are analysed to produce weekly reports of use of the network
Data	Full datasets of IP protocols, application types and packet size distributions

Table 3.1: Internet 2 NetFlow data source description

[1] The public data set of measurements is no longer publicly available. To know more details, please contact the author of the thesis

3.4.2 Measurement results

Internet users

The total number of Internet users was gathered from the Internet World Stats site [54]. We can observe in Figure 3.8 that the number of users has been increasing constantly, especially in the late 2000s. The compound annual growth rate (CAGR) from 1995 untill 2013 is 33% but this indicator is conditioned by the spectacular increase during the first years. If we only focus on the last five years, a period when the Internet has been (almost) extended globally the CAGR is 12%, a good figure taking into account that in 2013 only 39% of the worldwide population was connected. This increase is due to the introduction of broadband services in households and due to the global expansion of the Internet. The motivation for gathering this indicator is to discover the correlations with other indicators like the traffic demand or the predominance of some type of Internet traffic.

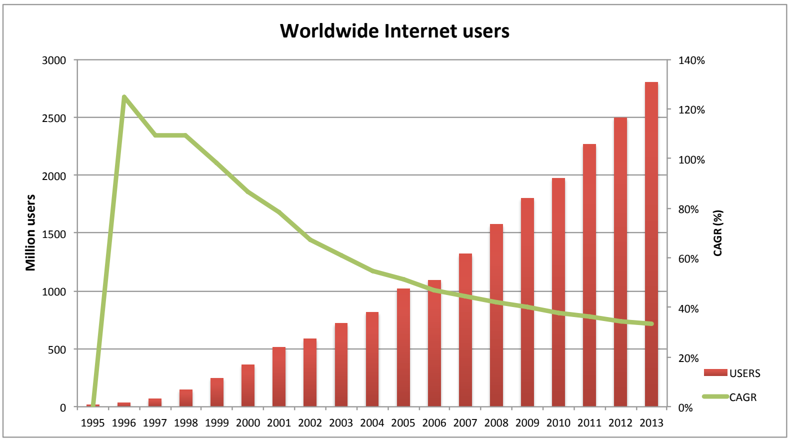

Figure 3.8: Worldwide Internet users growth. Source: [54]

Protocols distribution

The protocol distribution is obtained by processing the data gathered from the Internet2 project [55]. The protocol distribution has suffered a minimal variation during the studied period (2002 to 2010), as seen in Figure 3.9. TCP is the most used protocol, with almost 90% average usage during the observed period. UDP obtains an average of approximately 10%, whereas the sum of the rest of the protocols totals approximately 1 or 2%. This data demonstrates that TCP is the predominant protocol and that upper application protocols rely on it.

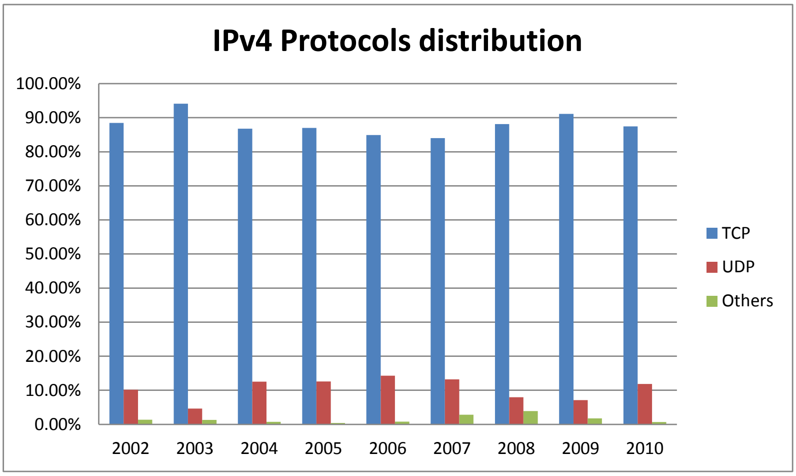

Figure 3.9: IPv4 protocols distribution

The protocols stability and the TCP predominance along such period of time surprised us because this period was highly proliferative in terms of new protocols and video codecs addressing real-time communications. The theory stated that protocols like TCP are not suitable for transporting low-latency services like video or audio and this gave room for improvement. This period of time witnessed the emergence of the early commercial audio/video conference software and the former video streaming services. Along that decade many research groups focused on implementing protocols for transporting multimedia data in an efficient way and adapting the streams to the different network requirements. An example of protocol is the Real Time Transport (RTP) protocol which runs over UDP. The intention of the RTP protocol was to become the default protocol for carrying real-time applications because of its efficiency and its management of real-time traffic, but the reality was completely different from its expectations. According to the obtained results, the adoption of such optimal protocols has not been reflected in the real world: the industry and the end-users are not adopting such efficient but complex protocols. The reasons include a mix of different factors. Real-time services arrived at the same time that the backbone networks suffered a huge upgrade and the access technologies reached the Mbps rate. Therefore just over-provisioning the network capacity any real-time service was able to run over the TCP protocol leaving aside the efforts put in efficiency with that specialised protocols. Other factors are that Internet industry did not support it with the required vehemence and moreover, users found many barriers because in most of the cases it was required to install a ”not-for-dummies” desktop application to enjoy the new services. Finally, the simplicity arises and the simple HTTP over TCP provides the versatility needed to carry any type of content.

Application protocol evolution

The Internet2 project identifies the application protocols based on the well-known transport ports. We are mostly interested in the evolution of HTTP and P2P traffic. To better understand the shown data, the different P2P technologies (BitTorrent, FastTrack, eDonkey2000, Shoutcast, Gnutella, Hotline, WinMX, Audiogalaxy, Blubster, Neo-Modus, Carracho and Freenet) are grouped into a single dataset called File Sharing. Furthermore, as Karagiannis et al. mentioned, P2P applications usually use random transport ports. Accordingly, this study considers that it is interesting to track the evolution of unidentified traffic, assuming that most of it could be also P2P traffic. Karagiannis et al., Parish et al. and Archibald et al., also mentioned other techniques to detect the unidentified traffic, including deep packet inspection, which seeks deterministic character strings in the packet payload.

As Figure 3.10 shows, HTTP traffic has undergone an important increase, especially after the mid-2000s, whereas File Sharing traffic has suffered a constant decrease in percentage values (though not in absolute values). HTTP has increased from a 6% to a 42% share of the total traffic, while File Sharing traffic has fallen from 31% to less than 2%. Conversely, the unidentified traffic, which we conjecture is hidden P2P traffic, has doubled from 20% to 44% during the analysed period (the 2003 sample is neglected here, as we do not have more in-depth information about what happened in that period). Therefore, we could assume that P2P is not decreasing, but it is doing a migration from well-known ports (file sharing classification) to random ports (unidentified classification) and increasing its presence in the total percentage of the Internet traffic. .

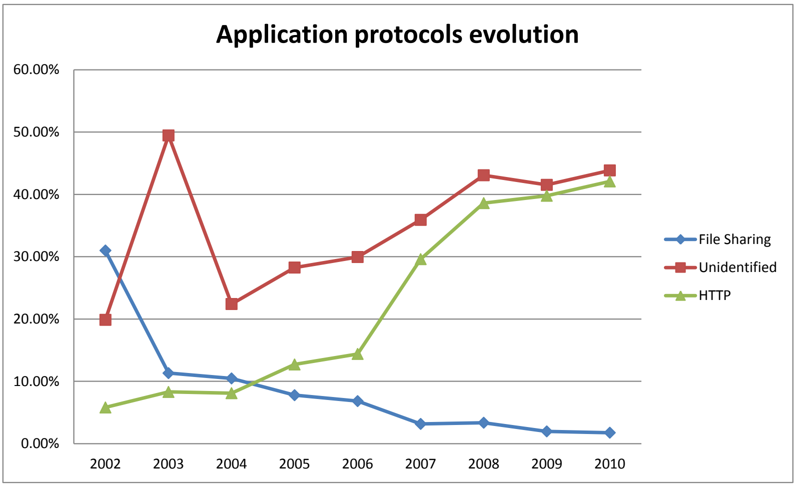

Figure 3.10: Application protocols traffic evolution

According to the obtained results, HTTP traffic has became one of the majors contributors of Internet traffic. This observation is in line with the evolution of the Internet users that we previously presented. The democratisation in the use of Internet has motivated the emergence of a vast quantity of web contents (looking for a business opportunity in the Internet market). But, why the web protocol? In previous sections we gave some hints: in the PC era, almost all operating systems have a browser and browsers are simple to use and the main clients for the consumption of Internet content. This simplifies the adoption of the technology and facilitates the proliferation of new (web)services. Another factor that has motivated the increase of the HTTP traffic is the *webification* of the applications. For example, many services like e-mail, file downloading or online video have adopted the web protocol. This concept is called *service-on-top-of-the-internet* (e.g. video-on-top-of-the-Internet). Therefore, HTTP will play an important role in the Internet ecosystem as it will agglutinate a wide variety of services/applications that in the past would have their own protocol and it will have a deep impact in the Internet traffic as the projections point to that HTTP will be the most relevant protocol for content transportation.

Packet size evolution

The Internet2 project analyses the data stream packet lengths and classifies them by size. Parish et al. and Archibald et al. found that applications can be identified using statistical techniques based on packet size distribution. For example, YouTube streaming has a typical traffic pattern that can be identified using these techniques.

The Maximum Transfer Unit (MTU) is a parameter that restricts the IP packet size. Classical Ethernet technologies offer a typical MTU of 1500 bytes. New Ethernet implementations offer larger MTUs up to 9000 bytes (usually called Jumbo Frames), but they are not fully implemented in the entire scope of the Internet.

As we can observe in Figure 3.11, the small packets (less than 100 bytes) have oscillated between a 35% and a 45% share of the traffic during the analysed period, whereas the share of the large packets (between 1401 and 1500 bytes) decreased during the first 4 years and began to increase during the last period until leading the classification with 42%. The last increase of the share of large packets has been accompanied by a relative decrease in the number of medium packets (between 100 and 1401 bytes). Jumbo frames are not observed on a large scale, and they have a residual usage. Real-time applications such as VoIP are closely related to small UDP packets to minimise the impact of packet loss. Small packets are also present in a large number of TCP connections due to the data to transmit is too small and it does not allow the TCP window to increase. Classic applications such as file downloading or multimedia applications such as video streaming (with buffering) are more connected with large packets when the TCP windows is sufficiently large and stable.

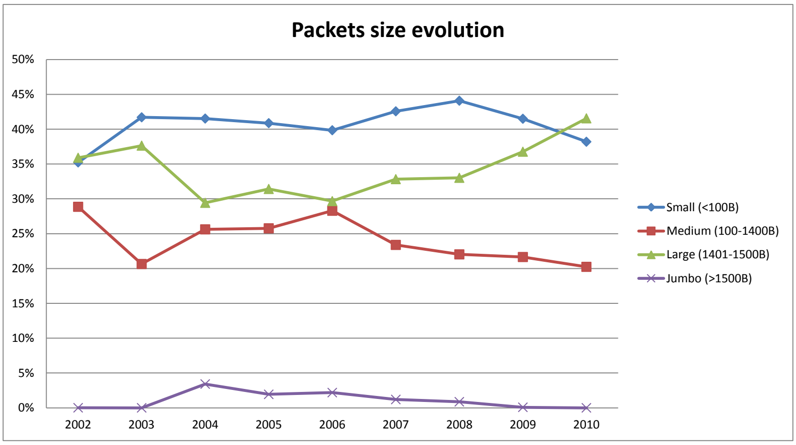

Figure 3.11: Packets size evolution

Internet application trends

The Internet application trends are obtained from the Cisco and Sandvine reports. Sandvine offers a snapshot of the application distributions in the fall 2011. Sandvine is not completely strict with its protocol definitions, but we can easily associate the service name with its application protocol (i.e., Netflix, YouTube or Facebook appear as applications when running on top of the HTTP protocol).

Focusing on the aggregated traffic in Table 3.2, applications that use the HTTP protocol almost exclusively make up the top 10 positions in both 2011 and 2014 snapshots. Only BitTorrent, RTMP, and MPEG-TS appear as applications that do not use the HTTP protocol. Table 3.2 also shows the Internet users' preference for multimedia applications, especially video streaming. Video applications such as Netflix, YouTube or Flash Video accumulated more than 40% of the total download traffic in 2011 whereas the group formed by Netflix, YouTube, Amazon Video and Hulu accounted more than 46% in 2014. This observation reflects two main points: applications are moving to the HTTP protocol instead of using particular protocols and video over HTTP is by far the most contributor to the Internet traffic.

Rank	Application (2011)	Share (2011)	Application (2014)	Share (2014)
1	Netflix	29.03%	Netflix	31.09%
2	HTTP	16.59%	YouTube	12.28%
3	BitTorrent	13.47%	HTTP	11.84%
4	YouTube	9.90%	BitTorrent	5.96%
5	Flash Video	3.04%	SSL	3.80%
6	RTMP	2.81%	iTunes	3.33%
7	iTunes	2.69%	MPEG	2.62%
8	SSL	1.96%	Facebook	1.83%
9	Facebook	1.84%	Amazon Video	1.82%
10	MPEG	1.49%	Hulu	1.58%
		82.83%		74.58%

Table 3.2: Application protocols distribution in fixed network (North America). Fall 2011 and 2014. Source: Sandvine

In Figure 3.12, we can observe the aggregation of successive CISCO forecast reports. The figure shows the traffic consumption per application segment for the period 2005-2015. Internet video (video streaming, video-on-demand, video calling) will see the largest increase and will be the predominant application inside the IP network. This trend is consistent with the growth curve of Internet users. Web, e-mail and data segment will have a slight but continuous increase. In other segments like online games and file sharing, the predictions do not match with the reality in terms of growth trends.

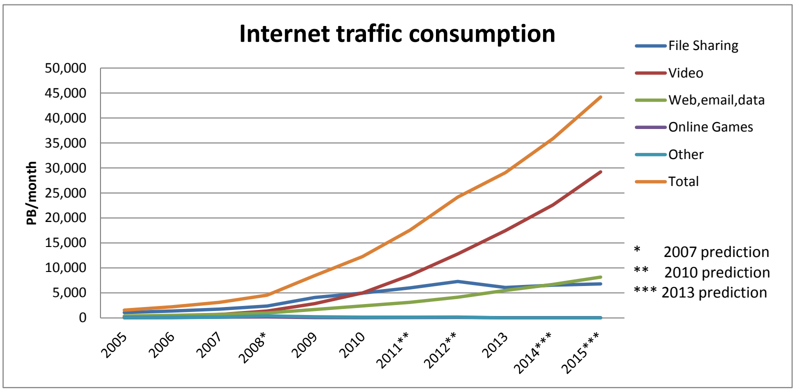

Figure 3.12: Internet traffic consumption. Source: CISCO

Network technologies

Network transmission speeds have exponentially increased in the last fifteen years. The evolution of the network speeds differ depending on the analysed part of the network infrastructure. Backbone networks have higher speeds than access networks, as the core network must support the traffic load inbound from the access networks. Tables 2 and 3 show the theoretical maximum speeds of the different network technologies. Furthermore, the tables include when a given technology was standardised. A standard specification is usually revised during the following years, and commercial deployments of such technology take several years after standardisation.

Backbone networks are based on optical multiplexing technologies that permit allocating more than one optical channel in the same optical fibre, including devices that multiplex up to 160 optical channels with speeds between 2.5 and 40 Gbps (modern optical multiplexers reach up to 100 Gbps to transmit the standard 100 Gigabit Ethernet [33]). Table 3.3 shows the most popular technologies: Dense Wavelength Division Multiplexing (DWDM) and Coarse Wavelength Division Multiplexing (CWDM).

As seen in Table 3.4, access network technologies have evolved from

Technology	Number of Channels	Lambda Speeds	Standard	Year
DWDM	40,80,160	2.5 Gbps, 10 Gbps, 40 Gbps	G.692	1998
CWDM	18	2.5 Gbps	G.694.2	2002

Table 3.3: Optical multiplexing technologies

some Mbps to one hundred Mbps during the last decade. We have only considered xDSL (ADSL, ADSL2, ADSL2+, VDSL and VDSL2), HFC (based on the DOCSIS standard) and 3G Wireless technologies (GPRS, UMTS, HSDPA, HSDPA+ and LTE) to analyse their speed evolution. The table shows theoretical maximum speeds, and, in practice, operators only assure a minimal bit rate to end users that is typically 10% of the contracted speed. The obtained maximum speed depends on different factors, including the state of the line or distance to the central office. 3G technologies typically share the spectrum. The maximum speed is therefore shared among the users of cell.

Technology	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
xDSL	8			12	24	55	100			
HFC	5	10	30				120			
2.5/3/3.5G	0.4	2		14					42	100

Table 3.4: Speed of access network technologies (speed in Mbps)

We have excluded FTTH networks based on xPON and new 4G access technologies, as we have considered that they were not widely deployed during the 2000s. Nevertheless, xPON and 4G technologies will enhance access speeds during the next years, and they will allow running next-generation services such as high-definition video on top of them.

3.5 Internet traffic analysis

Figure 3.8 shows how the number of Internet users has been increasing, especially during the last five years. Considering Figure 3.9, TCP remains the predominant transport protocol, even when streaming applications running over UDP offer better performance. Akhshabi et al. [7] review different strategies for video streaming and state that the TCP congestion mechanisms do not necessarily affect video streaming performance, as video players have adapted to throughput variations. Video services can use simpler solutions based on HTTP over TCP due to the HTTP advantages, instead of the more specialised RTP over UDP.

In Figure 3.10, we see that HTTP traffic has increased considerably during the last five years. Unidentified traffic has a similar development, while file sharing traffic (P2P application well-known ports) has continuously decreased during the same period. In Figure 3.11, we can also observe a significant increase in large packets during the latest period. As stated above, large packets are associated with video and file downloading applications. This observation matches Brownlee et al. [15], who have already detected this traffic pattern in their measurements.

If we combine these findings with the data in Table 3.2 and Figure 3.12, we can affirm that HTTP has become the predominant protocol due to the emergence of multimedia and file downloading global applications such as Google, YouTube, Yahoo, Facebook or Dropbox. We believe that the HTTP protocol flexibility and network upgrades attracted new services to the Web (*webification*), while the benefits of these services triggered a spectacular increase in the number of Internet users.

Video applications migrated to HTTP (over TCP), and new services are being implemented to run over the same Web protocol. Some P2P traffic, such as file sharing or P2P video streaming services, also migrated to HTTP. Labovitz et al. [63] observed the decline of P2P traffic in comparison to HTTP based on their measurements. However, they are likely not considering that there is a large amount of hidden P2P traffic using unknown ports. This share of P2P traffic is difficult to quantify, as Karagiannis et al. and Parish et al. observed.

The boom in web applications during this period is also related to the deployment of new network technologies in both the backbone and access networks of the operators. As seen in Tables 3.3 and 3.4, Optical Multiplexing (WDM) technologies and the upgrade of the DSL, HFC and 3.5G access speeds facilitate the appearance and spread of new multimedia applications. Services such as Netflix or YouTube require broadband access networks to offer an optimal user experience. Table 3.5 shows that high-quality video streaming requires download speeds between 1 and 5 Mbps depending on resolutions. This means that access networks should assure at least these minimum speeds in their Service Level Agreements (SLA). This kind of bandwidth-consuming application could thus not appear until the network technology was ready.

Video Resolution	YouTube	Netflix
1920x1080	3.75	5
1280x720	2.25	3
854x480	1.25	1.5
640x360	0.768	nd
480x270	0.768	nd
400x226	0.384	nd
176x144	0.1	nd

Table 3.5: Bitrates for different video resolutions (in Mbps)

3.6 Impact on the overall Internet

The impact of this study is related to both protocols and network infrastructures. This chapter identifies three main areas of impact related to upgrading the access networks, core traffic management and network economics, all under the assumption that the trends we identified in this study continue for the next few years.

3.6.1 Impact on access networks

If the consumption of multimedia applications such as video streaming continues increasing at the current rate, access networks based on copper technologies will soon reach their physical limit. There are two possible solutions to handle the growth in consumer traffic.

The first is continuing to overprovision access capacity, necessitating an upgrade to optical access technologies such as FTTx. As mentioned above, copper technologies can already offer higher rates. However, these higher rates depend on the line quality and the distance between the subscriber and the operator premises. Considering the evolution of Internet video from Figure 3.12, we can see that video consumption will be multiplied by 4 between 2011 and 2015, a factor that current access technologies can no longer support, as we have shown. Therefore, deploying the optical access networks should be the way forward to ensure the operation of future multimedia services.

The second option for accommodating the increased demands of consumer traffic is to optimise protocols and applications to increase data transmission efficiency. The pressure to employ such optimisations did not yet played a major role, according to our findings, as the bandwidth growth kept pace with or was even faster than the bandwidth consumption growth. However, if deploying new access technologies cannot keep pace with consumer demands in the near future, it might become necessary to squeeze the maximum out of the current network deployments. This can be done by applying traffic engineering techniques, including TCP ACK starvation, QoS, channel reservation procedures or implementing new high-speed TCP versions or the second version of HTTP. New video compression codecs and implementing video techniques such as Scalable Video Coding (SVC) or Multiple Description Coding (MDC) should also reduce the bit rate adding complexity to the subscriber devices (PC, mobile phones, set-top boxes or media players).

3.6.2 Impact on the core network

Similar to the access network, we might see two different developments in the Internet core, depending on whether bandwidth overprovisioning remains feasible. If it does, we expect to see the same protocols deliver content, i.e., HTTP over TCP. Under this assumption, traffic engineering techniques based on application ports are almost useless, as most applications run the default HTTP ports. Therefore, we believe that traffic management techniques based on IP source address filtering to identify Web server activity and URL filtering to identify Web sites could be the most appropriate. With these strategies, network managers can apply traffic differentiation, deciding that if a packet comes from a CDN network IP or URL, it would likely be multimedia traffic and will require different processing or routing. We consider that these techniques are simpler and more scalable than other more sophisticated methods based on deep packet inspection or statistical traffic classification. A possible solution is using the SDN OpenFlow switching technology [14] that allows filtering inbound flows based on different predefined rules.

Even if the core bandwidth is no longer sufficient to support the growing traffic, we doubt that an effect similar to the access networks might occur. QoS architectures that have not been widely implemented due to their complexity and a lack of pressure might again seem to be an interesting solution under these circumstances. Conversely, as we have seen, the traffic that will lead to resource exhaustion will likely be video traffic, i.e., traffic that needs one of the highest service guarantees. Prioritising streaming traffic over best-effort traffic will thus not necessarily be able to achieve better network utilisation, as the preferred traffic class alone will likely cause overloads. Under the assumption that Network Neutrality will be mantained, there is probably less room for optimisation in the network unless bandwidth overprovisioning or CDN solutions are applied as it will be discussed in the next section.

3.6.3 Impact on network economics

Although the backbone is currently sufficiently provisioned and will be easily upgraded by adding new channels to the fibre (increasing its capacity), the massive emergence of multimedia services requires new network strategies to satisfy the user demand. The Internet is formed by multiple network operators and the exchange of such amount of traffic require the establishment of many interconnection agreements to provide efficient services.

CDNs appeared as main actors for distributing content between the content providers and the subscribers located at the access networks. However, the asymmetry of the subscriber connections (downstream is higher than the upstream) generates imbalanced interconnections between operators that could end in disputes [26]. Access ISPs are in a dominant position as they have the keys to open the door of the customers. Therefore, when a transit ISP, CDN or large content provider wants to deliver its traffic into an access ISP, it could happen that the latter require a counterpart. Traffic senders claim that traffic exchange will be done using a settlement-free agreement (peering) as the access ISP is a termination and the only way to reach users but the access operator can refuse to do so if it has enough bargain power. Then, new interconnection agreements such as paid-peering may be required to solve these issues. However there are some situations in which it is not clear who may pay whom. The content value is sometimes higher than the delivery and vice versa. Access operators are in a privileged position, as they are directly connected to subscribers and can exert force on other ISPs. In this case, the same Internet market or a competent regulator body could again restore the equilibrium. From the point of view of a large content provider or CDN company it is commonly accepted to pay a *toll* for delivering traffic, but from the point of view of a transit provider this situation is odd because they are supposed to be the ones who are paid, not the ones paying. This new scenario is changing the status-quo of some Internet players such as transit providers that are being required to re-organise their business activities to offer new services like hosting or CDN services instead of only focusing on the transit business.

We observe these disputes through some proven examples in which access ISPs have discriminated some traffic over other, degrading the quality of service perceived by the end users (WhatsApp messaging servicer and Skype's VoIP service were blocked by some cellular ISPs in Europe [2] and Netflix's video service was degraded by some access ISPs in US). Access operators have acted in this way to put pressure on transit ISPs for exceeding the limits of the peering agreements and not paying for the extra traffic via a paid-peering arrangement and because some of these services/applications rivals with some of their services such as IPTV. These disputes between Internet players with conflicting positions raised debates in many countries about network neutrality and the open-

[2] See BEREC `http://ec.europa.eu/digital-agenda/sites/digital-agenda/files/TrafficManagementInvestigationBEREC_2.pdf`

ness of the Internet. In Europe and USA, regulators have long listening to the different parties and analysing the situation in order to ensure innovation and competition in the Internet. Recently, the FCC, the US regulator, presented a new directive [3] in favour of network neutrality in which presentation text we can read:

An Open Internet means consumers can go where they want, when they want. This principle is often referred to as Net Neutrality. It means innovators can develop products and services without asking for permission. It means consumers will demand more and better broadband as they enjoy new lawful Internet services, applications and content, and broadband providers cannot block, throttle, or create special "fast lanes" for that content. The FCC's Open Internet rules protect and maintain open, uninhibited access to legal online content without broadband Internet access providers being allowed to block, impair, or establish fast/slow lanes to lawful content.

3.7 Conclusions

This chapter presented the results of the analysis of a full data set of IP traffic measurements. Specifically, we have analysed the evolution of IP traffic between 2002 and 2010. Our main contribution is identifying a correlation between the deployment of new network technologies and the distribution and evolution of the IP protocols and their packet sizes. In particular, we have found that the network upgrade occurring during the mid-2000s caused a large adoption of the HTTP over TCP protocol for multimedia purposes such as video streaming. We state that the large-scale adoption of the HTTP protocol is related to its simplicity and flexibility. Furthermore, we suggest that P2P traffic is not declining but increasing, hidden in the unidentified traffic.

The rise of the HTTP protocol is based on how easily Internet applications have emerged worldwide. A simple but not optimal protocol (HTTP

[3] See FCC `https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf`

over TCP), in combination with a broadband network, has won the battle over other more specialised protocols. The results provide sufficient information to conclude that the HTTP protocol has reached the sufficient maturity and the Internet industry no longer needs to debate whether to create new protocols to provide new applications to the market. In this sense, we believe that new applications will be implemented using HTTP and, in the near future, using the HTTP version 2 that includes new capabilities to deal in a efficient way with the next generation of Internet services.

Finally, the chapter offered avenues to continue working in different areas, including analysing traffic influence on network architecture and traffic management. If applications continue consuming resources at the same increasing rate, we consider that new access network upgrades will be required. Furthermore, as we have observed from the obtained statistical data, Internet traffic is migrating to HTTP. Simpler traffic management techniques based on IP or URL source address filtering could therefore be more appropriate for managing traffic flows. Providing overprovisioned resources to the backbone networks is a strategy that could be more effective than implementing traffic management techniques.

Conversely, distributing Internet content using CDNs is a hot topic that is causing several changes in the Internet landscape. As we stated, video is the killer application, and new high-definition video services will require new investments that will possibly affect the interconnection landscape. Next chapter deals with the emergence of the CDNs and presents the implications of such content delivery solution in the Internet ecosystem.

Chapter 4

THE EMERGENCE OF CONTENT DELIVERY NETWORKS

In the previous chapter we have analysed the evolution of the Internet traffic from the point of view of the Internet protocols used to transport content. We witnessed an exponential increase in the demand of Internet content and we showed that although during the last years there have been many technical improvements, the use of transport and application protocols is quite stable and it is basically relied on TCP and HTTP. We have also observed that a large amount of traffic must be accommodated and we wondered whether *traffic overprovisioning* and *Content Delivery Networks* (CDNs) would be the best candidates to optimise the traffic distribution of high-demanding content like Internet video. In this chapter we discuss about the emergence of CDNs, their different typologies and how actors are using these solutions (see Figure 4.1). We put the CDNs in context with the asymmetric pattern of the current Internet traffic (difference between download and upload flows) which originated many disputes between ISPs. We also analyse how CDNs have influenced other Internet players making them to evolve their business activities. This chapter helps to understand the motivation and impact of CDNs: the high

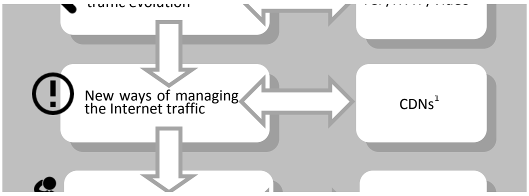

Figure 4.1: Thesis outline: the emergence of CDNs.

demand of video and new web applications required new technical solutions to accommodate the highly asymmetric and cacheable traffic. CDNs appeared as an efficient solution to deliver these traffic patterns and new specialised companies entered into the market while legacy ISPs reacted acquiring and re-selling CDN services from these specialised companies or by designing their own CDNs. All these facts together motivated a reconfiguration of the interconnection ecosystem promoting the competition and the increase of relationships between ISPs.

4.1 Introduction

Quality of Experience (QoE) for end users is a key factor in the success of any Internet service. Consequently, Internet applications require new strategies for distributing their content while offering the best user experience. Content Delivery Networks (CDNs) are technical solutions for providing high-performance content distribution. CDNs mirror the most popular and demanded web content to a set of distributed cache servers located at the edge of the network close to the end users. CDNs act as bypasses around the network that attempt to redirect users to the most suitable server based on performance criteria while avoiding saturated links. The main consequence of this bypass is that Internet Service Providers (ISPs) and content providers (CPs) are more efficiently interconnected, resulting in better performance and in a considerable reduction of traffic in the core of the Internet.

In this chapter we perform an extensive study about what the CDNs are and their implications in the Internet ecosystem. First, we introduce the Internet context in which CDNs appear and then we present the different CDN mechanisms and the differences versus traditional web delivery (hosting). Second, we conduct an analysis of different large and long-tail CPs with the objective of seeing whether they are using CDNs. Third, we provide a classification of CDNs based on the type of Internet player that offers this service. Fourth, we discuss about the policy implications of CDNs in the new Internet ecosystem.

4.1.1 Related work

Several related works have previously analysed how CDNs are changing the way content is distributed. Clark et al. [26] focus on the economic impact of CDNs on the interconnection ecosystem and observe a high asymmetry between the traffic flows from content providers to end users in comparison with the flows going in the reverse direction from users to content providers. This asymmetry has inspired many disputes between CDNs and access ISPs about their interconnection agreements, which have led to a change of their terms. On one hand, access ISPs are the destination of large amounts of traffic originated by CDNs. Thus, they believe that they must demand fees for delivering this traffic to end users using their network resources. On the other hand, CDNs consider that this is the only way to reach the end user because access ISPs are in a strategic position. Clark et al. discuss all these disputes in a net neutrality context and conclude that CDNs may pay access ISPs even if it could be considered a risk to competition, because they believe the market will provide enough competitive transit prices to sustain the ecosystem.

Dimitropoulos et al. [38] and Gao [45] have also analysed operator interconnections from a more technical perspective. They use a methodology to quantify the type of inter-Autonomous System (AS) relationships that exist in the Internet and classify them into three groups based on the state of the Border Gateway Protocol (BGP) messages: customer-to-provider, peer-to-peer, and sibling-to-sibling relationships. They found that more than 90.5% of the relationships are customer-to-provider, less than 8% are peering and less than 1.5% are sibling relationships (agreements between ASes of the same organisation). Other authors like Siganos et al. [97] have focused on generating a global Internet topology using a power law methodology. They have obtained results about the percentage of nodes that a node can reach in each hop, demonstrating that more than 99% of the Internet nodes can be reached within a maximum of six IP hops.

