UNIVERSIDAD AUTÓNOMA DE MADRID

ESCUELA POLITÉCNICA SUPERIOR

Ph.D. Thesis

Characterizing the spatial and temporal diversity of the Internet traffic:

A capacity planning application to the RedIRIS network

Author:
José Luis García Dorado

Supervisors:
Prof. Javier Aracil Rico
Dr. Jorge E. López de Vergara Méndez

Madrid, March 2010

DOCTORAL THESIS:	Characterizing the spatial and temporal diversity of Internet traffic: A capacity planning application to the RedIRIS network
AUTHOR:	José Luis García Dorado
SUPERVISORS:	Prof. Javier Aracil Rico Dr. Jorge E. López de Vergara Méndez

The committee for the defense of this doctoral thesis is composed by:

PRESIDENT:	Prof. Enrique Vázquez Gallo
MEMBERS:	Dr. Carlos García Rubio
	Dr. Marco Mellia
	Dr. Daniel Morató Osés
SECRETARY:	Dr. Sergio López Buedo

TO MY PARENTS

Abstract

Both Internet service providers and the research community have understood the importance of traffic measurements. On the one hand, the research community uses these measurements to characterize the network traffic in order to improve the understanding of the Internet dynamics. On the other hand, operators use them to ensure the Quality of Service, detect problems in their networks, invoice clients, and perform capacity planning, among other applications. In the development of all these tasks, network operators must handle representative information about the traffic volumes traversing their network. This entails the following issues:

First, the representativeness and generality of the traffic measurements still remains an open issue. It has been a general belief that certain internetwork traffic statistics show a similar behavior in networks with similar features, and the conclusions derived from the measurements of a given network can be extrapolated to a similar scenario. This work makes no starting assumption concerning this issue and undertakes a "spatial" analysis of network measurements. Our experiments, using measurements from the Spanish academic network, RedIRIS, show that although the frequency statistics of IP addresses and port numbers follow a Zipf distribution (as expected) the distributions' characteristic parameter values vary significantly in a spatial dimension, that is, across the individual networks.

Moreover, we model the distribution of the geolocation of Internet connections in an extensive set of campus networks using a Zipf-Mandelbrot law, and we further apply Analysis of Variance (ANOVA) to the measurements. Then, by comparing the distributions' characteristic parameter values and analyzing ANOVA results, we found that again the behavior between such campuses is far from being homogeneous.

In practical terms, this means that network designers, analysts, and operators should not assume that the statistics for applications usage for one network may accurately characterize other networks, even when those networks have similar user bases and environments. Furthermore, we show that experiment durations of more than 30 days are necessary for the traffic processes to display stationarity.

Second, the amount of traffic measurements available over which to perform any analysis, processing and storage is usually overwhelming. For this reason, the research community has paid attention to find an effective mechanism to reduce (or subsample) such huge amount of data, with minimum loss of information. In this light, this thesis proposes a new method that enables to subsample network measurements optimally using Multiresolution Analysis with wavelets: The "queueing equivalent" thresholding method.

Additionally, this thesis addresses the characterization of the Internet's traffic busy hour given that this metric is usually adopted by the operators for capacity planning. The results of analyzing six months worth of Netflow data and MRTG logs collected from RedIRIS, show that the traffic volume marginal distribution during the busy hour is Gaussian, and further show that there is no correlation in the busy-hour traffic volume over different working days. This implies that the traffic volumes in the busy hour can be modeled by a white Gaussian process, which can be used to derive the capacity required such that the traffic volume is met with certain probability.

Finally, we go one step further and examine the influence of the networks' intrinsic features, mainly population size and access link capacity, on the Internet busy-hour traffic. Well-known statistical methodologies, such as ANOVA and Analysis of Covariance, show that the network size in terms of number of users justifies most of the busy-hour traffic information. We further provide a linear-regression model that adjusts the amount of traffic that each network user contributes to the busy-hour traffic mean and variance values, with a straightforward application to the problem of link capacity planning of IP networks.

Keywords: Internet Measurements; University Networks; RedIRIS; Network capacity planning; Netflow; Spatial and Temporal Diversity; Heavy-hitters; Time-Series subsampling; Queueing equivalent thresholding method; Multiresolution

Analysis with wavelets; Internet traffic busy hour; ANOVA; ANCOVA; Intrinsic features; Zipf; Zipf Mandelbrot; Geolocation .

Resumen

Tanto los proveedores de servicio de Internet como la comunidad científica han comprendido la importancia de las medidas de tráfico de las redes de datos. Por un lado, la comunidad científica usa estas medidas para caracterizar Internet contribuyendo al mejor conocimiento de su dinámica. Por otro lado, los operadores de red las usan para asegurar la calidad de servicio de sus clientes, detectar problemas, facturar a los clientes y para dimensionar correctamente sus redes, entre otras aplicaciones. Al llevar a la práctica estas tareas, los operadores de red deben capturar información representativa del volumen de tráfico que atraviesa sus redes. Esto conlleva las siguientes complicaciones:

Primero, la representatividad y generalidad de las medidas de red es un problema por resolver. Típicamente se ha considerado que ciertas estadísticas del uso de Internet muestran un comportamiento equivalente entre redes con características similares, y que las conclusiones derivadas de estas medidas podían ser extrapoladas a escenarios semejantes. Este trabajo no asume esta premisa y analiza la diversidad "espacial" de las medidas de red. Nuestros experimentos, usando medidas de la red académica española, RedIRIS, muestran que, aunque la frecuencia de uso de las direcciones IP y puertos siguen una distribución Zipf (como era esperado), el parámetro característico de esta distribución varía significativamente en la dimensión espacial, esto es, entre un grupo significativo de redes aun cuando la población y características de estas redes son similares.

Del mismo modo, modelamos la distribución de probabilidad de la localización geográfica a donde se conectan un conjunto de redes universitarias mediante una distribución Zipf-Mandelbrot y, posteriormente, aplicamos Análisis de Varianza (ADEVA) a los datos. Entonces, al comparar las parámetros de este modelo y analizar los resultados de ADEVA, encontramos que, de nuevo, el comportamiento de las redes universitarias está lejos de ser homogéneo.

En la práctica, esto significa que los diseñadores de red, analistas y operadores no deberían asumir que las estadísticas de uso de aplicaciones de una red en particular pueden caracterizar otras redes aunque tengan similar población e infraestructura. Además se muestra que la duración de los experimentos debe ser como mínimo de 30 días para que ciertas estadísticas muestren un comportamiento estacionario.

Segundo, la cantidad de medidas de tráfico disponibles sobre las que desarrollar cualquier análisis, procesado o almacenamiento es frecuentemente desbordante. Por esta razón, la comunidad científica ha prestado especial atención a la búsqueda de mecanismos para reducir o submuestrear tal cantidad de datos, minimizando la perdida de información. Este trabajo propone un nuevo método que permite submuestrear medidas de red óptimamente usando Análisis Multiresolución con wavelets: "The queueing equivalent thresholding method".

Esta tesis también analiza la caracterización de la hora cargada de Internet debido a que esta medida es típicamente usada por los operadores para dimensionar los enlaces de sus redes. Los resultados de analizar seis meses de datos Netflow y MRTG de RedIRIS, muestran que la distribución de tráfico, en cuanto a volumen, durante la hora cargada es Gaussiana, y que, además, no existe correlación entre los diferentes días laborables de la semana. Esto implica que el tráfico puede ser modelado como un proceso blanco Gaussiano, el cual puede ser usado para estimar la capacidad requerida tal que el volumen de tráfico generado por los usuarios no exceda con cierta probabilidad esta capacidad.

Finalmente, examinamos la influencia de las características inherentes de las redes, principalmente la población y la capacidad de los accesos, en el tráfico de Internet durante la hora cargada. Metodologías bien conocidas como ADEVA y el Analisis de Covarianza, muestran que el tamaño de la red en términos de usuarios explica convenientemente la dinámica de la hora cargada. Con esto, proponemos un modelo de regresión lineal que predice el volumen de tráfico que genera una red durante la hora cargada. Esta propuesta tiene aplicación directa al problema de asignación de capacidad a las redes IP.

Palabras clave: Medidas de Internet; Redes universitarias; RedIRIS; Dimensionado de redes; Netflow; Diversidad temporal y espacial; Heavy-hitters; Submuestreo de series temporales; “Queueing equivalent thresholding method”; Análisis Multiresolución con wavelets; Hora cargada de Internet; ADEVA; ADECOV; Características inherentes; Zipf; Zipf Mandelbrot; Geolocalización.

Acknowledgments

I would like to deeply thank my supervisor, Javier Aracil, and co-supervisor, Jorge E. López de Vergara, for their constant support, inestimable contributions and patient supervisions in the development of this thesis. Thank you Javier for giving me the opportunity to become a researcher.

My sincere thanks to José Alberto Hernández for his guidance and contributions to this thesis.

I owe my deepest gratitude to Felipe Mata for his priceless participation on the geolocation characterization of the traffic measurements carried out in Chapter 4, as well as for his support on the statistical analysis of the traffic measurements.

I would also like to thank Víctor López for our fruitful discussions about all the topics of this thesis.

In the same way, let me thank my colleagues from the High Performance Computing and Networking research group: Sergio López, Iván González, Francisco Gómez, Luis de Pedro, Bas Huiszoon, Gustavo Sutter, Sergio Arroyo, Jaime Garnica, Pilar Holgado, Francisco Javier Ramos, Germán Retamosa, Alfredo Salvador, Pedro Maria Santiago, Víctor Lucas, Diana Carolina Araujo, Luis Miguel Gómez, Álvaro Jiménez-Peralo, Álvaro Katsu, Luis Moreno, Víctor Moreno, and Gonzalo Polo. Thank you for all this time at the laboratory and for making the lab. a nice place to work. This also definitely includes my other colleagues from the laboratory José Luis Añamuro, Walter Fuertes, Elías Todorovich, Juan González, and Eduardo Boemo. All my work would not have been possible without the support of the Universidad Autónoma de Madrid and the Departamento de Ingeniería Informática.

A great experience during this Ph.D. period has been my internship at the Telecommunication Networks Group of Politecnico di Torino, Turin, Italy. Thank you Prof. Marco Mellia for hosting me. I extend my thanks to all other members of the research group and Ph.D. students of the "sala dottorandi delen".

Appreciation is given to RedIRIS for providing us with the traffic measurements that have been fundamental to this thesis.

I wish to acknowledge Spanish Ministry of Education and Science that has funded this research under project *DIOR* (TEC2006-03246) and F.P.I. fellowship program.

Last but not least, I would like to thank my family for its constant support and understanding, and for always pushing me to succeed. Thanks for making our home such a wonderful place, full of love and affection.

Contents

List of Figures

List of Tables

Acronyms

ANCOVA Analysis of Covariance.

ANOVA Analysis of Variance.

ARIMA Auto Regressive Integrated Moving Average.

AS Autonomous System.

CCDF Complementary Cumulative Distribution Function.

CDF Cumulative Distribution Function.

CDN Content Distribution Networks.

CLT Central Limit Theorem.

CV Coefficient of Variation.

Diffserv Differentiated Services.

DNS Domain Name System.

DoS Denial of Service.

HTTP Hypertext Transfer Protocol.

IM Instant Messaging.

Intserv Integrated Services.

IPTV Internet Protocol Television.

ITU International Telecommunication Union.

IXP Internet Exchange Point.

MIB Management Information Base.

ML Maximum Likelihood.

MLE Maximum Likelihood Estimation.

MRA Multiresolution Analysis.

MRTG Multi Router Traffic Grapher.

MSE Mean Squared Error.

MVA MultiVariate Analysis.

NAT Network Address Translation.

NNTP Network News Transfer Protocol.

NREN National Research and Education Network.

P2P Peer-to-Peer.

PCA Principal Component Analysis.

PMF Probability Mass Function.

POPs Points of Presence.

POTS Plain Old Telephone System.

QoA Quality of Audio.

QoS Quality of Service.

QoV Quality of Video.

RFC Request for Comment.

SLA Service Level Agreements.

SNMP Simple Network Management Protocol.

TELNET Teletype Network.

TERENA Trans-European Research and Education Networking Association.

UAM Universidad Autónoma de Madrid.

VPNs Virtual Private Networks.

ZM Zipf-Mandelbrot.

Chapter 1

Introduction

This chapter provides an overview of this Ph.D. thesis, presents its objectives and hypotheses, and, finally, outlines its structure.

1.1 Overview and motivation

Network traffic measurements collected across the Internet provide very meaningful information for researchers, service providers, and other members of the Internet community [CK06, FK03, BC02].

On the one hand, network operators may benefit from such information in their goal of ensuring the appropriate Quality of Service (QoS) to their customers. Indeed, the ever-increasing user demands and wide variety of application requirements are forcing Internet Service Providers (ISP) to develop network capacity plans very carefully, not only to maintain the QoS provided, but also to reduce the need for investment. ISPs have not underestimated the benefits of traffic measurements, and have traditionally applied their potential to other related fields, namely the performance evaluation of networks, the detection of anomalies and Denial of Service (DoS) attacks, and even the generation of the costumers' invoices.

On the other hand, the research community has also found essential the use of real network measurements to improve the understanding of the Internet dynamics, and further apply this knowledge to the development of network models, with direct application to network operators' needs mentioned above.

However, the collection of representative traffic measurements is not a straightforward process. In light of this, the authors in [FP01] provide a detailed explanation of the problems that can be found on the simulation of the Internet, some of which also arise in the process of measuring networks. Examples of such problems include the large size and the heterogeneous nature of the Internet, the ever-increasing number of new applications being introduced to the network, the fast and unpredictable way the Internet changes, the size and date of the sample collected, and the handling of outliers in the measuring process. In addition, new problems are constantly emerging. For instance, the authors in [JTO09] explain the current legal limitations that the research community is finding to share actual traffic measurements or even to carry out new traffic measurement campaigns.

In this thesis, we pinpoint two additional difficulties: First, the *"spatial diversity"* of measurements, that is, whether the information obtained from measurements collected at diverse locations with similar features differs significantly or not; and secondly, the time required to capture stationarity, the *"temporal diversity"*, that is, the amount of measuring time needed to bring a sampled distribution which persists over time. Specifically, we have analyzed three important characteristics of the Internet, (i) the popularity of IP addresses, (ii) the popularity of port numbers, and (iii) the geolocation of the Internet connections.

To this end, this thesis has had access to traffic measurements of RedIRIS, the Spanish National Research and Education Network (NREN). RedIRIS connects more than 300 institutions and the measurement capture process has lasted more than 2 years, resulting in an overwhelming amount of data which is by itself a difficult problem to deal with.

This work also addresses this issue. Specifically, we further propose a mechanism to downsample traffic time-series using Multiresolution Analysis (MRA) with wavelets, and evaluate the optimal subsampling level based on comparing the queueing behavior of the subsampled and original time-series at the output of a router. This mechanism is more related to network performance than conventional downsampling methods, since queueing delay is a very representative QoS metric. In addition, we take advantage of the fact that the internetwork measurements show a strong periodicity due to their relation to the human activity patterns. Thus, by applying MRA with wavelets we take into account information not only in the time domain but also in the frequency domain.

Next, this work pretends to characterize the RedIRIS' busy hour given its importance for accurate capacity planning. This challenge is of fundamental importance for the ISPs since the quality that their users receive directly depends on the link bandwidth. In general, there are two approaches [vdM06] to meet the Service Level Agreements (SLA): (i) using Integrated Services (Intserv) or Differentiated Services (Diffserv) to limit either the number of users in the network or the resources that can be requested, and (ii) overprovisioning the network capacity such that all the users and applications' requirements are met. In this thesis, we pay special attention to this latter option.

We have found that the two main drawbacks of the most of the current approaches to the capacity planning problem are: (i) The temporal and spatial diversities are ignored, and (ii) such approaches are typically based only on a priori measurements of the demand for capacity. For instance, these approaches use dedicated measurement systems to obtain the bandwidth consumption and, then, the measurements are used to estimate the demand for bandwidth [PNvdMM09]. However, sometimes it is not possible to measure in a given network, for example, a new institution that joins RedIRIS or a subnetwork without a dedicated measurement system. In addition, sometimes the problem is to foresee "what happens if" a certain feature changes. That is, the problem is to deem the variation on the demand for network resources by adding new network users, when network topology changes or it is upgraded. If the estimation are based on previous measurements only, then, these questions cannot be addressed.

In this light, this thesis takes one step further and shows how the demand for bandwidth can be estimated by means of the intrinsic features of the networks. Basically, such features include the number of users and the network access capacity. Thus, given these features, network managers can estimate the demand for bandwidth in their networks.

1.2 Objectives and hypotheses

The overall aim of this work is to show how the internetwork measurements (sufficiently representative and appropriately captured, validated and reduced) can be useful to characterize a facet of the Internet's behavior as important as the Internet's traffic busy hour and how it can be estimated from the IP networks' intrinsic features. Consequently, the following specific hypotheses and objectives are defined:

1. *Hypothesis:* Traffic measurements gathered from a limited number of networks and limited duration cannot be considered to be sufficiently representative of the Internet.
   Objectives: We pretend to assess to what extend the traffic measurement campaigns must last to obtain stationary internetwork statistics of a given network.
   In addition, we pretend to assess if a homogeneous set of networks shows similar behavior with regard to several internetwork statistics.

2. *Hypothesis:* If the Internet traffic measurement campaigns must last for long periods of time, the volume of data that such campaigns entails can result by itself difficult to analyze, monitor, and store.
   Objectives: This objectives comprises several aspects:

   – To propose new techniques to subsample Internet traffic measurements. We focus on the fact that it is well known that such measurements follow the patterns of the human behavior, and, consequently, they show strong periodicity.

   – To define an automatic and objective mechanism to identify when the subsampled signal is not longer representative of the original one.

   – To validate the proposed approaches with real data that represents measurements and statistics of the Internet traffic using an extensive set of networks during a representative period of time.

   – To compare the proposed approaches with previous well-known methodologies to subsample time-series.

3. *Hypothesis:* The demand for bandwidth in the busy hour over mid-length periods can be characterized by a stochastic process.

Objective: The aim of this objective is to model the traffic volume exchanged during the busy hour over time by means of a stochastic process. This task comprises two subtasks:

- Visual inspection of the data in order to propose a model.
- Validation of such model with real data, in this case measurements from RedIRIS.

4. *Hypothesis:* The demand for bandwidth during the busy hour over long periods calls for a non-stationary process model.

   Objective: This objective includes the inspection of stationary of the demand for bandwidth in the long term, bearing in mind that it is expected that the demand increases over time. Thus, we pretend to assess if such increment is either at a constant rate or, conversely, the demand changes as a staircase function (that is, as a set of consecutive stationary processes).

5. *Hypothesis:* The demand for bandwidth in low-utilized networks are not polluted by their access capacities. As RedIRIS networks' utilizations are typically low, we support the hypothesis that access capacities are not "capping" the demand of the users.

   Objective: This objective tries to evaluated if the demand for bandwidth is correlated to the access capacity of the networks in an extensive set of network scenarios.

6. *Hypothesis:* The parameters of the process that models the demand for bandwidth over time can be estimated by means of the networks' intrinsic features. Consequently, the demand for bandwidth in given a network can be estimate in an objective and fairly fashion, and, even, avoiding to carry out a dedicated measurement campaign.

   Objective: Once the busy hour process is modeled by a stochastic process, this objective intends to infer the parameters of such process by means of explanatory variables, specifically, a set of network intrinsic features.

The reader may notice that these specific objectives are treated throughout this Ph.D. thesis and their evaluation is reported at the end of this document.

1.3 Thesis structure

This thesis is organized as follows:

– Chapter 2 presents the state-of-the-art. Firstly, we show literature that analyzes NRENs. Secondly, we focus on network monitoring tools that we have used along the thesis. Next, we present some different Internet traffic models that the research community has proposed to characterize the Internet. Finally, we show how some of these models have been applied to the capacity planning problem.

– Chapter 3 analyzes the popularity of the IP addresses and port number focusing in two aspects of the Internet's characterization: the spatial and temporal diversities, which have usually been ignored. In this light, this chapter compares the behavior of universities with similar characteristics during several months. Moreover, we also explain why both spatial and temporal diversities occur. In addition, this chapter presents the measurement scenario that is analyzed throughout this thesis. Essentially, it includes a description of the RedIRIS' architecture and measurement systems as well as the validation of the data.

– Chapter 4 analyzes the geolocation of the Internet traffic destinations paying also special attention to its spatial and temporal diversities.

– Chapter 5 deals with the traffic measurements reduction problem and proposes a new method to subsample network measurements over time based on the Multiresolution Analysis with wavelets.

– Chapter 6 analyzes the dynamics of the Internet traffic busy hour due to its importance in the capacity planning problem. Several goodness-of-fit tests are applied to an extensive set of RedIRIS' measurements, concluding that a Gaussian distribution can model the Internet traffic busy hour. In addition, we further study the impact of the intrinsic features on the demands for bandwidth at the Internet traffic busy hour.

– Chapter 7 presents the conclusions that can be drawn from this thesis and proposes some directions for future research lines.

Figure 1.1 shows the relation between the different contents that make up this work and the layout of the chapters.

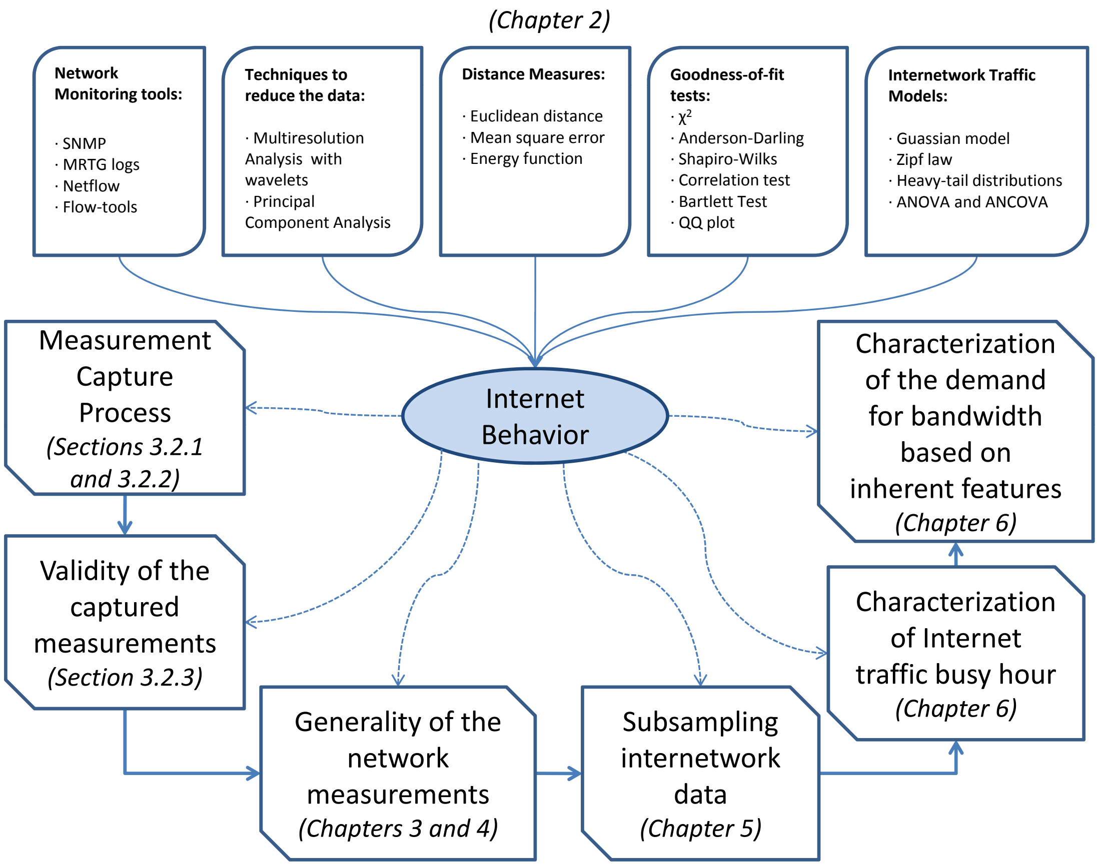

Figure 1.1: Structure of the document

Chapter 2

State of the art

This chapter presents a general background of the most important issues that this work deals with and it is useful to justify the decisions taken along the research process. Specifically, we focus on the following items:

- *Studies about other academic networks. In this section we analyze research on networks similar to RedIRIS since its measurements are an essential part of this thesis.*
- *Network monitoring tools, which serve to provide an accurate capture process and deal with the huge amount of data that the traffic measurements imply. We have paid special attention to the Cisco's Netflow and Multi Router Traffic Grapher logs because RedIRIS is currently using these tools.*
- *Data reduction. We present different proposed approaches that aim at reducing the volume of data for further processing and storage while still preserving relevant information.*
- *Traffic characterization. In this section we present several models to characterize the Internet traffic that the research community has proposed.*
- *Network Dimensioning. In this last section we study proposed methodologies to dimension networks as well as techniques to predict the evolution of the Internet.*

2.1 National Research and Education Networks

The academic networks have received relatively much attention by the research community. Basically, this is due to the current regulation that usually prevents the research community from sharing actual traffic measurements and even capturing traffic in commercial networks [JTO09]. As a consequence and due to the greater accessibility of the academic networks, the researchers have typically used measurements from their own institutions which are normally academic networks. As this thesis analyzes measurement from an extensive set of campus networks, this section presents some previous studies based on academic networks' measurements.

On the one hand, there are many research studies that have used measurements from academic networks to evaluate new algorithms, systems, architectures or applications' performance. Examples of this are the MEHARI and MIRA projects [LASP+99, RGMG+02]. In particular, MIRA is a distributed and scalable real-time measurement platform and MEHARI is a low-cost programmable and scalable tool for the analysis of IP/ATM traffic. The authors tested both projects using RedIRIS' network and some details of the RedIRIS' infrastructure can be found in these works. Moreover, measurements from academic networks have proven useful to validate traffic classification methodologies. Such methodologies try to identify the applications that compose the traffic generated by a network. Examples of that are the community research's efforts to detect Skype [BMM+07, BMM+09] and to classify the traffic into different application groups [CEBRSP].

In addition, other studies have shown the current state of a given academic network in terms of bandwidth, number of nodes and number of centers connected; however, these studies neither explain the link capacity planning nor compare the behavior between the institutions that make up the network. For instance, in [Lia93, MKYK89, Mat87] the characteristics of the Japan and China's academic networks are shown. However the authors do not provide any conclusion or methodology for capacity planning. Similarly, the academic network of Slovenia is explained in detail in [JBOK03]. Specifically, the authors introduce the network topology, network technology, national and international connectivity, and services that are available in the network.

However, few studies have compared in detail the measurements from an extensive set of well-characterized networks. Only the Trans-European Research and Education Networking Association (TERENA) [Terb] has supported the communication and the flow of ideas between the network manager of each of the academic networks in Europe. Consequently, TERENA (especially by means of TERENA networking conference) has presented some results of comparing several European academic networks. The authors in [Tera] compare several characteristics of a set of academic networks. This includes the core capacity of each network, the capacity to Europe and to USA/rest of the world, network size, technology used in the core network, traffic per student, services provided, staffing, investments, budget per university student and country, and population.

It should come as no surprise that the networks with the greatest capacity, staff, and investments show the highest ratio between the traffic volume and the number of students per academic network. Such ratio ranges from lower than 0.6 Mb/s per student (in Albania) to more than 4 Mb/s per student (in Finland). When the budget per student between both countries are compared it becomes evident that Finland exceeds notably ones Albania. As a conclusion, the traffic generated by user strongly depends on the "intrinsic" features of the academic networks and their relation is evident.

In the case of this thesis, all the institutions under study are part of the same operator, the Spanish academic network. Therefore, the most of these features are not applicable. However, by comparing a set of institutions we assess if they share similar features in terms of connectivity, provided services, population, etc. as will be explained in detail in the forthcoming of this thesis (Chapter 3).

2.2 Network monitoring tools

A great deal of monitoring applications has been found in the literature. Such applications enable to estimate, among others, the available bandwidth, the router queueing occupancy, service times, and link one-way delays. For example, the authors in [ALML06] show an application that estimates many of these metrics, transparently to users. In [LdVBF04] it is shown how these metrics can be estimated, in a university network, using only Netflow information. A work that shows an extensive background about this topic is [PMDC03]. However, we note that the implementation of network monitoring tools are out of scope of this thesis. Actually, this work must adapt to the available measurements from RedIRIS. Specifically, RedIRIS collects Netflow records, which consist of a summary record per flow that traverses its routers, and logs from the Multi Router Traffic Grapher (MRTG) tool.

2.2.1 Netflow

The analysis of Netflows seems to be the best option to study the RedIRIS traffic characteristics. A Netflow is a sequence of packets that share the same resource and destination IP addresses, port numbers and protocol. More specifically, a Netflow summary includes: aggregated traffic volume in bytes, number of packets, port numbers, source and destination IP addresses, type of service, input and output interface indices, together with time-stamps for the flow beginning and end.

Initially, the Netflow protocol was implemented in Juniper and Cisco's routers to reduce the router workload and subsequently standardized [Lei04]. As a matter of fact, Netflow is the name that Cisco gave to the network flow records that its routers generate, however, currently Netflow refers to any network flow records. Additionally, the protocol that enables to export these records to other machines is usually named Netflow as well. Routers using these Netflow records could forward a packet without examining the routing table which implies an important computational cost reduction [Cis]. Since then, Netflow records have been used for many other tasks as we will show in this chapter and in the rest of this thesis. Such tasks include, among others, the intrusion-prevention systems, the client's invoicing, traffic characterization, capacity planning, and estimations of the available bandwidth capacity.

There are numerous tools and applications that enable to capture, store and even visualize the Netflow records. One of these applications is the Flow-Tools software package [RFL00], which has been used in this thesis. This application allows the capture of the Netflow records that a router generates and stores them compressed. Furthermore, it can be also filter the Netflow according to the IP addresses, port numbers, size (in bytes) and length (in seconds) among others.

In [LSLY05] a tool is presented that anonimizes the stored Netflow records in order to fulfill the privacy policy requirements.

Netflow records have been extensively used to provide security to the networks. There are many applications that enable to visualize a set of flows just in a graph (note that it can involve to depict at least six dimensions). Thus, at first glance, a network operator can detect problems in its network. In [OGK06] and [Yur06] several of these applications are shown, specifically they show how Netflow can be used to detect viruses, worms and attacks (like Denial of Service (DoS) attacks) in the traffic that traverses a certain network.

However, both the Netflow capture process and Netflow statistics creation process have two limitations. On the one hand, a flow may last days, nonetheless the routers have a limited memory capacity and they cannot maintain an infinite number of open flows awaiting for a hypothetical new packet. In general, a flow is said to be finished [Cis]:

- When a rejected or an end-of-connection packet (FIN/RST flags) is found (in TCP connection case).
- When a router does not detect traffic of an open flow for a reasonable period of time (say 30 seconds).
- When a flow remains open for a long time (usually 30 minutes).
- When the flow statistics table is completely full and the router needs to free some records to store new ones.

On the other hand, the traffic that goes through a router can be so intense that the time a router takes to update the flow statistic table is larger than the packet interarrival time; thus, the router cannot create all the Netflow records accurately. These limitations are explained in detail in [EV03]. The solution that has been accepted by the industry to avoid this latter problem is based on sampling the flows, taking only a percentage of the packet that a router forwards. Then, from this smaller set of packets the whole flow statistics are estimated. Obviously, this mechanism adds inaccuracy to the Netflow records.

In [EV03] two new sampling algorithms are proposed that improve the deterministic sampling performance (i.e., to take one out X packets, where X is a fixed number). These algorithms are based on the fact that the traffic is composed by a small set of large flows and a large set of small flows. Thus, paying special attention to the set of the largest flows results in a better traffic characterization than analyzing all the flows regardless of their sizes. However, these algorithms, in the case of this thesis, are unlikely to prove useful since they are a replacement of the algorithms currently installed in the network routers, and this is unfeasible. In contrasts to this, the above-mentioned work shows interesting performance comparison with the Cisco sampling algorithm (deterministic sampling). They show results that depend on the flow size: for those flows that represent more than 1% of the link capacity, the deterministic sampling algorithm has a error of 9% while in the case of smaller flows the error is much higher.

The authors in [DLT05] also assess the traffic characteristics after the sampling process and present a comparison between the deterministic sampling and the random sampling (i.e., to chose a packet with probability p). Reference [CB05] performs a similar analysis. Both studies conclude that, in practical terms, the flow statistics resulting from both subsampling algorithms are very similar, although the authors show that the results are not statistically identical because the algorithms can be distinguished using goodness-of-fit tests.

Cisco's Netflow records are used in [FGL$^+$01] to characterize a backbone network as well as to estimate its traffic matrix. The obtained results improve previous results clearly. Such previous results were based on civil engineering approaches where the researchers know the number of vehicles in a point of the road network but they do not know the destination of each vehicle. Therefore, these previous works estimated the traffic matrix using traffic load measurements in the network links without information from either source or destination addresses of the traffic [BR02]. Conversely, Netflow data comprises information both source and destination addresses, thus, the authors in [FGL$^+$01] estimate the traffic matrix of a backbone network analyzing only Netflow records in its peer links (and some extra internal links to estimate the traffic between clients of the same network). Additionally, they compare the traffic volume measurements that can be obtained using Netflow records with the estimations obtained using other techniques such as the Simple Network Management Protocol (SNMP) [CFSD90]. Similar comparisons were performed in [LPC$^+$04]. In this latter study it is estimated that the difference, regarding network load, ranges from 1% to 5%.

Another interesting work about traffic characterization is presented in [WP05]. In this case, the authors focus on the optimization of the traffic data storage process since some operators must retain information about their clients' traffic to comply with certain regulations. They compare different aggregation levels that fulfill such regulations: storing all the packets, only IP headers, Netflow records and pre-analyzed Netflow records. They conclude that Netflow is the most economical way to store traffic information while preserving the information required by the regulations. Specifically, the best option is to store pre-analyzed Netflow records; this pre-analysis consist of removing redundant and useless Netflow data. That is, the source port of a connection is sometimes chosen randomly so this information is pointless, however the destination port is usually important since it is often bounded to specific services and applications. They also propose to assemble similar flows in a single one, for instance, flows of the same application. However, this pre-analysis has an evident problem: the computational cost. Obviously, the process of assembling flows by application saves storage resources but the process of identifying application, by itself, is a challenge. From this point of view, storing Netflow without pre-analysis is the best option. According to [WP05] the ratio between the traffic volume and the space that Netflow records demand is in the range of 0.4% and 0.04% depending on the pre-analysis process.

2.2.2 Multi Router Traffic Grapher logs

MRTG [Oet98, Shi08] is a software tool that reports the amount of traffic forwarded by a router interface. The data collection process is performed by means of polling SNMP-enabled network devices in order to obtain the value of a given variable Management Information Base (MIB). Specifically, it performs polling of the *ifInOctets* and *ifOutOctets* counters of the Interfaces MIB [MK00]. The time between two consecutive readings is configurable but it is typically set to 5 minutes, as it is the case of RedIRIS. MRTG stores such information in its own text-format (MRTG log files) and then generates graphs and web pages for the given time interval. It includes statistics such as maximum, minimum and mean values. The MRTG log files format is presented in detail in [MRT]. Basically such files comprise five columns that represent:

- The Unix timestamp for the point in time the data on this line was measured.
- The average incoming transfer rate in bytes per second. This is valid for the time between the timestamp of the current line and the timestamp value of the previous line.
- The average outgoing transfer rate in bytes per second since the previous measurement.
- The maximum incoming transfer rate in bytes per second for the current interval. Note that MRTG files comprise not only records per two consecutive pollings of the Interfaces MIB, but also the aggregation of several pollings. Consequently, this column accounts for the maximum value between some consecutive aggregated pollings.
- The maximum outgoing transfer rate in bytes per second for the current interval.

2.2.3 Data reduction

One of the major challenges of the current monitoring systems is that the volume and diversity of measurements that such systems are able to collect is so humongous that makes it difficult to process, visualize and even to store (both in capacity and speed terms). In this light, several methods have been proposed aiming at reducing the amount of data while preserving the most of its relevant information. In this section we present some of them, for further details the authors in [CK06, Chapter 6] give a detailed review. Such methods include:

- Counters. This approach forms time-series of counts of internetwork traffic statistics which are aggregated over certain time intervals. That is, as first step this approach measures some traffic statistics, typically at fix time intervals, then, it stores the data, and, finally, an aggregation process is applied to reduce the data in the time domain. This is the behavior of the MRTG tool, presented in previous section, which counts the traffic volume traversing a given network using SNMP and then aggregates the data per day, week, month and year. In this thesis, we show how such time-series can be even reduced applying the fact of that the time-series that represent the traffic statistics are strongly related to the human activity patterns and, consequently, they show periodicity.

– Flow aggregation. This approach consists of summarizing sequences of packets into flows, thus, only information about the flows is stored instead of information of each packet. This method was previously explained in Section 2.2.1 because the RedIRIS' routers are able to export Netflow records of the traffic that they handle.