Shavitt and Weinsberg [94] recently discussed the topological trends of content providers. They create a snapshot of the AS-level graph from late 2006 until early 2011, then analyse the interconnection trends of the transit and content providers and their implications for the Internet ecosystem. AS graphs are built by traversing IP traceroutes and resolving each IP address to its corresponding AS. Shavitt and Weinsberg proved that large content providers like Google, Yahoo!, Microsoft, Facebook, and Amazon have increased their connectivity degree during the observed period and are becoming key players in the Internet ecosystem, strengthening the idea that the Internet is becoming more meshed.

4.2 Bypassing the Internet

The Internet architecture implemented until the early 2000s was presented as a multi-tier hierarchic structure [63]. Tier-1 ISPs were on top of the hierarchy followed by the tier-2 regional ISPs and the access ISPs at the lower part of the hierarchy connecting the end users. In this scheme, tier-1 ISPs were highly connected to other ISPs and offered transit services to other ISPs in lower layers. Content was distributed through access ISPs or, in the best cases, through ISPs located at advantageous points. Traffic flows were required to go up and then down in the hierarchy to reach end users (see Figure 4.2).

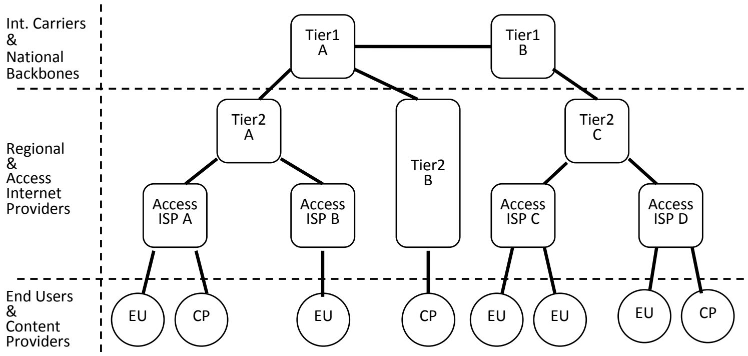

Figure 4.2: Traditional Hierarchic Internet Structure. Adapted from Labovitz[63]

Currently, Internet architecture is evolving and introducing new peering agreements, and most of the ISPs have increased their connectivity level. CDNs have emerged and have modified the interconnection paradigm as they move the servers from which end users download content closer to them, bypassing transit networks for these connections. In addition, apart from the new CDNs, some telco operators have also adapted and upgraded their networks and systems to offer CDN-like services (see Figure 4.3).

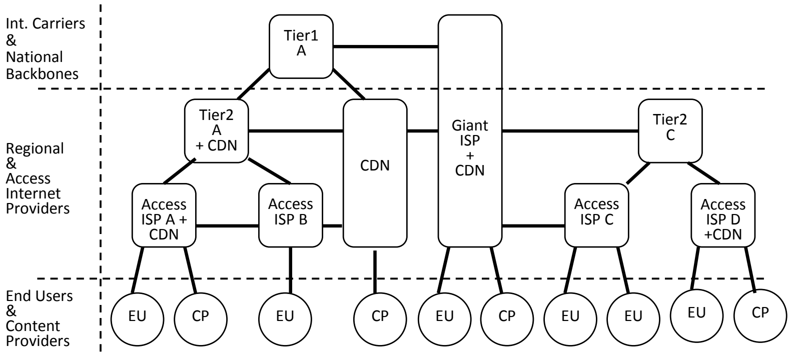

Figure 4.3: Current Internet Structure. Adapted from Labovitz[63]

Figure 4.2 depicts how content hosted in the access ISP A must be transmitted up and then down in the hierarchy passing through the different operators to reach end users of the access ISP D. In Figure 4.3, we can observe a different Internet topology in which content hosted on CDNs ”bypasses the network” and easily reach the end users thanks to direct and faster connections. This bypass means a breakthrough in network performance and, at the same time, it optimises the network resources thanks to the use of cache servers near the end user. The content of a CDN does not need to be transferred through the entire network each time an end user makes a request. CDNs cache content close to the end users and deliver a cache copy without affecting the rest of the network.

Figure 4.4 shows the difference in the network resource utilisation between the traditional hosting scheme and the CDN paradigm. The traditional scheme (a) requires transmitting a flow with the content for each user request, while in the CDN alternative (b) the CDN data centre (the CDN element that first receives the content from the content provider) transmits a flow to each cache server and then each user request is served by the cache server. This comparison helps to illustrate the traffic savings in the backbone made possible by the CDN. However, we note that this diagram is a simplification of a more complex scenario. For example, we do not consider the content expiration that would require a new data replication from the CDN data centre to the cache servers.

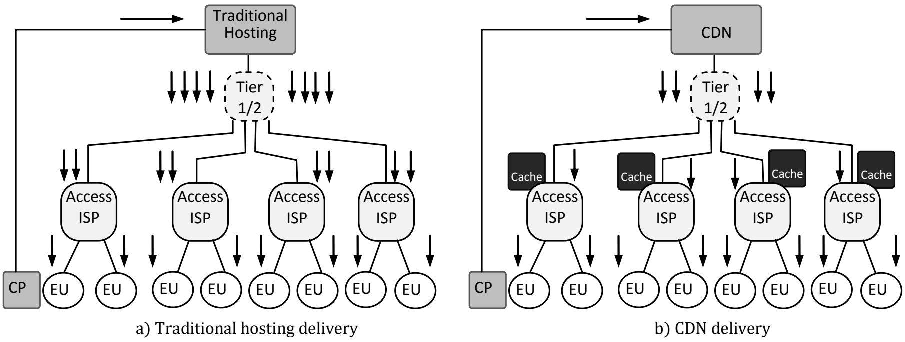

Figure 4.4: Traditional Hosting vs. CDN Delivery.

4.3 The taxonomy of content delivery networks

The term Content Delivery Network is a concept of network structure that can lead to different interpretations. A content delivery network is defined as a technical solution that deploys many servers in many distributed locations to offer the best QoE to the end user. However, this definition is too generic and each provider implements its network solution in a different way. Huang et al. [53] identify two main types of CDN. On one hand, there are CDNs that build large content distribution centres in few but strategic locations, and that connect these centres with high speed links to the ISPs (see Figure 4.5.a). In this strategy, CDNs try to place the distribution centres at vantage points (VP) that are simultaneously close to many large ISPs. On the other hand, there exist CDNs that deploy their cache servers inside the ISPs. Level3 is an example of the first type of CDN, whereas Akamai is an example of the second type (see Figure 4.5.b).

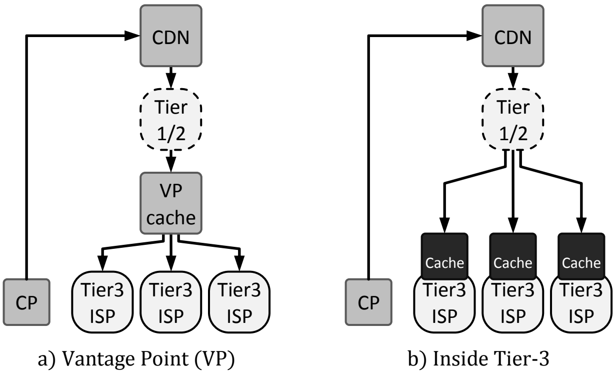

Figure 4.5: CDN Strategies.

Both types of CDN strategies have technical advantages and disadvantages. The first strategy has a simpler management overhead, but at the expense of increasing the response time. The second strategy has a better performance in terms of response time due to its location being very close to the end user. However, the management and server deployment are more complex. CDNs must analyse what is the most profitable strategy depending on the content: investing in distributed network equipment and management, or investing in few vantage points and buying network capacity.

In addition to these two types, there are hybrid architectures that combine inside-ISP-server-collocation with large CDN data centres, other more hierarchic structures that do not have direct connection with the target Access ISP, or CDNs that manage caches using P2P structures [6]. Another emerging paradigm is the CDN federation [87], i.e. the interconnection of different and heterogeneous CDNs to obtain a greater worldwide presence assuring a predefined Service Level Agreement (SLA). There are two models of CDN federation [105]: Figure 4.6.a shows the *bilateral model* where every member directly interconnects with other CDN members of the CDN federation. Figure 4.6.b shows the exchange-based model where every member of the federations interconnects with a centralised hub that performs some of the interworking functions (e.g. billing and routing) on behalf of all members.

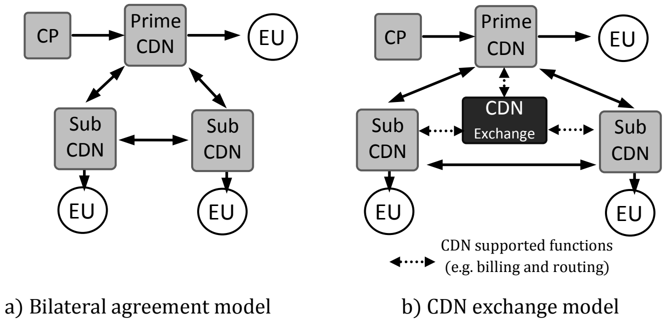

a) Bilateral agreement model

b) CDN exchange model

Figure 4.6: CDN Federation models. Adapted from CISCO[105]

4.3.1 Content delivery

Technically, a CDN is formed by a content server network and a specialised Domain Name Sever (DNS) network that redirects end users to the most appropriate content server based on a sophisticated algorithm. The basic mechanism of a CDN consists of the following steps (see Figure 4.7):

- An end user (EU) queries its local DNS (LDNS) to resolve the IP address of a web resource (e.g. http://images.example.com).
- The LDNS connects to the authoritative DNS server of ”example.com” which returns the CNAME ”server1.cdn.com” in response.
- The LDNS connects to the authoritative DNS server of ”server1.cdn.com” which finally returns the IP address (usually, servers return two or more IPs to allow client-side load balancing) of the most convenient CDN content cache server that hosts the example.com content.

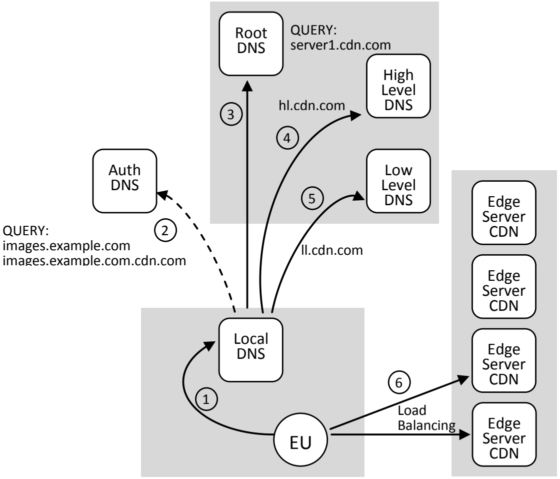

Figure 4.7: CDN DNS Redirection.

Most of the CDNs follow the delivery procedure illustrated in Figure 4.7. However, each CDN adds its own mechanisms to improve the performance of the content delivery. Some of these mechanisms are based on deploying the content and the DNS servers at strategic locations, implementing sophisticated cache algorithms at the content servers or implementing advanced load balancing algorithms and monitoring tools to redirect users to the most appropriate content server.

Table 4.1 shows a summary of the different content delivery strategies and their advantages. The content providers wanting to publish non-multimedia content will choose a traditional hosting service as it is the cheaper connectivity solution and covers the minimum user requirements sufficiently. Content providers that want to publish real-time resources, large amounts of data (photos, videos, audio files, etc.) or simply want to offer a better QoE to the end user will choose a CDN service. The performance of a CDN inside an ISP or located at a vantage point is quite similar although the former is slightly better. The content provider must evaluate if it is necessary to pay more for an extra level of quality. Therefore, using a CDN solution with cache servers inside the ISP is a good choice for content providers who believe that the quality of the services provided to their customers (e.g. minimising delay) is crucial to their value chain.

Content delivery solution	Quality of Service	Network infrastructure costs	Occupation of backbone resources	Network management (number of DNS and cache servers)	Applications	Target market
CDN inside an ISP	Extremely high	Very high	Low (if content is highly requested)	Very high	VoD and IPTV, photos, and large files	Large CPs
CDN located in a vantage point	Very high	High (depending on the number of PoP)	Medium (one flow per user request from the vantage point)	High	VoD and IPTV, photos, and large files	Large CPs and any small or medium CP whose business model is focused on delivering some type of QoE sensible content
Traditional hosting	Medium or low	Medium or low	High (one flow per user request)	Medium or low	Web hosting	Small and medium CPs

Table 4.1: Summary of the Different Content Delivery Strategies

4.4 Inspecting the anatomy of content delivery

In this section we perform an in-depth analysis to determine the content delivery solution each site is using and to classify them. The objective of this study is to confirm the use of CDNs and to measure the degree of CDN penetration across web sites of different sizes (from Hyper-Giant content providers to regional web sites) and to create a taxonomy with the different type of content delivery solutions.

Most of the large Internet companies use some kind of CDN solution for distributing their Internet content. However, this section aims to identify which strategies the companies are using. CDNs were designed to transport and cache large amounts of Internet content, such as HTML code, JavaScript, large files, images, audio, and video. The most appropriate way to identify whether a web site is using a CDN service is to inspect its HTML code looking for URLs linked to CDN providers. Web sites using CDNs map their resource URLs using names like *cdn1.mydomain.com* or *cdn1.mycdnprovider.com*. The redirection of a link to a well known CDN provider is clear evidence that a particular web site could be using a CDN solution. Otherwise (e.g. a link to an unknown third party domain), it is not, unless we prove that the domain is a CDN.

To perform this task, we have carefully analysed 43 web sites (see Table 4.2) available from Spain. We have selected 21 sites from ”the most visited web sites rank” created by Google [1] and a random selection of 22 long tail web sites. We distinguish between popular and long-tail sites because we want to observe the different traffic delivery models between them. We consider this sample of web sites sufficiently representative because the selected large content providers represent most of the Internet traffic according to [91]. Therefore, with only few web sites it is possible to create an accurate classification of the different CDN types.

After the selection of the web sites, we inspected their HTML code looking for image URLs. Then, we input these URLs in the *nslookup* tool to obtain the IP addresses of the servers hosting the images. Additionally, *nslookup* returned Canonical Names (CNAMEs) for those URL aliases of the web sites that use CDN solutions. Finally, for each IP address was searched in a database [103] to identify its owner ISP and, in combination with the interpretation of the CNAME, we can conclude the type of content delivery solution used by the different sites.

We conclude after analysing the selection of web sites that top-ranked sites tend to use CDN solutions while less-visited sites (usually classified as long tail sites) use traditional Internet hosting. The selection of one content delivery solution over others depends on the web site’s requirements.

Table 4.3 shows a classification of the content delivery services used by the analysed sites in which can be seen a wide variety of solutions. In the following list we describe each type of delivery service:

[1] Google provided a service that ranked web sites for customers who use its Ad Planner service. This service is no longer available. Last access June 17, 2013

Type of site	Sites
Large Content Provider sites	google.com, youtube.com, facebook.com, live.com, msn.com, yahoo.com, tuenti.com, wikipedia.org, marca.com, microsoft.com, as.com, elmundo.es, ask.com, eltiempo.es, elpais.com, lacaixa.es, wordpress.com, rtve.es, softonic.com, twitter.com, bing.com
Long Tail sites	chipspain.com, onabarcelona.com, ovellanegra.com, lacasadelosdisfraces.es, fibratel.es, www.planafabrega.com, km77.com, hotelmajestic.es, enriquetomas.com, viladepiera.cat, autoprint.es, labraseria.es, floristeriasnavarro.com, restaurantegorria.com, copiservei.com, inoutalberg.com, corchosgomez.com, barcelonavinos.es, dondisfraz.com, forocoches.com, www.condis.es, hoteles-catalonia.com

Table 4.2: Web sites Selection

- **Hosting**: operators whose business models are focused on hosting content in data centres located in vantage points.

- **Pure CDN**: network operators whose business models are focused on delivering content through their cache servers and renting link capacity to connect directly to the vast majority of Access ISPs.

- **Tier-1 CDN**: tier-1 operators that are capable of offering CDN services using their network infrastructures and deploying cache servers at vantage points.

- **Telco**: access operators that are capable of offering CDN or hosting services.

- **CP**: large content providers that deliver their content using their own network infrastructures which might include their own CDNs.

Type of delivery service	ISPs
Hosting	Acens, Arsys, IDH Telvent, Dinahosting, Easynet, OneAndOne, OVH, Softlayer
Pure CDN	Akamai, Edgecast, Amazon Cloudfront
Tier-1 CDN	Level 3, NTT, Telia, Cogent
Telco	Telefonica, Colt, BT, Digiweb
CP	Google, Microsoft, Yahoo, El Pais, Softonic

Table 4.3: Content delivery classification.

According to our dataset of web sites, large content providers are more likely to use pure CDNs, tier-1 CDNs, and their own CDN solutions in the same proportion. On the other hand, long tail sites prefer hosting services as a first option and telco solutions as a second option. From the 43 analysed web sites we found that the content delivery market is quite diverse. Nevertheless, it is important to note that a pure or tier-1 CDN will offer a better user experience although at a higher economic cost than a hosting solution. In addition, the total amount of Internet traffic delivered by large content providers is at a different scale compared with the long tail content providers. Therefore, each site must select the best solution based on its requirements.

Table 4.4 examines the case of Akamai in detail, based on the analysis of the IP addresses returned by the DNS servers of different Spanish access ISPs. We can observe that almost all large access ISPs in Spain have Akamai cache servers inside their premises or at a distance of one hop. Akamai is the world's leading CDN, present in more than 1900 networks in 70 countries, serving nearly 30% of global Internet traffic [5], making it the prime example of this CDN paradigm. Other ISPs like Level3 or NTT have located their servers in data centres close to the access ISPs. Google, the largest content provider in the world, which is estimated to contribute between 6% and 10% of global Internet traffic [31], also uses vantage points to host its servers as close as possible to the access ISPs.

Another interesting aspect of this analysis is that tier-1 ISPs are competing in two markets. They continue offering tier-1 transit services to interconnect ISPs and at the same time they use their own networks to offer CDN services to content providers. Moreover, we can observe that there are some access ISPs with international backbones that compete in both transit and CDN businesses and also offer Internet connectivity to residential users.

Access ISP	CDN Strategy	Returned IP(s)	Details of the AS
Telefonica	Akamai inside the ISP	194.224.66.42, 194.224.66.19	AS6813 Flexnet Com. Int. de Telefonica
ONO	Akamai at one hop through Telia	213.248.113.32, 213.248.113.24	AS1299 Telia Net
Jazztel	Akamai inside the ISP	212.106.219.176, 212.106.219.130	AS12715 Jazz Telecom Global Spanish ISP
Orange	Akamai inside the ISP	90.84.53.16, 90.84.53.59, 90.84.53.64, 90.84.53.81, 90.84.53.9	AS5511 Opentransit. France Telecom, Orange IP Backbone
Vodafone	Akamai at one hop through Espanix IXP	92.123.73.59, 92.123.73.81	AS20940 AKAMAI-ASN1 Akamai Technologies European AS

Table 4.4: Access ISPs using Akamai CDN.

4.5 Content delivery implications

ISPs are focused on offering new delivery services because (1) they need to satisfy the content providers' demand for faster connectivity solutions; and (2) they need to compensate the decrease of profitability for of transit services [113]. ISPs can offer faster solutions in many ways. One of these ways is to apply QoS policies to prioritise the IP packets from the content provider. As an alternative, ISPs can deploy content cache servers inside their networks.

Applying QoS mechanisms like traffic prioritisation could be seen as the natural strategy for ISPs to accelerate some content traffic over the rest, as they have already deployed mature technologies such as MPLS. While it is fairly easy to implement DIFFSERV to provide preferential treatment to certain traffic flows in the same ISP, ensuring QoS across several ASes using INTSERV is much more challenging. The large-scale deployment of these mechanisms would be more a problem of coordination between operators rather than a technological issue. However, prioritising some content could degrade other traffic flows and thus violate the non-discrimination principle of the Internet. The implicit degradation in the transmission quality of one type of content in comparison to other types could be seen as an anti-”network neutrality” practice. In this context, regulators like FCC in United States [29] and the European Commission [73] are taking action on the matter.

In contrast, CDNs accelerate the delivery of content without directly degrading the rest of the IP traffic. In the previous section we observed that large content providers use CDN solutions to deliver their services. The use of this strategy has a huge impact on the core of the Internet, as most of the information is bypassed to the Access ISP. According to the most recent Sandvine report, YouTube as distributed by the Google CDN and Netflix as distributed by different CDNs (Akamai, Level3, LimeLight, and its own solution) represent 50% of total Internet traffic [91]. Such an amount of traffic bypasses the core of the Internet, decreasing the utilisation of the network. This also has other implications as it disrupts the traditional tiered hierarchy leading to a mesh model in which all ISPs want to increase their interconnections.

It is important to note that both strategies, applying traffic prioritisation and deploying CDNs, have the same main goal in common: accelerating the delivery of specific content. Large content providers want to reduce latency in their content delivery and it does not matter whether this is done using prioritisation or CDNs. Also, from the user’s perspective no explicit difference may be observed because users may perceive similar QoE. The interesting part is how different regulators are dealing with the different strategies. QoS mechanisms, which are the natural and in some cases the more cost-effective option for some ISPs, could be considered an anti-net neutrality practice as they prioritise some traffic over other. In contrast, CDNs are not being considered as violating net neutrality principles, although they offer ”faster lanes” for those content providers who can afford it, possibly also leading to a two-class (or more) Internet. In this context, one can argue that although CDNs are not degrading the rest of the traffic, how can a long tail video web site compete against a ”hypergiant” whose content is distributed using high speed connections?

The impact of preferential interconnection agreements between large content providers and ISPs could affect competition in the rest of the Internet ecosystem. Content providers with significant market power to hire or even implement their own content-accelerating services would obtain a better user perception, which can translate into higher popularity. End users would only access video web sites from dominant content providers as the rest of video sites would offer a poorer user experience. This scenario would have the consequence that long tail web sites could hardly compete as they would not be able to pay for the prioritised traffic or CDN services [27]. Consequently, the Internet would be formed by few but hyper-giant content providers that would aggregate most of the Internet's content (an oligopoly), and ISPs and CDNs that only provide fast pipes to those who are willing to pay for them. In this situation, small web sites would be forced to delegate part of their core business (e.g. the marketing channel) for being aggregated into the ecosystem of the hyper-giants [2] unless they could contract affordable content delivery solutions.

Thus, we encourage policymakers to address how CDNs affect network neutrality because their effects on the Internet ecosystem are potentially the same as those of traffic prioritisation. Both strategies offer similar user perception in which some content is delivered with better quality than others, and where one must pay for this increase in quality. Therefore, we consider that the debate over network neutrality should also include CDNs.

4.5.1 Becoming part of the CDN business

CDNs have also had a large impact in the interconnection ecosystem as they have changed the way ISPs offer connectivity solutions. CDNs have influenced in a noticeable price reduction in transit costs and this fact has generated new business opportunities. The previous section showed that heterogeneous content providers can choose between different content distribution solutions, confirming that the supply of services is wide and that operators are investing in this type of network solution.

[2] This is similar to the apps ecosystem in which two platforms, Android and Apple, control the distribution of applications through their exclusive online stores

Hence, this study confirms what Wulf et al. [114] previously observed and reveals that many carriers like Level3, NTT, Telia, or AT&T have deployed, acquired, or resold CDN solutions to compete in this emerging market. They are transforming their businesses and compensating for the price reduction of transit service offering CDN services where prices are considerably higher than for pure transit services. With this strategy, carriers are taking advantage of their existing networks and are deploying CDNs to obtain higher revenues. Figure 4.8 illustrates the carrier strategy to enter into the CDN market.

Moreover, the emergence of ISPs into the CDN market may foster the deployment of new enhanced services such as cloud computing or advanced security applications for both residential and enterprise markets. Therefore, the CDNs are more than accelerating or caching content, they are a new business opportunity that will open new doors to innovation and to commercialise new services.

Spanish ISPs are aware of this tendency, but only a few of them are offering CDN solutions. Only Telefonica and Orange are explicitly offering CDN services, though with different strategies. Telefonica has recently decided to deploy its own CDN while Orange has chosen to resell Akamai services. The rest of the operators are probably already assessing the possibility of offering CDN services and in the coming months there will possibly be new announcements. Table 4.5 below shows the current status of CDN solutions offered by the largest Spanish operators. Nevertheless, all these operators are connected to the Spanish Internet eXchange Point (IXP) Espanix[3], which means that those who do not yet have a CDN solution are in an optimal position to resell a third-party CDN service.

4.6 Conclusion

In this chapter we analysed and identified the different content distribution solutions used by various content providers, with a focus on the use of CDNs. The study analyses the URLs of different web resources in a set of

[3] See Espanix, Peering entre miembros, http://www.espanix.net/esp/peering.htm.

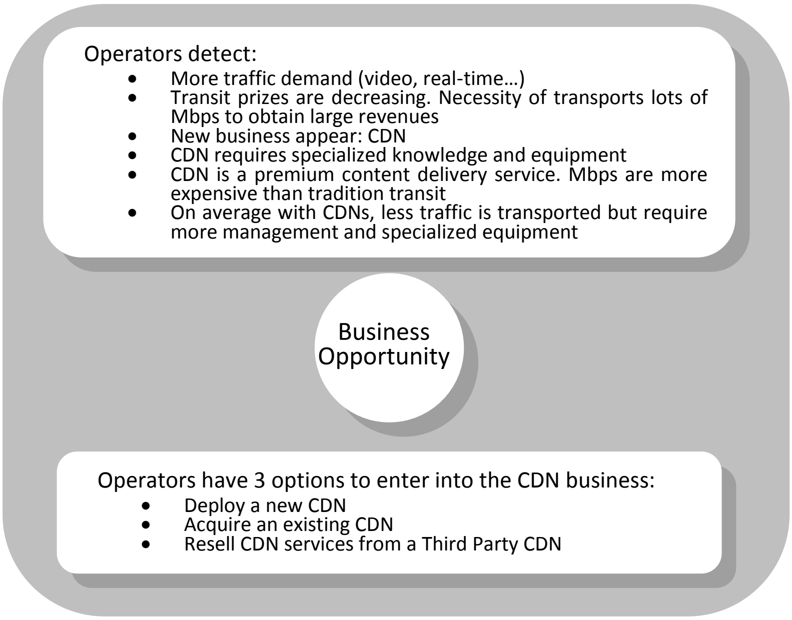

Operators detect:
- More traffic demand (video, real-time…)
- Transit prizes are decreasing. Necessity of transports lots of Mbps to obtain large revenues
- New business appear: CDN
- CDN requires specialized knowledge and equipment
- CDN is a premium content delivery service. Mbps are more expensive than tradition transit
- On average with CDNs, less traffic is transported but require more management and specialized equipment
Business Opportunity
Operators have 3 options to enter into the CDN business:
- Deploy a new CDN
- Acquire an existing CDN
- Resell CDN services from a Third Party CDN

Figure 4.8: ISP Strategy to compete in the CDN Business.

Carrier	CDN/Hosting Provider	Services	Type of co-operation	References
Telefonica	Telefonica	Web optimisation, video services, smart download, private delivery	In-house development	`http://www.telefonica.com/cdn/en/`
ONO	ONO	Traditional hosting	In-house development	`http://www.ono.es/empresas/productos/internet/alojamiento/`
Orange	Akamai	Web optimisation, video services, smart download, private delivery	Reselling	`http://www.akamai.com/html/about/press/releases/2012/press_112012_1.html`
Vodafone	Vodafone	Cloud services and traditional hosting	In-house development	`http://www.vodafone.es/empresas/es/soluciones-unificadas/servicios-en-la-nube/hosting/`
Jazztel	Jazztel	Traditional housing	In-house development	`http://empresas.wholesale.jazztel.com/servicios/wholesale`

Table 4.5: CDN Strategies of Spanish ISPs.

web sites, and the *nslookup* tool confirms network redirections pointing to CDNs. The results obtained from the study reveal that large content providers tend to use some type of CDN solution to spread their content over the Internet while web sites from the long tail use traditional hosting.

Among the different CDN solutions, two types of architecture can be distinguished. The first type of CDN solution (mainly used by Akamai) deploys its cache servers inside the different target ISPs. The second type (Level3, NTT, Cogent, etc.) deploys its cache servers in vantage points very well connected (at only one hop) to the target ISPs. The first type obtains a better performance in terms of latency and throughput at the expense of a greater investment in the number of servers and management, whereas the second type obtains slightly worse performance but with lesser maintenance and a smaller number of servers while offering a more competitive price.

The content delivery sector is an expanding market with a high level of competition, as can be observed by the heterogeneity of content distribution solutions used by content providers. This study finds that content providers use many different content delivery solutions and provides a classification depending on the nature of the solution. The study identifies content delivery solutions from pure CDNs, from tier-1 ISPs, from access ISPs, from large content providers, and from traditional hosting ISPs. Some of the conclusions after the analysis are that traditional tier-1 ISPs have evolved from simple carriers and now offer content delivery services close to the end user. Access ISPs have also started to deploy CDNs to distribute their own and external services, and at the same time they continue offering classic hosting solutions. Pure CDN ISPs have emerged and have deployed their data centres and networks at vantage points close to the end user. Large content providers are also involved in the CDN business because they have found that offering an optimum QoE is the key to success. Long tail content providers choose hosting solutions from third parties located in vantage points that also offer the possibility of upgrading their services through cloud and CDN solutions.

Finally, the chapter discusses the impact of preferential interconnection agreements between large content providers and ISPs, and the potential consequences for the rest of the Internet ecosystem. The article suggests that policymakers should analyse whether CDNs are susceptible to affecting network neutrality. The study also discusses how ISPs have begun to take part in the content delivery business by deploying their own CDNs, acquiring existing CDNs, or reselling CDN services from third parties.

4.6.1 Future Implications of CDNs in the New Internet Ecosystem

As this chapter concludes, large content providers commonly use CDNs in the distribution of their content while long tail content providers tend to use traditional hosting solutions. By analysing the ”hyper-giants” (Google, Yahoo, Microsoft, Amazon, etc.), one can observe that they are not just content providers, but they also are basically content aggregators. This means that these hyper-giants are attracting content from smaller sites or individuals and publish it via their high-speed infrastructures. Somehow a cannibalisation process has begun in which the hyper-giants are absorbing content from the long tail, entering fully into the niche of the traditional hosting companies.

Taking as examples a user who wants to publish his own blog, or an SME (small and medium-size enterprise) willing to create its webpage, the hyper-giants provide many easy and specialised solutions to cover these necessities. In the first example, users can create their own blogs, using for example Blogger from Google, and take advantage of the Google CDN and its excellent network performance at zero cost. In the second example, an SME can create its own web site using Google Sites for plain HTML5 pages or Google App Engine and Amazon Web Services for more sophisticated web sites. So an SME can publish its web site for free until the site exceeds a traffic limit or network resources, which is when the hyper-giants begin to charge the SME. Therefore, hyper-giants offer free minimum services (or not-so-minimum depending on your site requirements) to test their environment and when the free quota is exceeded, they begin to charge you for the consumed resources with competitive prices.