– Sampling. In this approach only a subset of the total set of packets is measured which implies a reduction of the amount of data collected as well as a cut in the computational load of the monitoring system. Note that the flow aggregation methodology typically also applies sampling to limit the load that the flow aggregation process implies. The use of packet sampling by forming Netflow records was introduced in Section 2.2.1. As introduced, the sampling process can be done in deterministic or random intervals or, alternatively, applying adaptive [CPZ04] and stratified sampling [EV03]. The former approach adapts the sampling rate to the traffic characteristics, for instance, reducing the sample rate at the busy hour or increasing it when the traffic load is low. On the other hand, stratified sampling divides the packets into subsets and then sampling is applied within each of such subsets.

– Packet truncation. Another approach to reduce the storage requirements is to collect only a fraction of the packets that make up the traffic. Basically, such methodology only collects a fixed number of bytes of each packet or a variable number of bytes, for instance, only the TCP/UDP headers. Additionally, this latter option circumvents legal concerns about privacy by avoiding to record potentially sensitive packet payloads [JTO09]. Albeit this is at the expense of extra computation load, because it requires to inspect the packet to remove such payload. The main drawback of this approach is that the amount of data to be stored, in spite of the packet payload being removed, can still be excessive as was shown in Section 2.2.1.

– Other approaches proposed to deal with the data reduction problem are the Bloom filter [BM03], and the dimensionality reduction. We note that an overview of the applicability of this latter methodology is given in Section 2.3.1 due to its importance in the characterization of the Internet traffic.

2.3 Traffic Characterization

As the Internet characterization has turned out to be useful both from a commercial and a technical standpoint [CK06], the research community have recently started to present works whose aim is to characterize different aspects of the Internet's dynamics.

There are a number of studies that have analyzed Internet characteristics, such as web domains [BYCE07], web servers' workload [AW97], web performance [MA98], popularity of the port numbers and IP addresses [GDHA+08], Peer-to-Peer (P2P) applications [IUKB+04], Internet Protocol Television (IPTV) applications [MM09], online games [Bor00], trends and tendencies in the traffic [Pax94], just to mention some of them.

The authors in [CK06], [NP08] and [FP01] present extensive overviews to many of the well-known topics and the major challenges of the Internet measurement characterization.

Specifically, in this latter reference the authors define an Internet invariant as a facet of behavior that is empirically shown to persist for some time in a wide set of measured samples. Thus, each invariant can be considered as an Internet traffic characteristic and can result useful as foundation for future characterizations. Some of the invariants are:

– Daily traffic pattern [TMW97]. The Internet follows daily patterns clearly with direct relation to human behavior. According to this, the traffic during the working days is "similar" from one day to another, but they are different to the holidays and weekends that are, in turn, similar between them. Likewise, the traffics during a week, month and year are similar between them. Similarly, within the 24 hours in a day the traffic also shows clear trends. We have found in RedIRIS' traces that the traffic during the night is low, it increases about 8-9 a.m., then it shows a dip during the lunch-time, after this resting time traffic reaches a peak at 3-4 p.m. and finally it begins to decrease until 8-9 p.m. when the traffic becomes stable again, following the night traffic pattern. However this pattern cannot be extrapolated to any scenario, it depends on the population, the kind of network and the protocols among other factors.

Regarding the population, obviously, the traffic patterns depend on the population's regular lunch and work schedules.

Regarding the kind of network, we note there are many types of networks, for instance, academic networks (as RedIRIS), commercial networks and private networks. The traffic patterns in these networks can be very different. In the case of academic networks the traffic during the weekends (especially, incoming traffic) is very low, however in a commercial network the traffic during the weekends can be as high as during the working days since the users usually have more spare time during the weekends.

Furthermore, there are some protocols that are independent of the human behavior, such as the Network News Transfer Protocol (NNTP) and automatic software updates. Figure 2.1 shows the Netflow traffic that RedIRIS generates, the traffic patterns described above turn out to be evident.

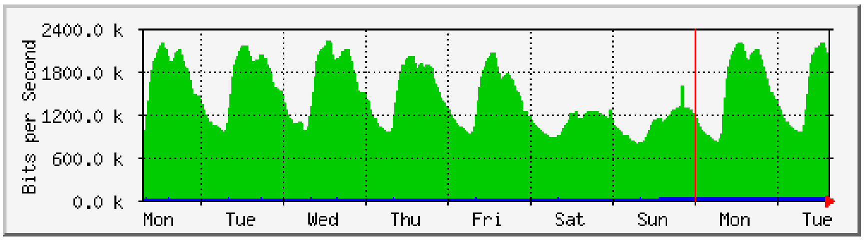

Figure 2.1: RedIRIS' Netflow traffic for one week

- The interrarival time between packets show a self-similar behavior, in contrast with previous approaches that assume that the traffic follows a Poisson model.

- The network "user session" arrivals can be modeled by a Poisson distribution. A "user session" interrarival is defined as the time between users decide to use the network for a certain task and they decide to begin another one.

– The connection sizes follow a log-normal distribution, however there are many other Internet network activities that behave following a heavy-tail distribution like Pareto (including the case when Pareto's parameter α is lower than 2 which implies infinite variance). Such distributions play an important role in this thesis and they will be explained in detail later.

– Finally, there are other minor Internet invariants such as those related to certain protocol behaviors like Teletype Network (TELNET), or even regarding the Earth's topology: it never changes since the cities and the countries are always in the same place.

On the Internet characterization task the research community has always taken into account the Internet invariants since they are a very useful a-priori knowledge that facilitates the understanding of the Internet.

In general, any characterization can be addressed from the parametric and non-parametric point of view, the difference lies in that the former assumes that the data follows a model or distribution, for instance, the Normal distribution. In the following two sections we give an overview to both parametric and non-parametric approaches that are typically used by the Internet community. We note that in the forthcoming chapters some of such approaches are explained in a deeper detail whereas other ones, not so typically used in the Internet measurement field, are only introduced when we apply them to RedIRIS' data.

2.3.1 Non-parametric analysis of Internet Traffic

The Internet traffic measurements can be analyzed as a mere set of data. Therefore, any of the well-know data mining techniques can be applied. Such techniques include the cluster analysis, the MultiVariate Analysis (MVA), neural networks, Principal Component Analysis (PCA) and MultiResolution Analysis (MRA) with wavelets. These techniques enable to find and group patterns of a given traffic trace as well as filtering and finding the parameters that characterize the traffic in an optimal and formal fashion.

The main problems that arise in the analysis traffic measurements are: (i) The data usually covers very long time periods, (ii) the patterns are not well-characterized, and (iii) the number of variables or dimensions is usually too high.

All these problems, specially the high dimensionality of the data, make the task of applying any of the data mining methodologies difficult. Thus, typically two approaches have been applied, on the one hand those techniques that transform the original variables into new ones; this transformation usually consists of projecting the old variables into new dimensions and removing the lowest informative dimensions. An example of these techniques is PCA. On the other hand, those techniques that subsample the information, i.e., techniques that reduce the number of samples of a signal while retaining most of the information. One example of this technique is the MRA with wavelets. The following research studies show how these techniques can be used in practice.

The authors in [LPC+04] capture traces of a backbone network and apply PCA to them. They find out that the Internet traces can be divided into three classes: deterministic traffic, bursty traffic and Gaussian noise. Moreover, they find periodic trends in the traffic each 12 hours, 24 hours, week and month. Finally, they show that by using PCA the traces under study can be characterized by only a few dimensions. This point is interesting since the volume of data that we are collecting from RedIRIS makes it difficult to apply any of the mining data algorithms.

However, the characterization of universities not only includes bandwidth measurements but also other variables such as the number of students, the university staff, types of degrees, etc. PCA transforms these variables, projecting the original variables into new axes that are optimal (they maximize the variance of the data) but leaving the new variables without meaning. Note that after applying PCA the variables do no further represent the number of students or the university location, they are only coordinates. This drawback can be important when the traffic measurement variables are well-known as in the RedIRIS' network.

The authors in [dOSVP06] follow the above-mentioned article's guidelines adding some interesting contributions. Basically, they repeat the same experiments (obviously, they use their own traces), resulting that the dimension reduction, by means of PCA, is not so powerful. In addition, they show how to apply cluster algorithms to the Internet user's traffic. Specifically, they cluster the clients of a commercial operator using 1-week-traffic measurements. The results show that the population can be grouped into three or four clusters with a high clustering coefficient between the members of each cluster. Then, they apply PCA to the data and repeat the clustering process, resulting in important differences (specially when the number of groups is high, say 4 or 5). However, these experiments were performed over 1-week data in which each pattern represented the traffic of a single user with a total set of 6,000 users (mostly residential clients). Hence, as pointed out in this paper, the results can be influenced by the high measurement variance caused by the low aggregation level. The RedIRIS' measurements are much more aggregated since the data comes from universities with thousands of Internet users.

Another technique to face the high dimensionality problem is shown in [PTZD05]. They propose the use of MRA with wavelets to smooth and subsample temporal series, i.e., consecutive bandwidth traffic measurements. Such measurements comprise MRTG logs, which implies a sample per 5 minutes that represents the average throughput during this time interval. As they find periodic trends each 12 and 24 hours, they state that subsampling interval must be multiple of these trends. To achieve this goal they use the analysis with wavelets [Dau90]. This analysis consists of dividing a signal into two sub-signals: the first one, named Details, comprises the high frequency terms and the second sub-signal, named Approximations, comprises the low frequency components. If the Approximation sub-signal is divided, again, into Approximation and Details signal (and so forth), finally, the signal that we obtain is a subsampling of the original one since by each step the number of samples is divided by two. In the MRTG measurements case the original signal has a sample per 5 minutes, after the first subsampling process one sample per 10 minutes, then per 20 minutes and so on. Finally, it is necessary to have a distance measure that determines whether the subsampling process is optimal regarding the signal size and the loss information. In general, the most used distance measurement is the Euclidean distance [CFY03] and the Analysis of Variance (ANOVA) [PTZD05].

Additionally, wavelet functions have also been used to characterize the traffic in different aggregation scales [RRCB00], to estimate of the Hurst parameter [AV98] and to detect anomalies in the Internet traffic [BKPR02]. MRA with wavelets technique will be applied in Chapter 5, consequently, in such chapter a more detailed description of this technique is given.

2.3.2 Parametric analysis of the Internet Traffic

Nowadays it is usually assumed that the Internet traffic exhibits normal-distributed characteristics. In [vdMMP06, KN02] an analysis is presented that assesses whether the used bandwidth follows such a distribution in a short-term fashion. Thus, an operator can estimate the link load in the following time instant by calculating the traffic mean and variance. These two studies assess if the Gaussian model is still true when the traffic timescale is in the millisecond range (horizontal aggregation). Additionally, they also analyze if this Gaussian behavior depends on the network level aggregation (vertical aggregation). Both studies find Gaussian behavior regardless of the timescale (at least up to 5 milliseconds) and that the aggregation does not have to be large (in the order of tens of users). To check the Gaussianity of the data they propose to correlate the original signal and the Quantile-Quantile plot function [DS86, Chapter 2] instead of using some well-known goodness-of-fit tests. Which are usually too demanding for traffic measurements. Note that the networks are often upgraded, they suffer malfunctions and the traffic patterns change from working days to weekends or holidays making the use of Gaussian models difficult.

In [AH02] the Zipf distribution is used to model some of the Internet's characteristics. Among others, the webs' incoming links usually follows a Zipf law, the distribution of P2P networks or, for instance, the mails between Internet users. In Zipf distributions a few elements represent most of the distribution, and many elements explain a very small part of the distribution. Therefore, a small set of webs are linked by a large set of networks (think about Google, for instance) and there are many webs that do not have incoming links. In [FGL+01] it is shown that Zipf law fits the most popular IP addresses distribution, consequently a small set of IP address is visited by millions of users and an extensive set of IP addresses are hardly visited. We note that Zipf law will be applied in the chapters 3 and 4 where we will provide a more extensive review.

Similarly, the research community have noticed that most of the Internet traffic is generated by a small fraction of network users [Bro02, PTB+02], often referred to as heavy-hitters [FGL+01]. A heavy-hitter is typically defined as a user whose use of the network resources has a significant impact in the aggregated traffic of the whole network. According to this, a user is considered to be a heavy-hitter:

- If the traffic generated is higher than a certain threshold, for instance, 100 times the mean traffic/user or, in absolute terms, more than 1 GB.

- A user can be also considered as a heavy-hitter if such user accounts for a certain percentage of the total traffic; in this case, depending of the total number of users, a user whose traffic represents more than 1% or 0.1% of the traffic volume.

- Another way to detect heavy-hitter users is to consider as such to the set of users that represent a high percentage of the total traffic, namely 80% or 90%.

Finally, in this thesis we will use the ANOVA and Analysis of Covariance (ANCOVA) [DC74, Jai91] as a way to compare and contrast traffic measurements from the extensive set of networks under study. However, the Internet community has not paid much attention to either ANOVA and ANCOVA methodologies. In this light, they will be explained in the chapters 4 and 6, respectively, in detail. The authors in [PTZD05] apply ANOVA, as mentioned in the previous section, to reduce the high dimensionality of Internet measurements. Basically, they quantify the amount of variability accounted for by each term in a multiple linear regression model (in this case each term represents a wavelet detail signal) and they select only those that comprise significant information. Additionally, in [Ber06] ANOVA is applied to identify those factors that have a significant impact on the Quality of Video (QoV) and Quality of Audio (QoA) of multimedia applications over the Internet, such factors include the transmission delay, jitter, and packet loss ratio among others.

2.4 Network Dimensioning

In [vdMMP07, PNvdMM09] the "Smart Network Dimensioning" concept is introduced. This is defined as the lowest capacity (C) that a link must have to meet the Quality of Service (QoS) requirements. The aim of such study is to assess the validity of a general dimensioning rule for any network scenario: $C = M{\cdot}d$, where M is the traffic mean (during the busy hour, for instance) and d is a network constant. It turns out that the d value strongly depends on the network scenario (infrastructure, population, configuration, among others) and the level of aggregation; consequently, a universal d parameter does not exist. As the former formula is not general they proposed to use the following one (assuming that traffic is Gaussian): $C = M + \sigma/T\sqrt{V(N)}$. Where $V(\cdot)$ is the traffic variance, T is the time period, $\sigma = \sqrt{-2\log\epsilon}$ and ϵ is the probability that the required bandwidth is lower than the available bandwidth during T temporal units. In a previous work [vdBMvdM+06], the authors defined the following equation that did not take into account the traffic variance, it is more simple but more dependent of the kind of the traffic: $C = \rho + \alpha\sqrt{\rho}$ being ρ the traffic mean and α an empirically estimated parameter ranging from 0 and σ according to the kind of traffic in the link under study. They estimate that the d parameter should be 3 at the most, and, however, in practice this constant can be larger than 30 [Odl03]. This shows the excessive overprovisioning of the current networks.

2.4.1 Relation between Traffic and Population

Another important aspect regarding network dimensioning is to find relations between the characteristics of a network, its population and the traffic measurements. In the literature, relatively few attention has been given to this issue. In [AB97, BAD99] the "service factor" is introduced. They propose a direct relation between any characteristic of the network under study and the generated traffic. For instance, if the population doubles, then the operator should double the traffic bandwidth. However, in [GMS+07] it is shown that this perfect linear relationship cannot be assumed using traffic measurement from several Spanish universities.

In [BBD05] a methodology is shown to allocate certain available capacity bandwidth of a generic backbone network to the network's users. This methodology is based on dividing the traffic in classes according to its priority and its QoS requirements. They propose an algorithm that, in execution time, estimates if the bandwidth per traffic class is correct, otherwise it allocates more resources to any inadequately dimensioned traffic class.

The authors in [MACL06] presents an empirical model to estimate the demand for bandwidth according to the number of users in high speed networks in a short term view. They first characterize the users' demand of 190 ADSL subscribers as a

Gamma distribution in a 1-sec scale during one day. Then, they repeat their study focusing only on the busy hour. The results show how that they can extrapolate the results for such amount of subscribers to any given number of users applying a simple formula. However, the authors do not take into account that the behavior of one network with respect to others is far to be homogeneous in order to extend the results to other scenarios. In addition the impact that the access capacity can have in the demands is also ignored.

RedIRIS gives service directly to an large fraction of the total set of Spanish universities. However there are some other universities, which make up regional networks, that do not receive service from RedIRIS directly. Actually, such regional networks are a single Internet Exchange Point (IXP) for RedIRIS; this implies that the traffic inside such networks is not captured by RedIRIS (RedIRIS' architecture and measurement system are explained in the following chapter). Regarding to this, the authors in [MKYK89] analyze whether this traffic between universities, in this case between several Japanese universities, is high. In particular, they assess the requests for using online computation resources, resulting that the most of these requests come from members of a university to their university itself. Another important part of these requests linked close universities and, finally, the traffic between the rest of the universities was low. However, the ratio between intra-universities traffic and external traffic shows that the traffic volume of the former is almost marginal. This implies that, very likely, the RedIRIS' traffic that is not captured is very low compared to the total traffic.

2.4.2 Bandwidth variability and demand

The Internet users and network infrastructures are always changing, which entails a high variability of the user's demand for bandwidth and a challenge to the network operator. The authors in [dFV02] address these problems. They analyze the investment that an operator must perform to upgrade its infrastructures in order to meet client's needs. Specifically in [dFV02] it is shown how to estimate the increase of demand for bandwidth in a fixed period under uncertainty. That is, from an estimated growth rate they add a factor that estimates the volatile part of the demand. They support that this factor can be modeled by means of a Wiener stochastic process. Such a process has been used to model the fluctuations of the share prices, for instance, in [PP00] is used to model the Madrid stock market exchange. Moreover, the authors in [dFV02] intend to estimate the "optimal timing investment into new capacity". Once they model the traffic variability, using traces from a university network, they estimate the time in which the operator should upgrade its infrastructure given the required QoS. Finally, they show, again, that the current networks are excessively overprovisioned. Only in non-typical situations (such as a very large increase in the demand for bandwidth or in the case of absurd volatilities) an operator should upgrade its networks at use-rates of 50%. However operators often think that they must upgrade their networks at much lower rates. Nonetheless, some of the estimations that this model requires, such as the estimation of initial bandwidth demand or the estimated growth, by themselves, are very challenging to obtain.

In [All08] the author proposes to forecast the capacity for web servers in a peak-driven fashion. This approach consists of the characterization of how the peaks of demand change over time and requires to set a time window in which the peak is measured. In this specific case, the demands refers to web servers resources as the bandwidth, storage capacity, CPU load, among others. The analysis focuses on the mid-term, therefore the temporal window is set to a week. Once a 4-month measurement campaign is carried out, the author shows how the number of processes in a set of Apache servers (measured as the weekly peaks) of the photo-sharing site Flickr.com increases at constant rate because the number of stored photos also does. Finally, the author shows how the upper limit capacity of the Flickr.com infrastructure will be reached in three months assuming that the rate remains unchanged.

In the literature we have found some works that predict the demand for petrol, gas and electricity among others things. Although the premises and constrains do not match our scenario, an Internet operator, the models and methodologies that the authors show can be useful in this thesis. In particular, the authors in [BM95] deal with the prediction of the demand for petrol from a local gas company's point of view. The demand fluctuates according to the month, holiday periods, temperature and the demand for other industrial and agrarian products. This study proposes to use linear regression and neural networks to perform such estimations; the results show that neural networks can predict the demand with more precision. Similarly, in [EC96] it is intended to predict the electricity demand for a much larger scenario: the Spanish electricity market. However, the forecasting is just for short term, since the electricity operators need to know the demand from one day to the following one. Specifically, they propose to use an Auto Regressive Integrated Moving Average (ARIMA) model. They show a detailed study of the factors that can influence the demand. In this case the most influential factor is, above all, the temperature. They find a non-linear relation, just as expected, since when the temperature ranges within a comfortable interval (say 15-25°C) the air conditioning equipments and heaters are usually switched off, and consequently, the demand for electricity remains low and constant. Nevertheless, outside the comfortable intervals the demand suddenly increases, but when the devices are working close to their limits the demand for energy becomes constant again. Obviously all these considerations cannot be extrapolated to our scenario: an Internet operator. It does not seem reasonable that the temperature can have an effect on the demand for bandwidth. However both types of demands fluctuate according to the time (holidays, weeks, months, among others).

The authors in [PTZD05], as previously seen in this chapter, show how to use MRA with wavelets to subsample signals. Nonetheless, the final goal of such article is to estimate the optimal moment to upgrade a backbone network (similar goal to [dFV02]). In this case, they model the traffic using ARIMA (as in [EC96]). They have traffic measurements for several years, and use this data to predict the traffic for the following six months.

2.5 Conclusions

In this chapter we have shown the state of the art of some mathematical models, techniques and tools that will be used along this thesis.

These include (i) monitoring and capturing tools such as MRTG, Netflow and Flow-tools. (ii) Data reduction methods such as MRA with wavelets. (iii) Distance measurements such as the Euclidean one. (iv) Goodness-of-fit techniques such as QQ-plot function. (v) Several internetwork traffic models, for instance, Gaussian, Zipf and Gamma distributions among others. In addition, we have shown different approaches to problems such as the capacity planning problem, optimal timing investment into new capacity, traffic characterization, traffic prediction, etc.

As a result, we have mainly found three aspects that have not been analyzed in-depth:

1. The generality of the captured measurements. That is, we have found many studies that analyze their new algorithms, architectures and methodologies using measurements from university networks, residential or office environments. However the authors in these studies have not taken into account whether these measurements were representative of the Internet, or, otherwise, these measurements only represented the behavior of a small set of the Internet's users. We believe that this lack of analysis is due to the difficulties in obtaining measurements from an extensive set of networks over large periods of times. This analysis is performed in Chapter 3 and Chapter 4 of this thesis.

2. The research community has presented different techniques to reduce the number of samples of a traffic measurement. These techniques are of paramount importance since such measurements usually involve overwhelming amounts of information. However these well-know methods are based on mathematical distances and they do not take into account the meaning of the information, i.e., traffic measurements. In addition, the research community has shown that some aspects of the Internet's behavior follow well-know patterns. In light of these premises, Chapter 5 of this thesis shows how to make the most of this a-priori knowledge in the traffic measurement subsampling process.

3. The research community has shown the importance of the busy hour in the capacity planning problem. Essentially, this measure is usually employed by the Internet operators to estimate the bandwidth of their links. Despite busy hour's importance, we have not found any characterization of its dynamics nor explanatory factors in an extensive set of network over a long period of time; consequently, this thesis addresses this issue in Chapter 6.

Chapter 3

Generality of the network measurements and measurement scenario

Often, Internet measurement-based studies have followed a three-step procedure: (i) Collection of network measurements, (ii) measurement-based model inference, and (iii) generalization of the results obtained to other scenarios. Indeed, it has been a general belief that certain internetwork traffic statistics, such as the mostly used IP addresses and port numbers, show a similar behavior in networks with similar features, and the conclusions derived from the measurements of a given network could be extrapolated to a similar scenario.

This study makes no starting assumption concerning this issue and undertakes a "spatial" analysis of network measurements. The measurement set comprises a six-month trace collected by RedIRIS (the Spanish National Research and Education Network) at different monitored points across the country during 2007.

Our experiment shows that, although the frequency statistics of IP addresses and port numbers follow a Zipf distribution (as expected), the distributions' characteristic parameter values vary significantly in a spatial dimension, that is, across the individual university networks, *even when the profile of the networks' user base are similar. Furthermore, we show that experiment durations of approximately one month are necessary for the traffic processes to display stationarity. Hence, in order to obtain accurate statistics on traffic characteristics of large internetworks using state-of-the-art measurement techniques, long and spatially diverse experiments may be necessary.*

The structure of the present chapter is as follows: The first section presents the introduction and motivation of the work carried out in this chapter. Section 3.2 is devoted to the description of the measurement scenario that will be used in this chapter as well as in the rest of this thesis. Then, Section 3.3 presents the universities selected to perform the study. The experiments performed and results obtained both in the time and space dimensions are shown in Section 3.4, and, finally, Section 3.2 summarizes the conclusions of this chapter.

3.1 Introduction

Collecting representative traffic measurement is a fundamental task in the Internet traffic characterization. However, as stated in Chapter 2, it is an involved process due to the heterogeneous nature of the Internet, its continuous expansion in speed and size and the different application requirements, among other reasons.

In this chapter and the following, we pinpoint two additional difficulties: First, the "spatial diversity" of measurements, that is, whether the information arisen from measurements collected at diverse locations with similar features differs significantly or not; and secondly, the time required to capture stationarity, the "temporal diversity", that is, the amount of measuring time needed to bring a sampled distribution which persists over time. Essentially, we try to answer the following two questions:

Can the conclusions derived from a measurement experiment in a given network be further applied to a similar network/scenario? And, how long should the measurement experiments last until stability in the metrics under study is reached?

Throughout this work, the term *similar* networks shall refer to networks which share certain common intrinsic features. In this light, the research community has generally accepted that the conclusions derived from a given network are valid for a scenario with similar characteristics, such as population size, bandwidth capacity and filtering policy. Therefore, measurements have been taken from links that are believed to be sufficiently representative of the Internet, typically university, residential or even smaller networks.

To answer the questions above, this chapter focuses on the distribution of the most popular IP addresses and port numbers (often bound to specific services/applications), and the following one analyzes the geolocation of the Internet connections in a set of university network access points nationwide.

It is worth noticing that this study is not focused on the measurement results themselves, which have been partially reported elsewhere, but instead on the representativeness of network measurement experiments, in terms of spatial and temporal diversity. Temporal diversity is related to the concept of "horizontal aggregation", as introduced in [KN02], whereby the authors study the necessary timescale such that aggregated traffic follows a Gaussian distribution. However, in this work we follow a rather different approach: the problem is not to estimate the timescale to reach Gaussianity but to rather find the time horizon above which the distribution parameters remain stable. Such time horizon is typically in the range of days or weeks, a much coarser time-scale than the ones often considered in such horizontal aggregation studies (seconds or milliseconds). Other works have aimed at ranking the top traffic generators in a network scenario, often known as "heavy-hitters" (see Section 2.3.2) and their persistence over time in such ranking [WF06].

Concerning spatial diversity, this has received little attention from the research community. For instance, the authors in [MKYK89] make a comparison study of the inter and intra-use of mainframes between seven Japanese regions in the late 1980s, but nonetheless the spatial diversity of the measurements was not analyzed. We believe that such lack of spatial diversity related studies is due to the difficulties in capturing traffic from a large number of distant networks and over large periods of time.

In fact, our work analyzes an extensive set of measurements (Netflow records) collected from a large number of university networks kindly donated by RedIRIS, the Spanish National Research and Education Network (NREN) [Red]. RedIRIS spans more than 70 universities whose size, user population and organization is well documented in central repositories by the Spanish Ministry of Education for statistical purposes [Con]. Therefore, it is possible to group universities by similar features, for instance, number of users, bandwidth, traffic filters (e.g., restrictions on Peer-to-Peer (P2P) applications such as music file sharing), and proceed with the analysis to check whether or not, university networks with similar intrinsic characteristics produce similar traffic patterns.

3.2 Measurement Scenario

In this section we presents the measurement scenario that will be used in this chapter as well as in the rest of this thesis. The set of traffic measurement was kindly donated by RedIRIS for research purposes[1]. First, a brief summary of RedIRIS' architecture is given, secondly the data collection infrastructure is explained, and, finally, we assess the accuracy and validity of the data.

3.2.1 RedIRIS' architecture

RedIRIS serves more than 300 institutions, mainly universities and research centers, and comprises 18 Points of Presence (POPs) across the country, as shown in Figure 3.1 (right). Each node represents a set of communication equipments that concentrate the backbone transmission media and access lines of the institutions of each region. However, it is worth remarking that RedIRIS does not give service inside five regional networks, namely Andalucía, Cataluña, Galicia, Madrid and Pais Vasco. This implies that only the traffic that enters/leaves these regional networks is captured by RedIRIS. Moreover, it has external links to other European academic networks such as GEANT, and Internet Exchange Point (IXP) with TELIA, COGENT, LEVEL3, among others. For the experiments, RedIRIS provided the traffic measurements at the access routers of a large number of universities connected RedIRIS, typically with a bandwidth ranging from 100 Mb/s to 1 Gb/s.

[1] The data is stored in isolated servers and never treated at the individual flow level, in full compliance with the Spanish regulation concerning privacy of the electronic communications

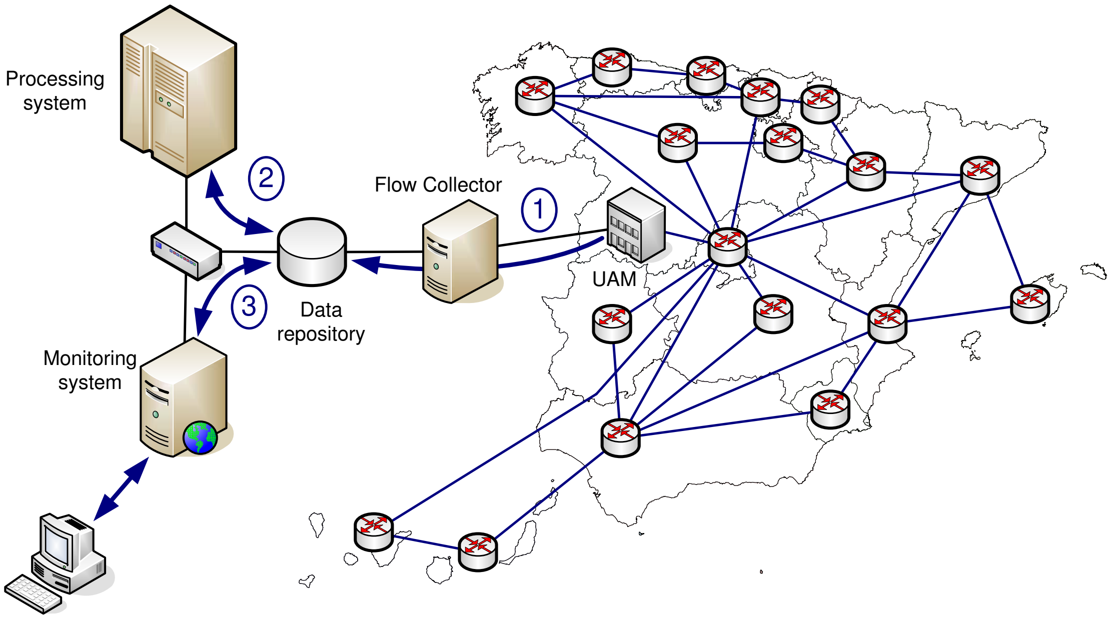

Figure 3.1: Measurement system architecture and RedIRIS' network topology

3.2.2 Data collection infrastructure

The flow summaries provided by RedIRIS are sent to a central repository, located at the Universidad Autónoma de Madrid (UAM) campus. The average input rate to the repository was 2 Mb/s (flow compressed records), over a three year period (April 2007-2009). Figure 3.1 (left) shows the measurement system architecture. First, the *Flow-Tools* software package was used for data collection at the repository. Then, a number of statistics were obtained by the processing subsystem, which included total bandwidth consumption, most active IP addresses and port numbers, busy-hour bandwidth requirements, heavy-hitters users (see Chapter 6), and geolocation information (see Chapter 4) per university. In the rest of this thesis we labeled the data of each university as U_1, U_2, ..., U_N, where N is the number of RedIRIS institutions, due to privacy reasons. Finally, the Monitoring System provides a graphical interface, whereby such processed information can be accessed via web and properly visualized (this is the third stage).

Additionally, RedIRIS have provided us with Multi Router Traffic Grapher (MRTG) logs of eight universities, four regional networks and five external links

/IXPs from February 2007. In the following we refer to such links, for privacy reasons, as U_a, U_b, ..., U_h in the case of university networks, RN_a, ..., RN_d to the set of regional networks and, finally, we have labeled the set of external link and IXPs as EL_a, ..., EL_e.

Let us introduce some of the traffic definitions that are used in this chapter and the following ones. We shall denote "incoming traffic" as the traffic volume, flows or packets that are sourced by a host located somewhere in the Internet and destined for a host located in one of the RedIRIS' institutions, and we shall denote "outgoing traffic" as the converse, i.e., the traffic volume, flows or packets that are sourced by a host in one of the institutions of RedIRIS and destined for a host in the Internet (see Figure 3.2). Note that with these definitions, inter-RedIRIS' insitutions traffic is neither incoming nor outgoing traffic, *indeed we did not include such traffic in our experiments* in order to homogenize the set of measurements.

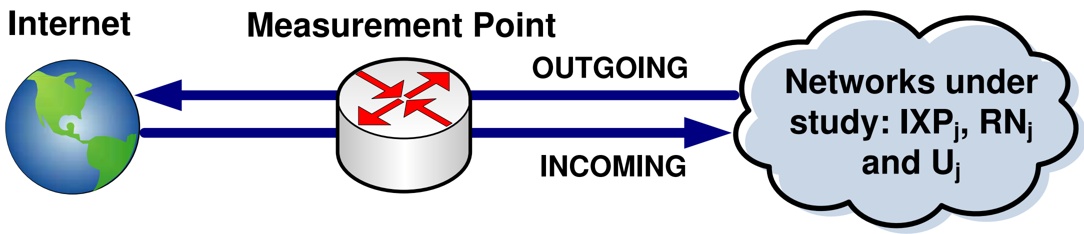

Figure 3.2: Incoming and outgoing directions of network traffic

3.2.3 Traffic Measurement validation

The RedIRIS' data validation involves three main tasks:

1. In the process of generating Netflow records, as stated in Section 2.2.1, each flow is subsampled at packet level in order to reduce the router's workload. In the case of RedIRIS' routers the sampling ratio is configured to be either 1:100 or 1:200. The first question that arises is whether these statistics, calculated multiplying the traffic by the inverse of the sampling ratio, are accurate enough to perform any subsequent analysis.

2. The flow-tools application stores all the Netflow records that we are receiving without any source distinction. Thus, the traffic per institution is estimated filtering by IP address ranges.

3. The RedIRIS' Netflow records traffic uses UDP as transport layer, so we could lose some packets and consequently, the statistics' accuracy diminishes.

To ensure that the packet loss ratio is low, that we can filter the traffic by university and regional router, and that the influence of the packet sampling process is negligible, we compare the traffic measurement estimated using the Netflow records and the measurements obtained using other tool such as MRTG. In this light, we compared the MRTG logs of 8 universities and the sampled Netflow records of these same 8 universities after inverting the sampling ratio and filtering by university. To perform the filtering process RedIRIS informed us about regional routers and universities network's IP addresses that compose the network. Figure 3.3 shows this comparison for three campus networks. The used bandwidth calculated using Netflow records is plotted as a solid line; the bandwidth according to MRTG is plotted as a full area. The accuracy can be visually evaluated in such figure, the rest of universities showed similar results. This results are consistent with those from other studies [FGL$^+$01, LPC$^+$04, SF02, MCS$^+$06].

In conclusion, the statistics obtained from Netflow records fit the actual traffic with accuracy in spite of the above-mentioned three limitations at least at 5-minute aggregates. This thesis focuses on the characterization of traffic measurements at a much larger granularity, thus, as we have shown, we can ignore such three limitations.

3.3 University networks under study

The collected traffic sample comprised more than 70 universities, with different user base populations, access link capacities, filtering policies (P2P applications), proxies and Network Address Translation (NAT) capabilities. Clearly, such intrinsic features have an impact on the traffic pattern. For instance, if NAT services or proxies are available it is very possible to find that most traffic comes from a single IP address, but the truth is that a large number of traffic sources are shar-

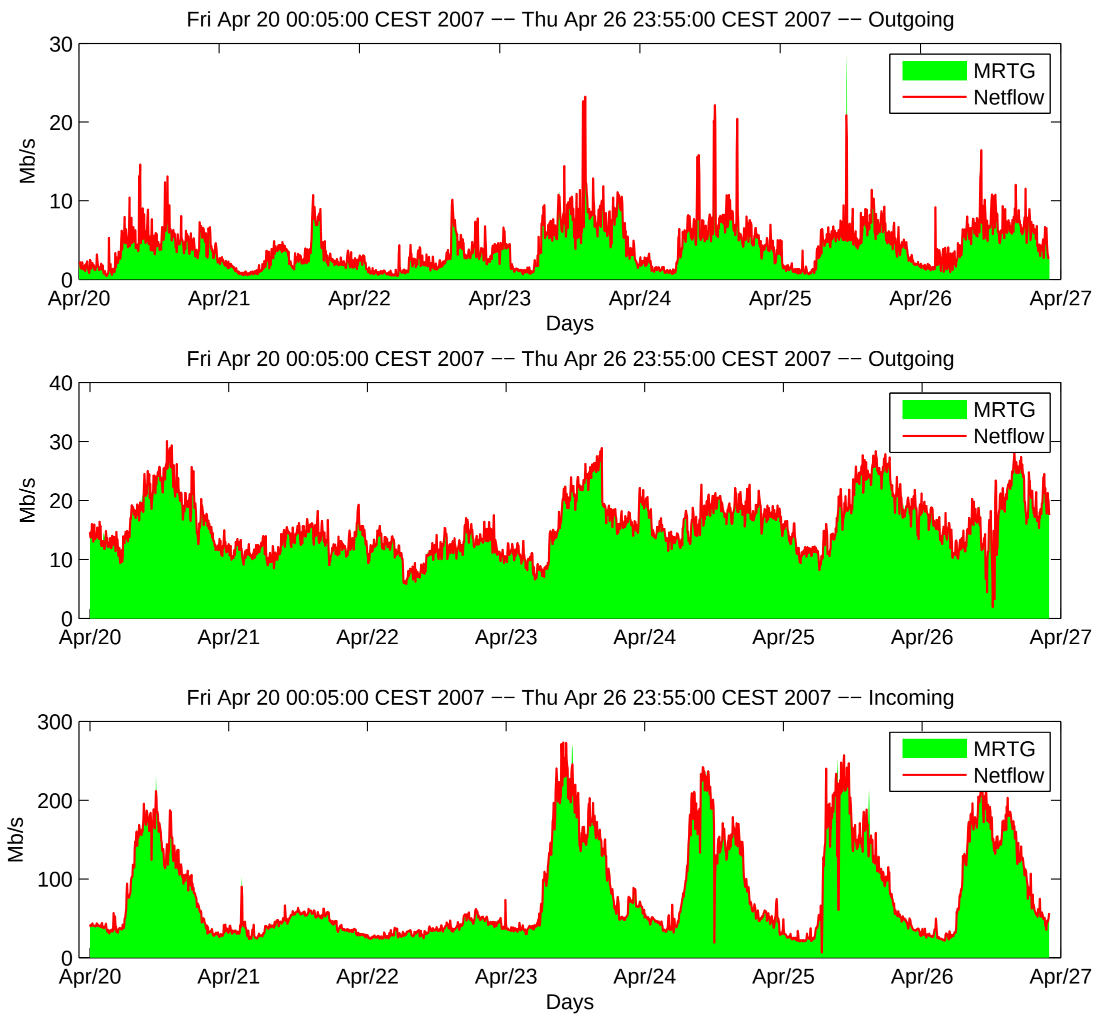

Figure 3.3: Traffic for a set of 3 different campus networks for 7 days, as a full area the bandwidth estimated using MRTG, as a solid line using Netflows (outgoing direction in the first two examples and incoming in the last one)

ing the same IP address. In the same way, NAT not only affects IP addresses but also port numbers, since every traffic source under the same IP address is given a different port number.