This new scenario has many advantages for end users and SMEs. They obtain free hosting and a free subdomain (example.blogger.com, example.appspot.com, etc.), several enterprise tools (e-mail, calendars, etc.), search engine optimisation (SEO) capacities, and the high-speed connectivity and high availability of a CDN. However, the hyper-giants offer these ”free services” as closed products and they impose their operating methodology through their own Application Programming Interfaces (API) and Content Management Systems (CMS). In addition, they do not usually offer a technical service, but they offer FAQs and technical forums where specialists and current users share their knowledge about a specific issue. The reason for this approach is that their cloud environments are completely ready to publish content and that user problems can be solved by consulting online forums. Therefore, the degree of freedom of customisation (web technologies, system performance, etc.) is limited in comparison with the hosting solutions, but they are in constant evolution offering more and better adapting services every day.

So, what is the benefit for the hyper-giants? The hyper-giants offer these free services because they obtain more potential users for other services, whom they can impose to use their integrated accounts (e.g. Google Accounts), and because these long tail sites can be used to position advertisements from third parties, which usually is the core business of the hyper-giant. Other secondary arguments for offering such services is that they provide an economic benefit when the sites exceed the minimum quotas, and that the hyper-giants can enhance the life cycle of their environments thanks to the contributions (suggestions, error detections, etc.) of the long tail users through Web 2.0 tools such as working groups, wikis, and forums.

Finally, there are many open questions to be answered about the future of the Internet ecosystem. Will content providers tend to choose hyper-giant CDN solutions instead of traditional hosting? Will this situation lead to a scenario featuring an overlay content network formed by the hyper-giants’ CDN solutions, in which traditional hosting ISPs could be removed from the market? Will the hosting ISPs be converted to resellers of the hyper-giants or carrier ISPs? Could this emerging situation generate oligopolies that could affect the competition?

This chapter has provided information about the implications of the CDNs in the Internet ecosystem and has presented a classification of the content delivery strategies. Next chapters aim to go further in the research of these network solutions and their interaction with other players and purposes a measurement tool that facilitates the analysis and discovery of different types of network interconnections.

Chapter 5

DESIGN OF AN INTERNET MEASUREMENT PLATFORM

In the previous chapters we have studied the impact of the increasing demand of Internet traffic and the implications of the CDNs in the Internet ecosystem. With the motivation of going deeper in our research of how the different Internet players are organised, we require to gather measurements and evidences to prove their relationships. And for this reason we have developed a new tool to measure and study more in detail the interactions between Internet players.

This chapter aims to present the measurement platform that we have developed to gather interconnection data from different Internet players (see Figure 5.1). First, we introduce the need of such tool and we compare it over other existing measurement platforms to research the Internet structure. Then, we test it to analyse many hidden interconnections between access ISPs and content providers which are crucial to understand the evolution of the Internet topology in the last years. One of our objectives is to discover many elusive interconnections like when a particular player places their content servers within the access ISPs. This type of observations shows us that these Internet players are really interested in getting closer to the end users to optimise their content delivery and this model of interconnection is transforming the global topology of Internet.

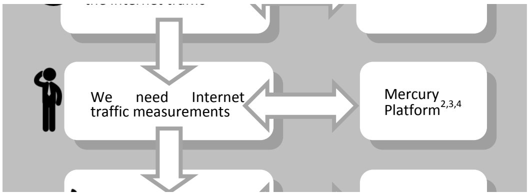

Figure 5.1: Thesis outline: design of a measurement platform.

5.1 Introduction

Discovering the hidden interconnections between access Internet Service Providers (ISPs) and Content Providers (CPs) is a challenge for researchers due to the lack of data available to the general public and due to the several weaknesses of the existing tools and methodologies to extract them.

From the point of view of the ISP business relationships, these direct interconnections are critical because of the increasing demand for multimedia content requiring an optimal quality of experience (QoE) for the end users. The main reason why these interconnections are elusive is because they are usually based on peering agreements (p2p) using the Border Gate Protocol (BGP) which is the responsible of configuring the interconnection links between different ISPs. When BGP is running, the edge routers of an ISP exchange BGP announcing messages which are only visible between peers and to their customers and this mode of operation hinders the propagation of these link establishments. Furthermore, we also focus on other type of hidden interconnection between access ISPs and CPs which consists in placing the CP servers infrastructure inside the ISP. This can only be seen combining an analysis of the server IP addresses returned by the DNS Servers and translating the IP addresses along the path between the end-host and the server to AS numbers.

Several efforts have been made to discover the Internet topology at the Autonomous System (AS) level. Historically, two approaches have been used: the analysis of BGP paths and the *traceroute* traces. The BGP approach discovers the AS topology using a set of distributed monitors that sniff the BGP announcing messages. This methodology uses the AS paths included in the BGP messages to infer the interconnections between the different ASes. However, BGP announcements describe the control plane which does not necessary correspond to the real path of the Internet traffic and this is one of the drawbacks of using that method. Also, it is not effective for discovering the hidden interconnection and requires much more monitors to obtain a global topology. In contrast, the traceroute approach has the advantage that can discover real path of the Internet traffic between two end points. As a disadvantage, the traceroute tool works at the IP level and requires an IP-to-AS mapping to infer the traversed ASes. The results collected with the traditional Van Jacobson traceroute are affected by the multi-path diversity, essentially preventing their use to mapping the AS interconnections.

Providing a global map of the Internet topology has been researched for a long time [45, 41], but is out of the scope of this chapter, our motivation is only focused on discovering the hidden interconnections between CPs and access ISPs. Toward this end we introduce *Mercury*[1], a platform for discovering the AS-level interconnections between content providers and content consumers. Mercury enables users to visualise the AS topology of access ISPs when they connect to other organisations and to identify how CPs organise their server infrastructure to reach their target audience. To perform this, Mercury combines the two approaches presented previously. On one hand, project participants run a desktop client in their hosts that collects traceroute targeting multiple URLs. We have extended a special version of traceroute, called Paris traceroute [12], which attempts to mitigate the multi-path problem for routers that implement per-flow load balancing. We rely on our own BGP datasets, collected from several common data sources, to perform the IP-to-AS translation. This helps us sanitise inconsistencies between neighboring hops. The client uploads the measurements to the Mercury platform which presents all this information and statistics using a web site and allows expert users to dig more in depth via a built-in API.

[1] The project takes the name from the mythologic Roman god Mercury, the god of the messages, communication and commerce. More information about Mercury Platform is found at Appendix A

By leveraging this data, Mercury discovers the AS paths and interconnection relationships, while putting a special emphasis on the detection of interconnections between CPs and access ISPs. Mercury offers information about the number of AS hops to reach selected CPs, the type of AS relationships and the existence of IXPs in the path, and statistics related with the AS geographic distribution. Furthermore, Mercury allows researchers to consult aggregated statistics of different AS paths from multiple geographical locations to the same destination. This analysis is particularly useful for the identification of the server infrastructure used by Content Distribution Networks (CDNs). In addition, the platform itself is highly flexible, making it easy to add future extensions via its API. In the following, we present several of its functionalities:

- **Identification of the hidden interconnections:** Mercury identifies many direct interconnections between CPs and access ISPs as the traceroute measurements are mostly done from commercial access ISPs with destination the most popular web sites. We consider that the release of such information is central to prove that CPs are getting closer to the access ISPs to offer a better user experience.

- **Identify the architecture of CDNs:** We identify the interconnection for CDNs, by revealing different paths to the same web resource from many points of origin. This helps researchers to analyse and classify existing CDN deployments, and to propose novel caching and data transfer techniques. At the same time, content providers can evaluate network performance for reaching their target customers.

- **Deployments for small-scale content providers:** Mercury facilitates the benchmarking of the interconnection agreements between access ISPs and hosting ISPs. With this feature, small web sites have the opportunity of discovering the best way to reach their target audience (i.e. a small web site whose target customers are from a single country could be only interested in hosting its contents in a hosting provider directly interconnected with the national operators of this country).

- **Modeling the AS topology and the taxonomy of AS interconnections:** Mercury aggregates data of Internet paths, facilitating the generation of AS graphs. The interconnection degree improves with the addition of new measurements. Such a dataset can serve for the evaluation of interconnection models, taxonomy of research and commercial networks, and to determine their interconnection strategies.

- **Optimising the interconnection agreements:** Mercury provides information about the AS neighbors and their interconnection agreements. This information could be useful for those network operators who want to optimise their interconnections and traffic routing policies based on the observation of other AS topologies. Thanks to this database, an ISP can discover peering partners and avoid using transit (bypass) connections [26].

5.1.1 Related work

In the last years many authors have noticed substantial changes in the Internet topology and the interconnection models. Gill et al. [46] identify that large content providers began to deploy their own wide-area networks (WANs) allowing them to have end-users closer in detriment of the tier-1 ISPs usage. Faratin et al. [43] not only identified the emergence of the large content providers and CDNs but also observed an expansion of the access ISPs with a progressive upgraded of their international backbones. Labovitz et al. [63] confirmed the evolution of the Internet topology from hierarchical to a more flattened and meshed model where large content providers and CDNs tend to concentrate most of the Internet traffic. More recently, Shavitt and Weinsberg [94] have analysed the interconnection trends during the 5 years for large CPs concluding that there is an exponential increase and diversification of interconnections among actors (using IXPs), confirming a loose of presence of tier-1 providers.

The previous literature used measurements processed at AS level targeting overall trends in interconnection models and techniques. Most of the related studies are based on performing extensive *traceroute* measurements combined with *BGP data* from public datasets like RouteViews [71]. Despite their pioneering work, researchers realised that measuring the Internet topology from few vantage points leads to partial results [20]. To address this, Mao et al. developed a traceroute tool at the AS-level, creating a database of IP-to-AS mappings based on the observation of both BGP announcements in combination with IP traceroute measurements [67, 68]. However, the authors noticed the difficulty of detecting IXPs and sibling relationships, as well as mapping mismatches due to measurements in a limited geographical region.

Recognising the benefit of measuring the Internet from the edge, Shavitt et al. developed DIMES [92], a measurement infrastructure using a large number of software agents to obtain the Internet graph at the AS and IP level. DIMES collects traceroute and ping traces from a set of specific agents and process the IP addresses assigned to AS numbers. Meanwhile, Dimitropoulos et al. [41] focused on modeling and generating synthetic but realistic AS topologies using BGP data from RouteViews. The need for more extensive measurements has been conducted by efforts like the CAIDA ARK project [16] as a large-scale infrastructure that coordinates large-scale traceroute-based topology measurements including both IP and AS levels. Another example is RETRO, implemented by He et al. [50] that uses public traceroute servers to collect measurements from many diverse locations. One of the challenges in Internet measurements methodlogies is to deal with partial information and possible missing peering links, most of them located at IXPs [50]. Chen et al. [21] used a plugin called ONO embedded in a BitTorrent client to perform random traceroute measurements from end-users. ONO is centred in the detection of hidden peering links combining data from public IXP and interconnection information from CAIDA datasets [18]. Augustin et al. provided a new approach for detecting IXPs and inspecting their AS participants based on various databases and traceroute measurements [11]. They detected 223 out of 278 IXPs and demonstrated that most of the remaining IXPs are invisible to tracerouting. In late 2010, the Regional Internet Registry (RIR) in Europe RIPE, released a collaborative tool called ATLAS [89] that consists in a global network of probes that measure Internet connectivity and reachability. This tool provides thousands of end hosts to execute distributed traceroutes and a public database of existing measurements. ATLAS is a great tool for executing distributed topology measurements, however it does not include the IP-to-AS translation and user cannot control the number of attempts in case of failure. In parallel, other projects such as PeeringDB [81], PCH [78], and EuroIX [42] have been developed to facilitate the interconnection between ASes. They put a special emphasis on peering interconnections, such that participant ASes can register and advertise their IP prefixes. Based on this data, in the recent years there have been many new contributions augmenting the AS topology with relationship information [77, 76]. To dig more into measurement initiatives go to Table 2.3 in Chapter 2.

Other studies analyse the hosting infrastructures and facilities of content providers. Huang et al. [53] measured the number of servers used by the main CDN providers Akamai and Limelight. Adhikari et al. analysed the infrastructures of the two major video content providers: Youtube and Netflix. This analysis, based on Planetlab servers [83], showed that YouTube [2] accumulates up to 80% of the analysed IP addresses whereas the rest belong to other ISPs like Comcast or Bell Canada. In contrast, Netflix [1] bases its video delivery services combining three different CDNs. Calder et al. [19] looked at the Google infrastructure determining the geographic location of the cache servers based on an approach called client-centric geolocation that consists in geolocating front-end servers by the geographical mean of client locations combined with the use of the EDNS-client-subnet extension. Ager et al. [4] also used a new methodology based on BGP snapshots and DNS queries for detecting web content infrastructure, and realising that few hosting infrastructures (e.g. Akamai and Google) are serving a large number of hostnames.

Although the research in interconnection models based on measurements is vast, there are no studies dealing with the strategies and models targeting the differences between global and regional (access and content) providers. Some articles aim to infer the AS interconnections while other studies are more case-centred addressing an specific provider (Google, YouTube, Akamai or Netflix). To address these shortcomings, in this chapter and the following we present a new measurement platform called Mercury that combines elements from prior work and focuses on discovering interconnection models from an end-to-end (access ISP to CP) point of view. We use this measurement methodology to discover the CP interconnections based on a simple taxonomy and to characterise the different interconnection models of a large data set of CPs.

5.2 Objectives and Methodology

Mercury is a measurement platform dedicated to discovering the interconnection between content providers and content consumers. Fig. 5.2 illustrates the AS information that one can discover using Mercury. The figure shows an snapshot with the interactions between different ASes and the types of interconnection in their links. In the left part of the figure appear the access ISPs and their links to their sibling core networks (e.g. Orange Spain is part of France Telecom that has an international backbone to with the rest of Internet). Then, the core networks directly connect to the ASes of the content providers or intermediary trasit ISPs (e.g. Level3). The figure also shows the interaction between two ASes through an IXP (e.g. see the interconnection between Vodafone and Microsoft).

On the consumer's side, Mercury can create a taxonomy of how access ISPs reach the content providers. This result should help us to answer the question on whether access ISPs use different interconnections strategies, depending on their size. On the provider's side, Mercury facilitates the identification of content distribution architectures, depending on the determined geographical, IP or AS destinations.

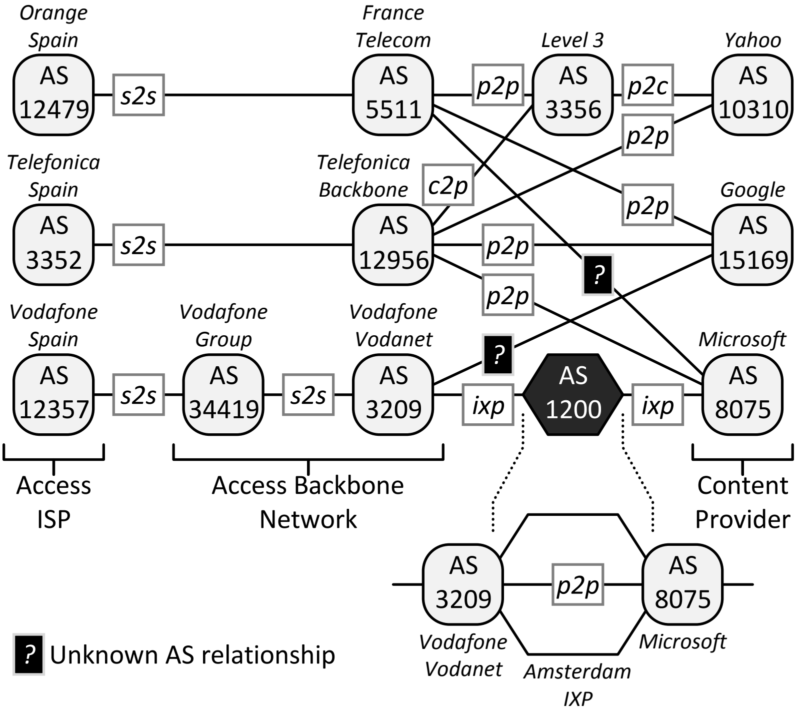

Figure 5.2: Example of AS interconnections.

Our objective is transform Mercury into a comprehensive public dataset of end-to-end Internet paths. In addition to accessing the data, users can contribute, sending their own traceroute measurements using a desktop client that can be downloaded from the project's web site. Although we evaluated the option of implementing a web-browser plugin or applet, we discarded this option due to the inability of accessing to certain restricted functions of the operating system such as opening raw sockets. Therefore we finally implemented a stand-alone application inspired by the DIMES and ONO projects. As a future research, we will investigate the possibilities of upgrading the user interface in order to make it less intrusive for the user.

Mercury features a user-friendly web environment to navigate through the available datasets, query for certain data entries and to visualise interactive plots. Expert users can alternatively use the provided API to send traceroute measurements and to query for stored information. The API is open and based on the REST Web Service protocol. It allows third parties to contribute with new extensions.

5.2.1 Measurement Methodology

As illustrated in Fig. 5.3, Mercury collects topology data from two main sources: (i) traceroutes from end-points clients across the Internet, and submitted via the REST API, and (ii) Internet topology information from several trusted databases. The measurement methodology is the following.

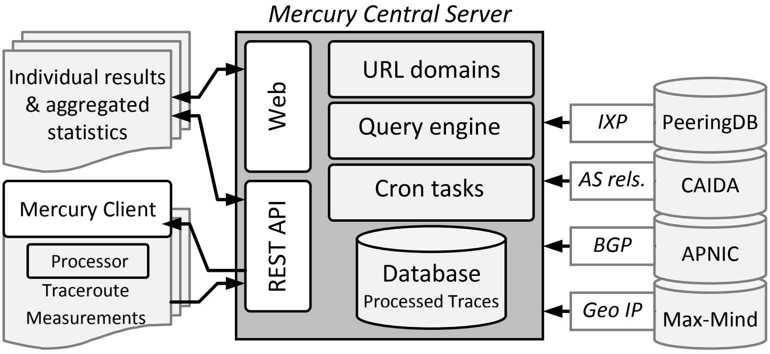

Figure 5.3: The Mercury architecture.

Content Destination Selection: we selected the top 100 most popular web site destinations from Alexa Top Sites [8] for Spain (the system supports up to 18 different countries). This includes large CPs like Google, Microsoft, Facebook, Yahoo!, YouTube or Amazon and popular web sites from Spain, e.g., ElMundo, ElPais or Softonic. Once, we have the top 100 destinations, we parse each web site (using web scraping of the HTML code) in order to extract all the URL resources (links, images, videos, gadgets, etc.). Extracting URLs from embedded resources, we determine all web servers that contribute content to the rendering of an individual web page. This step is important because some of these URLs are likely to point to CDNs. Finally, we obtain between 700 and 800 unique URL destinations per country that we store in the Mercury main server and the desktop clients download their URL list according to their country using the REST-API.

Traceroute Collection: Mercury platform receives IP traceroute paths from the desktop clients and stores them into a database. Our client extends and implements Paris *traceroute* [12] in order to prevent multipath anomalies across routers that perform load balancing. Because the correctness of traceroute data is paramount to the inference of real AS relationships, we resort to multiple tests per destination, at least 50 in most cases, which helps us to filter correct traceroutes. Using a desktop client, as opposed to a specially deployed traceroute servers in vantage points, guarantees a greater topological and geographical diversity [21]. Processed traceroute information is uploaded via the REST API to the Mercury web platform. Client participants probe a set of web site destinations automatically provided according to the region where the client is located. The rationale for this approach rests on the belief that access IPS and CPs have incentives to directly connect when the corresponding content is popular in their country.

AS Resolution: Upon executed a traceroute measurement, the client translates the IP addresses to the equivalent AS number requesting mapping data to the Mercury platform. A scheduled task in the Mercury platform collects periodically every 24 hours IP prefixes from the APNIC/RouteViews from the daily BGP announcements in [82] and updates its own database with the new IP-to-AS mappings. All this IP-to-AS data is also available through the Mercury public REST-API.

Geolocation: In parallel to the AS resolution, the Mercury client obtains, when it is available, the geographical location of IP addresses (country and city). The Mercury platform has a scheduled task that requests the IP-to-geo dataset from the MaxMind geolocation service [69]. Although we offer this information, we have to note that sometimes these mappings are not precise, specially those related to IP address that belong to Hyper-

Giant content providers or CDNs.

IXP Identification: We search the IP addresses matching an IXP from the PeeringDB database [81]. At this step, we both detect the IXP in the routing path, and the participating ASes. Fig. 5.2 illustrates the principle of IXP interconnection, where the IP addresses within the IXP subnet and revealed by the traceroute, actually belong to the connecting ASes.

AS Relationships: Finally, we examine the interconnection relationships between ASes using data from CAIDA [18]. Thanks to this dataset, we identify the peering (*p2p*), customer-to-provider (*c2p* or *p2c*) and sibling (*s2s*) relationships (see Fig. 5.2). In addition to this, we provide an extra dataset with (*s2s*) relationships extracted from analysing the business owners of the ASes.

5.2.2 Internal design

Mercury is formed by two main software instances: (i) the *desktop clients*, geographically distributed around the world and used by the project participants, and (ii) the *Mercury Central Server* or *MCS*. The MCS is the responsible of formatting and storing the information from the external datasets and the responsible of storing and publishing the processed measurements sent by the clients. The desktop client performs the traceroute probes and processes the results for obtaining the IP-to-AS translation, the IXP detection and the AS relationships of the end-to-end path. To better understand how Mercury Platform works please refer to Appendix A.

Mercury Central Server (MCS):

The MCS is a web server that aggregates information from multiple sources. The MCS has an updated database of IP-to-AS and IP-to-Geo mappings and AS relationships. The IP-to-AS translation is provided after the processing of BGP data (IP-prefixes from APNIC and RouteViews) obtained from [82] on a daily-basis. To improve the IP-to-AS database we include IP-prefixes present at the most important IXPs using PeeringDB [81]. The

IP-to-Geo service geolocates the origin and destination IP addresses using data from MaxMind [69]. The AS relationships provide interconnection information about two AS neighbors based on the data from CAIDA [18]. All this data is available via WS-REST API and is consumed by the MCs to process the traceroute paths. The MCS also stores the processed measurements from the distributed MCs and offers the possibility of inspecting the resulted data using a web interface or the built-in API. In addition, the MCS provides multiple filters to query data based on different criteria (e.g. origin or destination IP/AS/domain/location).

Mercury Client (MC):

The MC is the responsible of performing distributed traceroute measurements, data-path processing and the upload of the collected traces to the MCS. The MC is a desktop application that performs topology measurements using a modified version of Paris traceroute [12]. MC executes Paris traceroutes using ICMP and UDP protocols in order to reduce the multipath nature of IP forwarding. Then, MC translates the IP addresses from each trace to ASes using the MCS API and it identifies and corrects (when is possible) corrupted traces (e.g. missing hops or loops). Once MC has the AS-level trace, it uses the MC API to aggregate metadata information about the interconnection relationships between consecutive AS neighbors (AS hops) and it also adds the geographic location of the origin and the destination IP. Finally, MC uploads the processed measurement to the MCS database.

```
CLIENT operations and interaction with MCS

1. GET URLs $\Leftarrow$ MCS
2. foreach $\boxed{\text{URL}}$
    (a) GET $\boxed{\text{destIP(s)}}$ $\Leftarrow$ DNS Lookup
3. foreach $\boxed{\text{destIP}}$
    (a) Execute $\boxed{\text{parisTR}}$ $\{$50 probes per $\boxed{\text{destIP}}\}$
    (b) Translate $\boxed{\text{IP-2-AS}}$ $\Leftarrow$ MCS
    (c) Sanitise and Identify VALID traces
4. Aggregate FLOWs in $\boxed{\text{AS-PATH(s)}}$
5. foreach VALID $\boxed{\text{AS-PATH}}$
    (a) GET $\boxed{\text{AS-Rels}}$ for adjacent ASes $\Leftarrow$ MCS
    (b) Generate $\boxed{\text{STATs}}$
    (c) GET $\boxed{\text{IP-2-Geo}}$ $\Leftarrow$ MCS
    (d) Send PROCESSED $\boxed{\text{AS-PATH}}$ $\Rightarrow$ MCS
```

5.3 Case Studies

In order to prove the capabilities of Mercury towards discovering the hidden interconnections between CPs and access ISPs, we conducted an experiment using a set of traceroute measurements from various end-points located in the major Spanish access ISPs. The objective of this experiment is twofold. Firstly, we test the detection and aggregation of interconnection relationships for a set of given URL destinations, showing the types of connectivity for a given end-point. Secondly, we leverage the known relationships between end-points and the Internet AS graph to show how we can discover the architecture of complex distributed systems, such as Content Distribution Networks (CDNs).

Toward this end, we selected a set of 100 web destinations from the Alexa Top 100 list in Spain as content providers. From these 100 sites we extract the embedded URLs from each site giving us more than 700 URLs. We parse the embedded URLs in order to identify content (images, audios, videos, gadgets, etc.) hosted in other servers than the main web site. This will help us to identify content distribution solutions.

On the side of content consumers, we have a set of volunteers that run the Mercury client from their home ISPs. For this experiment we use 5 participants from Barcelona, each located in each of the 5 major Spanish access ISPs (Telefonica, Orange, ONO, Jazztel and Vodafone). Each participant uses the Mercury client for tracerouting the corresponding set of destinations, and upload the results automatically to the Mercury platform using the REST API. We require clients located at the commercial access ISPs in order to see how these ISPs interconnect with the CPs. Running the MCs from commercial ISPs, rather than using platforms like PlanetLab [83] which nodes are mostly located within research networks, allows us to determine how the content reach end-users of access ISPs. The clients were executed in the area of Barcelona (Spain) and the measurements were done from each access ISP during the first week of September 2014.

Upon receiving the measurement data, Mercury distinguishes between *completed* and *inconsistent* traceroutes. Reasons for the latter include (i) the origin node being placed behind a firewall blocking ICMP traffic, (ii) destinations behind a firewall preventing the completion of a traceroute, and (iii) incomplete databases keeping Mercury from performing an IP-to-AS mapping. We use a heuristic algorithm to identify the traceroutes that do not yield a path between the origin and destination ASes. Finally we collect both correct and incorrect measurements and we store them with a flag that identifies the different inconsistences of a path.

Although the Mercury platform can be used for multiple purposes, in this study we only focus on detecting the hidden interconnection between CPs and access ISPs in Spain. Therefore we use Mercury for detecting direct interconnections between CPs and access ISPs without intermediaries and CPs that place their server infrastructure inside the access ISP network. For instance, Fig. 5.4 illustrates an example where Google has a direct interconnection with the Spanish access ISP Jazztel and also places servers inside the Jazztel network.

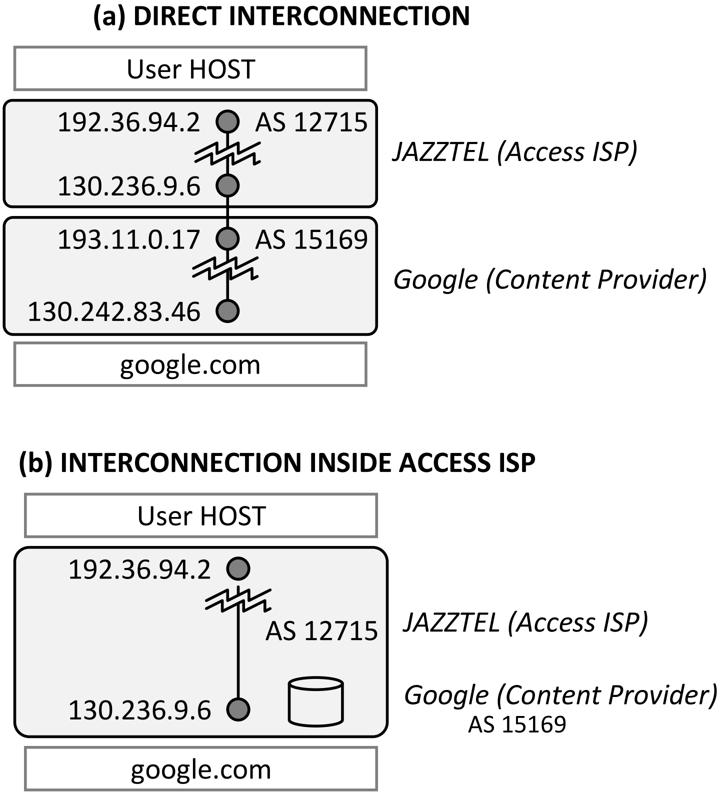

Figure 5.4: Processed traceroute measurement

Although Mercury was implemented to collect data from participants around the world, in this thesis we focus exclusively on the Spanish interconnection case because is a mature post-monopoly market with multiple access ISPs that gives a broad picture of one of the largest and more dynamic Internet markets in Europe with up to 50 million inhabitants and 71.6% of Internet users [2]. Therefore, we expect that this experiment will obtain similar results in the larger European countries because they have a similar ISP market structure. This experiment does not require too many participants as we consider that most of the users of a certain access ISPs will be routed, at the AS-level, using the same policy [62]. Hence, at least one participant in each one of the five major Spanish access ISPs will be enough to draw conclusions on their interconnections, as shown in Fig. 5.5.

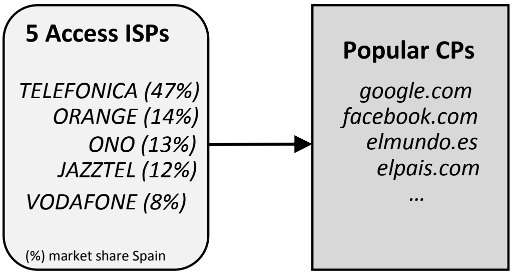

Figure 5.5: Major access ISPs probing popular CPs [28]

5.3.1 Revealing Hidden Interconnections

We can leverage Mercury to discover the interconnection between a particular access ISP and a particular content provider. This takes advantage of the aggregated statistics for traceroutes from multiple origin access ISPs. Table 5.1 shows the identification of direct interconnections for a subset of popular web sites that include global and local content providers.

[2] Source: WorldBank Indicators Database. 2013. `http://data.worldbank.org/indicator/IT.NET.USER.P2`

It compares the existence of direct interconnection relationships, either physically direct or across a sibling AS of the same organisation, with the AS relationships from the the CAIDA dataset [18].

In these results, we use checkmarks to emphasise the matches between Mercury and CAIDA dataset. We note that there are many instances where the direct interconnections are not present in the CAIDA dataset (see crossings in Table 5.1). Dashes are used to illustrate that in this case we cannot compare with CAIDA datseet because Mercury have not found a direct interconnection. As demonstrated by Calder et al. [19], Google's infrastructure is particularly hard to localise. Mercury shows us that Google has direct connections with all the major spanish ISPs.

	Google	Facebook	Yahoo	Twitter	Amazon	MSN	Wikipedia
Telefonica	Sibling ✗	Sibling ✓	Sibling ✓	No –	Sibling ✓	Sibling ✓	Sibling ✓
Orange	Sibling ✗	Sibling ✓	No –	No –	No –	Sibling ✗	Sibling ✓
ONO	Direct ✗	No –	No –	No –	No –	No –	No –
Jazztel	Direct ✗	Direct ✗	Direct ✗	IXP ✗	IXP ✗	IXP ✗	No –
Vodafone	Sibling ✓	Sibling ✓	Sibling ✓	IXP* ✓	IXP* ✓	Sibling ✗	Sibling ✓

Note: IXP* is a relationship where a sibling AS is connected to an IXP

Table 5.1: Identification of direct interconnections.

Although Mercury sometimes is not capable to identify the AS relationship type of a direct interconnection, it at least detects it, making possible to focus on this link in future studies in order to detect the relationship type. However, we conjecture that most of the direct interconnection are based on peering or paid-peering relationships based on the analysis of the peering policies of both access ISPs and CPs. One can observe that access ISPs and large content providers find more attractive this type of interconnection than using an intermediary AS (see Google and Microsoft with Jazztel). They find the direct interconnection mutually beneficial. The content provider can be closer to its target audience and can offer a better QoE while the access ISP obtains an economic compensation from the paid-peering agreement.