Consequently, we made a choice of universities with similar features, and compared the resulting most popular port and IP addresses distribution. In this light, we have carefully selected 9 universities out of the total set, for which the above intrinsic features are very much alike. Let us labeled such universities as $U_1 \ldots U_9$ for privacy reasons.

Firstly, regarding the filtering policy, we have chosen universities in which most non-educational traffic is allowed with no rate control except for well-know P2P application ports. Additionally, it is worth noticing that the analyzed measurements comprise traffic to the Internet only, not between campuses. Thus, such inter-university traffic from supercomputing or grid facilities is explicitly not included. Furthermore, we also performed an inspection of the most active flows, in order to ensure that no outliers were present in the sample.

Secondly, concerning the use of NAT, we focus on most frequently accessed IP addresses and ports on the Internet side, i.e. destination IP addresses and port numbers of outgoing flows from campus networks, and origin IP addresses and port numbers of incoming flows (see Figure 3.4). Such measurements provide a more meaningful and representative portrait of the user behavior browsing Internet content, rather than pursuing a characterization of the Internet users that access hosts in the university campuses.

The population size of the universities under study ranges from $20,000$ to $40,000$ members with a similar proportion between subpopulations (strata), i.e. students, faculty and administration, thus favoring the representativeness of the aggregated traffic (see table 3.1). Furthermore, this table shows the number of collected Netflow summaries for the selected universities, along with the number of active IP addresses in the busy traffic hour. The latter gives a hint of the population activity, to reinforce the fact that the sample is representative in terms of number of active users.

In addition to this, the access bandwidth capacity in all universities under study is exactly 1 Gb/s and they are connected to the Internet through a single Exchange Point, located in Madrid.

University	Population (ratio students / staff)	No. of Flows per day	Different IP addresses in the busy hour (University/ to Internet)
U_1	40,000 (9.6)	1,400,000	4,000 / 23,000
U_2	28,500 (10.8)	1,300,000	4,000 / 22,800
U_3	20,000 (12)	2,000,000	3,200 / 30,000
U_4	31,500 (11)	5,870,000	5,000 / 90,000
U_5	30,500 (10.3)	3,000,000	4,300 / 66,000
U_6	36,000 (11.2)	4,000,000	5,600 / 66,000
U_7	33,500 (12.2)	3,500,000	4,500 / 58,000
U_8	26,500 (11.2)	2,400,000	6,500 / 30,000
U_9	28,000 (10.5)	2,500,000	2,000 / 30,000

Table 3.1: User-base population size, average number of flows collected per day, and average IP addresses in the busy hour per day for all universities under study

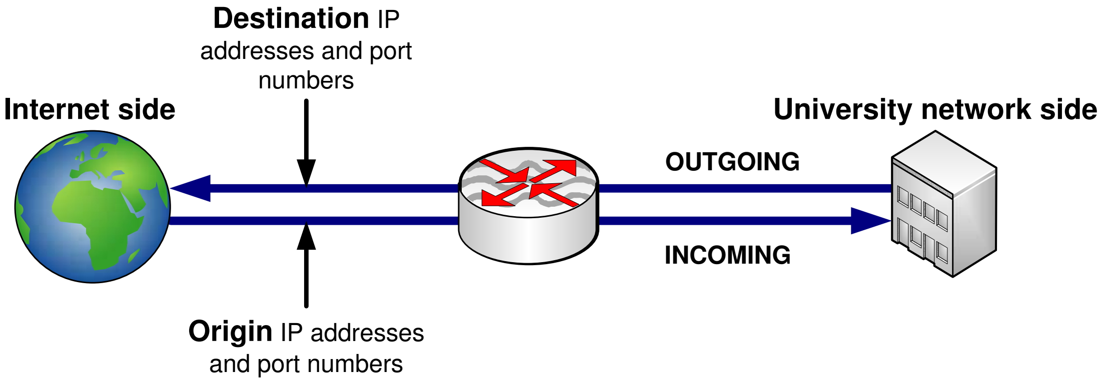

Figure 3.4: Analysed data: Destination IP addresses and port number for outgoing flows, and origin IP addresses and port numbers of incoming flows

We conclude that the selected universities are similar in terms of user base populations, access link capacity, filtering policies (P2P applications) and availability of proxies and NAT services. It is finally worth remarking that the measurements were collected over the same time period, thus avoiding any contamination of the spatial diversity by temporal factors.

3.4 Experiments and results

The following presents a measurement analysis from the spatial diversity point of view, that is, whether or not equivalent universities share similar behavior. It also shows the timescale for which the observed behavior becomes stable, i.e. the sampling distribution does not significantly change as the sample size increases.

A typical invariant that can be observed from measurements of a university network concerns the IP addresses and port numbers most widely found in the traces. It is well-known that, although the amount of possible destination IP addresses of flows and port numbers is huge, most users typically connect to the same sites and use the same services [FGL+01]. Moreover, the amount of traffic either sourced or destined to the most popular IP addresses and port numbers follows a Zipf distribution. Zipf-like phenomena have been observed in the past in internetwork traffic traces [AH02], and often appears in other disciplines, such as economics, sociology and linguistics.

The Zipf Cumulative Distribution Function (CDF) is given by:

$$F(k) = \frac{\sum_{n=1}^{k} 1/n^s}{\sum_{n=1}^{N} 1/n^s}, \qquad k = 1, \ldots, N \tag{3.1}$$

where $s > 0$ characterizes the Zipf distribution, N is the number of most popular IP addresses or port numbers included in the study, and k refers to their rank.

In our spatial analysis, we shall study the most popular (namely comprising most exchanged traffic in bytes) IP addresses and port numbers. Thus, we shall use $F(k)$ to represent the cumulative fraction of traffic (in bytes) over the total that are sent to the k^{th} most popular IP address or port number $k = 1, \ldots, N$ in the Internet.

For example, in Zipf distributions with $s = 1$, the most popular port number ($k = 1$) or IP address comprises as much as twice the traffic exchanged by the second ($k = 2$) most popular one, and thrice the traffic of the third ($k = 3$) popular one, and so on. For $s > 1$, the percentage of total traffic of the most popular one with respect to the others is even larger, and viceversa, i.e., if $s < 1$, such percentage is smaller. Hence, the s parameter is related to the tail decay of the Zipf distribution.

The purpose of the following experiments (spatial diversity) is to check whether or not university networks with similar intrinsic features, as discussed in the previous section, show the same behavior, in terms of the s parameter of the Zipf law.

However, prior to any spatial analysis, it is first necessary to find a time-scale at which the parameters under study are stable. This is the purpose of next section.

3.4.1 Temporal diversity analysis

This section examines the temporal aspect of the measurement set over which we perform the spatial diversity analysis in the next sections. In other words, this section aims to check that the measurement set under study shows stationarity features, i.e., distributions that do not change with time. To do so, we evaluate the number of days worth of data required until the s parameter of the Zipf distribution for the most popular IP addresses and port numbers remains stable.

Figures 3.5 and 3.6 shows the most active destination IP addresses and port numbers of outgoing flows for universities U_1 and U_2 (for 1-day and 1-month time slot), together with its most-likely Zipf distribution, obtained following the least squares linear regression technique described in [Nic87]. The accuracy between the measured data and the theoretical Zipf fit can be visually checked in the figures. Note that only the fifteen ($N = 15$) most popular IP addresses and port numbers were taken into account in the estimation of the Zipf parameter s. We remark that similar behavior was observed for $N = 8$ and $N = 20$, although such results have not been included for the sake of clarity.

These figures also shows that the estimated s values vary for different time-scales. Hence it is necessary to consider a large traffic sample until the s parameter

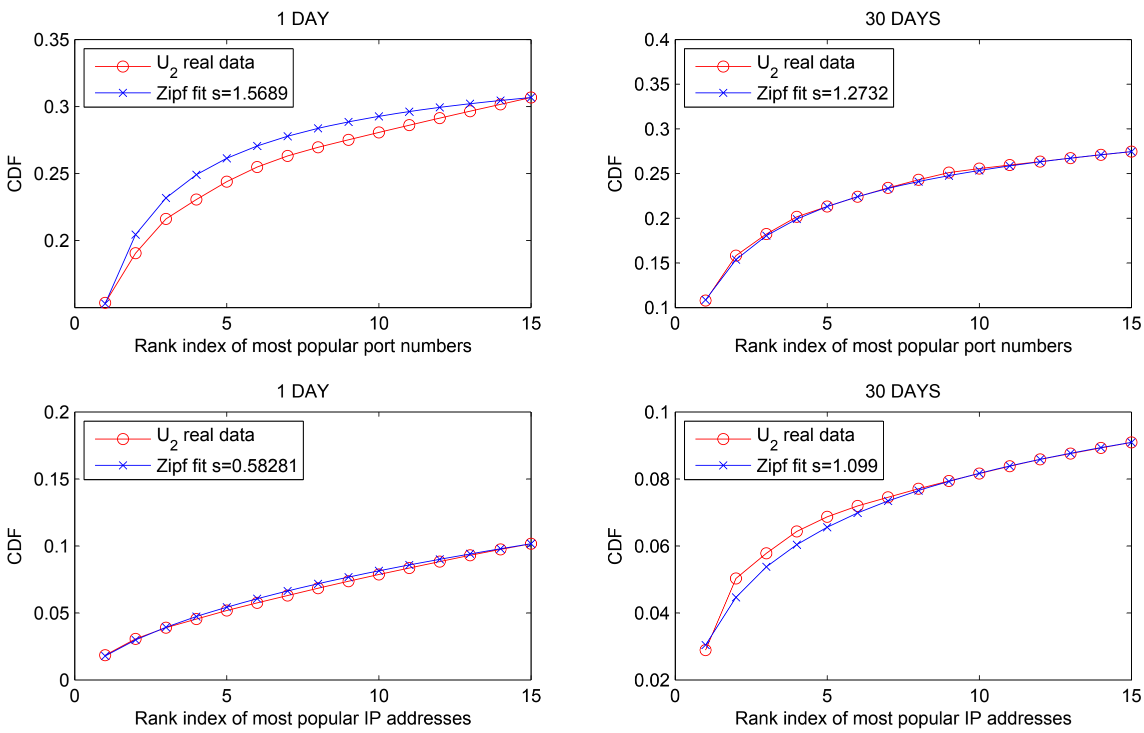

Figure 3.5: CDF of most popular port numbers and IP addresses (outgoing) for U_1 and its Zipf distribution fit, assuming 1 day worth of data (left) and 30-days worth of data (right)

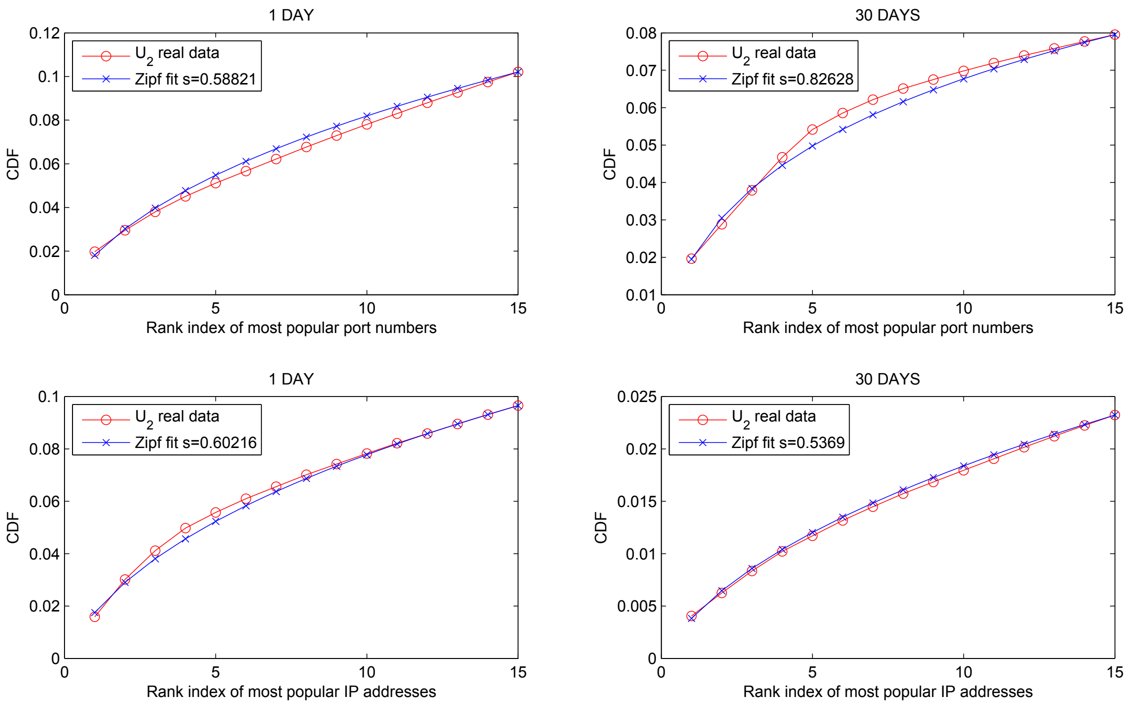

Figure 3.6: CDF of most popular port numbers and IP addresses (outgoing) for U_2 and its Zipf distribution fit, assuming 1 day worth of data (left) and 30-days worth of data (right)

becomes stable. Following this, Figure 3.7 shows the estimated s value assuming several days of measurements. As shown, the s parameter estimate becomes smoother as we increase the trace length, bringing a stable value after 30 days of data. We consider a s estimate is stable if it varies less than 5% after five consecutive days.

It is also worth noticing that the s estimate after 30 days of data is different for all networks under consideration. This issue is analyzed in the next section.

3.4.2 Spatial diversity of most popular IP addresses and port numbers

Figure 3.8 shows the CDF of the fifteen most popular IP addresses (on the right) and port numbers (on the left) for all universities under study, both in the outgo-

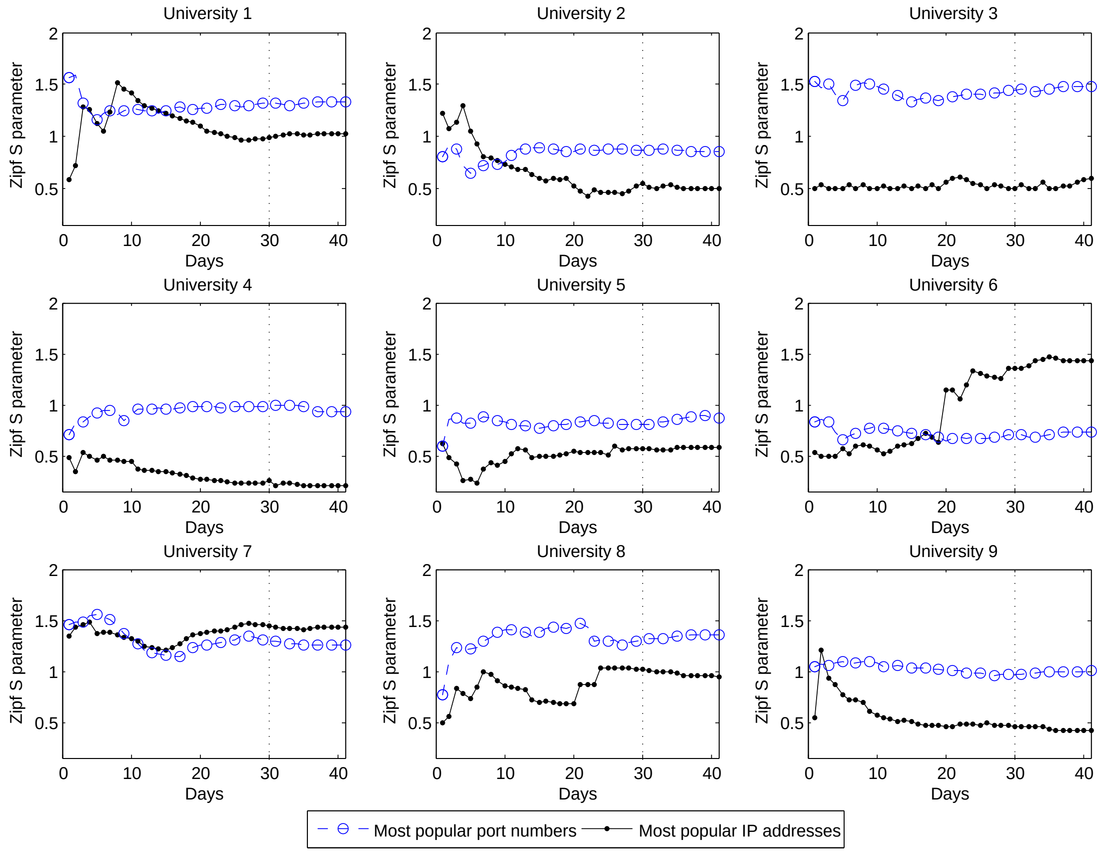

Figure 3.7: Most-likely Zipf distribution s value for the 15 most popular port numbers and IP addresses for all university networks (only outgoing flow direction) for various time-scales of traffic statistics (from 1 day to 40 days worth of aggregated data)

ing (top) and incoming (bottom) direction from the Internet side. The numbers shown refer to the cumulative ratio of transferred bytes over the total in the trace. Following the results of the previous section, we have used 30 days worth of data in order to obtain a reliable estimate of the CDF.