We also observe that not all CPs have direct interconnections to all access ISPs. This could be for different reasons: there are some CPs that only allocates the cacheable content using direct interconnections and

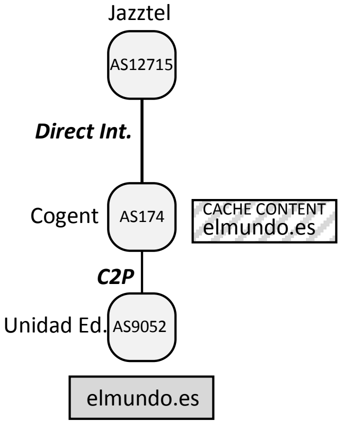

Figure 5.6: Multiple content delivery strategies for web sites.

there are also some CPs that have only agreed direct interconnections with some certain access ISPs and require an intermediary to reach the rest of access ISPs. This can be seen in Fig. 6.4, where one of the major Spanish media groups (ElMundo.es) maintain its own AS, but they contract a third party for delivering their cached content.

5.3.2 Revealing Interconnections Inside the Access ISP: The CDN

Thanks to Mercury we observe that many Content Providers use different strategies to place their contents close to the end users. When seeking evidence of whether a content provider uses a CDN solution, Mercury provides a number of statistical indicators that are adequate to this task: (i) the number of AS hops to reach a server, (ii) the list of destination countries, and the list of (iii) IPs and (iv) ASes for a same URL. The number of AS hops to reach a CDN destination is a weak parameter of CDN existence, but it reveals the location of the content server relative to the user (when zero, the server is insider the user ISP).

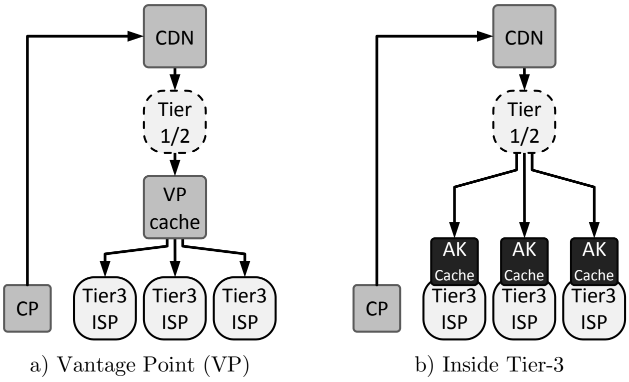

Figure 5.7: Types of CDN architectures.

The list of destination countries is a better indicator of CDNs when the geolocation service is accurate, giving us the distribution of the server infrastructure. However, in practice, we must be careful because global content providers do not publish the location of their servers. For example Google servers are geolocated only in the USA, based on the registration of their AS. Currently, a lot of research effort is invested into improving the geolocation of IP addresses [95, 19, 4].

The list of destination IPs and ASes are the strongest indicators of a CDN. They show that a web site is deployed along multiple distributed servers, confirming the existence of some type of load balancing or caching technique, which are intrinsic to the use of CDN solutions. In addition, we can use them to determine the taxonomy of CDN strategies. As illustrated in Fig. 5.7, we can distinguish between two main types: (i) CDNs that host their cache servers inside the Tier 3 ISPs (Akamai and Google strategy) and (ii) CDNs that locate their cache servers in vantage points (VPs), near to the access operator (Cogent or Level3 strategies).

Therefore Mercury helps us to detect where the content is hosted by leveraging these indicators in combination with the AS-path to reach them. Our results indicates that most of the web sites we analysed resort to some type of CDN.

Table 5.2 summarises the Mercury data for several CDN destinations.

We observe that many global content providers like Facebook or Microsoft use Akamai, which deploys servers in both access ISPs and Tier 1 carriers, in addition to their own VPs. Furthermore it is also interesting that Elpais, the second major press group in Spain, also uses Akamai. The number of servers deployed by Akamai in other ASes stands well above the other CDNs, something observed by previous research publications [99]. Google uses a similar strategy and has direct interconnections with most of the access ISPs and has servers inside some access ISPs (we detected Google cache servers pointing to Jazztel IP addresses in the range 212.106.221.0/24).

CP	Google	Facebook	Yahoo	Amazon	MSN	Instagram
CDN	Google	Akamai	Yahoo	Amazon	Akamai	Amazon
# servers	43	5	4	85	4	12
Inside ISP	✓	✓	–	–	✓	–
VP	✓	✓	✓	✓	✓	✓
Multi-vendor	–	–	–	–	–	–

CP	Elmundo	Elpais	LinkedIn	20minutos	Wikipedia	
CDN	Cogent Interroute	Akamai	Limelight Level3 ...	Level3 Fujitsu ONO	Wikipedia	
# servers	41	12	4	4	1	
Inside ISP	–	✓	–	–	–	
VP	–	✓	–	–	✓	
Multi-vendor	✓	–	✓	✓	–	

Table 5.2: CDN strategies for different content providers.

The remaining content providers use VPs to host their cache servers, based on different strategies. LinkedIn uses a multi-vendor CDN solution formed by the mentioned Level3, LimeLight, EdgeCast and others. This case is quite interesting because LinkedIn diversifies the spending in content delivery services. Similarly, the Spanish press groups ElMundo and 20Minutos use the same multi-vendor strategy. For Amazon CDN, we highlight the large number of servers. Whereas Akamai and Google manages the CDN for its own services, Amazon sells CDN services to other web sites like Instagram.

5.4 Conclusion

In this chapter, we introduced Mercury, an Internet measurement platform that aggregates traceroute measurements from multiple locations and analyses the AS interconnection relationships along a network path. Mercury stands-out over other solutions because it discovers the end-to-end network path at the AS-level, while including information about the AS relationships, detecting IXPs and adding geolocation. We evaluate Mercury for a set of web sites using clients located at major Spanish access ISPs. Our results reveal the existence of many direct AS interconnections between access ISPs and content providers that are hidden for other methodologies. This suggests that some access ISPs and content providers find this interconnection strategy more attractive over using an intermediary and it confirms the trend that content providers are increasingly closer to the end users. In addition, our results emphasise the interconnection degree of the Spanish market, which is relatively high due to the small number of AS hops between access ISPs and content providers. Finally, Mercury provides indicators for discovering the architecture of CDNs and we successfully identify different content delivery strategies used by many web sites. Mercury detects Akamai and Google servers inside some Spanish ISPs which demonstrates the interest of these companies in offering a high quality content delivery.

In the next chapter we will exploit the power of Mercury Platform with the objective of studying the different models of interconnection that content providers follow to distribute their services.

Chapter 6

INTERCONNECTION MODELS OF CONTENT PROVIDERS

In this chapter we evaluate a large sample of content providers with the purpose of observing which type of interconnection model they are using to reach the end users located at the access ISPs (see Figure 6.1). This chapter presents different models of interconnection and uses the data gathered from Mercury platform (see Chapter 5) to research how Internet players interconnect to deliver content. The obtained results give us interesting results as we observe a correlation in the methods of interconnection that similar Internet players follow. Thanks to the information extracted from the developed tool and the lessons learned about how Internet traffic has evolved in the last years, we are able to give solid arguments about why a large content provider seeks the direct peering interconnection while smaller ones tend to use more affordable ways of interconnection like using hosting ISPs. Following we will go deeper analysing the drivers behind the interconnection models of different Internet players.

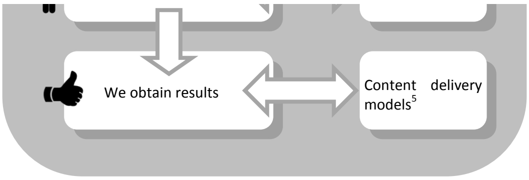

Figure 6.1: Thesis outline: interconnection models of CPs.

6.1 Introduction

The Internet topology is continuously evolving and in the last years we have witnessed a spectacular increase in the consumption of multimedia content [63] shifting the traditional roles of most of the Internet players. The needs for having faster delivery pipes, together with the irruption of content delivery and cache mechanisms, have impacted the interconnection models between the Internet Service Providers (ISPs) and the content providers (CPs). The balances between global and local content, and larger and smaller service providers brings a plethora of interconnection models within the same Internet market that depends on a set of strategies followed by competing Internet actors.

Global CPs, with presence in most worldwide Internet markets (like the *big five* Google, Facebook, Yahoo, Microsoft and Amazon), need to design and invest in interconnection models to reach end-users in a cost-efficient way. Those players, with strong market position, are able to make extensive use of costly direct (and peer) connections to deliver their services with higher quality and performance (peak and average rate, stable jitter and low delay). Local players, with less finantial muscle need to use other interconnection strategies mostly based on shared connections (by using third-parties or Internet eXchange Points) to provide similar performance with much affordable investment. In this chapter we analyse those strategies using recent measurements to figure out the Internet topology for each type of actor (content or service providers; with global or local presence).

The emergence of content delivery networks (CDNs) influences the evolution of the traditional interconnection models, by having specialised actors that concentrate the heaviest traffic, being able to deliver it globally and changing their business models [43]. This dynamic and continuously evolving market forces some actors to reinvent themselves and adapt to the fierce competition for traffic delivery. Tier-1s are a clear example of those dynamics, moving from mere traffic transit services to provide hosting and CDN interconnection services as a central value in their offers.

In this chapter, we analyse the interconnection models followed by local and global providers by measuring the current Internet topology at the Autonomous System (AS) level. The analysis is done from both sides: how CPs are interconnected to ISPs to reach end-users; and how ISPs allow their customers to reach global and local contents. The research done combines partial and fragmented publicly available information with a deep knowledge of the Internet market to process the measurements. Although the results obtained are for the Spanish Internet market, where global (Telefonica, Vodafone, Orange) and local (Jazztel, ONO) actors coexist targeting global (Google, Facebook, Yahoo) and local (ElPais, UnidadEditorial, Softtonic) contents, it is reasonable to expect similar results for other European and U.S. Internet interconnection markets. The same analysis can be easily extended to other markets by using the same measurement tool[1] to collect exhaustive measurements in other areas.

The chapter identifies different interconnection models depending on the provider's profile: while the *big five* content providers tend to be directly connected with most ISPs, content delivery actors and hosting providers have extensive presence to provide efficient alternatives to reach local and global content.

The chapter is organised as follows: following the introduction, we present a literature review in interconnection models and measurements tools. Then, we present the methodology used and the analysis of collected data to show the results. Finally, we explain the conclusions of the analysis done.

[1] BETA version of Mercury Platform available at `http://mercury.upf.edu/mercury`

6.2 Methodology

The methodology used in this analysis is based on active traceroute measurements. The data collected allows to explore the existing interconnections among the Internet actors and snapshot the existing topology. The methodology used consist in analysing the measurements to get the Autonomous Systems (ASes) involved in each traceroute path to then discover the interconnection models of the affected players. The analysis allows the detection of direct connections between actors, the existence of CDN servers in between or the use of Internet eXchange Points (IXPs). The tool used also detects corrupted traceroute paths and the analysis is only base on successful routes. In this section (and in Chapter 5) we provide more details regarding the measurement process and the data collected.

6.2.1 The measurement platform

Towards this end we use the measurement platform called Mercury that we introduced in the Chapter 5 and [79]. Mercury is a multi-purpose platform consisting of a central server (MCS) and downloadable clients (MC) that perform AS-level traceroute measurements and upload them to the MCS (see Figure 6.2.a). Once the MCS has stored the processed measurements from the distributed MCs, it offers the possibility of inspecting the resulted data using a web interface or via its API.

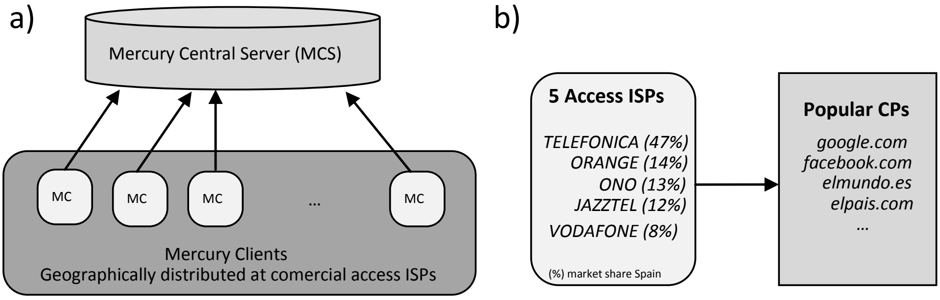

Figure 6.2: Mercury platform: a)General view, b)MCs at Spanish Access ISPs targeting popular CPs

6.2.2 Data collection

Mercury Platform crosses information from existing databases to get the Internet paths between providers: CAIDA [18] to get neighbour relationships; PeeringDB [81] to detect interconnections at IXPs; and routing data from [82] to collect BGP information and translate IP addresses into ASes. More information about how Mercury deals with different anomalies (AS-loops or missing hops) along the internet path process can be found in Chapter 5 and [79].

The data to carry on the analysis has been collected from the Spanish Internet market. As we previously mentioned, we run 5 clients, each located in each of the 5 major Spanish access ISPs (Telefonica, Orange, ONO, Jazztel and Vodafone). (see 6.2.b). Spain shows a mature Internet market in terms of competition (number and size of ISPs), usage (most visited places and penetration) and networking infrastructures (tier-1s, IXPs and fibre-based access networks), comparable to other western countries. Therefore, the methodology can be directly extended to other countries and areas expecting quite similar results in terms of interconnection models.

6.2.3 Data analysis

Once all the traceroute measurements are stored in the MCS database, we use the MCS API to execute queries that provide us the required data. For each URL, we analyse the different ISPs along the AS-path and its interconnection relationships. Then, we compare different traces from different CPs to find interconnection similarities.

To facilitate the analysis of the different content delivery strategies, we define different metrics based on the valid traces (platform algorithms also sanitise when possible corrupted traces from missing hops and loops). We locate where CPs have their hosting infrastructures: within the access ISPs, within a tier-1, within a commercial CDN or within their own network. Finally, we analyse whether CPs tend to interconnect at IXPs and whether CPs tend to direct-connect with access ISPs bypassing tier-1s when possible.

6.2.4 Definition of the interconnection models

To better understand the diversity of interconnections between CPs and access ISPs, we have defined three different interconnections models based from the point of view of the CP: using an ISP as intermediary; using a direct connection; or by means of an IXP. In addition, each model may include the use of a CDN provider to speed up some of the contents. Here we provide more details for each one of the models:

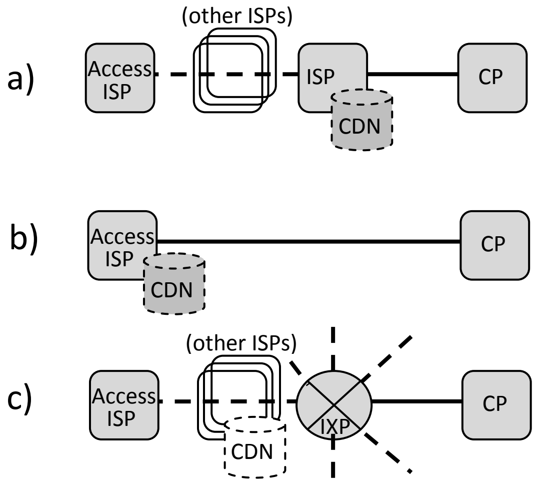

Figure 6.3: Interconnection Models

A) Interconnection through an intermediary transit provider:

In this model of interconnection (see Figure 6.3.a) the CP uses a transit provider as intermediary to carry its traffic. CPs usually use the same or a set of providers to reach end users for each particular access ISPs. This interconnection is commonly based on a transit agreement where the CP contracts from the transit ISP a transport service. There is also the option to use a CDN solution to deliver the traffic.

B) Direct interconnection:

In this model of interconnection CP has a direct link to the access ISP (see Figure 6.3.b). The direct model, which avoids intermediaries ISPs in contrast to the previous one, tends to use (paid) peering agreements to improve the performance of the traffic delivery. As in model A, the access ISP can deploy its own CDN service or use a third party CDN, like Akamai[2], to speed up the traffic. This situation is extremely difficult to detect, because it does not affect the interconnections and the CDN service is deployed within the access ISP. Therefore, we are required to apply other techniques based on DNS resolutions to detect the originating CPs, e.g. when we detect that a targeted URL does not go away from the access ISP (0 AS hops) we inspect the URL using the tools like *nslookup* to identify whether CDNs like Akamai or Google deploy their servers within the access ISP.

C) Interconnection through an IXP:

In this model of interconnection (see Figure 6.3.c) the CP is directly connected to an Internet eXchange Point (IXP) where it has the possibility to interconnect with other actors (both transit or access ISPs) within the same location. At the IXPs, ISPs interconnect based on both transit and peering arrangements. Strictly speaking, we could consider the IXP interconnection model together with the previous two models. However, the broad flexibility that provides the use of an IXP in terms of number of interconnections incline us to treat it as a separate case.

6.3 Results

In this section we use the collected data to analyse the interconnection models between the Spanish ISPs and the top CPs. All the measurements are processed to group the traces belonging to the same organisation by checking the Autonomous System Number. Following that grouping rule, traces from providers such as YouTube belong to Google ASN or MSN to Microsoft ASN, having all those measurements united in the same set.

[2] We consider Akamai part of this model (B) only when it uses the IP-address space belonging to the access ISP

6.3.1 Content delivery interconnection models

The measurements analysed allow us to identify the interconnection models for each content provider according to the models defined in the previous section. A first observation of the collected data gives interesting insights that complement other databases such as CAIDA. An example is Google that appears directly connected to either the Spanish access ISPs or a tier-1 (Level3), while CAIDA only reports direct links to Cogent, Telia, NTT and Tinet. These results do not affect the final conclusions that are based on interconnection models rather than in connections between particular peers.

The results are presented in four groups: i) CPs with their own networking infrastructure, ii) CPs using specialised CDNs, iii) CPs using CDN/hosting solutions of carrier ISPs and iv) CPs using hosting or cloud solutions. The motivation for each CP to use one model or other depends on different factors such as transit or hosting costs, the status of their current infrastructure or their needs to offer an enhanced QoS. The analysis is done for each CP type depending on the models previously defined in Section 6.2.4: A)interconnection through tier-1, B)direct interconnection between CP and access ISP and C)interconnection through an IXP. We present results using tables that show the interconnection models (A, B or C) used by different content providers to reach the five access ISPs.

Interconnection models of CPs with network infrastructure:

Table 6.1 crosses the interconnection model used for each pair of (global or regional) CPs and access ISPs. The results show how the *big five* providers (Google, Facebook, Microsoft, Yahoo and Amazon) tend to use direct interconnections (B model) with the access ISPs, by-passing tier-1s (A model). These large global CPs have extensive networking infrastructures, that facilitate more efficient interconnections instead of relying only on tier-1 providers. In addition, the more strict low-latency requirements for most Internet services, makes more cost-effective peering interconnections with better performance than regular transit services. Leading Spanish local CPs, such as ElMundo or ElPais, use direct links (B model) with the two largest operators (Telefonica and Orange). The rest of the CPs tilt the balance towards transit services (A model) avoiding the costly direct connections.

ASN (name)	15169 (Google)	32934 (Facebook)	8075 (Microsoft)	10310 (Yahoo)	16509 (Amazon)	20049 (LinkedIn)	13414 (Twitter)	19679 (Dropbox)	43821 (Wikimedia)	33612 (Tumblr)	15224 (Adobe)	43996 (Booking)	11643 (Ebay)	12678 (Badoo)	47195 (Gameforge)	51773 (Softonic)	50974 (ElPais)	9052 (Unidad Ed.)
Telefonica	AB	AB	B	B	AB	A	A	A	A	A	A	A	A	B	A	A	B	B
Orange	B	AB	B	A	AB	A	A	A	A	A	A	A	A	A	A	A	B	B
Jazztel	B	B	C	BC	ABC	A	BC	A	AC	A	A	C	A	A	C	A	A	A
Vodafone	B	B	B	ABC	BC	A	BC	A	BC	A	A	A	A	A	A	A	C	A
Ono	B	A	A	A	B	A	A	A	A	A	A	A	A	A	A	A	A	A

Table 6.1: Interconnection models of CPs with network infrastructure

Interconnection models of CPs hosted in a specialised CDN:

Table 6.2 presents the interconnection models for CPs using specialised CDNs to deliver their traffic. CDNs like Akamai, Edgecast, CDNetworks, Limelight or CloudFlare, combine multiple models prioritising the use of direct interconnections (B model) to connect with the access ISPs. This model provides the best performance which is an interesting incentive to reduce interconnection costs. However, in some sporadic cases, these CDNs also use tier-1s (A model) or IXPs (C model) less frequently.

Interconnection models of CPs hosted in hosting and cloud companies:

Table 6.3 shows the interconnection models used for a list of CPs that use hosting and cloud providers to deliver their services. These CPs are mostly regional and local web sites without a strong networking infrastructure and require more affordable interconnection models. The analysed hosting companies mostly use IXPs (C model) to deliver their services. This solution is mainly adopted by companies with presence in

ASN (name)	20940 (Akamai)	15133 (Edgecast)	36408 (CDNetworks)	22822 (Limelight)	13335 (CloudFlare)
Telefonica	B	A	A	B	B
Orange	B	B	A	B	A
Jazztel	BC	B	AC	B	C
Vodafone	B	B	AC	C	C
Ono	A	BC	A	BC	A

Table 6.2: Interconnection models of CPs hosted in specialised CDN

local markets and with the aim to optimise resources having most interconnections in a single point. The majority of the measured IXP interconnections use ESPANIX (located in Madrid, Spain) and LINX (London, UK) IXPs. It is noteworthy that, although measurements are done from Barcelona, there are few CPs using the CATNIX IXP located in the same city which it could be more efficient. However, these companies also combine this interconnection model with the use of tier-1s (A model) and rarely use direct interconnections (B model). This last observation is completely reasonable because smaller CPs do not prioritise their investments in their own infrastructure as the direct connections require (ISPs must physically connect their networks and operate their links).

ASN (name)	8220 (Colt)	3324 (Fujitsu)	36351 (SoftLayer)	13768 (Peer1)	16276 (OVH)	8560 (1and1)	16371 (Acens)	20718 (Arsys)	39020 (Comvive)	45037 (Hispaweb)	24592 (Nexica)	24931 (DediPower)	39743 (Voxility)
Telefonica	C	C	B	A	A	A	B	C	A	A	A	A	A
Orange	C	C	A	A	B	A	C	C	A	C	C	A	nd
Jazztel	C	C	C	BC	C	C	A	C	C	C	A	C	C
Vodafone	C	C	AC	A	BC	A	C	C	C	C	A	A	A
Ono	A	nd	A	AB	A	A	C	A	C	A	A	A	A

Table 6.3: Interconnection models of CPs hosted in hosting and cloud companies. (n.d.: not defined)

Adapt or perish - Interconnection models of CPs hosted in carriers:

Table 6.4 shows the interconnection models of a list of CPs using carrier ISPs to deliver their traffic. Tier-1 ISPs like Level3, Cogent, NTT or Interoute have detected that their transit services are losing share in front of the direct interconnections (peering) and the CDN solutions. As a response, these carrier ISPs have decided to take advantage of their large backbone networks to extend their services porfolio adding CDNs or hosting services to CPs. The motivation for CPs to host their contents using carriers is because these ISPs are in a good market position to build a new offer consisting of an affordable connectivity performance and highly cost-efficient to middle-size CPs. Typically, CPs hosted by these tier-1s have direct interconnections (B) with the access ISPs although some of them are transit agreements rather than (paid) peering. Access ISPs with sibling international networks (Telefonica, Orange and Vodafone) tend to have peering agreements while regional access ISPs (Jazztel and ONO) have transit agreements. On one hand, large access operators such as Telefonica, Orange and Vodafone move significant amounts of Internet traffic through their international backbones and consequently, comply the conditions to establish direct (peering or paid-peering) connections with tier-1s. On the other hand, Jazztel and ONO only operate at the regional level and the volume of traffic that they move is not still eligible to be exchanged using peering agreements as they are dependant from these transit providers to reach the ”rest of the Internet”. This case is an example where size (traffic volume) does matter.

ASN (name)	3356 (Level3)	174 (Cogent)	2914 (NTT)	8928 (Interoute)
Telefonica	B	B	B	BC
Orange	B	B	B	BC
Jazztel	B	B	B	B
Vodafone	B	B	B	BC
Ono	B	B	B	A

Table 6.4: Interconnection models of CPs hosted in carriers

Placing servers inside access ISPs:

Another interesting result from the measurements is to observe an interconnection model where CPs decide to locate the servers within the access ISP network. Based on the *akamaized*[3] URLs of some traces, we detected that some of the CPs are using Akamai to deliver their static content. Akamai delivers this content using servers at their vantage points or using servers located within the access networks. Table 6.5 shows that a large number of CPs are using the Akamai servers located in Telefonica, Orange, Vodafone and Jazztel. The first three ISPs use the Akamai servers located in their sibling backbones (Telefonica Backbone, OpenTransit and CableWireless) while Jazztel use their local network. We have observed that global CPs like Microsoft, Facebook, Yahoo, Ebay or Apple and local media CPs like RTVE, Antena3, Telecinco, ElPais, ElPeriodico or LaVanguardia require enhanced content delivery solutions like Akamai to ensure the better quality of service.

	Akamai	Access ISP hosting/CDN
Telefonica	Apple, Facebook, eBay, Microsoft, LinkedIn, Yahoo, ElPais, As, Badoo, LaVanguardia, Mundodeportivo, RTVE, Abc, Antena3, Telecinco, Mediaset, Sport, ElPeriodico	ElCorteIngles, BancSabadell, Europapress, Fotocasa, Segundamano
Orange		-
Jazztel		-
Vodafone		-
Ono	No detected Akamai servers within Ono	20minutos

Table 6.5: CPs using Akamai servers and hosting solutions within access ISPs

In addition, we have also observed that some measurements targeting Google never leave some access ISPs (e.g. Jazztel or the Spanish research network RedIRIS when we made the test measurements from Pompeu Fabra University). We identify a similar behavior to the one used by Akamai, where Google places its content servers inside the access ISPs. These results are consistent with Calder et al. [19] who previously observed this new trend. It may seem reasonable to expect this interconnection model

[3] An Akamaized URL is an URL which contains lexical references to be part of the Akamai CDN, e.g. *s-static.ak.facebook.com*

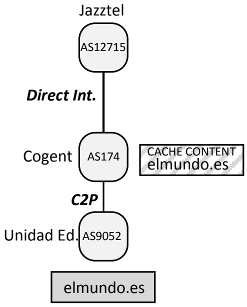

Figure 6.4: Multiple content delivery strategies for web sites

from Google to optimise its transit costs and reduce latency in some services like YouTube.

Multiple interconnection model:

There is also a specific case in which a CP with its own AS requires one or multiple intermediary ISPs (tier-1s) to reach the access ISPs. The particularity of this case is that these CPs host their dynamic content at their own AS but they delegate its cached content to the CDN service of an intermediary ISP. This can be seen in Fig. 6.4, where one of the major Spanish press groups (ElMundo.es) maintain its own AS, but they contract a third party for delivering their multimedia content. The main motivation of this model is to reduce transit costs using cache servers for static content while boosting the user experience.

6.4 Conclusions

In this chapter, we have explored the Internet paths between the most visited CPs and the major access ISPs in Spain. The results obtained are consistent with the literature and provide answer to the research questions that we formulated at the beginning of this thesis.

The results confirm that large CPs tend to bypass tier-1 networks and prefer direct interconnections with access ISPs. This behaviour is exploited by global content providers like Google, Facebook, Microsoft, Yahoo or Amazon that have developed networks with international presence. Those global players move huge amounts of Internet traffic, therefore they require advanced and large network infrastructures to accommodate and distribute their content. Then, they have a clear preference for direct connections instead of using transit intermediaries: a major control over the delivered content. Direct connections based on (paid) peering agreements obtain better performance than transit services [112] but not all CPs are in conditions of bargain this interconnection type. The fervent belief of these specific players in expanding their own infrastructures it is not a mere coincidence. In addition to their core Internet content business, these players also explore new business opportunities based on offering virtual cloud services (Google Cloud, Amazos WS, Microsoft Azure) to other companies that are not feasible without the tenure of a global and high-developed network.

The increasing popularity of direct interconnections [34, 60, 59] is a fact that is threatening tier-1s. In the past tier-1s played a key role because they were at the top of the hierarchy of the Internet topology and they were indispensable to transport Internet content from one part of the globe to another. Currently, this privileged position has changed and CPs have multiple alternatives to distribute their contents, e.g. direct (paid) peering with access ISPs or using third party CDNs. Furthermore, many tier-1s have been involved in many disputes because access ISPs are requiring them to pay a fee for delivering traffic into their networks [58] when their established peering agreement is exceeded (tier-1s introduces substantially more traffic into the access ISP than in the other direction). Therefore, tier-1s are really conditioned by this new situation where large CPs are moving a significant fraction of the total Internet traffic/business using alternative methods to avoid passing (and paying the toll) through the tier-1 networks. Accordingly, tier-1s are losing market power in this sector formed by large CPs but they are reacting rapidly: tier-1s like Level3, Cogent or NTT are leveraging their large networks to offer CDN and hosting services to those CPs (a large majority) that do not consider strategic to deploy their own content delivery infrastructure and prefer to outsource these network resources. This confirms that tier-1s are evolving their business beyond the traditional IP transit and they are offering a new catalogue of cost-effective services to smaller and local CPs. This change of course in the tier-1 original philosophy have re-positioned them in the Internet ecosystem *"from the transport to the content"* bringing them more business opportunities and providing more content delivery alternatives to the market.

Along this chapter we have laid special emphasis on analysing specialised CDNs. These type of Internet players provide featured and low-latency content delivery services to CPs which require to place their networks strategically. Therefore, they prioritise the use of direct interconnections as this model allows them to provide the required low-latency level service agreement to those CPs that demand advanced content delivery services that assure low-latency transmissions. We also notice that there are different categories of CDN addressed to different audiences: while more developed CDN companies tend to use direct (paid) peering and are focused on high demanding CP customers, there are other CDNs that combine the use of transit intermediaries and their presence at IXPs to offer an affordable service to general CPs. The former CDNs follow the same model as large CPs, putting performance above all, while the latter CDNs follow the hosting companies model which prioritise the use of shared infrastructures (probably some of the analysed CDNs could be also considered as a featured hosting company). In addition to this, we highlight the Akamai case. The measurement processing allowed us to identify Akamai servers within the networks of 4 of the 5 surveyed access ISPs which outstands the strategy of this company over others in favour of reducing the latency response.

We have also identified a large number of CPs that rely on hosting companies which mainly interconnect at IXPs. Hosting companies offer affordable solutions (including cloud and cdn solutions) that comply with the requirements of the majority of content providers (web sites with few -or not strategic- multimedia resources). Hosting companies business is based on *serving the content* (using their server farms) and they delegate the transportation to other Internet players by interconnecting to them in strategic points. In this context, hosting companies have two interconnection main options: they can interconnect their data centres using at least two transit providers (to ensure redundancy) or they can interconnect at an IXP that offers more alternatives. IXPs offer shared infrastructures that facilitate the interconnection of multiple ISPs in a single point. This structure removes the dependency from a single vendor and provides more options in case the hosting company want to expand in the future. From the point of view of a hosting company, to be present in a IXP is synonym of flexibility and resource optimisation.

Finally, according to the obtained results, we conclude that the Internet market is in continuous evolution and there are sufficient content delivery alternatives to cover the different necessities of heterogenous CPs.

Looking at the whole picture of the Internet, this chapter has helped us to better understand the current roles of different Internet players. Active measurements is a good methodology to research the behaviour of Internet traffic and provides sufficient information to extract solid conclusions about the current state of the Internet ecosystem. However, the Internet is a dynamic environment and probably the observed roles of the analysed players will evolve constantly demonstrating that the Internet is alive. At this point, our goal is fulfilled although we do not close the door to continue exploring the evolution of the Internet players in future research. The next chapter finally develops the conclusions drawn after the completion of the different studies along this thesis.

Chapter 7

CONCLUSIONS

In this final chapter, we summarise the main conclusions of the overall research work conducted throughout the thesis. This chapter compiles the observations and analysis from each previous chapter and notes the fundamental aspects with the objective of better understanding the Internet ecosystem from a traffic analysis perspective and clarifying the interactions among the different Internet players. At the end of the chapter we include a set of future research lines.

7.1 Concluding remarks

This thesis began highlighting the dynamics of the Internet ecosystem. Now, in concluding the research, we can provide a set of reasons to confirm this statement. The evolution of the technology and the rapid adoption of Internet by end users have affected most of the players that form the Internet ecosystem and have forced them to evolve rapidly to continue competing in the market.

Throughout this thesis we have addressed different topics regarding Internet traffic and the interactions among the Internet players. We began analysing the evolution of Internet traffic with the objective of understanding whether network technology and protocols had evolved with influence from by the generalisation of low-latency and multimedia Internet content.

Then, we continued addressing how Internet content has affected the ecosystem and how it has reconfigured the roles of the different Internet players. In this context, some content providers gained significant market share and began to operate their own network infrastructures, something that was previously reserved to telco operators. This fact brought content providers (CPs) greater prominence in the ecosystem.

Moreover, a new technical solution appeared to deal with the huge amount of multimedia traffic: Content Delivery Networks (CDNs). CDNs motivated the emergence of a new Internet actor, which we called **specialised CDN**, focused on exploiting this solution and other featured content delivery services. From a technical perspective, CDNs appeared as a response to the network limitations of the classic client-server architecture. CDNs cache dynamically popular content and place it strategically close to the end user to optimise the network capacity and maximise the quality of experience, which is translated into a dramatic increase in web navigation performance. CDNs have become very popular, and almost all ISPs, regardless of their size and the type of content they manage, implement them in their networks to accelerate their content distribution. More specifically, large content providers use to deploy their own CDN infrastructure, and ISPs more focused on providing Internet connectivity tend to offer CDNs as a service to other players.

After this deep analysis of the state of the art, we decided to develop a measurement platform to actively discover the Internet paths that content follows to reach the end users. Using this platform we have been able to characterise the different interconnection models that Internet players use to distribute Internet traffic. The measurements allowed us to discover many interesting details about how ISPs interact. For example, we have determined that large content providers are trying to interconnect directly with the access ISPs, by-passing transit ISPs, while smaller content providers tend to find other affordable formulas such as contracting content delivery services provided by transit carriers or specialised hosting companies.

Thanks to the lessons learned during the analysis of the evolution of the Internet traffic during last years and thanks to the information extracted from the developed tool, we are able to offer solid arguments about how Internet players interact within the Internet ecosystem and to answer the initial research questions.

7.1.1 From the Internet traffic perspective

We began the thesis by formulating the following research question: "**What is the evolution of the Internet protocols behind the Internet traffic?**". After the analysis of the protocols in Chapter 3, we observed that while Internet traffic, specifically video and multimedia content, has spectacularly increased its demand and technology has incredibly improved the network capacity to absorb these huge amounts of traffic, the protocols responsible for encapsulating and transporting Internet applications have hardly changed. As we showed in Chapter 3, most of the Internet traffic is served using the TCP transport protocol, while UDP and others have a minority presence. Although TCP was not designed for transporting real-time or multimedia traffic, we have observed that using UDP and moving the flow and congestion control to the application protocols has not succeeded. Because TCP in combination with high capacity networks performs fairly well, there is no need to use the simple UDP protocol and to code the intelligence into the application layer, and therefore developers can simply focus on implementing the main purpose of the service and forget about the network functionalities. The TCP congestion and recovery mechanisms facilitate the proliferation of new services in the sense of "*implement what you really know how to do and forget about everything else*", but none of this development would have been possible without the network upgrade and the global adoption of the HTTP web protocol as the de facto standard. With the HTTP protocol something similar happened to TCP. HTTP was initially designed for transporting hypertext web pages but soon many developers exploited its flexibility to carry many types of applications (e.g., SMTP/POP3 for email, FTP for data transfers, RTP for video streaming). This concept, introduced in Chapter 3 and called *webi-*

fication, breaks with the idea of designing a specific application protocol for each service. Therefore, this non-optimal but adaptable protocol increased its popularity and its presence in the network pipes thanks to the ease of implementing almost any type of service and because there is a web browser in almost all devices with Internet connectivity. It is also important to note the support that the HTTP protocol received from the content providers industry, which was key to its expansion and its future improvements (e.g., Google is leading the working group for HTTP version 2). Finally, the consolidation of HTTP over TCP is a valuable experience for both researchers and industry, as it has created a new framework, ***the web platform***, that facilitates the implementation of new services on top of it and the evolution of hardware (devices and networks) in the lower layers of the TCP/IP stack.

7.1.2 From the Internet players' perspective

We continued the thesis by formulating the following research question: "**What are the different interconnection models among Internet players?**". We address this question by analysing how Internet players interact to deliver Internet traffic and we present different subsets of questions to separately address the peculiarities of each one.

First, we focus on the CDNs because we consider that they play an important role in the evolution of the Internet topology. As we explained in the previous section, the Internet experienced a boom of multimedia content, mostly video, and CDNs emerged as an efficient technical solution to deal with the particular requirements of latency and large traffic volume that content providers have. Certainly there are many specialised CDNs such as Akamai and Limelight that operate their networks globally but CDNs are not something exclusive to them, and there are other Internet players that operate their own CDNs (with different levels of complexity) for internal purposes or as a value-added service for third parties. CDNs bring Internet content closer to the end user and by-pass the existing transport networks, breaking with the hierarchic model that we learn from academia. This new paradigm of content distribution has a great impact in the Internet topology because a significant amount of traffic is distributed more efficiently from cache servers located close to the users, and this fact affects many Internet players whose business interests consist of exchanging Internet traffic because they are removed from the value chain. CDNs are simply an evolution in the way the traffic is distributed across the Internet, but this development has motivated the Internet players to rethink their strategies by exploring new business lines or by focusing on a specific market sector.

In parallel with the CDN phenomenon, we have observed the strategic positioning of access ISPs and large content providers. The former, principally ones that belong to the incumbent operators, have played hard ball, investing in expanding their network capacity and range and deploying their backbones towards the global scope. The latter have left their data centres, as birds leave a nest, and have evolved their networks to be closer to their target customers located at the access networks. Not all the content providers have developed their own networks because not all of them considered it strategic, but it is a common approach to be closer to the end users and they have found different formulas to do it. After analysing the interconnection models between access ISPs and large content providers, we have observed that there is a reduced group of content providers with highly developed networks, the Hyper-Giants such as Google, Amazon, Yahoo, Microsoft and Facebook that prioritise the use of direct links with the access ISPs. These content providers take profit from their network muscle (backbones and CDNs) to connect to the access ISPs under (paid) peering agreements, avoiding the use of intermediaries. In addition to the Hyper-Giants, we observed that there are some local media content providers (media groups) that also use this method with some access ISPs. We consider this approach reasonable, as media groups move a high volume of real-time information. Hence, this model of interconnection is restricted to content providers that meet the necessary conditions (traffic load, network and economic resources) for establishing a peering arrangement and that consider it part of their strategy. In this paragraph we have provided our vision regarding the research question "**Is direct interconnection a common trend among access ISPs**

and content providers?”.

On the other hand, there are other *medium-size* content providers that, although they have their own AS networks, use other interconnection alternatives because they are not at the stage of maturity to conduct peering with access ISPs or because the company strategy goes in another direction instead of investing in such complex agreements (they would need to negotiate individually with each access ISP, which requires a high management overload). This category of content providers tends to use the services of tier-1s (CDN and/or IP transit) or specialised CDNs to distribute their content to the end users. This strategy offers a good relation between performance and price, and it makes sense to use this approach when seeking for a scalable solution with global presence and trying to outsource the content distribution service.

The introduction of paid peering by access ISPs has conditioned the tier-1s (transit providers). Initially, tier-1s acted as intermediaries and were responsible for connecting distant networks. However, paid peering forced tier-1s to reformulate their business strategy. The Internet will continue to need IP transit services, but transit ISPs cannot rely wholly on this because there is an important fraction of the traffic that they can no longer serve (the paid peering and CDN shares). Thus, tier-1s decided to expand their services catalogue by providing new featured products around content distribution and taking advantage of their international backbones. Therefore, tier-1s have evolved from their original IP transit services and gone a step further, focusing on a new market sector with other target customers, the **distribution of Internet content** through CDNs (for middle-size content providers) and hosting services (for smaller content providers). This paragraph has addressed the research question ”**How are tier-1s reacting to attract and retain content providers?**”.

Smaller content providers do not require complex interconnection solutions; however, there are many ISPs that target this market sector by providing competitive and affordable alternatives. Normally, small content providers do not have their own network infrastructure and outsource this service to hosting companies. These content providers usually rent servers, Internet connectivity and featured services such as CDNs, load balancing or firewalls, and focus basically on their core business, which is the development and operation of services. Therefore, small CPs take advantage of the benefits that cloud technology has recently released to the market, such as the *virtualisation* that allows creating and scaling virtual machines and virtual network devices on demand and at a competitive price. Hosting companies are addressed to this sector, which forms the *long tail* of Internet content and they operate by grouping multiple small content providers to generate business volume. As we mentioned, *virtualization* has helped hosting companies to allocate resources efficiently, which has been translated into a price reduction. Regarding the interconnection of hosting companies, we have observed that they tend to be present at IXPs. Hence, hosting companies connect their data centres with the IXPs with the objective of maximising the possibilities of interconnecting with other networks.

Related to the research question "**What is the role of the Internet eXchange Points?**" we remark that IXPs have had a significant impact on the Internet ecosystem, as their shared infrastructures facilitate interconnection to any organisation in a single place by reducing costs. We can make an analogy between IXPs and business centres: one can do business with anyone, but someone needs to move to seal the deal, and this point has determined costs. In business centres, companies are close to others, and it is easy to contact others and to promote new deals. In the Internet context where connectivity degree is a competitive/technological advantage, the IXPs offer the possibility of interconnecting with a large number of ISPs in a single location. We see no reason not to be present in an IXP, unless an ISP does not intend to operate in the area where the IXP is located or, is ISP too small. In the latter case, we highly recommend contracting the services of someone present at the IXP. The strength of an IXP is derived from being placed in a strategic location and from having sufficient ISPs interested in participating in the node, i.e., an IXP requires the network effect, as without a critical mass of participants, its existence does not make sense. Therefore, it is necessary to have the involvement of an external agent (e.g., a governmental authority or private organisation) that takes charge of promoting the presence of ISP participants in the IXP node.

The research question **What are the interconnection alternatives for regional and local content providers?** was also addressed when we discussed the tier-1, CDN and hosting solutions. Regional or local does not mean small; it only means that they target a specific geographic location. It sounds crazy to talk about local when the Internet is generally considered *a global tool*, but there are some content providers that are focused only on a small area. Moreover, a small area does not necessarily mean less network resources. For example, media web sites are among the most visited sites in all countries and require specific content delivery solutions (CDNs) to distribute their up-to-date news and real-time video. In contrast, the e-commerce site of a small shop in your neighbourhood will surely be served using a 50 euro-per-year service from a hosting company. Accordingly, the geographic scope is covered by different Internet players.

Finally, we focus on the research question "**Is there any correlation or interconnection pattern for the different types of content providers?**". After analysing the collected data from the measurement tool, we observed that content providers within a similar profile tend to compete using similar interconnection methods. To summarise, we conclude that large CPs prioritise direct connections, medium CPs contract the services of a transit ISP or a specialised CDN, and small CPs place their servers in hosting companies that are mostly in IXPs where they can connect to transit ISPs. There are also some CPs that combine different solutions, e.g., some media groups use their own network to distribute their most recent content and at the same time contract CDN services to deliver their cacheable content. Hosting represents the cheapest and least complex solution, while paid peering represents the most complex and costly solution, not only for the fee for exchanging traffic but also for the investment in network resources and management. Lastly, CDN provides a simpler but more expensive solution than transit service; however, CDNs provide better performance than IP transit. The CDN structure is most likely more complex than establishing a transit agreement, but here we are analysing the situation from the perspective of the content provider, e.g., with a CDN, the CPs need only to deploy their content in the CDN servers, while with a transit service, the CP must manage the network resources. Thus, we can confirm that the Internet market provides sufficient interconnection alternatives to match the different profiles of content providers.

7.2 Future research lines

During the course of this thesis, we have realised that there are many topics that we could not address because they were beyond the scope of this research or because they deserve a whole thesis to themselves because of their relevance. Below, we discuss some future lines of research connected to this work.

7.2.1 Evolution of the current research

We consider this thesis to be a good starting point for further research, and we think that the current measurement platform can be upgraded to improve its capabilities. First of all, we realise that it will be necessary to promote the Mercury platform in multiple countries with the objective of obtaining a greater diversity of measurements and of identifying similarities among different national markets. For this purpose, it would be necessary to establish more partnerships with other universities and research centres and to initiate a more aggressive campaign to promote the platform worldwide (e.g., more presence in conferences, workshops, etc.). Moreover, it could be interesting to involve more researchers in the project to exploit the potential of the platform more intensively and to increase the measurement database. One of the short-term objectives of the initiative would be to execute periodic analysis of the interconnection to continue researching the evolution and dynamics of the Internet ecosystem over time. We expect that the release of new results will help the research community to understand the evolution of content providers, the expansion of access operators and the transformation of tier-1s and other actors.

The Mercury Platform is currently open and available in the GitHub code repository [1] and we invite other researchers and developers to contribute to improving the capabilities of the architecture. Some of the improvements that we propose are i)to implement the Mercury client for Linux/UNIX environments, ii)to implement a graphical user interface (GUI) of the Mercury client iii)to adapt the Mercury server source-code to accept traces from other platforms such as RIPE Atlas, iv)to add new web visualisations that facilitate data interpretation and the creation of self-generated plots, and v)to implement an extension for representing CDN servers localisation and topology AS graphs.

7.2.2 Impact of virtualisation

Virtualisation is a hot topic that consists of generating virtual entities of resources such as computer functions, operating systems, network functions and storage space. Virtualisation is having a notable impact on the Internet ecosystem, as it is reducing network resource costs, promoting the proliferation of a new industry around it and facilitating the implementation of new services. Virtualisation is strongly connected with the trending concept of the *Cloud*. Many companies are specialising in offering cloud services that include virtual private machines, hosting space, high performance computing, machine learning, virtual storage and network functions. Companies from the content sector such as Amazon with Amazon Web Services, Microsoft with Azure Cloud and Google with Google Cloud are exploring this new opportunity successfully. Other companies such as Dropbox, SugarSync, Box or Mega exist specifically to offer cloud storage. And many other ISPs (tier-1s, CDNs, access ISP and hosting companies) are upgrading their infrastructure to provide these virtualisation services. Hence, virtualisation has opened the door to a new market segment where different profile actors are competing, and it

[1] Mercury server (`https://github.com/manuelpalacin/mercury`) and Mercury client (`https://github.com/alexbikfalvi/Mercury`)

has revolutionised the industry by democratising the implementation of services that previously were restricted to certain specialised companies. Virtualisation is enabling the emergence of complex applications because it makes it possible to scale the infrastructure on demand while ensuring resource optimisation. In this context, virtualisation has changed the old pricing policy of pay-per-server for a new one based on pay-per-use. Finally, we note that virtualisation is not only having an impact on the network economy, but also on the network topology: virtualisation creates an overlay network abstraction that represents a complete network environment but shares resources with other entities, e.g., you can logically interact with different network elements in a virtualised environment, but physically, you could be interacting within the same physical resource that virtualises the different network functions. Hence, a network misconfiguration will have to be solved via software through powerful diagnostic tools.

7.2.3 Impact of HTTP/2 and web technologies

As previously mentioned, we expect the HTTP version 2 protocol (HTTP/2) to have a huge impact on the Internet market. *Webification* has caused many applications that previously implemented their own protocol to shift to the web protocol. This trend has facilitated the adoption of many services by end users and has increased the presence of the HTTP protocol in the networks. However, the new brand of multimedia applications requires empowering HTTPv1.1 with many mechanisms that boost the web experience. Some of these mechanisms are to manage caches at the application level or to tune the application server to multiplex several HTTP connections or to maintain the existing ones. With HTTP/2, these tasks (or tweeks) will no longer be needed, as the new protocol will deal with them in a transparent way. Therefore, HTTP/2 will require a *re-education* of web developers to avoid using these provisional workarounds (now considered bad practices) and to encourage them to exploit the new functionalities of HTTP/2. Moreover, HTTP/2 will motivate the update of both the client and server side: server vendors will need to implement the new protocol in the server code, server administrators will need to configure the new capabilities of the software, and web browsers will release new versions. According to recent experiences in the adoption of web technologies, the process will presumably be rapid, as there is a large content industry supporting the change, and the web browser update is simple to perform. HTML5 is an example of how technology has revolutionised the Web. Since the introduction of the HTML5 specification, there has been a spectacular proliferation of new services that rapidly reach their target audience (e.g., WebRTC is a specific but successful case of how real-time communications have landed into the Web). The Web environment facilitates the upgrade of an application from BETA status to a final release, something very strange in the IT industry, where it has traditionally taken years from prototyping to production. This short time-to-market decreases costs and promotes a sustainable environment where different entities (e.g., private companies, standardisation organisms, research centres, open-source communities) collaborate in the development of new specifications. Therefore, it could be interesting to analyse the impact of the adoption of the HTTP/2 protocol and the different web technologies, as they will involve many actors and will affect the organisation of the Internet ecosystem in the short term in many aspects, e.g., Internet browser updates, server adaptation, new programming skills for developers, new applications, and network optimisation.

7.2.4 Impact of IPv6 on the Internet ecosystem

According to recent studies [32], version 6 of the IP protocol (IPv6) currently represents a small but growing fraction of the total Internet traffic (0.64% in 2013), but we consider that the massive introduction of such a protocol will affect the Internet ecosystem (see also the Google online report about IPv6 adoption [2]). For over ten years, IPv6 has been the subject of extensive research, although it seems that its meaningful presence in Internet traffic has not been a reality until a couple of years ago. The

[2] Availability of IPv6 connectivity among Google users `https://www.google.com/intl/en/ipv6/statistics.html#tab=ipv6-adoption`

depletion of IPv4 addresses has motivated the adoption of version 6, but this migration is being long and costly. One of the drawbacks of IPv6 is that to maintain a communication, end-hosts and intermediate devices must understand IPv6, and herein lies the complexity of the migration process. Most IP routers and end device NICs are ready for the new version, but the vast majority of applications are still configured by default to use version 4. Fortunately, IPv6 was introduced in top-level DNS servers in 2008, and by the beginning of 2014, approximately 91% of these top-level servers were ready for version 6. Although security is not mandatory in IPv6, the protocol provides native support for IPsec and it recommends its implementation which requires extensive configuration to be properly secured. The mobile capability of IPv6 will also have a deep impact, as it will introduce new possibilities for roaming between wireless systems. Other elements such as NAT will have to be rethought, which will affect the address space management of millions of organisations, especially in countries that have had a dramatic adoption of the Internet in the recent years like China, India and other regions of Asia and Africa. The change is large and affects multiple actors, and therefore it will definitively have a huge impact on the Internet.

Appendices

Appendix A

MERCURY PLATFORM SPECIFICATIONS

The Mercury Platform is one of the main contributions of this PhD thesis. Thanks to this platform we can perform the Internet measurements and we can obtain and process the information.

A.0.5 Objective

The objective of the Mercury Platform is to provide to the research community a tool to easily discover the interconnection strategies of the different Internet players from an end-to-end point of view.

A.1 Overall Platform

The Mercury Platform is formed by two elements: the Mercury Central Server and the Mercury clients (see Figure A.1). The Mercury clients are software agents installed at the PCs of the participants located in residential networks (access ISPs). The clients perform multiple AS-level traceroutes to different content provider's locations and then they upload these processed path measurements to the Mercury Central Server. The Mercury Central Server is formed by a database and an application server with a web interface and API that allow researchers to find, view and filter the different stored paths and their interconnections.

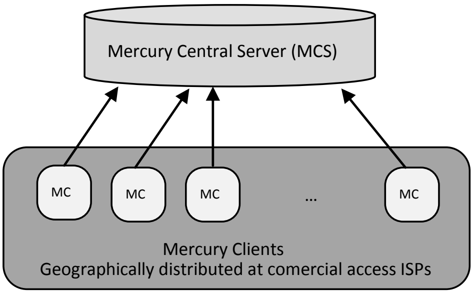
Mercury Central Server (MCS)
MC
MC
MC
...
MC
Mercury Clients
Geographically distributed at comercial access ISPs

Figure A.1: The Mercury Platform.

A.1.1 Mercury client

The Mercury client is a software component that collects AS-level traceroute measurements from the residential networks pointing to the most visited web sites. These AS-level traceroutes show the path that Internet packets follow and the ISPs (autonomous systems) that these packets traverse to reach the destination content provider.

Specifications

- Software written in C#
- Available for Windows Vista/7/8 with Framework v3.0
- It requires administrator privileges to be executed
- It requires Internet connectivity with ICMP blocking rule of firewall disabled
- It surveys the URLs of the top100 most visited web sites by Alexa Ranking [8] per selected country

- It surveys the URLs of the resources (images, js, css, etc.) of the previous top100

- It executes a customised Paris traceroute version to mitigate the multi-path nature of classic traceroute

- It can be configured the number of different flows and the number attempts for each flow

- It translates IP addresses to their corresponding AS

- It detects when an IP address corresponds to an IXP

- It detects and correct (when possible) anomalous traces and mark them using a scale of consistency

- It identifies the interconnection relationships between AS neighbors along the paths

- It generates statistical information (AS hops and AS relationship types)

- It uploads automatically the information to the Mercury Central Server

Implementation internals

The Mercury client is a desktop application written in C# that automatically downloads the list of URLs and executes the traceroute probes. Before tracerouting, the client executes a *nslookup* query for each URL in order to obtain the corresponding IP address. Sometimes it obtains multiple IPs for a same URL. This denotes the existence of load-balancing which is a technique of CDN architectures. Then, the client executes the traceroute measurements for the set of IPs of each URL. The traceroute is performed using a modified version of Paris traceroute [12] to minimise the effect of multiple paths during IP routing. That is, using Paris traceroute we modify the IP packets (IP and ICMP/UDP headers and payload) for generating traffic flows that follow the same paths across routers that implement per-packet load balancing [10]. In our implementation, for each destination IP address we generate 5 different flows with 5 attempts per flow, for both ICMP and UDP traffic, respectively. Therefore, we generate 50 different traceroutes for each destination. Then, the client translates each IP address from each IP hop to an AS number using the MCS-API. It is important to note that we are mostly interested in the AS level of a traffic path. This means that we are mostly concerned about detecting and correcting inconsistencies at edge routers between different ASes. This step is crucial because here is where we aggregate flows and we detect, correct (if it is possible) and discard multi-path behaviours that lead us to obtain loops and missing ASes. One of the inference heuristics that we use to solve these anomalies is using an algorithm that detects AS hop inconsistencies and corrects them based on the analysis of the previous and next AS hops (see Fig. A.2) and assuming a hot potato policy at the AS level. Therefore, the combination of Paris traceroute with this algorithm minimises the number of incorrect traces.

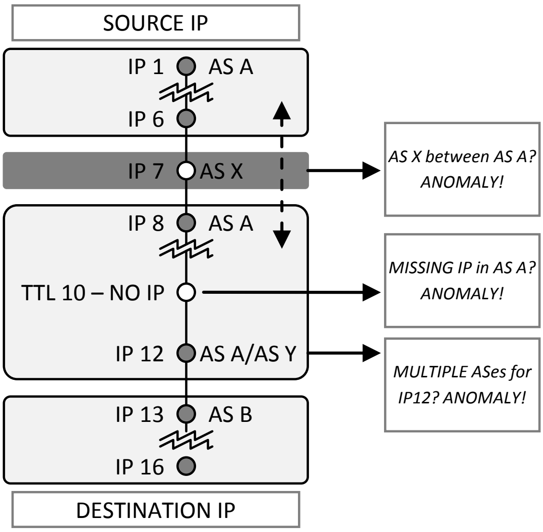

Figure A.2: Detecting AS anomalies.

Once the traceroutes are executed, the program annotates the traceroute data with information from the external datasets of the MCS. In this, step we include to the traces the AS relationships between adjacent ASes, the geolocation of the source and destination end-points and we generate an statistical summary. The client is the responsible of processing all the information releasing the MCS for doing only searching and publication tasks. Finally, when the client ends processing, all the data is sent to the MCS via the REST-API.

To store the measurement data, we use a high availability and low latency system formed by a MongoDB *database*. In conjunction with a *query engine* developed in Java2EE, it gives a great flexibility in uploading data, applying filters or aggregating statistical results. For an improved user experience, Mercury uses a cache system to boost performance for the most popular queries.

Installation and execution

The Mercury client does not require an specific installation because it just need to be placed into a Windows folder. However, it requires to be installed in a Windows Vista or newer with the Microsoft Framework v3.0 installed and it is necessary to enable de Windows firewall to not blocking ICMP packets. In addition, it requires to be executed using *administrator* privileges.

To execute the Mercury client you need to open a Windows command line terminal and call the ”MercuryTool.exe” executable (see Figure A.3). It automatically downloads a list of URLs according to your country (based on your operating system locale) to probe and it starts the measurements. Alternatively, you can pass a comma-separated list of URLs as an argument (e.g. google.com,facebook.com,twitter.com,upf.edu). While it executes the measurements, it shows the results and the autonomous systems between the client and the destination URL (content provider server). When it finishes, it asks you to press a button to close the application.

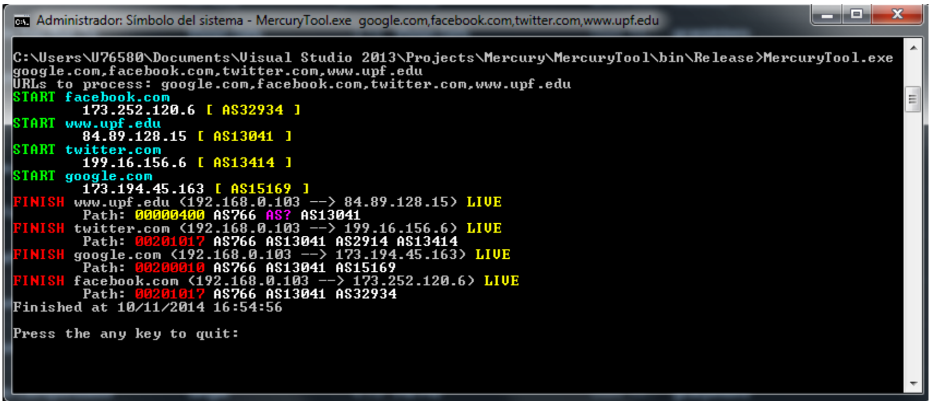
Administrador: Símbolo del sistema - MercuryTool.exe google.com,facebook.com,twitter.com,www.upf.edu

```
C:\Users\U76580\Documents\Visual Studio 2013\Projects\Mercury\MercuryTool\bin\Release>MercuryTool.exe
google.com,facebook.com,twitter.com,www.upf.edu
URLs to process: google.com,facebook.com,twitter.com,www.upf.edu
START facebook.com
        173.252.120.6 [ AS32934 ]
START www.upf.edu
        84.89.128.15 [ AS13041 ]
START twitter.com
        199.16.156.6 [ AS13414 ]
START google.com
        173.194.45.163 [ AS15169 ]
FINISH www.upf.edu (192.168.0.103 --> 84.89.128.15) LIVE
        Path: 00000400 AS766 AS? AS13041
FINISH twitter.com (192.168.0.103 --> 199.16.156.6) LIVE
        Path: 00201017 AS766 AS13041 AS2914 AS13414
FINISH google.com (192.168.0.103 --> 173.194.45.163) LIVE
        Path: 00200010 AS766 AS13041 AS15169
FINISH facebook.com (192.168.0.103 --> 173.252.120.6) LIVE
        Path: 00201017 AS766 AS13041 AS32934
Finished at 10/11/2014 16:54:56

Press the any key to quit:
```

Figure A.3: MercuryTool execution.

A part from the Mercury client, one can use the developed AS traceroute developed for Windows. This tool is a sub-component of the Mercury client that can be use for testing proposes. As we have previously explained, it implements Paris traceroute and can be configured for probing using the ICMP and UDP protocols and with different payload configurations. To execute it, you just need to open a Windows command line terminal and call the ”MercuryTraceroute.exe” executable (see Figures A.4,A.5,A.6). This tool can be used for detecting the packets path at the IP-level and at the AS-level and to discover whether it exists multi-path load balancing in the packet route.