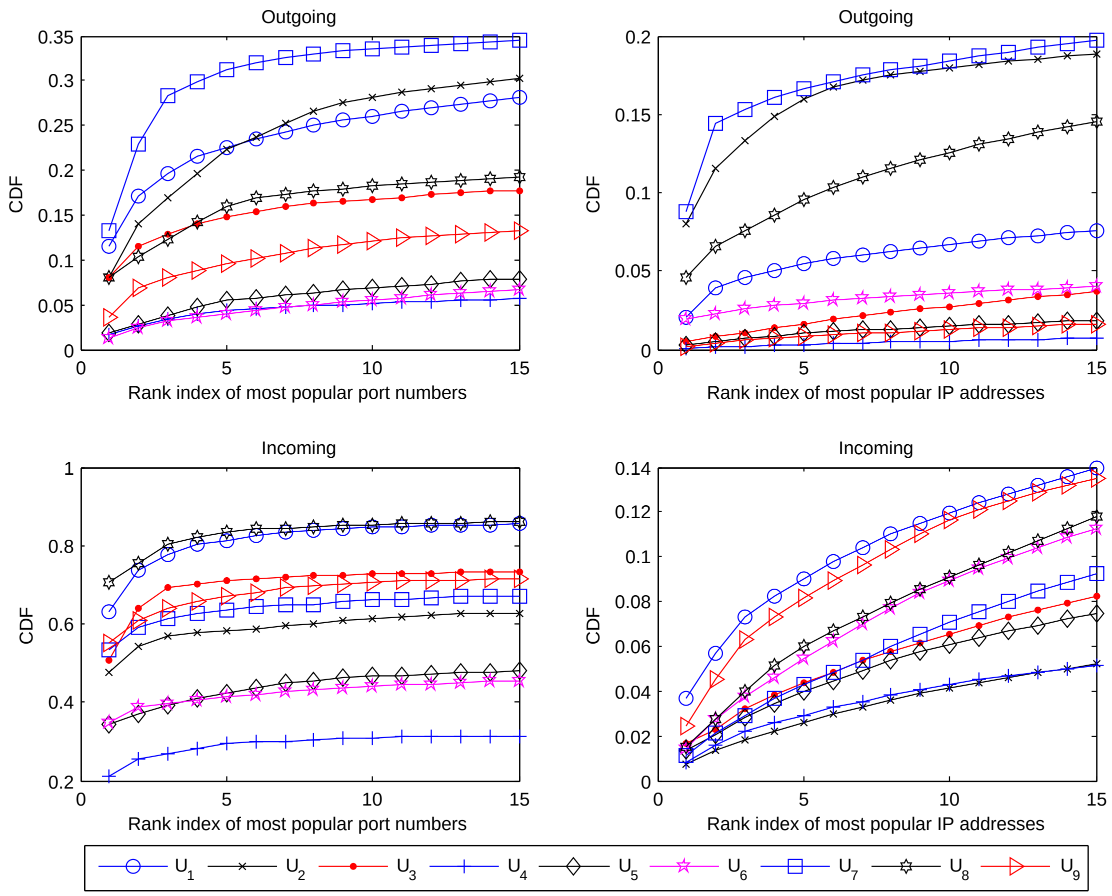

Figure 3.8: CDF of most popular IP addresses and port numbers for all universities under study

Surprisingly, although the networks under study were carefully chosen with similar intrinsic features (large aggregation level, filtering policy, access bandwidth, proxies, NAT and population size and strata), the observed traffic profiles, as measured by the s parameter values, are different from one another. It is worth noticing that the population sizes of all networks under study are large enough (more than $20,000$ Internet users) such that the CDF are expected to converge to the same distribution.

In conclusion, the most popular IP addresses and port numbers of each university network follow a Zipf distribution, but the spatial analysis has shown that the particular s parameter is different in each case (see Figure 3.7). Hence, measurements collected at one university are not generally valid to another, even if they have similar intrinsic features.

3.5 Explanation of the spatial and temporal diversity behavior

This section analyzes why the spatial and temporal diversity occurs. Basically, the possible reasons for this diversity can be divided into two groups: Those linked to the Internet traffic characteristics and those relate to the population that makes up a university.

First, some studies have shown that the size of the Internet connections and the files downloaded by the users follow a heavy-tailed distribution, as mentioned in Chapter 2. This fact may involve that the traffic directed to a port number or IP address never shows stability even though more and more traffic is aggregated.

Second, the population must be large enough to expect representative and comparable results. In this study the smallest university under study has more than 20,000 students which supports this hypothesis. Nonetheless, it is well-know that only a small set of the network users generate the most of the traffic, these users are know as "heavy hitters" (see Section 2.3.2). Therefore, from the total of university users we may only be analyzing and comparing the behavior of a much smaller set of users and, consequently, without representativeness. Thus, particular traffic patterns of a heavy hitter user may have impact on the results of the whole university and explaining the different behavior that the universities have shown.

– *Reason 1:* Heavy-tailed distribution of the Internet connections. To assess if the heavy-tailed distribution of the Internet connections can strongly affect the results, we simulated an scenario in which each connection follows a Pareto's distribution, with α ranging from 1.01 to 1.4. Then, we calculated the ratio between the received aggregated traffic by a certain port number or IP address and the sum of the remaining 15 ones for one year, assuming 1 million connections per day. We fixed the total number of IP addresses and port number to 15 since we use this number in our earlier experiments.

Figure 3.9 shows the results of these simulations. The conclusion that can be drawn from these simulations is that with low α values, about 1.1, the process attains the expected value, 1/15, and shows stationarity. This implies that, bearing in mind that the values for the α parameter for actual traffic traces is higher than 1.1, the heavy-tailed distribution of the Internet connections cannot explain the spatial diversity behavior.

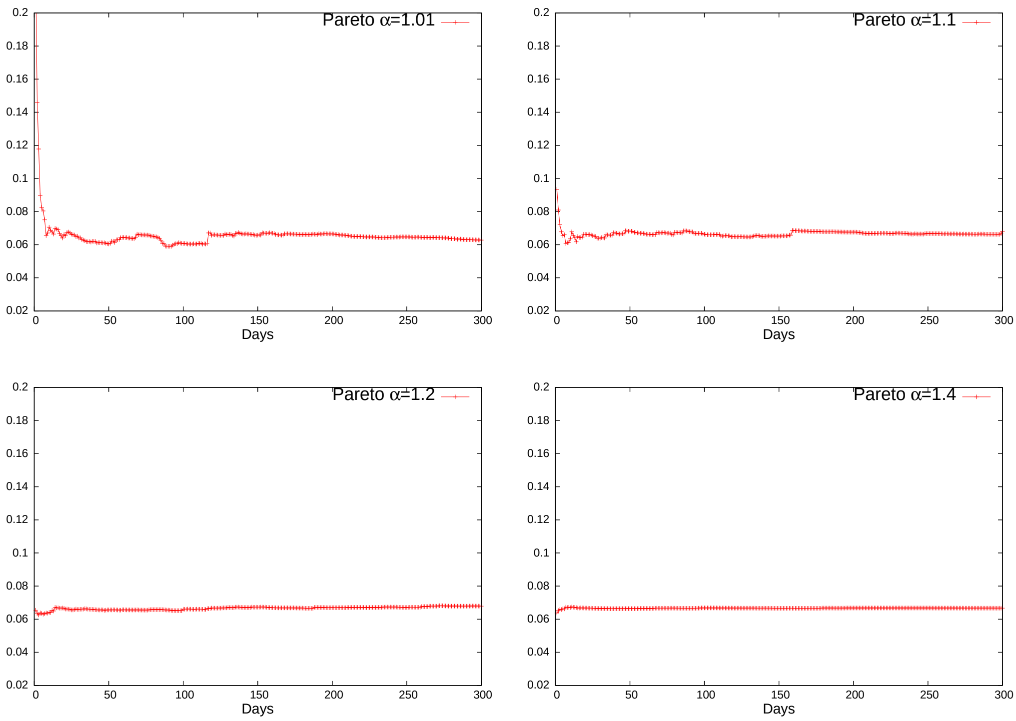

Figure 3.9: Ratio between the aggregated traffic received by one port and the sum of the other 15 ones during several days, assuming 3 million connections per day and the connection sizes follow a Pareto distribution with α ranging from 1.01 to 1.4

– *Reason 2: Heavy-hitters.* In this point we analyze the heavy-hitters' behavior. A heavy-hitter is defined as a user that generates such a huge amount of data that this volume is representative with regard to the total traffic. In RedIRIS' traffic we have detected some heavy-hitter users. For instance, we detected that one IP address of one the RedIRIS' universities sent 122 GBytes to a certain external IP address with destination port 22 in only one day. However, this IP address did not connect again to this university in the following 10 days. Obviously, this heavy-hitter user's behavior had an important impact in the network statistics. It is worth remarking that this connection started 20 days after the capture process did; and in spite of this, such external IP address was the top-traffic producer until that moment, and only in one day. After this, this external IP address was ranking lower successively, until disappearing from the most popular IP addresses ranking.

Figure 3.10 shows the minimal number of different IP addresses that acceded to some of the 15 most popular port numbers as well as the Zipf's s parameter to the ports number and IP addresses. Note that some ports are only used by a very low number of users, such users are heavy-hitters. However, obviously, there are heavy-hitters that usually hit the most popular ports and, therefore, they are more difficult to detect. Interestingly, note that all those distributions whose Zipf parameter needs more time to become stable show a low number of minimal IP addresses (say less than 10). However there are some other university networks whose Zipf parameter becomes stable in a few days and show a low number of minimal IP addresses. This happens because this minimal number of IP address refers to the 15th or 14th most popular port traffic and the volume of exchanged traffic was not representative.

In this section we have tackled the following issues: (i) the universities under study show very different behavior regarding certain network metrics, and (ii) the slow convergence rate to stability of the process of the most popular IP addresses and port number. With regard to the first point the heavy-hitters make the groups under comparison small and not representative of the whole population; in this light, there is no reason to expect similar behavior between the networks. Figure 3.10 showed that these groups can be smaller than 10 users. Regarding the second point, we have shown that the heavy-hitters' behavior has impact on the

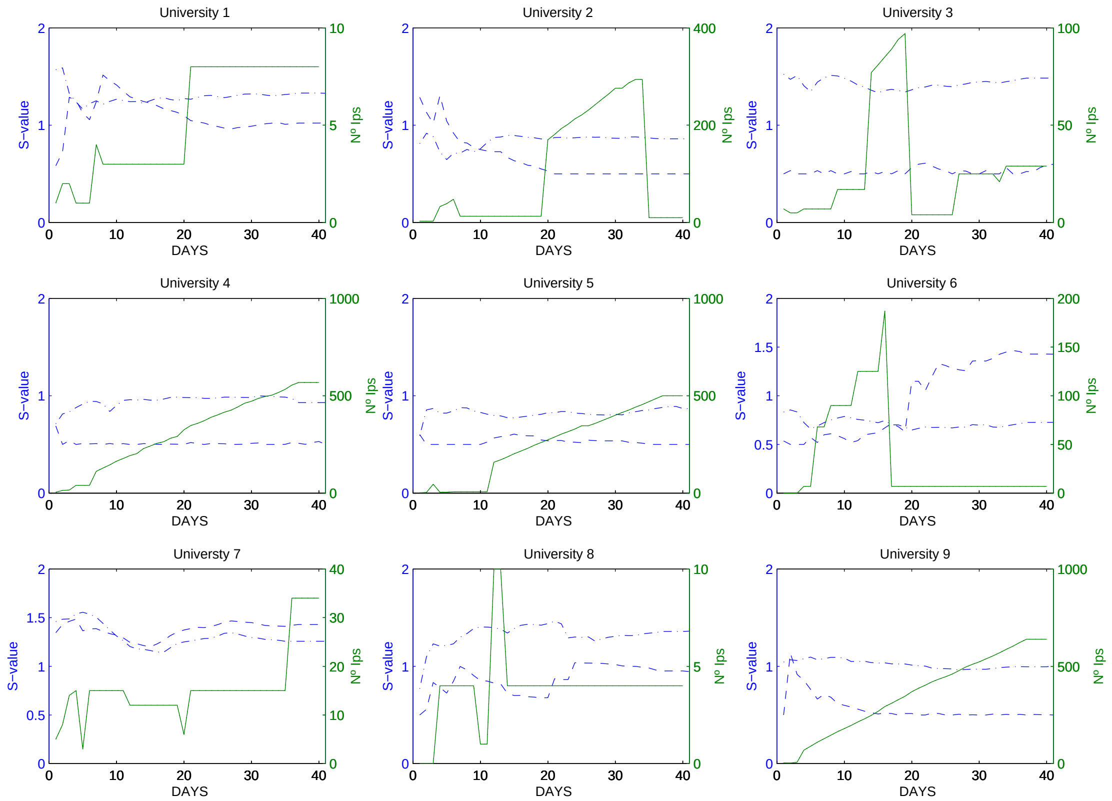

Figure 3.10: Minimal number of different IP addresses that access to some of the 15 most popular port numbers (solid lines) as well as the Zipf s parameter for the most popular ports number and IP addresses (dashed lines)

network metrics and that this behavior is not regular. Consequently, the whole university network metrics show this same behavior for days, until the aggregation level makes it impossible that a heavy-hitter enters into the most popular ranking.

3.6 Conclusions

This chapter provides a new point of view in the study of network measurements: the spatial analysis. Essentially, the spatial analysis aims to check whether or not the conclusions derived from the analysis of a given set of measurements gathered from a particular network scenario are valid to another but similar network scenario. The answer to this question is negative. Although a number of invariants have been identified to persist across different scenarios, our findings show that, when measurements from networks with similar intrinsic features are compared, the distribution of the most popular port numbers and IP addresses differ from one network to another.

Additionally, the experiments have shown that the distribution of the most popular IP addresses and port numbers experience high variability, and only reach some stability when long periods of measurements are considered, typically in the range of weeks. However, it is important to remark that, given the heterogeneous nature of the Internet and the fast and unpredictable way it changes, the results do not remain valid for long periods of time, thus requiring continuous monitoring and measuring, as noted in [FP01].

This involves two important consequences: Firstly, the duration of internetwork experiments must last until the measurements under study become stable, which involves a much longer traffic trace than usually believed; and, secondly, single-link measurements do not suffice for a meaningful analysis, hence a spatially diverse measurement experiment must be carried out. As a result, the required measurement infrastructure must be designed accordingly, and may involve sophisticated and costly equipment, both in terms of storage capabilities and number of probes.

Chapter 4

Analyzing the geolocation of the Internet connections

In this chapter we study the geographical location of end-hosts from connections sourced in similar Spanish campuses from RedIRIS. We remark that the geographical characterization of the Internet connections is of fundamental importance for the cost-optimization of peering agreements and it is important to estimate the most efficient intra-network route paths.

As done in the previous chapter, we specifically focus on the "temporal" and "spatial diversity". First, we obtain the distribution of connection destinations (per country) and fit it to a Zipf-Mandelbrot law. Our results show that the distribution of connection destinations has a very slow convergence rate to the domain of attraction of such distribution with the number of days added to the sample. More specifically, we conclude that traces of less than 30 days are not valid to detect the end-hosts geographical location pattern.

Then, factor analysis is performed to understand the relation between the response variable of averaged number of bytes, flows, and packets per day with dependent variables such as the source campus network, traffic direction and destination country. We also show that the distribution of geographical destinations is strongly dependent on the campus *network. Even though there are thousands of users in a typical campus network, it turns out that the aggregation level which is required to observe a stable end-hosts geographical location pattern is even larger. Therefore, the aggregation level to set up an efficient routing arrangement based on destinations is very large, possibly in the range of hundreds of thousands of users. This is far beyond the average customer population size of a small-to-medium Internet Service Provider.*

This chapter is organized as follows: Section 4.2 shows the preliminaries of this work, that is, we present the measurement set and briefly review the different options to geolocate IP addresses, together with the statistical methodology. Sections 4.3, 4.4 and 4.5 are devoted to results and discussion. Finally, Section 4.6 concludes this chapter with a summary of the main findings.

4.1 Introduction

The Internet research community has not underestimated the benefits of the geographical location analysis, and has documented its potential range of applications, such as traffic engineering [WGT95], peering agreements with other Internet Service Providers (ISP) [LS08, LMRT01] and location-aware applications [LBCM03]. We provide a different approach by looking at the connection pattern of similar user populations. Note that this analysis is not focused on the actual location of the destination hosts, which has received an extensive research effort (Section 4.2.4). More specifically, we perform a country-wide measurement campaign that comprises the whole Spanish academic network. Then, we focus on the per-campus destination pattern and focus on whether similar campus populations provide the same connection pattern or not.

This is particularly important for the routing and capacity planning arrangements of the ISP. In our case study, we make an analogy between campus network and ISP. Let us assume that a new ISP starts business with a population base which is similar to an existing population from a different ISP. Then, is the traffic destination pattern different, even though the customer population looks similar?. Note that if the destination pattern differs so does the routing strategy and possible peering agreements from the ISP.

In principle, being the per-campus population large (over 28,000 students and faculty/staff), one could assume that the overall connection pattern, when adding a new campus, would simply scale up by a given factor. This is a consequence of the large number of users, which basically drive the demand to the same limit distribution. Namely, the destination pattern should be similar if the number of aggregated users of the same type (students and faculty/staff) grows very large. However, our findings show that the distribution of exchanged traffic volume per country is far from being homogeneous. Interestingly, the latter result suggests that the type of traffic that campus networks exchange depends on the country in which the end-host is located, i.e., users tend to use certain services in certain countries.

The following issues are analyzed in this chapter:

1. We model the geographical location distribution of Internet end-hosts per country, with connections originated in Spain. Specifically, we show that such distribution can be effectively modeled by means of a *Zipf-Mandelbrot* function. This implies that a small set of countries accounts for the most of the traffic destinations. Actually, we perform a χ^2 goodness-of-fit test to validate such model.

2. We estimate the amount of measurement time which is required to reach stability in the model parameters. Namely, we address the issue of how long a network should be measured to infer a destination pattern and possibly set up a peering agreement.