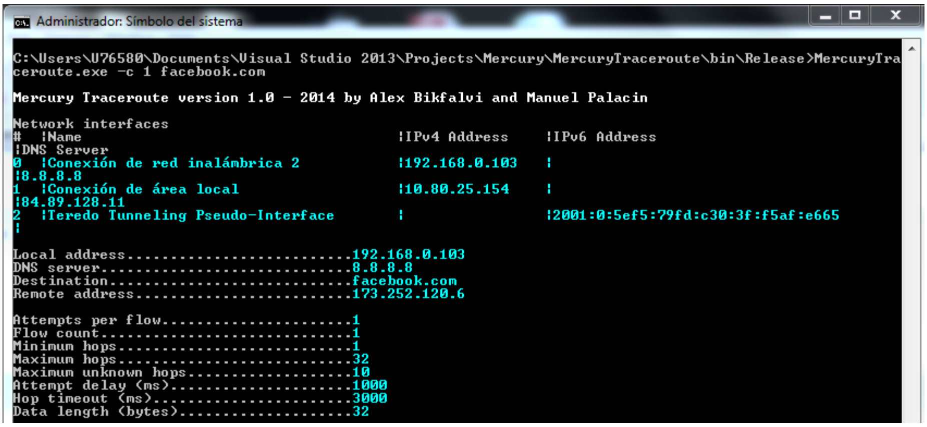

```
C:\Users\U76580\Documents\Visual Studio 2013\Projects\Mercury\MercuryTraceroute\bin\Release>MercuryTra
ceroute.exe -c 1 facebook.com

Mercury Traceroute version 1.0 - 2014 by Alex Bikfalvi and Manuel Palacin

Network interfaces
#  ¦Name                                    ¦IPv4 Address    ¦IPv6 Address
¦DNS Server
0  ¦Conexión de red inalámbrica 2           ¦192.168.0.103   ¦
¦8.8.8.8
1  ¦Conexión de área local                  ¦10.80.25.154    ¦
¦84.89.128.11
2  ¦Teredo Tunneling Pseudo-Interface       ¦                ¦2001:0:5ef5:79fd:c30:3f:f5af:e665
¦

Local address..........................192.168.0.103
DNS server.............................8.8.8.8
Destination............................facebook.com
Remote address.........................173.252.120.6

Attempts per flow......................1
Flow count.............................1
Minimum hops...........................1
Maximum hops...........................32
Maximum unknown hops...................10
Attempt delay (ms).....................1000
Hop timeout (ms).......................3000
Data length (bytes)....................32
```

Figure A.4: Mercury Traceroute1 execution: DNS information

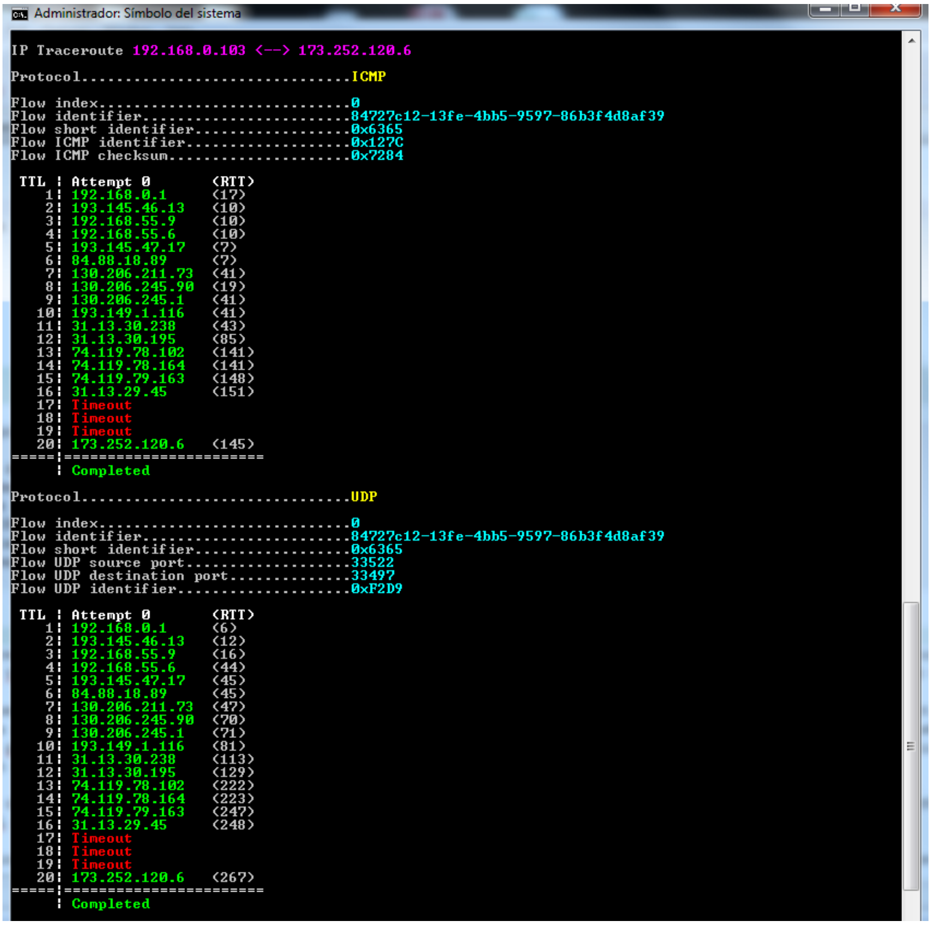

```
IP Traceroute 192.168.0.103 <--> 173.252.120.6

Protocol...............................ICMP

Flow index.............................0
Flow identifier........................84727c12-13fe-4bb5-9597-86b3f4d8af39
Flow short identifier..................0x6365
Flow ICMP identifier...................0x127C
Flow ICMP checksum.....................0x7284

 TTL ¦ Attempt 0       <RTT>
    1¦ 192.168.0.1     <17>
    2¦ 193.145.46.13   <10>
    3¦ 192.168.55.9    <10>
    4¦ 192.168.55.6    <10>
    5¦ 193.145.47.17   <7>
    6¦ 84.88.18.89     <7>
    7¦ 130.206.211.73  <41>
    8¦ 130.206.245.90  <19>
    9¦ 130.206.245.1   <41>
   10¦ 193.149.1.116   <41>
   11¦ 31.13.30.238    <43>
   12¦ 31.13.30.195    <85>
   13¦ 74.119.78.102   <141>
   14¦ 74.119.78.164   <141>
   15¦ 74.119.79.163   <148>
   16¦ 31.13.29.45     <151>
   17¦ Timeout
   18¦ Timeout
   19¦ Timeout
   20¦ 173.252.120.6   <145>
=====¦=======================
     ¦ Completed

Protocol...............................UDP

Flow index.............................0
Flow identifier........................84727c12-13fe-4bb5-9597-86b3f4d8af39
Flow short identifier..................0x6365
Flow UDP source port...................33522
Flow UDP destination port..............33497
Flow UDP identifier....................0xF2D9

 TTL ¦ Attempt 0       <RTT>
    1¦ 192.168.0.1     <6>
    2¦ 193.145.46.13   <12>
    3¦ 192.168.55.9    <16>
    4¦ 192.168.55.6    <44>
    5¦ 193.145.47.17   <45>
    6¦ 84.88.18.89     <45>
    7¦ 130.206.211.73  <47>
    8¦ 130.206.245.90  <70>
    9¦ 130.206.245.1   <71>
   10¦ 193.149.1.116   <81>
   11¦ 31.13.30.238    <113>
   12¦ 31.13.30.195    <129>
   13¦ 74.119.78.102   <222>
   14¦ 74.119.78.164   <223>
   15¦ 74.119.79.163   <247>
   16¦ 31.13.29.45     <248>
   17¦ Timeout
   18¦ Timeout
   19¦ Timeout
   20¦ 173.252.120.6   <267>
=====¦=======================
     ¦ Completed
```

Figure A.5: Mercury Traceroute execution: IP level.

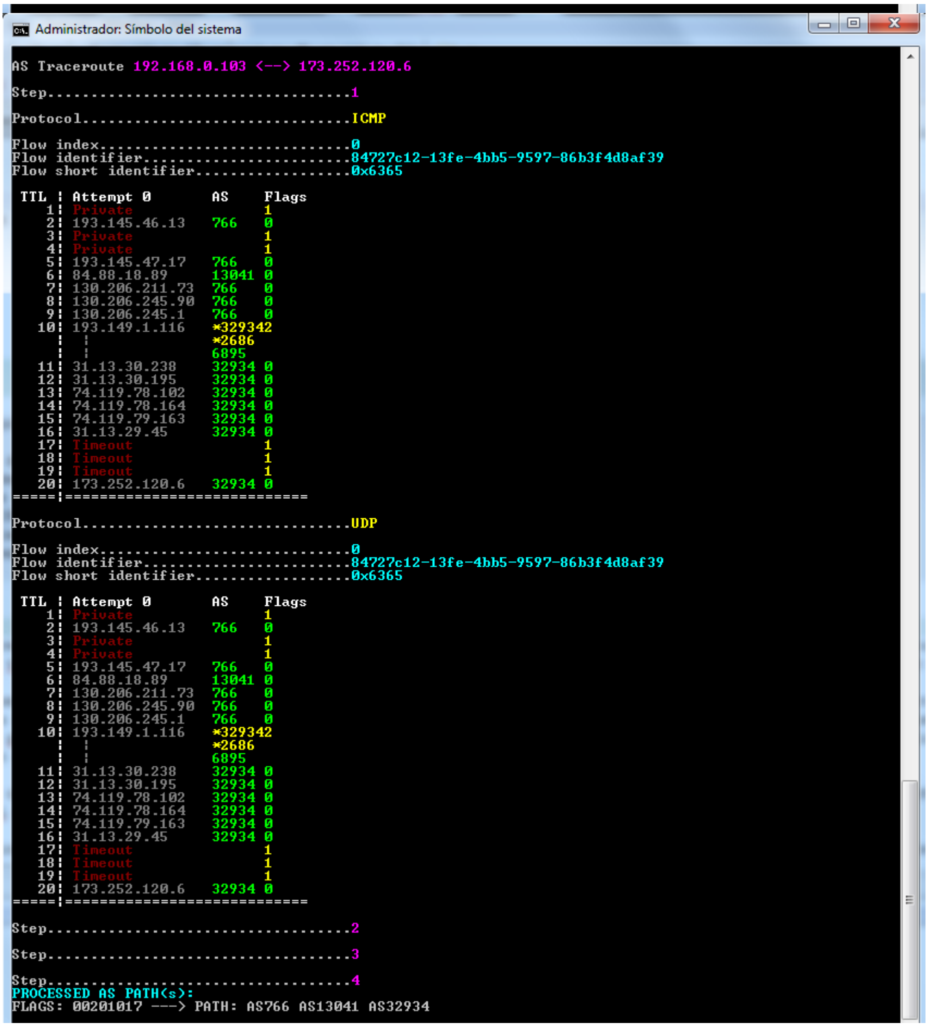

```
AS Traceroute 192.168.0.103 <--> 173.252.120.6

Step..................................1

Protocol..............................ICMP

Flow index............................0
Flow identifier.......................84727c12-13fe-4bb5-9597-86b3f4d8af39
Flow short identifier.................0x6365

 TTL ¦ Attempt 0         AS     Flags
    1¦ Private                  1
    2¦ 193.145.46.13     766    0
    3¦ Private                  1
    4¦ Private                  1
    5¦ 193.145.47.17     766    0
    6¦ 84.88.18.89       13041  0
    7¦ 130.206.211.73    766    0
    8¦ 130.206.245.90    766    0
    9¦ 130.206.245.1     766    0
   10¦ 193.149.1.116     *329342
     ¦   ¦               *2686
     ¦   ¦               6895
   11¦ 31.13.30.238      32934 0
   12¦ 31.13.30.195      32934 0
   13¦ 74.119.78.102     32934 0
   14¦ 74.119.78.164     32934 0
   15¦ 74.119.79.163     32934 0
   16¦ 31.13.29.45       32934 0
   17¦ Timeout                  1
   18¦ Timeout                  1
   19¦ Timeout                  1
   20¦ 173.252.120.6     32934 0
=====¦===========================

Protocol..............................UDP

Flow index............................0
Flow identifier.......................84727c12-13fe-4bb5-9597-86b3f4d8af39
Flow short identifier.................0x6365

 TTL ¦ Attempt 0         AS     Flags
    1¦ Private                  1
    2¦ 193.145.46.13     766    0
    3¦ Private                  1
    4¦ Private                  1
    5¦ 193.145.47.17     766    0
    6¦ 84.88.18.89       13041  0
    7¦ 130.206.211.73    766    0
    8¦ 130.206.245.90    766    0
    9¦ 130.206.245.1     766    0
   10¦ 193.149.1.116     *329342
     ¦   ¦               *2686
     ¦   ¦               6895
   11¦ 31.13.30.238      32934 0
   12¦ 31.13.30.195      32934 0
   13¦ 74.119.78.102     32934 0
   14¦ 74.119.78.164     32934 0
   15¦ 74.119.79.163     32934 0
   16¦ 31.13.29.45       32934 0
   17¦ Timeout                  1
   18¦ Timeout                  1
   19¦ Timeout                  1
   20¦ 173.252.120.6     32934 0
=====¦===========================

Step..................................2

Step..................................3

Step..................................4
PROCESSED AS PATH(s):
FLAGS: 00201017 ---> PATH: AS766 AS13041 AS32934
```

Figure A.6: Mercury Traceroute execution: AS level.

Source code

Anyone can contribute to the development and improvement of the Mercury client at `https://github.com/alexbikfalvi/Mercury`.

A.1.2 Mercury Central Server

The Mercury Central Server (MCS) is the software component that stores the AS-level traceroute measurements from the mercury clients. It also provides a web and an API interface to search, filter and view the stored information. Furthermore it also aggregates critical information from external databases that is required for the measurement processing.

Specifications

- Java J2EE application based on Spring Framework 3.0
- Tomcat 7 application server
- MongoDB Database: non-relational database
- Application server and database deployed in an Ubuntu 12.04LTS virtual server at Universitat Pompeu Fabra
- It uses daily BGP reports to have an updated ip-to-as database
- It can detect a large amount of IP addresses located in IXPs
- It detect the interconnection relationship type between two neighbour ASes
- It provides a web interface to search, filter and consult data
- It provides a REST-WS API

Implementation internals

The Mercury Central Server is the responsible of formatting and storing the information from the external datasets and the responsible of storing and publishing the processed measurements sent by the clients. The desktop client performs the traceroute probes and processes the results for obtaining the IP-to-AS translation, the IXP detection and the AS relationships of the end-to-end path.

The MCS obtains the IP-to-AS mappings from the BGP monitors through the Routing Report project [82]. It downloads daily the BGP report that contains the AS origin of the running IP prefixes. This information is structured and stored in a database and combined with IXP mappings from the PeeringDB project. In addition, it obtains the relationship type of each interconnected AS pair from the CAIDA AS relationships table [18] and the geolocation of the IPs from the MaxMind free service [69]. Finally the central server has the list of URLs to be examined for each country. All these datasets are stored in the MCS and the clients can download them using the REST-API.

The MCS uses a high speed database based on the MongoDB non-relational database. This database allows a high flexibility of storage and a high performance in storing, searching and retrieving records. We use this database in front of traditional SQL database because its performance its considerably better when we manage databases over 1 GB. Figure A.7 shows the different components of the MCS and its interactions with external data sources.

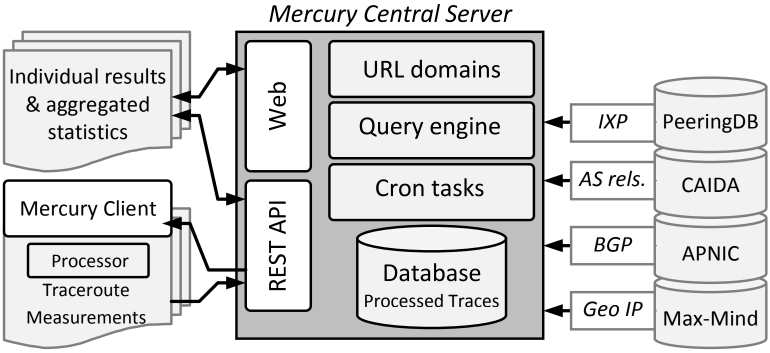

Figure A.7: The Mercury Central Server.

The *web interface* is developed using HTML5 standards. Furthermore, Mercury uses a responsive CSS template that adjusts to different browsers and devices and enriches the data visualisation using a set of intuitive JavaScript plots. All the functions offered by the web interface are offered using a *REST Web Service API* in JSON format.

All the stored measurements are publicly available using the Web interface and the REST-WS API. Researchers can obtain aggregated statistics for a set of paths. For example one can aggregate paths filtering by a single destination URL, origin/destination AS, origin/destination city and origin/destination country. These statistics include average, mean and standard deviation about the number of AS hops and the type of the AS relationships but also other network indicators like the number of destinations IPs and ASes that host a destination URL, e.g. *www.20minutos.es* is in average at 0.88 AS hops, points to two different IPs (193.148.34.26 and 89.140.253.190) and is hosted in two different ASes (AS3324 and AS6739).

External databases

The Mercury Central Server aggregates information from different external databases and re-structure this data into its own MongoDB database. Following we describe the used public databases:

- **Most popular web sites:** we collect the top 100 most popular web sites for each country from the ranking generated by Alexa [8].

- **BGP information:** we collect the BGP data that include the IP prefixes and their originating autonomous system from the BGP Routing Table Analysis [82]. We introduce the prefixes into our database creating an IP-to-AS mapping. The data can be originated from a router located in APNIC or from the the RouteViews project and is collected from the MCS in a daily basis.

- **IXP information:** in order to increase the number of IP-to-AS mappings and to detect ASes using IXPs, we collect information from the public database of IXPs, PeeringDB [81]. We adapt the information from this public dataset and we introduce it into our database.

- **AS relationships:** we collect the interconnection relationships between neighbour ASes using the AS Relationships dataset from CAIDA [18]. We use this dataset, to identify the peering (*p2p*), customer-to-provider (*c2p* or *p2c*) and sibling (*s2s*) relationships. CAIDA updates this dataset periodically based on their inferring algorithms (see State of the Art chapter for more information). All this information is introduced into the MCS database.

- **IP geolocation:** we collect the geolocation of the end points (source and destination hosts) using the free dataset of MaxMind [69]. This dataset is included into our database and we can return the geolocation data for a given IP address. This dataset is not as accurate as we wish (e.g. it geolocates a Google server from Madrid, Spain in Mountain View, California) but at least it can offer us the approximate position for most of the end points.

Web Interface

The web interface of the Mercury Platform is tool that users use to visualise the gathered and processed information. The web interface is available for everybody without any type of restriction at `http://mercury.upf.edu/mercury`. The web site is implemented using HTML5, a responsive CSS3 template and multiple javascript libraries that improve the user experience. The landing page shows you the multiple options that the web site offers (see Figure A.8). Then, the user can see the total list of processed measurements (see Figure A.9) and the details of each measurement (see Figure A.10). Inside the measurement details one can observe statistical information about the types of interconnection along the measurement path. Here we can also see descriptive information in JSON format about the autonomous systems placed between the source of the measurement and the destination content provider. In addition, the applications allows you to filter by source/destination city, country, IP, AS, URL (see Figures A.11,A.12,A.13). Finally, the web interface allows you to automatically call the API (see Figure A.14).

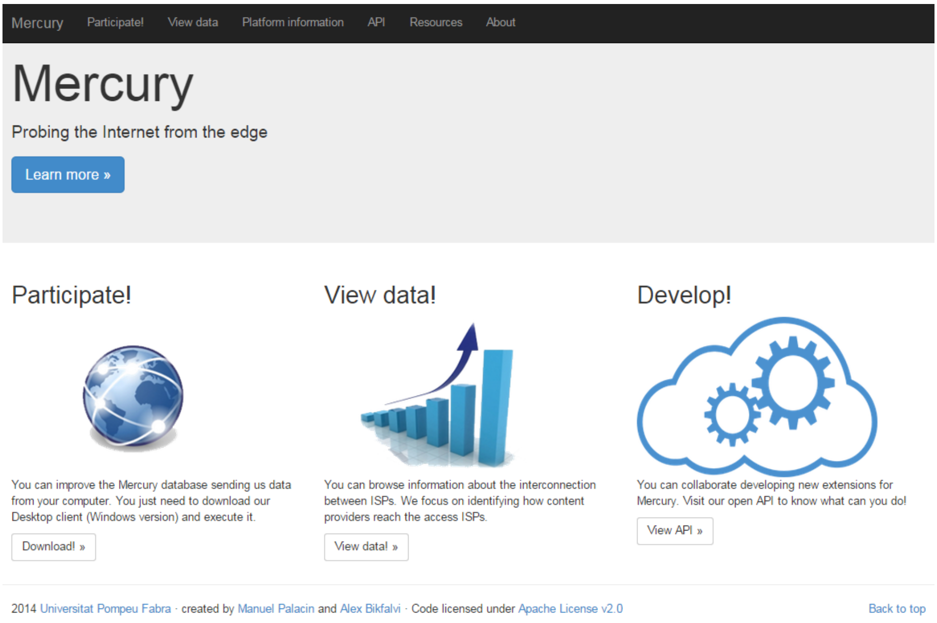

Mercury

Probing the Internet from the edge

Participate!

You can improve the Mercury database sending us data from your computer. You just need to download our Desktop client (Windows version) and execute it.

View data!

You can browse information about the interconnection between ISPs. We focus on identifying how content providers reach the access ISPs.

Develop!

You can collaborate developing new extensions for Mercury. Visit our open API to know what can you do!

2014 Universitat Pompeu Fabra · created by Manuel Palacin and Alex Bikfalvi · Code licensed under Apache License v2.0

Figure A.8: Home page of Mercury.

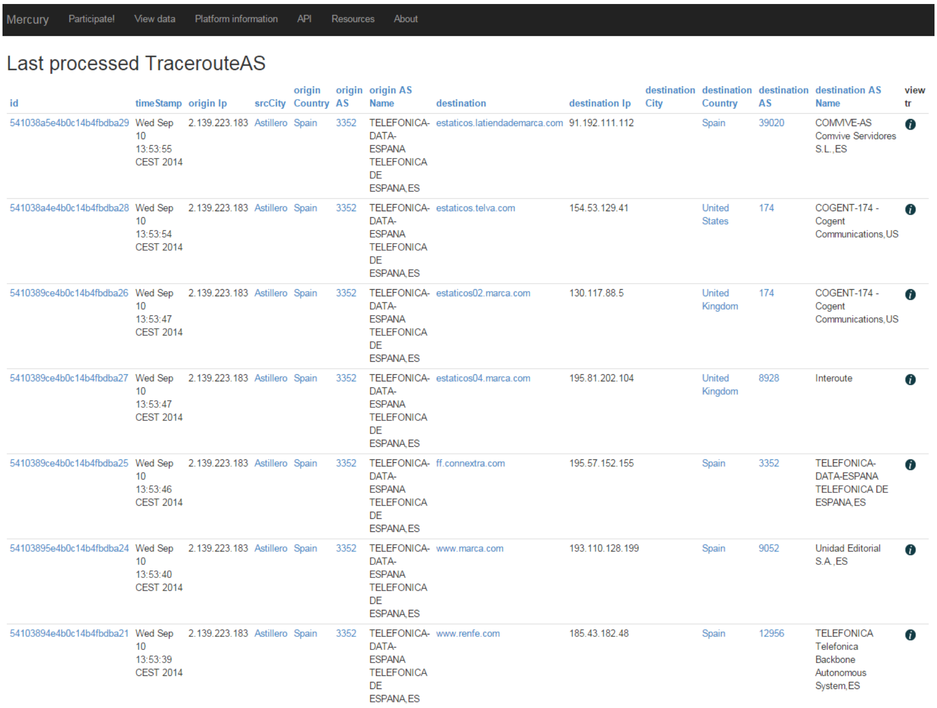

Last processed TracerouteAS

id	timeStamp	origin Ip	srcCity	origin Country	origin AS	origin AS Name	destination	destination Ip	destination City	destination Country	destination AS	destination AS Name	view tr
541038a5e4b0c14b4fbdba29	Wed Sep 10 13:53:55 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	estaticos.latiendademarca.com	91.192.111.112		Spain	39020	COMVIVE-AS Comvive Servidores S.L.,ES	ℹ
541038a4e4b0c14b4fbdba28	Wed Sep 10 13:53:54 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	estaticos.telva.com	154.53.129.41		United States	174	COGENT-174 - Cogent Communications,US	ℹ
5410389ce4b0c14b4fbdba26	Wed Sep 10 13:53:47 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	estaticos02.marca.com	130.117.88.5		United Kingdom	174	COGENT-174 - Cogent Communications,US	ℹ
5410389ce4b0c14b4fbdba27	Wed Sep 10 13:53:47 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	estaticos04.marca.com	195.81.202.104		United Kingdom	8928	Interoute	ℹ
5410389ce4b0c14b4fbdba25	Wed Sep 10 13:53:46 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	ff.connextra.com	195.57.152.155		Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	ℹ
54103895e4b0c14b4fbdba24	Wed Sep 10 13:53:40 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	www.marca.com	193.110.128.199		Spain	9052	Unidad Editorial S.A.,ES	ℹ
54103894e4b0c14b4fbdba21	Wed Sep 10 13:53:39 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	www.renfe.com	185.43.182.48		Spain	12956	TELEFONICA Telefonica Backbone Autonomous System,ES	ℹ

Figure A.9: Measurement list.

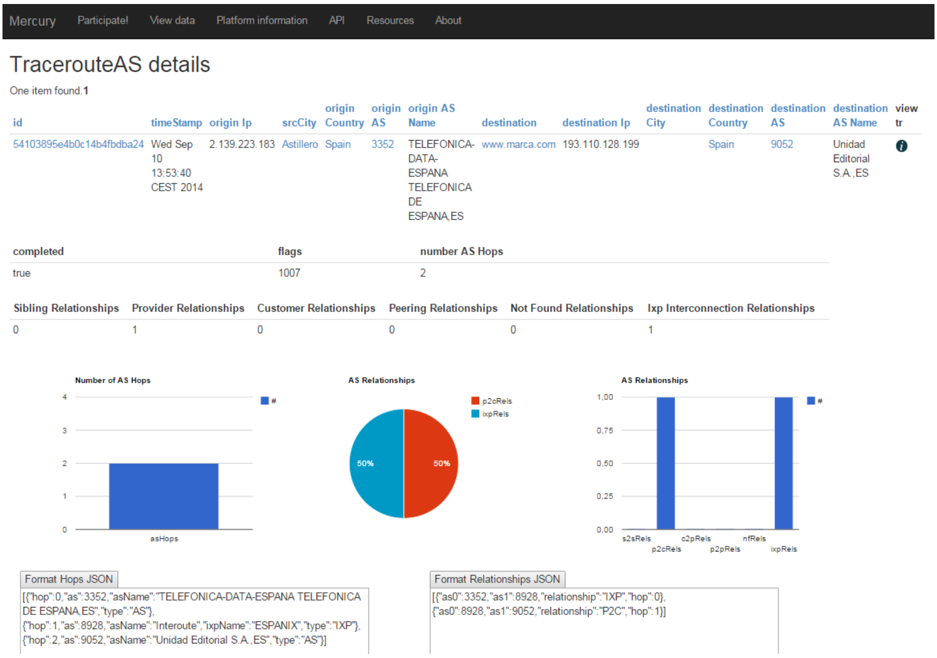

TracerouteAS details

One item found.**1**

id	timeStamp	origin Ip	srcCity	origin Country	origin AS	origin AS Name	destination	destination Ip	destination City	destination Country	destination AS	destination AS Name	view tr
54103895e4b0c14b4fbdba24	Wed Sep 10 13:53:40 CEST 2014	2.139.223.183	Astillero	Spain	3352	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES	www.marca.com	193.110.128.199		Spain	9052	Unidad Editorial S.A.,ES	ⓘ

completed	flags	number AS Hops
true	1007	2

Sibling Relationships	Provider Relationships	Customer Relationships	Peering Relationships	Not Found Relationships	Ixp Interconnection Relationships
0	1	0	0	0	1

[{"hop":0,"as":3352,"asName":"TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA,ES","type":"AS"},
{"hop":1,"as":8928,"asName":"Interoute","ixpName":"ESPANIX","type":"IXP"},
{"hop":2,"as":9052,"asName":"Unidad Editorial S.A.,ES","type":"AS"}]

[{"as0":3352,"as1":8928,"relationship":"IXP","hop":0},
{"as0":8928,"as1":9052,"relationship":"P2C","hop":1}]

Figure A.10: Measurement details in Mercury.

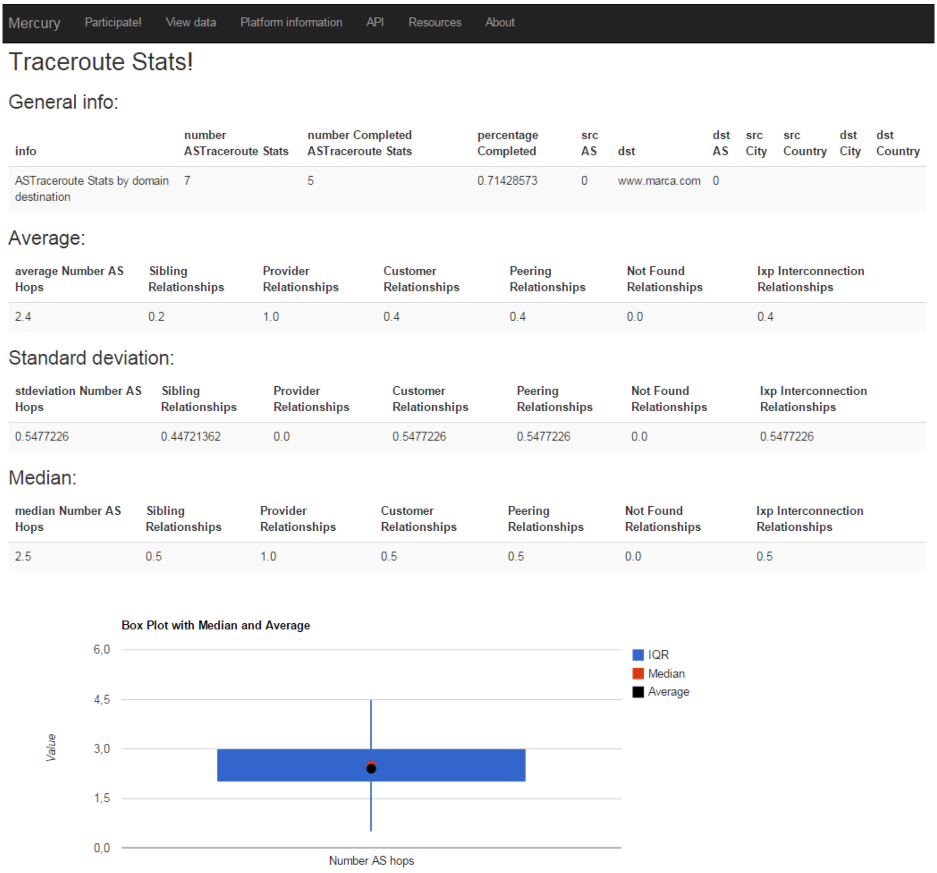

Traceroute Stats!

General info:

info	number ASTraceroute Stats	number Completed ASTraceroute Stats	percentage Completed	src AS	dst	dst AS	src City	src Country	dst City	dst Country
ASTraceroute Stats by domain destination	7	5	0.71428573	0	www.marca.com	0				

Average:

average Number AS Hops	Sibling Relationships	Provider Relationships	Customer Relationships	Peering Relationships	Not Found Relationships	Ixp Interconnection Relationships
2.4	0.2	1.0	0.4	0.4	0.0	0.4

Standard deviation:

stdeviation Number AS Hops	Sibling Relationships	Provider Relationships	Customer Relationships	Peering Relationships	Not Found Relationships	Ixp Interconnection Relationships
0.5477226	0.44721362	0.0	0.5477226	0.5477226	0.0	0.5477226

Median:

median Number AS Hops	Sibling Relationships	Provider Relationships	Customer Relationships	Peering Relationships	Not Found Relationships	Ixp Interconnection Relationships
2.5	0.5	1.0	0.5	0.5	0.0	0.5

Figure A.11: Filtering results in Mercury (1).

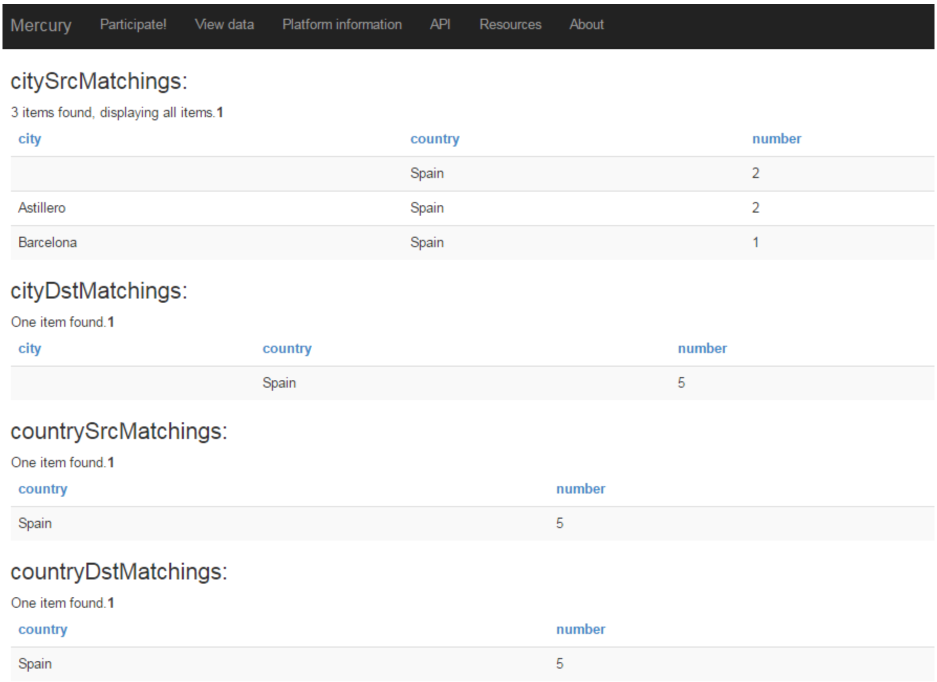

citySrcMatchings:

3 items found, displaying all items.**1**

city	country	number
	Spain	2
Astillero	Spain	2
Barcelona	Spain	1

cityDstMatchings:

One item found.**1**

city	country	number
	Spain	5

countrySrcMatchings:

One item found.**1**

country	number
Spain	5

countryDstMatchings:

One item found.**1**

country	number
Spain	5

Figure A.12: Filtering results in Mercury (2).

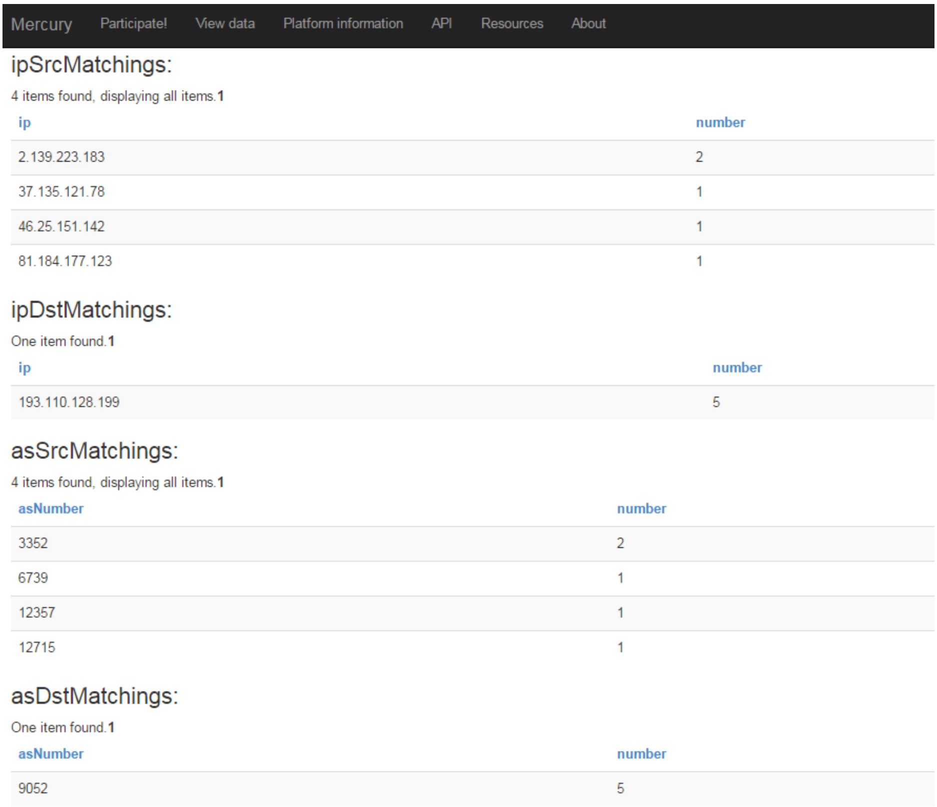

ipSrcMatchings:

4 items found, displaying all items.1

ip	number
2.139.223.183	2
37.135.121.78	1
46.25.151.142	1
81.184.177.123	1

ipDstMatchings:

One item found.1

ip	number
193.110.128.199	5

asSrcMatchings:

4 items found, displaying all items.1

asNumber	number
3352	2
6739	1
12357	1
12715	1

asDstMatchings:

One item found.1

asNumber	number
9052	5

Figure A.13: Filtering results in Mercury (3).

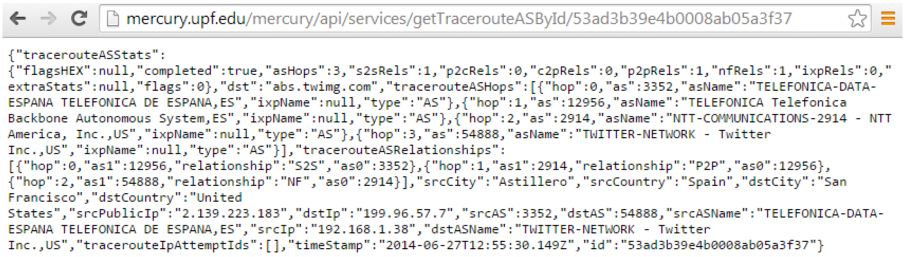
mercury.upf.edu/mercury/api/services/getTracerouteASById/53ad3b39e4b0008ab05a3f37

```
{"tracerouteASStats":
{"flagsHEX":null,"completed":true,"asHops":3,"s2sRels":1,"p2cRels":0,"c2pRels":0,"p2pRels":1,"nfRels":1,"ixpRels":0,"
extraStats":null,"flags":0},"dst":"abs.twimg.com","tracerouteASHops":[{"hop":0,"as":3352,"asName":"TELEFONICA-DATA-
ESPANA TELEFONICA DE ESPANA,ES","ixpName":null,"type":"AS"},{"hop":1,"as":12956,"asName":"TELEFONICA Telefonica
Backbone Autonomous System,ES","ixpName":null,"type":"AS"},{"hop":2,"as":2914,"asName":"NTT-COMMUNICATIONS-2914 - NTT
America, Inc.,US","ixpName":null,"type":"AS"},{"hop":3,"as":54888,"asName":"TWITTER-NETWORK - Twitter
Inc.,US","ixpName":null,"type":"AS"}],"tracerouteASRelationships":
[{"hop":0,"as1":12956,"relationship":"S2S","as0":3352},{"hop":1,"as1":2914,"relationship":"P2P","as0":12956},
{"hop":2,"as1":54888,"relationship":"NF","as0":2914}],"srcCity":"Astillero","srcCountry":"Spain","dstCity":"San
Francisco","dstCountry":"United
States","srcPublicIp":"2.139.223.183","dstIp":"199.96.57.7","srcAS":3352,"dstAS":54888,"srcASName":"TELEFONICA-DATA-
ESPANA TELEFONICA DE ESPANA,ES","srcIp":"192.168.1.38","dstASName":"TWITTER-NETWORK - Twitter
Inc.,US","tracerouteIpAttemptIds":[],"timeStamp":"2014-06-27T12:55:30.149Z","id":"53ad3b39e4b0008ab05a3f37"}
```

Figure A.14: API page of Mercury.

API

The Mercury application programming interface (API) is a complementary tool that interconnect the MCS with the Mercury clients and with the researches that want to test the capabilities of the platform. The API is implemented on top of the application server and is based on RESTFul Web Services (REST-WS). One can interact with the API using a simple Internet browser or using the command-line tools like cURL [35]. The Mercury API accepts HTTP-GET and HTTP-POST request messages and it returns the response messages using JSON. Following we describe the different API methods:

- **Get My info:** Get my Info provides you information about your public IP address and the corresponding Autonomous System (see Figure A.15).

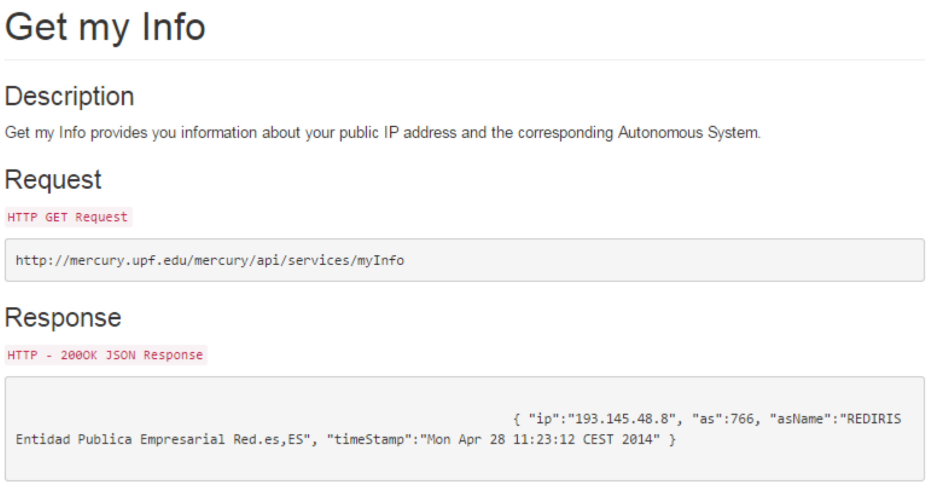

Get my Info

Description

Get my Info provides you information about your public IP address and the corresponding Autonomous System.

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/myInfo
```

Response

`HTTP - 200OK JSON Response`

```
                                                                { "ip":"193.145.48.8", "as":766, "asName":"REDIRIS
Entidad Publica Empresarial Red.es,ES", "timeStamp":"Mon Apr 28 11:23:12 CEST 2014" }
```

Figure A.15: Mercury API: Get My info.

- **Get AS Relationship:** Get AS Relationship provides you information about the interconnection relationships between two adjacent Autonomous Systems. AS relationships are extracted from CAIDA AS Relationships project (see Figure A.16).

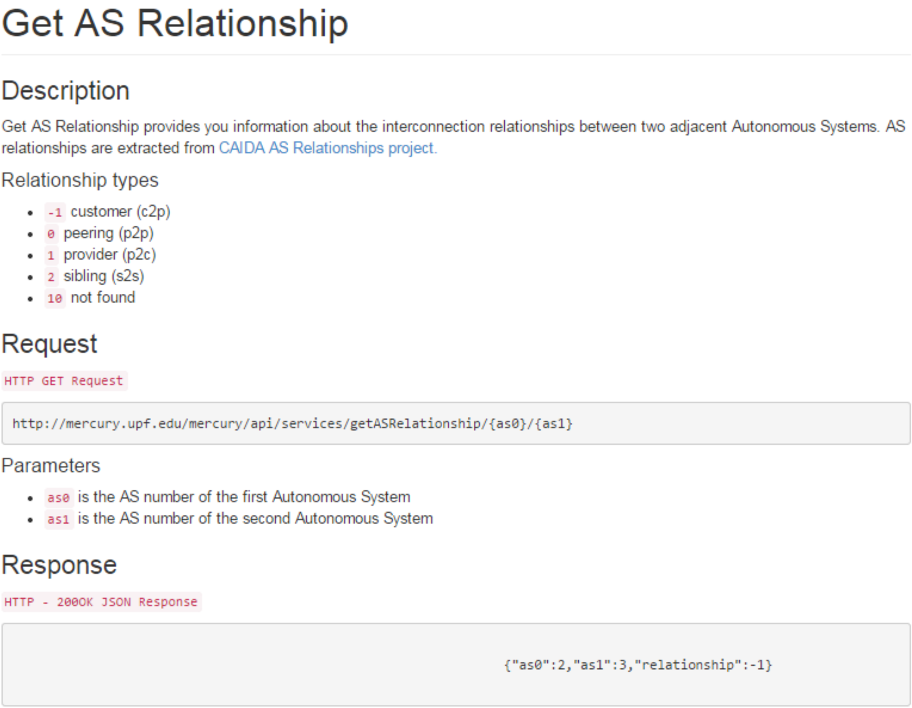

Get AS Relationship

Description

Get AS Relationship provides you information about the interconnection relationships between two adjacent Autonomous Systems. AS relationships are extracted from CAIDA AS Relationships project.

Relationship types

- `-1` customer (c2p)
- `0` peering (p2p)
- `1` provider (p2c)
- `2` sibling (s2s)
- `10` not found

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/getASRelationship/{as0}/{as1}
```

Parameters

- `as0` is the AS number of the first Autonomous System
- `as1` is the AS number of the second Autonomous System

Response

`HTTP - 200OK JSON Response`

```
{"as0":2,"as1":3,"relationship":-1}
```

Figure A.16: Mercury API: Get AS Relationship.

- **Get IP to ASN Mapping:** Get IP2ASN Mapping provides you information about the corresponding AS number for an IP address. IP to ASN translations are inferred from analysing BGP messages from BGP Routing Table Analysis (see Figure A.17).

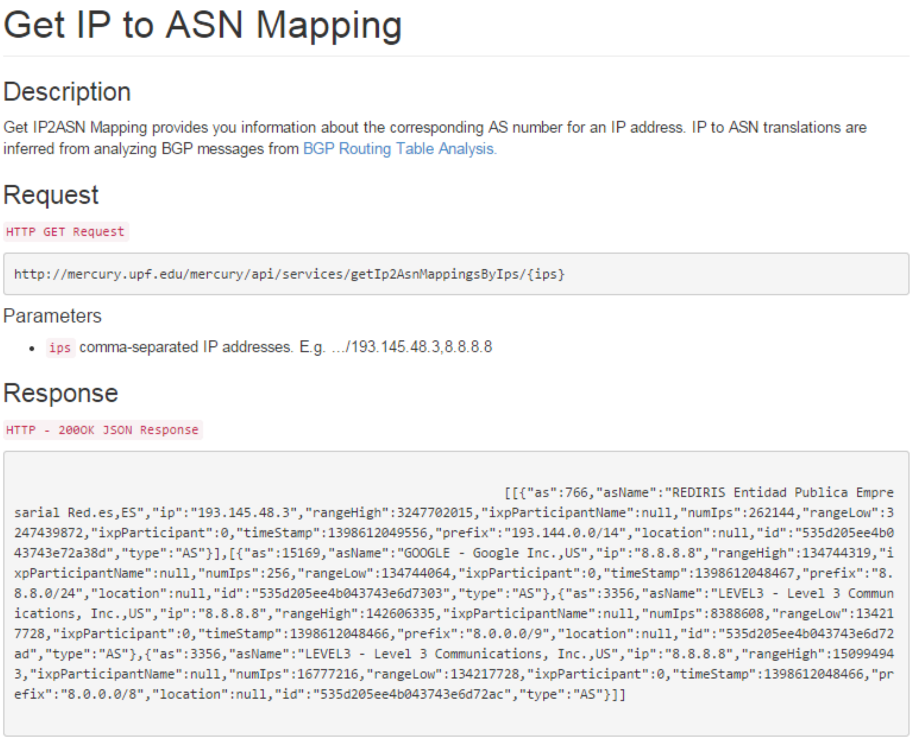

Get IP to ASN Mapping

Description

Get IP2ASN Mapping provides you information about the corresponding AS number for an IP address. IP to ASN translations are inferred from analyzing BGP messages from BGP Routing Table Analysis.

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/getIp2AsnMappingsByIps/{ips}
```

Parameters

- `ips` comma-separated IP addresses. E.g. .../193.145.48.3,8.8.8.8

Response

`HTTP - 200OK JSON Response`

```
                                                                  [[{"as":766,"asName":"REDIRIS Entidad Publica Empre
sarial Red.es,ES","ip":"193.145.48.3","rangeHigh":3247702015,"ixpParticipantName":null,"numIps":262144,"rangeLow":3
247439872,"ixpParticipant":0,"timeStamp":1398612049556,"prefix":"193.144.0.0/14","location":null,"id":"535d205ee4b0
43743e72a38d","type":"AS"}],[{"as":15169,"asName":"GOOGLE - Google Inc.,US","ip":"8.8.8.8","rangeHigh":134744319,"i
xpParticipantName":null,"numIps":256,"rangeLow":134744064,"ixpParticipant":0,"timeStamp":1398612048467,"prefix":"8.
8.8.0/24","location":null,"id":"535d205ee4b043743e6d7303","type":"AS"},{"as":3356,"asName":"LEVEL3 - Level 3 Commun
ications, Inc.,US","ip":"8.8.8.8","rangeHigh":142606335,"ixpParticipantName":null,"numIps":8388608,"rangeLow":13421
7728,"ixpParticipant":0,"timeStamp":1398612048466,"prefix":"8.0.0.0/9","location":null,"id":"535d205ee4b043743e6d72
ad","type":"AS"},{"as":3356,"asName":"LEVEL3 - Level 3 Communications, Inc.,US","ip":"8.8.8.8","rangeHigh":15099494
3,"ixpParticipantName":null,"numIps":16777216,"rangeLow":134217728,"ixpParticipant":0,"timeStamp":1398612048466,"pr
efix":"8.0.0.0/8","location":null,"id":"535d205ee4b043743e6d72ac","type":"AS"}]]
```

Figure A.17: Mercury API: Get IP to ASN Mapping.

- **Get IP to Geo Mapping:** Get IP2GEO Mapping provides you information about the geolocation for an IP address. IP to GEO translations are extracted from MaxMind GeoLite database (see Figure A.18).

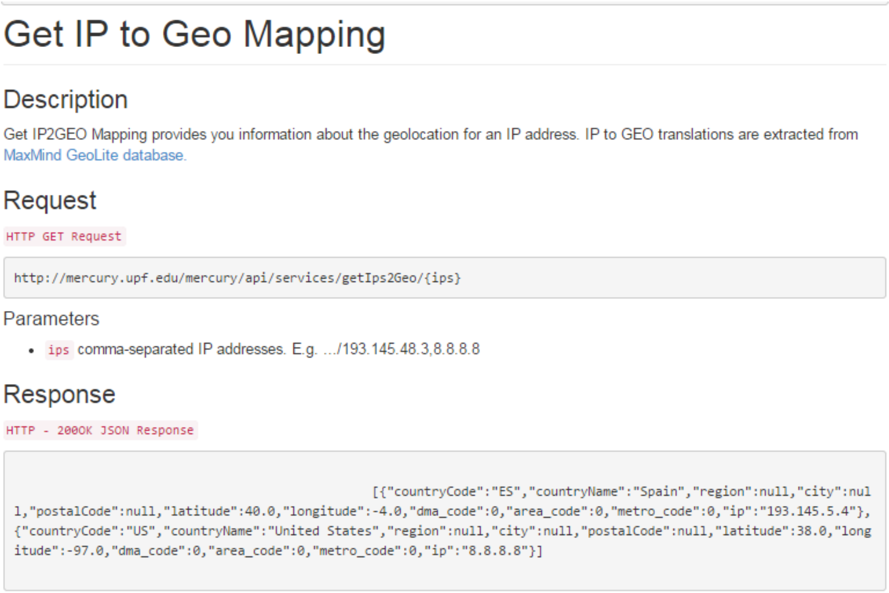

Get IP to Geo Mapping

Description

Get IP2GEO Mapping provides you information about the geolocation for an IP address. IP to GEO translations are extracted from MaxMind GeoLite database.

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/getIps2Geo/{ips}
```

Parameters

- `ips` comma-separated IP addresses. E.g. .../193.145.48.3,8.8.8.8

Response

`HTTP - 200OK JSON Response`

```
[{"countryCode":"ES","countryName":"Spain","region":null,"city":null,"postalCode":null,"latitude":40.0,"longitude":-4.0,"dma_code":0,"area_code":0,"metro_code":0,"ip":"193.145.5.4"},
{"countryCode":"US","countryName":"United States","region":null,"city":null,"postalCode":null,"latitude":38.0,"longitude":-97.0,"dma_code":0,"area_code":0,"metro_code":0,"ip":"8.8.8.8"}]
```

Figure A.18: Mercury API: Get IP to Geo Mapping.

- **Get IP to AS and Geo Mapping:** Get IP2ASNandGEO Mapping provides you information about the AS and the geolocation for an IP address. This method combines the previous two methods (see Figure A.19).

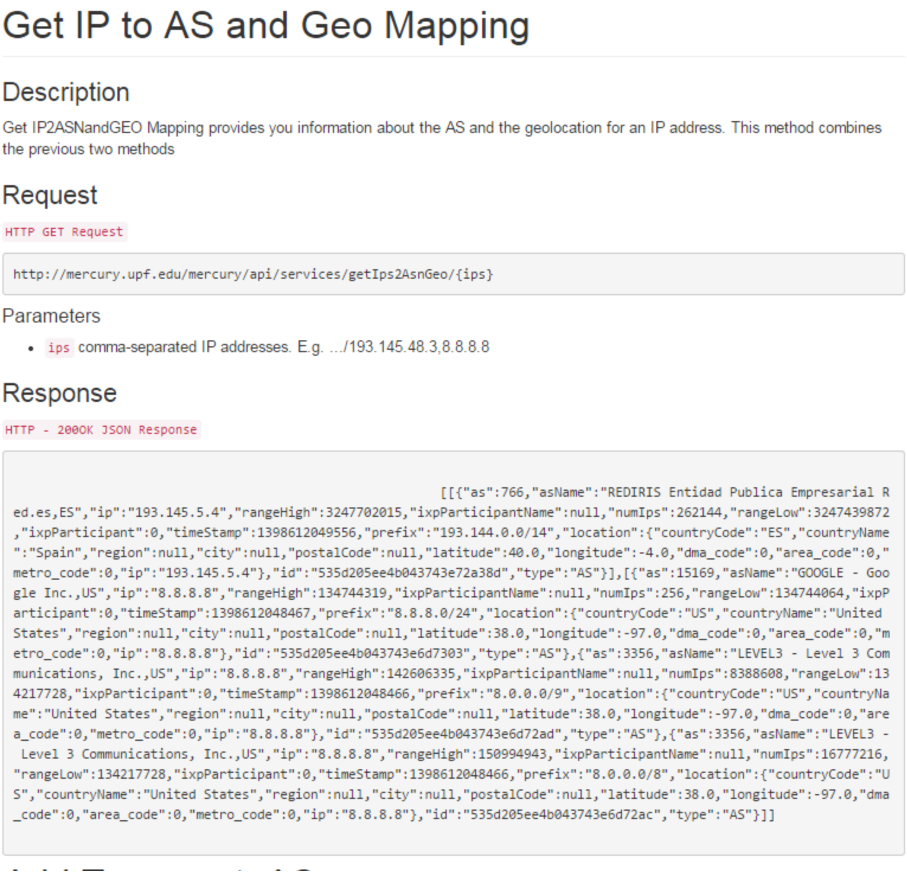

Get IP to AS and Geo Mapping

Description

Get IP2ASNandGEO Mapping provides you information about the AS and the geolocation for an IP address. This method combines the previous two methods

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/getIps2AsnGeo/{ips}
```

Parameters

- `ips` comma-separated IP addresses. E.g. .../193.145.48.3,8.8.8.8

Response

`HTTP - 200OK JSON Response`

```
                                               [[{"as":766,"asName":"REDIRIS Entidad Publica Empresarial R
ed.es,ES","ip":"193.145.5.4","rangeHigh":3247702015,"ixpParticipantName":null,"numIps":262144,"rangeLow":3247439872
,"ixpParticipant":0,"timeStamp":1398612049556,"prefix":"193.144.0.0/14","location":{"countryCode":"ES","countryName
":"Spain","region":null,"city":null,"postalCode":null,"latitude":40.0,"longitude":-4.0,"dma_code":0,"area_code":0,"
metro_code":0,"ip":"193.145.5.4"},"id":"535d205ee4b043743e72a38d","type":"AS"}],[{"as":15169,"asName":"GOOGLE - Goo
gle Inc.,US","ip":"8.8.8.8","rangeHigh":134744319,"ixpParticipantName":null,"numIps":256,"rangeLow":134744064,"ixpP
articipant":0,"timeStamp":1398612048467,"prefix":"8.8.8.0/24","location":{"countryCode":"US","countryName":"United
States","region":null,"city":null,"postalCode":null,"latitude":38.0,"longitude":-97.0,"dma_code":0,"area_code":0,"m
etro_code":0,"ip":"8.8.8.8"},"id":"535d205ee4b043743e6d7303","type":"AS"},{"as":3356,"asName":"LEVEL3 - Level 3 Com
munications, Inc.,US","ip":"8.8.8.8","rangeHigh":142606335,"ixpParticipantName":null,"numIps":8388608,"rangeLow":13
4217728,"ixpParticipant":0,"timeStamp":1398612048466,"prefix":"8.0.0.0/9","location":{"countryCode":"US","countryNa
me":"United States","region":null,"city":null,"postalCode":null,"latitude":38.0,"longitude":-97.0,"dma_code":0,"are
a_code":0,"metro_code":0,"ip":"8.8.8.8"},"id":"535d205ee4b043743e6d72ad","type":"AS"},{"as":3356,"asName":"LEVEL3 -
 Level 3 Communications, Inc.,US","ip":"8.8.8.8","rangeHigh":150994943,"ixpParticipantName":null,"numIps":16777216,
"rangeLow":134217728,"ixpParticipant":0,"timeStamp":1398612048466,"prefix":"8.0.0.0/8","location":{"countryCode":"U
S","countryName":"United States","region":null,"city":null,"postalCode":null,"latitude":38.0,"longitude":-97.0,"dma
_code":0,"area_code":0,"metro_code":0,"ip":"8.8.8.8"},"id":"535d205ee4b043743e6d72ac","type":"AS"}]]
```

Figure A.19: Mercury API: Get IP to Geo Mapping.

- **Add TracerouteAS:** Upload an TracerouteAS structure. This method is used by Mercury clients to add the measurements (see Figure A.20).

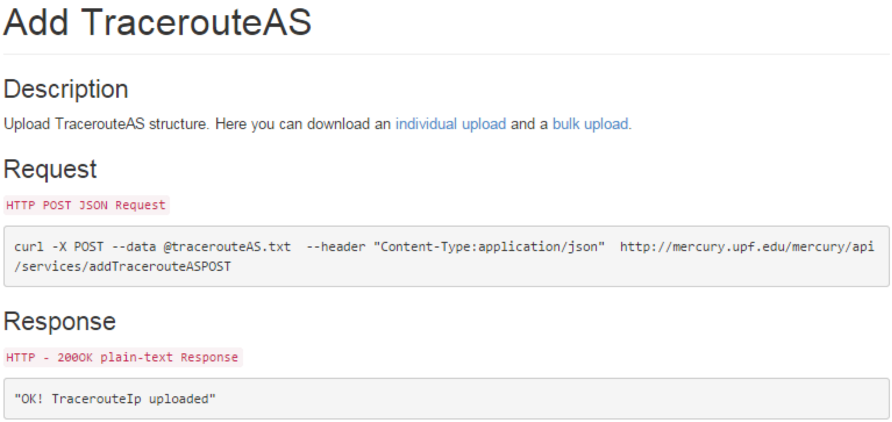

Add TracerouteAS

Description

Upload TracerouteAS structure. Here you can download an individual upload and a bulk upload.

Request

`HTTP POST JSON Request`

```
curl -X POST --data @tracerouteAS.txt  --header "Content-Type:application/json"  http://mercury.upf.edu/mercury/api
/services/addTracerouteASPOST
```

Response

`HTTP - 200OK plain-text Response`

```
"OK! TracerouteIp uploaded"
```

Figure A.20: Mercury API: Add TracerouteAS.

- **Get TracerouteASes:** Get TracerouteASes using custom Query using a POST request. In the request message you can query for the source/destination city, country, IP, AS and URL (e.g. ”dst”:”google.com”, ”srcCity”:null, ”srcCountry”:”Spain”, ”dstCity”:”Mountain View”, ”dstCountry”:”United States”, ”srcPublicIp”: ”46.25.151.142”, ”dstIp”: ”173.194.34.199”, ”srcAS”:12357, ”dstAS”:15169, ”srcASName”: ”COMUNITEL VODAFONE ESPANA S.A.U.,ES”, ”srcIp”:”192.168.0.2”, ”dstASName”: ”GOOGLE - Google Inc.,US”). See Figure A.21 for more details.

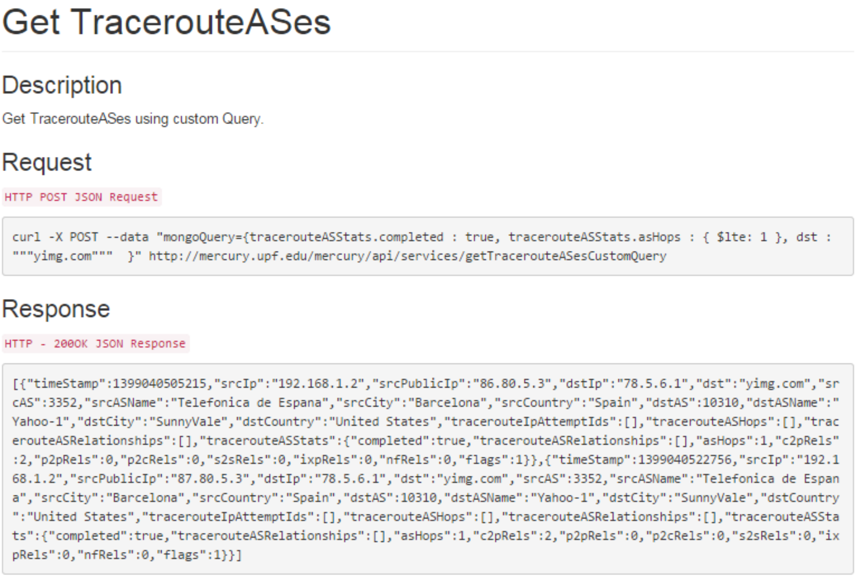

Get TracerouteASes

Description

Get TracerouteASes using custom Query.

Request

HTTP POST JSON Request

```
curl -X POST --data "mongoQuery={tracerouteASStats.completed : true, tracerouteASStats.asHops : { $lte: 1 }, dst :
"""yimg.com"""  }" http://mercury.upf.edu/mercury/api/services/getTracerouteASesCustomQuery
```

Response

HTTP - 200OK JSON Response

```
[{"timeStamp":1399040505215,"srcIp":"192.168.1.2","srcPublicIp":"86.80.5.3","dstIp":"78.5.6.1","dst":"yimg.com","sr
cAS":3352,"srcASName":"Telefonica de Espana","srcCity":"Barcelona","srcCountry":"Spain","dstAS":10310,"dstASName":"
Yahoo-1","dstCity":"SunnyVale","dstCountry":"United States","tracerouteIpAttemptIds":[],"tracerouteASHops":[],"trac
erouteASRelationships":[],"tracerouteASStats":{"completed":true,"tracerouteASRelationships":[],"asHops":1,"c2pRels"
:2,"p2pRels":0,"p2cRels":0,"s2sRels":0,"ixpRels":0,"nfRels":0,"flags":1}},{"timeStamp":1399040522756,"srcIp":"192.1
68.1.2","srcPublicIp":"87.80.5.3","dstIp":"78.5.6.1","dst":"yimg.com","srcAS":3352,"srcASName":"Telefonica de Espan
a","srcCity":"Barcelona","srcCountry":"Spain","dstAS":10310,"dstASName":"Yahoo-1","dstCity":"SunnyVale","dstCountry
":"United States","tracerouteIpAttemptIds":[],"tracerouteASHops":[],"tracerouteASRelationships":[],"tracerouteASSta
ts":{"completed":true,"tracerouteASRelationships":[],"asHops":1,"c2pRels":2,"p2pRels":0,"p2cRels":0,"s2sRels":0,"ix
pRels":0,"nfRels":0,"flags":1}}]
```

Figure A.21: Mercury API: Get TracerouteASes.

- **Get TracerouteASes by destination:** Get TracerouteASes filtering by destination domain (e.g. *upf.edu*). See Figure A.22 for more details.

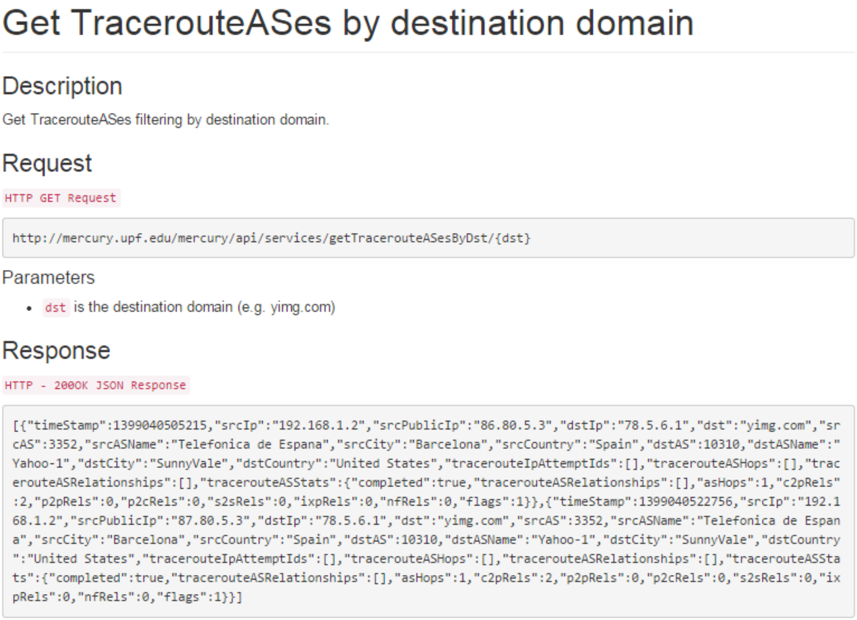

Get TracerouteASes by destination domain

Description

Get TracerouteASes filtering by destination domain.

Request

`HTTP GET Request`