3. Finally, we analyze the space diversity issue, i.e., whether similar campus networks produce similar traffic in terms of destinations. Namely, we would like to know to which extent there are invariants in destination patterns, that may serve to set up routing arrangements based on population groups. For example, traffic destination patterns should be similar for large Spanish campuses, due to the large aggregation level. Hence, we could predict a routing behavior on a per-campus basis and update the network accordingly when new users are added. However, we find that the destination patterns are indeed different.

Actually, Analysis of Variance (ANOVA) methodology is adopted to explain how the direction (incoming and outgoing), destination country and source network affect traffic destinations. The results show that the traffic destination per country heavily depends on the source campus network, despite of the large number of users. Furthermore, the distribution of traffic destinations does not seem to have a domain of attraction with the number of users in the source network. We note that the factor analysis has not been typically used by the research community to characterize and analyze Internet measurements.

This study is limited by the fact that we are considering academic users, who are different from residential users. However, we do not pursue the characterization of traffic destinations, but to which extent they are homogeneous if the user populations are similar. The methodologies presented in this chapter can also be applied to the case of residential networks, and provide valuable insight for a residential network operator.

4.2 Preliminaries

This section is devoted to describe the background of this study. First of all, we motivate the study of the geolocation of the Internet connections. Then, we present a description of the previous studies related to this one. Later, a description of the measurement capture process is presented. That is, we explain how we have modified the available data introduced in Section 3.2 in order to also include the geolocation information related with the measurements. Finally, we present a description of the main statistical techniques that have been applied in this study, namely Pearson's χ^2 test and ANOVA methodology.

4.2.1 Motivation

Regarding traffic engineering, it turns out that networks are typically connected at several exchange points that pass on the traffic as soon as possible, i.e., "hot potato routing" [NP08]. Such routing procedures make no geographical optimization whatsoever and lead to inefficient routing. Needless to say, knowledge of the destination patterns allows to perform a better routing, possibly decreasing the transit traffic in many intermediate networks [SPK02].

Second, we note that peering agreements are becoming increasingly popular [SS06]. Essentially, a peering agreement is a deal between two (or more) ISPs, such that both exchange traffic directly. On the other hand, a transit agreement is a deal between a local ISP and a transit ISP, such that the local ISP traffic to the rest of the Internet is carried through the transit ISP network. As a result of both kind of agreements, the infrastructures are shared and there is a cost reduction. Furthermore, latency can be also reduced if the most efficient path, in terms of transit ISPs, is selected. Therefore the traffic destination is a fundamental issue in the choice of peer and transit ISPs. Actually, ISPs use to serve certain geographical areas, most likely at the country level. Then, the choice of a peer ISP should be determined by the shortest path to the traffic destination [Nor01b, Nor01a]. However, the decision-making process for peering agreements is usually based on commercial reasons only, with less attention to the traffic volumes that the peering ISPs may exchange. This has caused well-known routing problems such as the path-inflation, i.e., end-to-end paths which are much longer than necessary [SMA03].

Third, knowledge of traffic destinations allows for a new class of location-aware applications that provide new functionalities to the Internet. As presented in [GZCF06] such functionalities include: targeted advertising on web pages adapted to the location where consumers live, content delivery control according to the local country policies and security features in electronic commerce and transaction verification.

4.2.2 Related work

Despite of the importance of factor analysis of traffic destinations, as shown in the previous sections, it turns out that the state of the art does not feature any similar study. We believe that such lack of research effort is due to the difficulties in capturing traffic from many geographically disperse source IP subnetworks. The authors in [AW97], more than ten years ago, presented a detailed workload characterization study of Internet web servers, on attempts to find invariants in the Internet behavior (as defined in [FP01], that is, characteristics that are common across an extensive set of networks for a significant period of time). One of those characteristics was the geographical distribution of document requests to several web servers. However, they only considered two possible options, whether the requests were local or remote to the web-server network, finding that most part of the requests were remote. In our case, we discriminate the traffic destination per country and do not restrict the analysis to the web service only. In [FCFW05] the authors analyze the traffic received by a certain online-game server. Among other characteristics, the players' location is included in the study. Their results indicate a clear geographical dispersion with only 30% of the clients placed close to the online-game server. However, the authors are exclusively focus on the online game service, i.e., the rest of the traffic is not considered and there is no comparison on a per-source IP subnetwork basis.

4.2.3 Measurement set description

In this chapter, we have used the Netflow records that RedIRIS is providing us with from each of its POPs as explained in Section 3.2. In this case, for the study of the geolocation of the Internet connections we have adapted the processing subsystem to include geolocation information. That is, we have upgraded the Netflow records with the country that the destination IP address belongs to. This geolocation information is obtained applying the methodology described in the next section.

As was defined in Section 3.2.2, in what follows let us denote incoming direction to the traffic that is destined to one of the universities. Conversely, the outgoing direction refers to traffic sourced from one of the universities.

To make this information more manageable, we have computed daily aggregates of the number of bytes, flows and packets (and their corresponding percentages over the total of the day), in {university, country, direction} triples. That is, for each day, we obtain the number of bytes, flows and packets and their corresponding percentages per source campus network to each country, i.e., for the outgoing direction (from campus to the rest of the Internet). We also obtain the same information for traffic sourced in each country and destined to each campus network, i.e., for the incoming direction (from the rest of the Internet to campus). In what follows, we use the term "measured items" to refer to bytes, flows or packets. The measurement set entries are presented in Table 4.1.

Following the methodology presented in the previous chapter, we have carefully

Field	Description
Source	University network/Country
Destination	Country/University network
Direction	Outgoing/Incoming
Bytes	Total number of bytes transferred from the source to the destination that day
Flows	Total number of flows transferred from the source to the destination that day
Packets	Total number of packets transferred from the source to the destination that day
Percentage of bytes	Percentage of the bytes transferred from the source to the destination that day
Percentage of flows	Percentage of flows transferred from the source to the destination that day
Percentage of packets	Percentage of packets transferred from the source to the destination that day

Table 4.1: Geolocation Measurement set summary

selected 12 universities out of the total set, for which the intrinsic network features, such as population size, bandwidth capacity, ratio students-staff, filtering policies (basically Peer-to-Peer (P2P) applications, etc.) are very much alike.

It is worth noting that Network Address Translation (NAT) capabilities and Content Distribution Networks (CDN) or proxies have no impact on our measurements. NAT groups the traffic of several different hosts in a single public IP address but this has no influence in the geographical location of hosts, neither remote nor local. Consequently, by relaxing the NAT constrain, we have extended the number of campus networks under study. In addition, we have refined the selected networks to comprise an even more similar set of networks. Thus, we have removed U_3 and U_8 from the set of universities and we have included information from other five ones. Table 4.2 provides some useful information about the selected universities, which are labeled according to previous chapter as U_1,

U_2, U_4, U_5, U_6, U_7, U_9, U_{10}, U_{11}, …, U_{14}.

It is worth remarking that local proxies inside the campus network do not have any influence in the results. They are accounted for as local hosts, that concentrate traffic. However, remote proxies will be accounted for as end-hosts because we have no means to locate the real end-host. Nevertheless, from the local ISP standpoint, this is traffic to an external host (the proxy) and accounts for in the peering agreement, the same way that non-proxy traffic does.

Finally, the sampling rate for Netflow records is the same throughout the measured routers, namely 1:100. We believe that the sampling error affects all measured campuses the same way and has no influence in our obtained percentages. Anyway, such sampling effect can be considered negligible for our analysis as shown in [MCS+06] (Section 3.2.3).

4.2.4 IP geolocation Methodology

There are several ways to find the physical location of an IP address. The most straightforward approach is to use a name resolver and make a Domain Name System (DNS) reverse query. Then, the address location is guessed by parsing the retrieved name. As a result, we can only obtain the location of the IP addresses whose domain location is known beforehand or that contains some localization information in the name itself. To circumvent these issues, we have used the geolocation database approach. We note that there are free and commercial versions of such databases. The differences between them are the accuracy of the geolocation and the number of different IP addresses that can be geolocated. In this study we have use the free version of the *GeoIP Country* database of *MaxMind*, i.e., *GeoLite Country* [Max]. This database has an accuracy of 99.5% and it is updated monthly. Actually, the accuracy of these databases has been studied and reported [GUF07, SGU08] and it seems to be adequate for our purposes. The *GeoIP Country* database has entries for the country code, country name and continent data. Recently, and keeping pace with the accuracy needed by the location-aware applications, there have been attempts to increase such accuracy [PS01, GZCF06]. We have discarded these methods because the database approach is simpler and it provides enough accuracy for our purposes (i.e., group destinations per country). For a better understanding of geolocation procedures, the reader is referred

University	Population	Ratio students/staff	No. of Flows (millions)
U_1	40,000	9.6	1.5
U_2	31,000	10.9	1.3
U_4	31,500	11.0	3.0
U_5	31,000	10.8	1.4
U_6	36,000	11.2	3.5
U_7	33,500	12.2	5.7
U_9	28,000	11.7	2.8
U_{10}	55,000	11.2	5.0
U_{11}	38,000	8	4.7
U_{12}	38,000	11.1	4.0
U_{13}	46,000	8.8	1.9
U_{14}	38,500	8.6	2.0

University	Capacity access	P2P filtering	Unique IP addresses acceded
U_1	1 Gb/s	No	23,000
U_2	1 Gb/s	No	24,000
U_4	1 Gb/s	No	66,000
U_5	1 Gb/s	No	22,800
U_6	1 Gb/s	No	58,000
U_7	1 Gb/s	No	90,000
U_9	1 Gb/s	No	30,000
U_{10}	1 Gb/s	No	70,000
U_{11}	1 Gb/s	No	55,000
U_{12}	1 Gb/s	No	66,000
U_{13}	1 Gb/s	No	40,000
U_{14}	1 Gb/s	No	27,000

Table 4.2: User-base population size, average number of flows collected per day, networks' bandwidth capacity, filtering policies and average Internet IP addresses acceded during busy hour per day for all universities under study (January 2009)

to [CK06, Section 5.3.6] and references therein.

4.2.5 Statistical Methodologies

In this section, we introduce the statistical techniques applied in this chapter and the following one. First, we present goodness-of-fit techniques that allow us to model the Internet end-hosts distribution and to assess the spatial diversity of the measurements and the temporal aggregation needed for stability. Second, we give a brief introduction to the ANOVA methodology, which allows us to measure the impact that factors such as the source network, country and direction have on the response variable, in this case the measured items (flows, packets, bytes).

Goodness-of-fit techniques

In order to find a suitable model for the traffic destinations, we perform visual inspection first (Section 4.3). This visualization can only give us some insight on the shape of the distribution, and this is not sufficient for hypothesis testing. To this end, we adopt a goodness-of-fit technique over a hypothesized distribution. In our case, the hypothesized distribution is a Zipf-Mandelbrot distribution [KV07, RB89] and the goodness-of-fit test is well-known χ^2 test [DS86, Chapter 3].

The underlying idea of Person's χ^2 test was to reduce the problem of goodness of fit of a general distribution to the testing fit to a multinomial distribution by dividing the support of the distribution in cells and comparing the observed values with the expected ones in each cell. In this way, to test if a random sample $X_1, X_2, \cdots, X_n$, with $n \geq 30$, follows a specified distribution with parameter $\Theta \in \mathbb{R}^r$, $F(x; \Theta)$ (null hypothesis), one starts by partitioning the support of the observations into M buckets, namely $C_1, C_2, \cdots, C_M$, where $M \geq 5$. Let us denote by O_i the number of observations that lie in the bucket C_i, then under the null hypothesis O_i follows a binomial distribution with parameters n and p_i, that is:

$$O_i \sim \mathscr{B}(n; p_i) \tag{4.1}$$

with p_i being the probability of choosing the bucket C_i under the null hypothesis

$$p_i = \int_{C_i} dF(x, \Theta) \tag{4.2}$$

It is recommended that approximately the same number of observations lie in each bucket, and also that each bucket has at least 3 observations. With these assumptions, it is reasonable to think that $O_i - E_i$, being E_i the expected number of observations in cell C_i ($E_i = np_i$), is a good measure of distance between the observed data and its theoretical distribution. It can be shown [DS86, Chapter 3] that under the null hypothesis, the statistic

$$\chi^2 = \sum_{i=1}^{M} \frac{(O_i - E_i)^2}{E_i} \tag{4.3}$$

follows a χ^2 distribution. The number of degrees of freedom depends on the number of constraints placed on the data. If the parameter Θ is known, then χ^2 has $d = M - 1$ degrees of freedom. In contrast, if Θ is estimated from the observations, the degrees of freedom of χ^2 are $d = M - 1 - r$. The estimation of the parameter Θ should be performed by Maximum Likelihood (ML) for the former assumptions to be true. Finally, the null hypothesis of goodness of fit is rejected if χ^2 has a large value, e.g. if $\chi^2 > \chi^2_{d,\alpha}$, being $\chi^2_{d,\alpha}$ the percentile (1-α) of a χ^2 distribution with d degrees of freedom.

Analysis of Variance

ANOVA is a widely used statistical methodology whereby the observed variance of a given response or dependent variable is split into explanatory factors. ANOVA provides a way to determine if such factors have any importance in explaining the variability of a response variable, and to which extent. ANOVA performs a contrast using the ratio between the adjusted sum of squares of samples that belong to each factor level, i.e., intra-level samples, and the total, inter-level samples. Such ratio is shown to follow a Snedecor's F distribution, provided that the samples fulfill the following hypothesis: First, they must be necessarily independent; second, they must be fairly Gaussian; and third, all of them must share the same intra-level variance (i.e., exhibit homoscedasticity). However, the results of ANOVA are generally accepted provided that the number of elements in each group are similar (Balanced ANOVA), and there is a non-excessive deviation from the homoscedasticity assumption [OA84, GPS72].

The null hypothesis supports the homogeneity of means within each factor. Basically, it contrasts, according to a given pre-defined significance level α (typically $\alpha = 0.05$), whether or not the intra-level variance values can be explained due to the randomness of measurements (generally, experimental errors) and not to differences in the population when grouped by categories (or levels). If so, the null hypothesis cannot be rejected and the factor used to build the groups is statistically non-significant. Otherwise, the factor explains enough variance and it is considered as significant.

Following this, the simplest ANOVA univariate model for a response variable y with an only significant factor α is given by:

$$y_{iu} = \mu + \alpha_i + \epsilon_{iu} \tag{4.4}$$

where y_{iu} represents the u^{th} observation on the i^{th} level ($i = 1, 2, \ldots, I$ levels), μ represents the overall mean response (or intercept). On the other hand, α_i refers to the effect due to the i^{th} level of factor α and ϵ_{iu} is the deviation, random or experimental error, in the u^{th} sample on the i^{th} level. We also note that $\sum_{i=1}^{I} \alpha_i = 0$.

The resulting model in case of two significant factor is:

$$y_{iju} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{iju} \tag{4.5}$$

and so forth in case of more than two factors. In this latter case, α_i and β_j represent the effect due to the i^{th} and j^{th} levels of factors α and β respectively. Similarly, $(\alpha\beta)_{ij}$ represents the interactions between i^{th} level of factor α and j^{th} level of factor β. Finally, ϵ_{iju} represents the deviation in u^{th} sample to the overall mean of the samples within i^{th} level of factor α and j^{th} level of factor β. Again, note that $\sum_{i=1}^{I} \alpha_i = 0$ and $\sum_{j=1}^{J} \beta_j = 0$ being J the total number of levels of factor β. The reader is referred to [DC74, Jai91] for further details on the ANOVA methodology.

4.3 Visual inspection

Following the common practice in data analysis, we first provide a visual inspection of the main descriptive statistics. We order the destination countries by descending value of the measured item. Then, replacing the name of the country by its rank in the ordered list and plotting the corresponding percentage of the measured item we observed a power law model (for example, see Figure 4.1(a) and Figure 4.2(a) for campus U_{10} and U_{11}). This observation is confirmed by the log-log plots of the same data (shown in Figure 4.1(b) and Figure 4.2(b)) where the values follow a straight line. However, the first value in the rank tend to look abnormal in the complete set of networks under study. This first ranked country is Spain for almost all campus networks under study, as expected. If we remove Spain from the former figures, the data shows a better fit to a power law distribution (figures 4.1(c) and 4.1(d), and figures 4.2(c) and 4.2(d)). Thus, we decided to fit the power law model to the data removing the first ranked country.

Concerning population aggregates, Figure 4.3 shows the top 15 countries for the aggregate of the networks under study in both directions after removing both Spain and USA for the three measured items for the sake of clarity since both countries account for the most of the traffic.

In this study, the majority of the bytes, around 40%, are sent and received within Spain. The United States comes in second place with 20% of the sent and received bytes, which is also expected because of web traffic. In the third place we find some of the most important countries of the European Union such as United Kingdom, Germany, France, to mention a few, see Figure 4.3(a). They account for a range between 2.5 and 6% of the total amount of bytes per country. In fourth place we find Latin American countries such as Argentina, Chile, etc., accounting for a range between 0.5 and 1.5% of the total share. These are Spanish-speaking countries and redirections to Web pages in Latin America are usual. Also, there are many researchers from such countries visiting Spanish universities. Finally, we find that there is traffic going to or from nearly all countries although their percentages of the total are negligible. In conclusion, a visual inspection of the data provides reasonable results.