```
http://mercury.upf.edu/mercury/api/services/getTracerouteASesByDst/{dst}
```

Parameters

- `dst` is the destination domain (e.g. yimg.com)

Response

`HTTP - 200OK JSON Response`

```
[{"timeStamp":1399040505215,"srcIp":"192.168.1.2","srcPublicIp":"86.80.5.3","dstIp":"78.5.6.1","dst":"yimg.com","sr
cAS":3352,"srcASName":"Telefonica de Espana","srcCity":"Barcelona","srcCountry":"Spain","dstAS":10310,"dstASName":"
Yahoo-1","dstCity":"SunnyVale","dstCountry":"United States","tracerouteIpAttemptIds":[],"tracerouteASHops":[],"trac
erouteASRelationships":[],"tracerouteASStats":{"completed":true,"tracerouteASRelationships":[],"asHops":1,"c2pRels"
:2,"p2pRels":0,"p2cRels":0,"s2sRels":0,"ixpRels":0,"nfRels":0,"flags":1}},{"timeStamp":1399040522756,"srcIp":"192.1
68.1.2","srcPublicIp":"87.80.5.3","dstIp":"78.5.6.1","dst":"yimg.com","srcAS":3352,"srcASName":"Telefonica de Espan
a","srcCity":"Barcelona","srcCountry":"Spain","dstAS":10310,"dstASName":"Yahoo-1","dstCity":"SunnyVale","dstCountry
":"United States","tracerouteIpAttemptIds":[],"tracerouteASHops":[],"tracerouteASRelationships":[],"tracerouteASSta
ts":{"completed":true,"tracerouteASRelationships":[],"asHops":1,"c2pRels":2,"p2pRels":0,"p2cRels":0,"s2sRels":0,"ix
pRels":0,"nfRels":0,"flags":1}}]
```

Figure A.22: Mercury API: Get TracerouteASes by destination.

Source code

Anyone can contribute to the development and improvement of the Mercury Central Server at `https://github.com/manuelpalacin/mercury2`.

Bibliography

[1] V. Adhikari, Y. Guo, F. Hao, M. Varvello, V. Hilt, M. Steiner, and Z.-L. Zhang. Unreeling netflix: Understanding and improving multi-CDN movie delivery. In *INFOCOM, 2012 Proceedings IEEE*, pages 1620–1628, March 2012.

[2] V. Adhikari, S. Jain, Y. Chen, and Z.-L. Zhang. Vivisecting YouTube: An active measurement study. In *INFOCOM, 2012 Proceedings IEEE*, pages 2521–2525, March 2012.

[3] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. *SIGCOMM Comput. Commun. Rev.*, 42(4):163–174, Aug. 2012.

[4] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Web Content Cartography. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '11, pages 585–600, New York, NY, USA, 2011. ACM.

[5] I. Akamai. Akamai and riverbed to accelerate applications over hybrid cloud networks. `http://www.akamai.com/html/about/press/releases/2011/press_051011.html`, 2011.

[6] I. Akamai. Akamai netsession interface. `http://www.akamai.com/client`, 2014.

[7] S. Akhshabi, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptation algorithms in adaptive streaming over http. In *Proceedings of the Second Annual ACM Conference on Multimedia Systems*, MMSys '11, pages 157–168, New York, NY, USA, 2011. ACM.

[8] Alexa. Alexa Top Sites. `http://www.alexa.com/topsites`.

[9] R. Archibald, Y. Liu, C. Corbett, and D. Ghosal. Disambiguating http: Classifying web applications. In *Wireless Communications and Mobile Computing Conference (IWCMC), 2011 7th International*, pages 1808–1813, July 2011.

[10] B. Augustin, T. Friedman, and R. Teixeira. Measuring Load-balanced Paths in the Internet. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, IMC '07, pages 149–160, New York, NY, USA, 2007. ACM.

[11] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped? In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 336–349. ACM, 2009.

[12] Augustin, Brice and Cuvellier, Xavier and Orgogozo, Benjamin and Viger, Fabien and Friedman, Timur and Latapy, Matthieu and Magnien, Clémence and Teixeira, Renata. Paris traceroute. `http://www.paris-traceroute.net/`.

[13] S. M. Besen. Advances in Routing Technologies and Internet Peering Agreements. *American Economic Review*, 91(2):292–296, 2001.

[14] A. Bianco, R. Birke, L. Giraudo, and M. Palacin. Openflow switching: Data plane performance. In *Communications (ICC), 2010 IEEE International Conference on*, pages 1–5, May 2010.

[15] N. Brownlee and K. Claffy. Understanding internet traffic streams: dragonflies and tortoises. *Communications Magazine, IEEE*, 40(10):110–117, Oct 2002.

[16] CAIDA. Archipelago Measurement Infrastructure. `http://www.caida.org/projects/ark/`.

[17] CAIDA. Skitter. `http://www.caida.org/tools/measurement/skitter/`, 1998.

[18] CAIDA. The CAIDA AS Relationships Dataset. `http://www.caida.org/data/active/as-relationships/`, July 2012.

[19] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 313–326, New York, NY, USA, 2013. ACM.

[20] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. Towards Capturing Representative AS-level Internet Topologies. In *Proceedings of ACM SIGMETRICS*, pages 280–281. ACM, 2002.

[21] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes From P2P Users. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, pages 217–228. ACM, 2009.

[22] Cisco Systems Inc. Cisco IOS NetFlow. White Paper. `http://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html`, 2007.

[23] Cisco Systems Inc. Cisco visual networking index: forecast and methodology, 2010-2014. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf, 2011.

[24] Claise, B. Specification of the IP flow information export (IPFIX) protocol for the exchange of IP traffic flow information. RFC 5101. https://tools.ietf.org/html/rfc5101, 2008.

[25] D. Clark. The design philosophy of the darpa internet protocols. *SIGCOMM Comput. Commun. Rev.*, 18(4):106–114, Aug. 1988.

[26] D. Clark, W. Lehr, and S. Bauer. Interconnection in the Internet: the Policy Challenge. In *Proceedings of the 39th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2011.

[27] L. Collins. In neutral [comms net neutrality]. *Engineering Technology*, 5(11):61–62, July 2010.

[28] Comision del Mercado de las Telecomunicaciones. Fixed broadband lines by technology and by operator in Spain. http://cmtdata.cmt.es/, 2013.

[29] F. C. Commission. Preserving the open internet; broadband industry practices. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf, 2010.

[30] T. N. Company. What Americans Do Online: Social Media And Games Dominate Activity.

[31] A. Craig Labovitz. How big is google? http://ddos.arbornetworks.com/2010/03/how-big-is-google, 2010.

[32] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey. Measuring ipv6 adoption. *SIGCOMM Comput. Commun. Rev.*, 44(4):87–98, Aug. 2014.

[33] Damas, et al. Extension-DNS (EDNS). http://tools.ietf.org/html/rfc6891.

[34] S. M. B. Dan Rayburn. Here's how the comcast and netflix deal is structured, with data and numbers. http://blog.streamingmedia.com/2014/02/heres-comcast-netflix-deal-structured-numbers.html, 2014.

[35] Daniel Stenberg and Jessica Ramos . cURL. http://curl.haxx.se/, 1997.

[36] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of the 6th International COnference*, Co-NEXT '10, pages 21:1–21:12, New York, NY, USA, 2010. ACM.

[37] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 1, pages 156–165 vol.1, March 2003.

[38] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley. AS Relationships: Inference and Validation. *SIGCOMM Comput. Commun. Rev.*, 37(1):29–40, Jan. 2007.

[39] X. Dimitropoulos, D. Krioukov, B. Huffaker, and G. Riley. Inferring AS relationships: Dead end or lively beginning. In *In Proceedings of 4th Workshop on Efficient and Experimental Algorithms (WEA '05*, pages 113–125, 2005.

[40] X. Dimitropoulos, D. Krioukov, G. Riley, et al. Revealing the Autonomous System Taxonomy: The Machine Learning Approach. In *Proceedings of Springer Passive and Active Measurement Conference (PAM)*. Springer, 2006.

[41] X. Dimitropoulos and G. Riley. Modeling Autonomous-System Relationships. In *Proceedings of the 20th Workshop on Principles of Advanced and Distributed Simulation (PADS)*, pages 143–149. IEEE Computer Society, 2006.

[42] European Internet Exchange Association. Euro-IX. `https://www.euro-ix.net/`.

[43] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. Complexity of Internet interconnections: Technology, incentives and implications for policy. In *in Proc. 35th Res. Conf. Commun., Inf. Internet Policy (TPRC*, 2007.

[44] G. E. for Network Innovations. Geni design principles. *Computer*, 39(9):102–105, Sept 2006.

[45] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, Dec. 2001.

[46] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of the 9th International Conference on Passive and Active Network Measurement*, PAM'08, pages 1–10, Berlin, Heidelberg, 2008. Springer-Verlag.

[47] Google Inc. SPDY: An experimental protocol for a faster web. `http://dev.chromium.org/spdy/spdy-whitepaper`, 2009.

[48] I. Grigorik. Making the web faster with http 2.0. *Commun. ACM*, 56(12):42–49, Dec. 2013.

[49] L. Guo, S. Chen, Z. Xiao, and X. Zhang. Analysis of multimedia workloads with implications for internet streaming. In *Proceedings of the 14th International Conference on World Wide Web*, WWW '05, pages 519–528, New York, NY, USA, 2005. ACM.

[50] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. A systematic framework for unearthing the missing links: Measurements and impact. In *Proceedings of 4th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 187–200. USENIX, 2007.

[51] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: A framework for discovering missing links in the internet topology. *Networking, IEEE/ACM Transactions on*, 17(2):391–404, April 2009.

[52] C. Huang, A. Wang, J. Li, and K. W. Ross. Measuring and Evaluating Large-scale CDNs. In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement*, IMC '08, pages 15–29, New York, NY, USA, 2008. ACM.

[53] C. Huang, A. Wang, J. Li, and K. W. Ross. Understanding Hybrid CDN-P2P: Why Limelight Needs Its Own Red Swoosh. In *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, NOSSDAV '08, pages 75–80, New York, NY, USA, 2008. ACM.

[54] Internet World Stats. Internet Growth Statistics. `http://www.internetworldstats.com/emarketing.htm`, 2011.

[55] Internet2 . Internet2. `http://www.internet2.edu/`, 2011.

[56] IRR. Internet Routing Registries. `http://www.irr.net/`.

[57] E. Jahn and J. Prüfer. Interconnection and Competition Among Asymmetric Networks in the Internet Backbone Market. Discussion Paper 2006-122, Tilburg University, Center for Economic Research, 2006.

[58] A. Jon Brodkin. Level 3 and cogent ask fcc for protection against isp tolls. `http://arstechnica.com/business/`

2014/03/level-3-and-cogent-ask-fcc-for-protection-against-isp-tolls/, 2014.

[59] A. Jon Brodkin. Netflix is paying comcast for direct connection to network. http://arstechnica.com/business/2014/02/netflix-is-paying-comcast-for-direct-connection-to-network-wsj-reports/, 2014.

[60] A. Jon Brodkin. Netflix may have gained direct connection to comcast network. http://arstechnica.com/information-technology/2014/02/netflix-may-have-gained-direct-connection-to-comcast-network/, 2014.

[61] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, and M. Faloutsos. Is p2p dying or just hiding? [p2p traffic measurement]. In *Global Telecommunications Conference, 2004. GLOBECOM '04. IEEE*, volume 3, pages 1532–1538 Vol.3, Nov 2004.

[62] R. Krishnan, H. V. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao. Moving Beyond End-to-end Path Information to Optimize CDN Performance. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '09, pages 190–201, New York, NY, USA, 2009. ACM.

[63] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. *SIGCOMM Comput. Commun. Rev.*, 41(4):–, Aug. 2010.

[64] S. Lukasik. Why the arpanet was built. *Annals of the History of Computing, IEEE*, 33(3):4–21, March 2011.

[65] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*, OSDI '06, pages 367–380, Berkeley, CA, USA, 2006. USENIX Association.

[66] Z. M. Mao, C. D. Cranor, F. Douglis, M. Rabinovich, O. Spatscheck, and J. Wang. A Precise and Efficient Evaluation of the Proximity Between Web Clients and Their Local DNS Servers. In *Proceedings of the General Track of the Annual Conference on USENIX Annual Technical Conference*, ATEC '02, pages 229–242, Berkeley, CA, USA, 2002. USENIX Association.

[67] Z. M. Mao, D. Johnson, J. Rexford, J. Wang, and R. Katz. Scalable and Accurate Identification of AS-level Forwarding Paths. In *Proceedings of the 23rd International Conference on Computer Communications (INFOCOM)*. IEEE Communications Society, 2004.

[68] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an Accurate AS-level Traceroute Tool. In *Proceedings of ACM SIGCOMM*, 2003.

[69] MaxMind, Inc. MaxMind GeoLite Databases. `http://dev.maxmind.com/geoip/legacy/geolite`.

[70] C. Metz. Ip over 2000: where have we been and where are we going? *Internet Computing, IEEE*, 4(1):83–87, Jan 2000.

[71] D. Meyer. University of Oregon Route Views Project. `http://www.routeviews.org/`.

[72] S. Murugesan. Understanding web 2.0. *IT Professional*, 9(4):34–41, July 2007.

[73] B. of European Regulators for Electronic Communications. Berec web site. `http://berec.europa.eu/`, 2010.

[74] U. of Oregon. Route Views Project. `http://archive.routeviews.org/bgpdata/`.

[75] U. of Washington. iplane. `http://iplane.cs.washington.edu/`, 2006.

[76] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. In Search of the Elusive Ground Truth: the Internet's AS-level Connectivity Structure. In *Proceedings of ACM SIGMETRICS Performance Evaluation Review*, volume 36, pages 217–228. ACM, 2008.

[77] R. V. Oliveira, B. Zhang, and L. Zhang. Observing the Evolution of Internet AS Topology. *ACM SIGCOMM Computer Communication Review*, 37(4):313–324, 2007.

[78] Packet Clearing House. PCH. `https://www.pch.net/`.

[79] M. Palacin, A. Bikfalvi, and M. Oliver. Mercury: Revealing Hidden Interconnections Between Access ISPs and Content Providers. In *Proceedings of the 20th EUNICE Conference*, number LNCS number 8846 in 20th EUNICE Conference. Springer's Lecture Notes in Computer Science, 2014.

[80] D. Parish, K. Bharadia, A. Larkum, I. Phillips, and M. Oliver. Using packet size distributions to identify real-time networked applications. *Communications, IEE Proceedings-*, 150(4):221–227, Aug 2003.

[81] PeeringDB. PeeringDB. `https://www.peeringdb.com/`.

[82] Philip Smith, Cisco Systems. BGP Routing Table Analysis. `http://thyme.apnic.net/`.

[83] Planet Lab. Planet Lab. `http://www.planet-lab.org/`.

[84] I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. Improving content delivery using provider-aided distance information. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC '10, pages 22–34, New York, NY, USA, 2010. ACM.

[85] I. Poese, B. Frank, B. Ager, G. Smaragdakis, S. Uhlig, and A. Feldmann. Improving content delivery with padis. *Internet Computing, IEEE*, 16(3):46–52, May 2012.

[86] Princeton University. CoDeen: A Content Distribution Network for PlanetLab. http://codeen.cs.princeton.edu/.

[87] D. Rayburn. Edgecast says cdn federation taking hold, details how operators are exchanging traffic. http://blog.streamingmedia.com/2012/05/edgecast-says-cdn-federation-taking-hold-details-how-operators-are-exchanging-traffic.html, 2014.

[88] J. Rexford and C. Dovrolis. Future internet architecture: Clean-slate versus evolutionary research. *Commun. ACM*, 53(9):36–40, Sept. 2010.

[89] Réseaux IP Européens Network Coordination Centre (RIPE NCC). RIPE Atlas. https://atlas.ripe.net/, 2010.

[90] Sandvine Inc. Global Internet Phenomena Report. Fall 2011. https://www.sandvine.com/trends/global-internet-phenomena/, 2011.

[91] Sandvine Inc. Global Internet Phenomena Report. Fall 2013. https://www.sandvine.com/trends/global-internet-phenomena/, 2013.

[92] Y. Shavitt and E. Shir. DIMES: Let the Internet Measure Itself. *ACM SIGCOMM Computer Communication Review*, 35(5):71–74, 2005.

[93] Y. Shavitt, E. Shir, and U. Weinsberg. Near-deterministic Inference of AS Relationships. In *Proceedings of the 10th International Conference on Telecommunications (ConTEL)*, pages 191–198. IEEE, 2009.

[94] Y. Shavitt and U. Weinsberg. Topological Trends of Internet Content Providers. In *Proceedings of the Fourth Annual Workshop on Simplifying Complex Networks for Practitioners*, SIMPLEX '12, pages 13–18, New York, NY, USA, 2012. ACM.

[95] Y. Shavitt and N. Zilberman. A Geolocation Databases Study. *IEEE Journal on Selected Areas in Communications*, 29(10):2044–2056, 2011.

[96] G. Shrimali and S. Kumar. Paid Peering Among Internet Service Providers. In *Proceeding from the 2006 Workshop on Game Theory for Communications and Networks*, GameNets '06, New York, NY, USA, 2006. ACM.

[97] G. Siganos, M. Faloutsos, and C. Faloutsos. The evolution of the internet:topology and routing.

[98] K. Sripanidkulchai, B. Maggs, and H. Zhang. An analysis of live streaming workloads on the internet. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*, IMC '04, pages 41–54, New York, NY, USA, 2004. ACM.

[99] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante. Drafting Behind Akamai: Inferring Network Conditions Based on CDN Redirections. *IEEE/ACM Trans. Netw.*, 17(6):1752–1765, Dec. 2009.

[100] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 618–627 vol.2, 2002.

[101] P. Szegedi, S. Figuerola, M. Campanella, V. Maglaris, and C. Cervello-Pastor. With evolution for revolution: managing federica for future internet research. *Communications Magazine, IEEE*, 47(7):34–39, July 2009.

[102] P. Szegedi, J. Riera, J. Garcia-Espin, M. Hidell, P. Sjodin, P. Soderman, M. Ruffini, D. O'Mahony, A. Bianco, L. Giraudo, M. De Leon, G. Power, C. Cervello-Pastor, V. Lopez, and S. Naegele-Jackson. Enabling future internet research: the federica case. *Communications Magazine, IEEE*, 49(7):54–61, July 2011.

[103] Team Cymru. IP-to-ASN Service. http://www.team-cymru.org/Services/ip-to-asn.html.

[104] Tel Aviv University. NetDIMES. http://www.netdimes.org/, 2009.

[105] C. D. N. C. F. H. S. C. W. the Battle for Content-Hungry Consumers. Scott puopolo, marc latouche, françois le faucheur, and jaak defour. https://www.cisco.com/web/about/ac79/docs/sp/CDN-PoV_IBSG.pdf, 2011.

[106] The Measurement Lab (MLab). The Measurement Lab. http://www.measurementlab.net/.

[107] Tim O'Reilly. What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. http://www.oreilly.com/pub/a/web2/archive/what-is-web-20.html, 2004.

[108] Tony Bates, Philip Smith, Geoff Huston. CIDR Report. http://www.cidr-report.org/.

[109] J. Turner. A proposed architecture for the geni backbone platform. In *In Proc. Architecture for Networking and Communications Systems*, pages 12–2006, 2006.

[110] N. University. Ono. http://www.aqualab.cs.northwestern.edu/projects/129-ono-developer-instructions, 2007.

[111] F. Wamser, R. Pries, D. Staehle, K. Heck, and P. Tran-Gia. Traffic characterization of a residential wireless internet access. *Telecommunication Systems*, 48(1-2):5–17, 2011.

[112] William B. Norton. The 2014 Internet Peering Playbook. `http://drpeering.net/FAQ/U.S.vs.European-Peering-Models.php`, 2014.

[113] D. William B. Norton. Internet transit prices: Historical and projected. `http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php`, 2010.

[114] J. Wulf, R. Zarnekow, T. Hau, and W. Brenner. Carrier activities in the cdn market - an exploratory analysis and strategic implications. In *Intelligence in Next Generation Networks (ICIN), 2010 14th International Conference on*, pages 1–6, Oct 2010.

[115] J. Xia and L. Gao. On the evaluation of AS relationship inferences. In *Global Telecommunications Conference, 2004. GLOBECOM '04. IEEE*, volume 3, pages 1373–1377 Vol.3, Nov 2004.

List of acronyms

.NET: is a software framework developed by Microsoft that runs primarily on Microsoft Windows.
3G: is the third generation of mobile telecommunications technology.
3.5G: is a grouping of disparate mobile telephony and data technologies designed to provide better performance than 3G.
4G: is the fourth generation of mobile telecommunications technology.
ACK: a signal, message or packet to acknowledge the receipt of data.
ADSL: asymmetric digital subscriber line is a type of DSL technology.
ADSL2: is an improved extension of ADSL with download rates up to 12 Mbps.
ADSL2+: is an improved extension of ADSL2 with download rates up to 24 Mbps.
AJAX: is a group of interrelated Web development techniques used on the client-side to create asynchronous Web applications.
API: application programming interface is a set of routines, protocols, and tools for building software applications.
ARP: is a telecommunication protocol used for resolution of network layer addresses into link layer addresses.
ARPANET: the Advanced Research Projects Agency Network (ARPANET) was an early packet switching network and the first network to implement the protocol suite TCP/IP.
AS: Autonomous System is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators on behalf of a single administrative entity or domain that presents a common, clearly defined routing policy to the Internet.
ASN: autonomous system number.
BGP: Border Gateway Protocol (BGP) is a standardised exterior gateway protocol designed to exchange routing and reachability information between autonomous systems (AS) on the Internet.
CDN: Content delivery network or content distribution network is a system of computers on the Internet that delivers content transparently to end users.

CERN: The European Organisation for Nuclear Research is a research organisation that operates the largest particle physics laboratory in the world and also the place the World Wide Web was first implemented.

CNAME: a Canonical Name record is a type of resource record in the Domain Name System (DNS) used to specify that a domain name is an alias for another domain.

CP: content provider

CSS: Cascading Style Sheets is a style sheet language used for describing the look and formatting of a document written in a markup language.

CSS3: CSS version 3 introduces new modules to add new capabilities or to extend features defined in previous versions, preserving backward compatibility.

CWDM: Coarse or conventional wavelength division multiplexing in contrast to conventional WDM and DWDM uses increased channel spacing to allow less sophisticated and thus cheaper transceiver designs.

DARPA: The Defense Advanced Research Projects Agency is an agency of the U.S. Department of Defense responsible for the development of emerging technologies for use by the military.

DIFFSERV: Differentitated Services is a model for providing QoS in the Internet by differentiating the traffic.

DOCSIS: Data Over Cable Service Interface Specification is an international telecommunications standard that permits the addition of high-bandwidth data transfer to an existing cable TV system.

DOM: The Document Object Model is a cross-platform and language-independent convention for representing and interacting with objects in HTML, XHTML, and XML documents.

DSL: digital subscriber line is a family of wide-area technologies that are used to transmit digital data over telephone lines.

DWDM: dense wavelength-division multiplexing refers originally to optical signals multiplexed within the 1550 nm band using different wavelengths of laser light but with denser channel spacing than WDM or CWDM.

FLASH: is a multimedia and software platform used for creating vector graphics, animation, browser games, rich internet applications, desktop applications, mobile applications and mobile games.

FLEX: is a software development kit (SDK) for the development and deployment of cross-platform rich Internet applications based on the Adobe Flash platform.
FTP: File Transfer Protocol (FTP) is a standard network protocol used to transfer computer files from one host to another host over a TCP-based network, such as the Internet.
FTTH: Fibre to the x (FTTX) is a generic term for any broadband network architecture using optical fibre to provide all or part of the local loop used for last mile telecommunications.
Gbps: Gigabit per second
GET: Request method of the HTTP protocol. Requests a representation of the specified resource.
GPRS: General packet radio service is a packet oriented mobile data service on the 2G and 3G cellular communication system's GSM.
HFC: Hybrid fibre-coaxial (HFC) is a telecommunications industry term for a broadband network that combines optical fibre and coaxial cable.
HSDPA: High-Speed Downlink Packet Access is an enhanced 3G mobile-telephony communications protocol which allows networks based on Universal Mobile Telecommunications System (UMTS) to have higher data speeds and capacity.
HSDPA+: Evolved High-Speed Packet Access, is a technical standard for wireless, broadband telecommunication with higher speeds than HSDPA that are comparable to the newer LTE networks.
HTML: HyperText Markup Language is the standard markup language used to create web pages.
HTML5: HTML5 is a core technology markup language of the Internet used for structuring and presenting content for the World Wide Web.
HTTP: Hypertext Transfer Protocol is an application protocol for distributed, collaborative, hypermedia information systems.
HTTP/2: is the second major version of the HTTP network protocol used by the World Wide Web.
INTSERV: Integrated services is a model for providing QoS in networks by building a virtual circuit in the Internet using the bandwidth reservation technique.

IP: Internet protocol.
IPFIX: Internet Protocol Flow Information Export standard defines how IP flow information is to be formatted and transferred from an exporter to a collector.
IPsec: Internet Protocol Security is a protocol suite for securing Internet Protocol communications by authenticating and encrypting each IP packet of a communication session.
IPTV: TV over IP. Service offered by some access ISPs.
ISP: Internet Service Provider.
IXP: Internet eXchange Point.
JS: JavaScript web programming language.
JSON: JavaScript Object Notation, is an open standard format that uses human-readable text to transmit data objects consisting of attribute-value pairs.
JSP: JavaServer Pages is a technology that helps software developers create dynamically generated web pages based on HTML, XML, or other document types.
LAN: Local area network.
LSA: Level service agreement.
LTE: Long term evolution is a standard for wireless communication of high-speed data for mobile phones and data terminals.
Mbps: Megabit per second.
MDC: is a coding technique that fragments a single media stream into N substreams referred to as descriptions.
MPEG: is a motion standard for ”the generic coding of moving pictures and associated audio information”.
MPLS: Multiprotocol Label Switching is a mechanism in high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table.
MTU: the maximum transmission unit of a communications protocol of a layer is the size (in bytes) of the largest protocol data unit that the layer can pass onwards.
OS: Operating system.

P2P: Peer-to-peer is a computing or networking distributed application architecture that partitions tasks or workloads among peers

PERL: is a family of high-level, general-purpose, interpreted, dynamic programming languages.

PHP: is a server-side scripting language created in 1995 and designed for web development but also used as a general-purpose programming language.

PON: A passive optical network is a telecommunications network that uses point-to-multipoint fibre to the premises in which unpowered optical splitters are used to enable a single optical fibre to serve multiple premises.

POST: method of the HTTP protocol designed to request that a web server accepts the data enclosed in the request message's body for storage.

QoE: Quality of experience is a measure of a customer's experiences with a service.

QoS: Quality of service is the overall performance of a telephony or computer network, particularly the performance seen by the users of the network. In computer network trafficking refers to resource reservation control mechanisms.

RARP: The Reverse Address Resolution Protocol is an obsolete computer networking protocol used by a client computer to request its Internet Protocol (IPv4) address from a computer network, when all it has available is its Link Layer or hardware address, such as a MAC address.

REST: Representational State Transfer is a software architecture style consisting of guidelines and best practices for creating scalable web services.

RFC: A Request for Comments is a publication of the Internet Engineering Task Force (IETF) and the Internet Society, the principal technical development and standards-setting bodies for the Internet.

RSH: the remote shell is a command line computer program that can execute shell commands as another user, and on another computer across a computer network.

RSS: Rich Site Summary or Really Simple Syndication, uses a family of standard web feed formats to publish frequently updated information: blog entries, news headlines, audio, video.

RTMP: Real Time Messaging Protocol is a TCP-based protocol which maintains persistent connections and allows low-latency communications.

RTP: Real-time Transport Protocol is a network protocol for delivering audio and video over IP networks.

RTSP: Real Time Streaming Protocol is a network control protocol designed for use in entertainment and communications systems to control streaming media servers.

RTT: Round-trip time.

SDN: Software-defined networking is an approach to computer networking that allows network administrators to manage network services through abstraction of lower-level functionality.

SMTP: Simple Mail Transfer Protocol is an Internet standard for electronic mail (e-mail) transmission.

SPDY: is an open networking protocol developed primarily at Google for transporting web content. SPDY manipulates HTTP traffic, with particular goals of reducing web page load latency and improving web security.

SSH: Secure Shell is a cryptographic network protocol for initiating text-based shell sessions on remote machines in a secure way.

SSL: Secure Sockets Layer is a set of cryptographic protocols designed to provide communications security over a computer network.

SVC: Scalable Video Coding standardises the encoding of a high-quality video bitstream that also contains one or more subset bitstreams.

TCP: Transmission Control Protocol (TCP) is a core protocol of the Internet Protocol Suite. TCP provides reliable, ordered, and error-checked delivery of a stream of octets between applications running on hosts communicating over an IP network.

TELNET: is an application protocol used on the Internet or local area networks to provide a bidirectional interactive text-oriented communication facility using a virtual terminal connection.

UDP: User Datagram Protocol is one of the core members of the Internet protocol suite. UDP uses a simple connectionless transmission model with a minimum of protocol mechanism.

UMTS: The Universal Mobile Telecommunications System is a third generation mobile cellular system for networks based on the GSM standard.

URI: in computing, a uniform resource identifier is a string of characters used to identify a name of a resource.

URL: a uniform resource locator is a reference to a resource that specifies the location of the resource on a computer network and a mechanism for retrieving it.

VDSL: very-high-bit-rate digital subscriber line is a technology providing data transmission faster than ADSL over a single flat untwisted or twisted pair of copper wires up to 52 Mbit/s.

VDSL2: is an enhancement to very-high-bit-rate digital subscriber line (VDSL), and most advanced currently deployed standard of digital subscriber line (DSL) broadband wireline communications. VDSL2 is designed to support the wide deployment of triple play services such as voice, video, data and high-definition television.

VoIP: is a methodology and group of technologies for the delivery of voice communications and multimedia sessions over Internet Protocol (IP) networks, such as the Internet.

WDM: wavelength-division multiplexing is a technology which multiplexes a number of optical carrier signals onto a single optical fibre by using different wavelengths of laser light

WEB2.0: Web 2.0 describes World Wide Web sites that emphasise user-generated content, usability, and interoperability.

WEBRTC: Web Real-Time Communication is an API definition drafted by the World Wide Web Consortium (W3C) that supports browser-to-browser applications for voice calling, video chat, and P2P file sharing without the need of either internal or external plugins.

WS: a Web service is a method of communication between two electronic devices over a network. It is a software function provided at a network address over the Web with the service always on as in the concept of utility computing.

WWW: the World Wide Web is an information system of interlinked hypertext documents and other digital resources that are accessed via the

Internet.
xDSL: set of technologies related to the DSL family.
XML: Extensible Markup Language is a markup language that defines a set of rules for encoding documents in a format which is both human-readable and machine-readable.
xPON: set of technologies related to the PON family