In order to further inspect the data set, we mapped each country with gray intensities according to the value of the measured items (Figure 4.4). We have

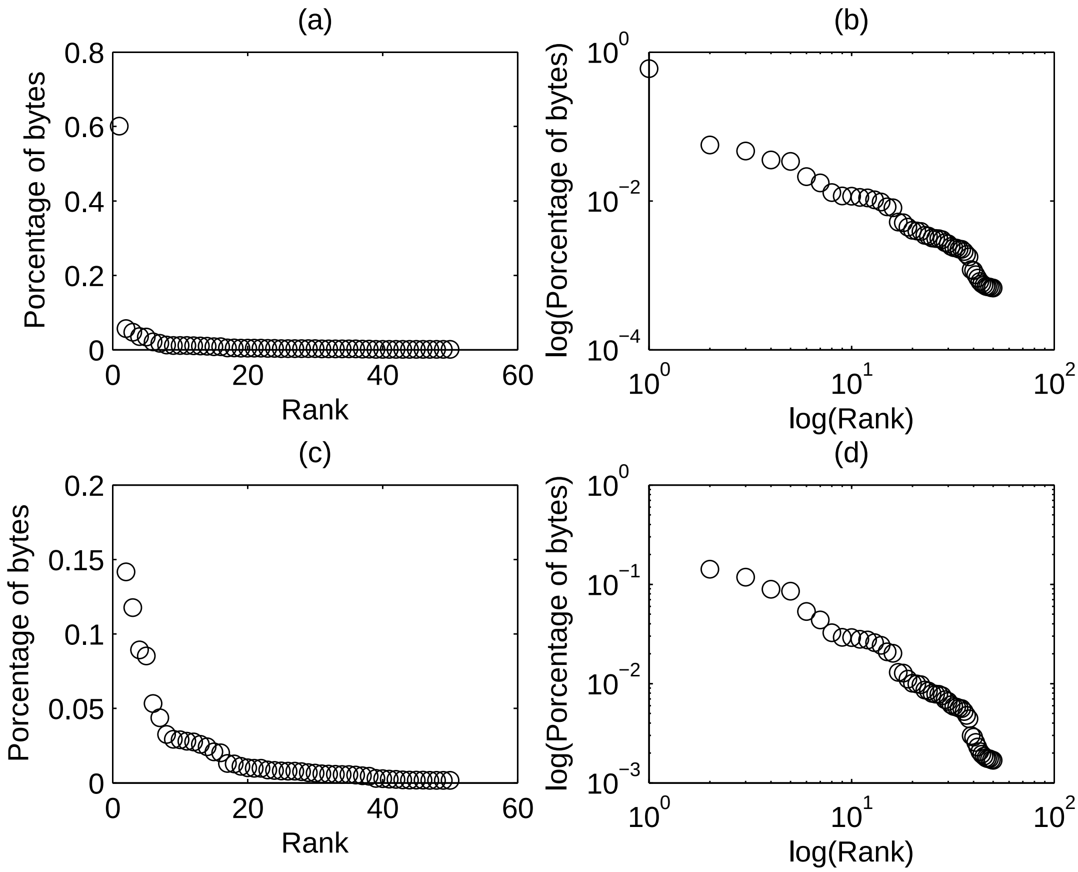

Figure 4.1: Visualization of the top 50 ranked countries for a day worth of measurements in the outgoing direction of U_{10}: (a) Percentage of bytes vs. rank. (b) Percentage of bytes vs. log of the rank. (c) Percentage of bytes vs. rank without the first ranked country. (d) log of the percentage of bytes vs. log of the rank without the first ranked country

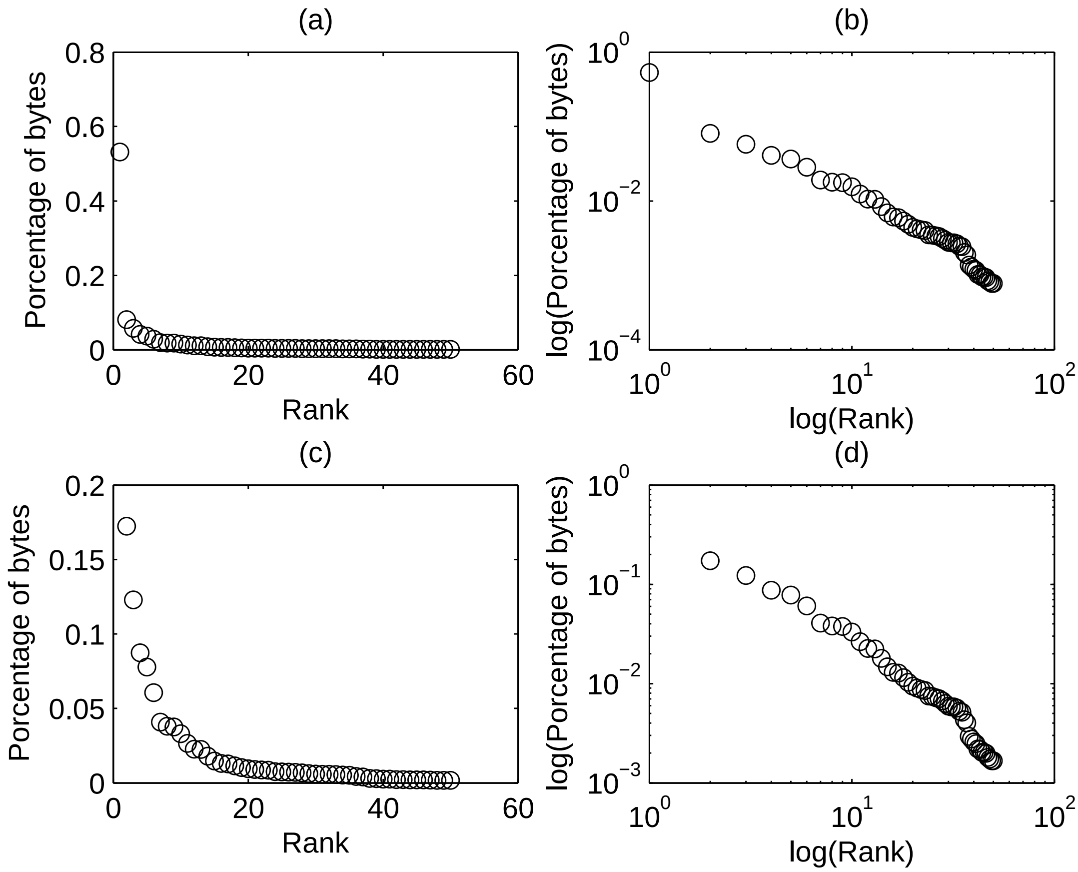

Figure 4.2: Visualization of the top 50 ranked countries for a day worth of measurements in the outgoing direction of U_{11}: (a) Percentage of bytes vs. rank. (b) Percentage of bytes vs. log of the rank. (c) Percentage of bytes vs. rank without the first ranked country. (d) log of the percentage of bytes vs. log of the rank without the first ranked country

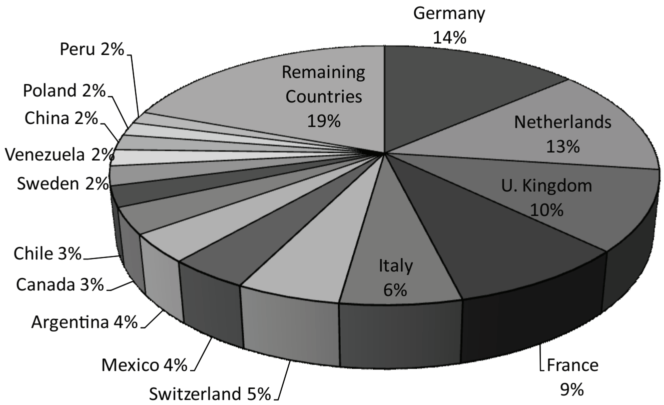

(a) Bytes

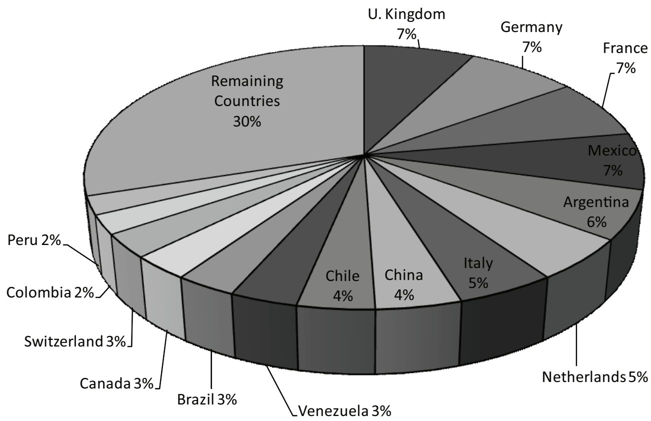

(b) Flows

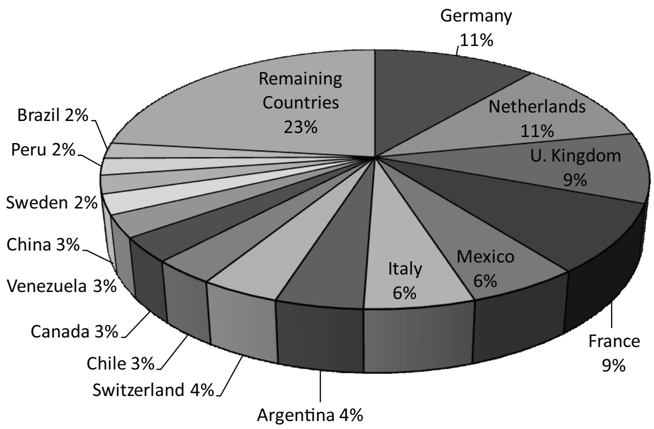

(c) Packets

Figure 4.3: Percentages of traffic sent and received from top-15-contributing countries excluding Spain and USA

excluded Spain and US for the reasons stated above and also those countries that account for less than $10^{-3}\%$ of the measured item. To draw the maps, we used the Google's Visualization API [1] that can be used directly as a gadget from Google docs [2].

4.4 On the characterization of end-hosts location

The visualization of the dataset presented in figures 4.1 and 4.2 evidences the need of a power law distribution to model the measurements. The most popular power law distribution with discrete support is the Zipf's law, whereby the probability mass function of the element whose rank is k, z_k, is proportional to an inverse power a of k, i.e.,

$$z_k = \frac{q}{k^a}, \tag{4.6}$$

where $a > 1$ and q is a normalization constant [JKK05]. Although this distribution has been widely used in Internet studies (including the study carried out in the previous chapter), this time we have chosen the Zipf-Mandelbrot (ZM) distribution, which is a generalization of the Zipf's law. The ZM distribution has three parameters instead of two, and shows better or equivalent performance in terms of goodness-of-fit.

The ZM Probability Mass Function (PMF) p_k is given by

$$p_k = \frac{c}{(k+b)^a}, \tag{4.7}$$

where $a > 0$, $b > -1$ and c is a normalization constant which is not necessarily equal to q.

When we used standard Zipf's law in the previous chapter, we showed that the goodness-of-fit between data and standard Zipf's law was correct. The explanation of this is that b parameter was equal to 0. In such a case, Eq. (4.6) and Eq. (4.7) become the same and, consequently, ZM and standard Zipf's law are equivalent.

[1] http://code.google.com/intl/es-ES/apis/visualization/documentation/gallery/intensitymap.html
[2] http://docs.google.com/

(a) Bytes

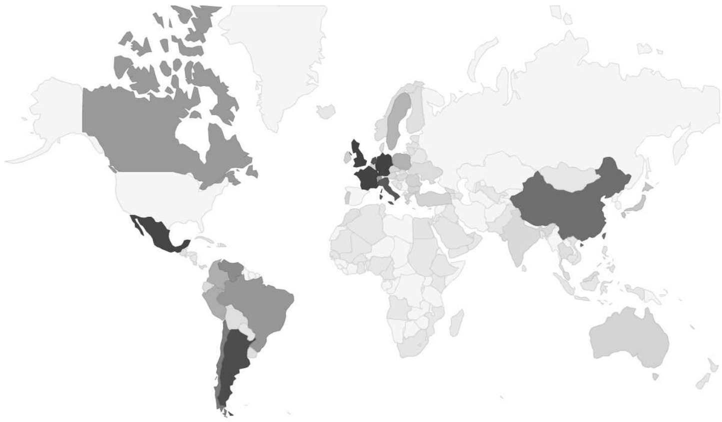

(b) Flows

(c) Packets

Figure 4.4: Percentages of traffic sent and received from top-15-contributing countries in graduate gray scale (excluding Spain and USA)

The Maximum Likelihood Estimation (MLE) procedure for the ZM distribution finds the parameters a and b that maximize the likelihood function for a random sample X of size n and it is given by

$$l(X; a, b) = \frac{n!}{n_1! n_2! \ldots n_N} \prod_{k=1}^{N} \left(\frac{c}{(k+b)^a} \right)^{n_k}, \tag{4.8}$$

where n_k is the number of instances of the element in the k^{th} order. The numerical optimization of this function is a very challenging task, and several procedures have been studied to compute the multinomial coefficients involved in the likelihood function $l(X; a, b)$ in a precise and fast way. One option to circumvent the computation of the multinomial coefficients is presented in [Izs06], whereby coefficients are obtained through the probability mass function of a Binomial distribution:

$$l(X; a, b) = \prod_{i=1}^{N-1} B_{n^j}^{p_{a,b,j}/p^j_{a,b}}(n_j), \tag{4.9}$$

being

$$B_s^t(r) = \binom{s}{r} t^r (1-t)^{s-r}. \tag{4.10}$$

In [Izs06] the calculation procedure for $p_{a,b,j}$ and $p^j_{a,b}$ is described, and this method can be easily implemented in a mathematical software package like Matlab, where the coefficients of Eq. (4.10) are optimally computed.

After obtaining the MLE parameters, we applied the χ^2 test in order to measure to which extent our model fits the data. As the ZM distribution is a discrete distribution, the buckets in the χ^2 test are defined by such discrete support. Even though it is recommended to have all buckets filled up with the same number of observations, this is unfeasible with our dataset, due to its power-law nature. However, we merged buckets with small number of samples in the tail of the distribution on attempts to have all the buckets with at least 5 samples on them.

We pursue a twofold objective in our analysis. On the one hand, we would like to assess the validity of the ZM distribution to model the data. Furthermore, we wish to find the smallest period of time such that the ZM parameters remain stable. It is worth noting that this stability check also provides hints about the trace length which is required to obtain meaningful results.

4.4.1 Goodness-of-fit tests

Table 4.3 shows the results of the χ^2 test for the 12 campus networks for a period of 90 consecutive days between December of 2008 and March of 2009. Similar results were obtained with the other two measured items and are not presented here for the sake of clarity.

The first column shows the (anonymized) university name as described in Section 4.2.3. The second column shows the direction of the traffic, either incoming or outgoing the campus network. The accuracy in the third column is defined as the percentage of days in the sample for which the χ^2 test null hypothesis of goodness-of-fit cannot be rejected at the significance level $\alpha = 0.05$. Finally, the last column shows the average p-value from all the performed χ^2 tests. We show this average only for those pairs university-direction where the accuracy was 100%, and it gives an estimate on how good the goodness-of-fit was, the larger the better. As can be seen in the table, except for a small number of university-direction pairs, the null hypothesis of goodness of fit cannot be rejected for a 50% of days or more, which supports our initial ZM distribution assumption. Remarkably, in most cases the ZM distribution fits the measurements better in the outgoing direction than in the incoming direction. We hypothesize that it can be due to the asymmetry of the Internet applications and services. A further analysis of this issue is performed using factor analysis (see Section 4.5).

To assess the stability of the estimated parameters, we form a time series of aggregated days and measure the relative error in the parameters for all the universities. The relative error $re_p(t)$ for a time series $p(t), t = 1, \dots, N$ is defined as follows:

$$re_p(t) = \frac{p(t+1) - p(t)}{p(t)}, \qquad t = 1, \dots, N-1. \tag{4.11}$$

In our case, t stands for the number of days used in the estimation and $p(t)$ is the estimated parameter a, b, c in Eq. (4.7), for all the university-direction pairs showed in Table 4.3 using Eq. (4.11). Figure 4.5 shows the evolution of the relative error in time of a and b parameters (note that c is function of a and b), for U_{10}

University	Direction	Accuracy(%)	Mean p-value
U_1	Outgoing	100	0.5917
	Incoming	42.17	–
U_2	Outgoing	50.57	–
	Incoming	62.50	–
U_4	Outgoing	68.29	–
	Incoming	100	0.5332
U_5	Outgoing	100	0.7304
	Incoming	49.12	–
U_6	Outgoing	100	0.8780
	Incoming	57.53	–
U_7	Outgoing	100	0.8452
	Incoming	100	0.8778
U_9	Outgoing	98.84	–
	Incoming	45.88	–
U_{10}	Outgoing	100	0.9974
	Incoming	44.71	–
U_{11}	Outgoing	100	0.8562
	Incoming	60.23	–
U_{12}	Outgoing	0	–
	Incoming	100	0.7847
U_{13}	Outgoing	100	0.9864
	Incoming	17.44	–
U_{14}	Outgoing	100	0.7956
	Incoming	8.05	–

Table 4.3: Results of the goodness of fit tests

and U_{11} networks.

As can be seen in Figure 4.5, at least one month worth of aggregated data are necessary to make the parameter estimation stable, i.e., to make the relative error be close to zero. This implies that the parameter estimation does not change if we add one more day worth of data. According to it, Table 4.4 shows the stable parameter values for all the universities under study. It turns out that the parameter values differ from one university to another, even though some of them are similar. Consequently, we find differences between campus networks, which, in principle, are similar in terms of population, access bandwidth, etc. This is the motivation for the factor analysis presented in the next section, which takes into account source network, direction and destination country.

Finally, Figure 4.6 shows the percentage of connections per destination country together with the fitted ZM distribution for U_{10} and U_{11}, which shows remarkable goodness-of-fit.

4.5 Factor analysis

In this section we apply ANOVA to our measurement set. The aim of such analysis is to assess the impact that the traffic direction (both incoming and outgoing), the campus network under study and the country in which the end-host is located have in the response variable. In this study the response variable refers to the three measured items introduced in Section 4.2.3 specifically to the percentages of bytes, packets and flows. Note that we could also perform ANOVA analysis using the absolute value of the measured items instead of percentages. However, this analysis provides misleading results because small differences between the campus networks have a large influence in the response variable. For instance, the traffic load in absolute terms within the set of universities is not identical even though the universities are similar. To avoid this overshadowing effect, we choose traffic percentages.

ANOVA allow us to assess the impact of a specific country (i.e., knowing exactly which country is) instead of using only the rank order as in the previous analysis –Zipf-Mandelbrot characterization–. This permits to contrast:

– Whether or not the set of countries under study generate equivalent volumes

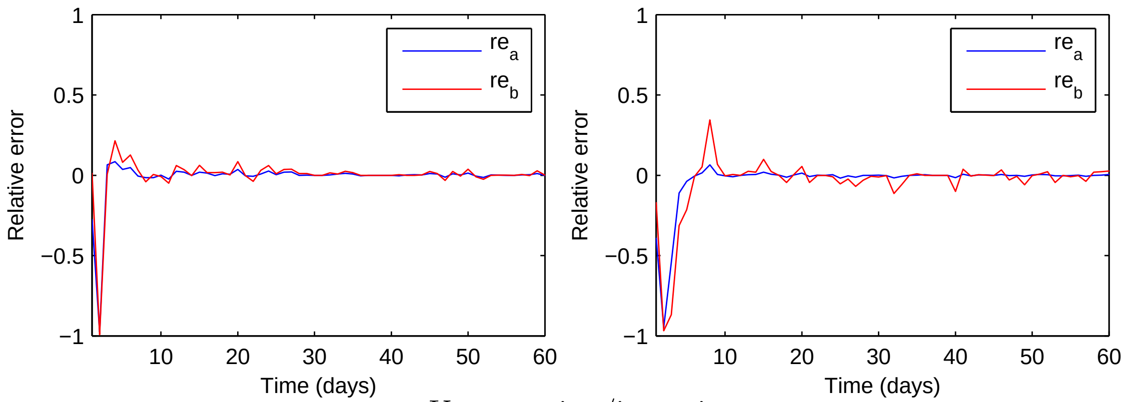

(a) U_{10} outgoing/incoming

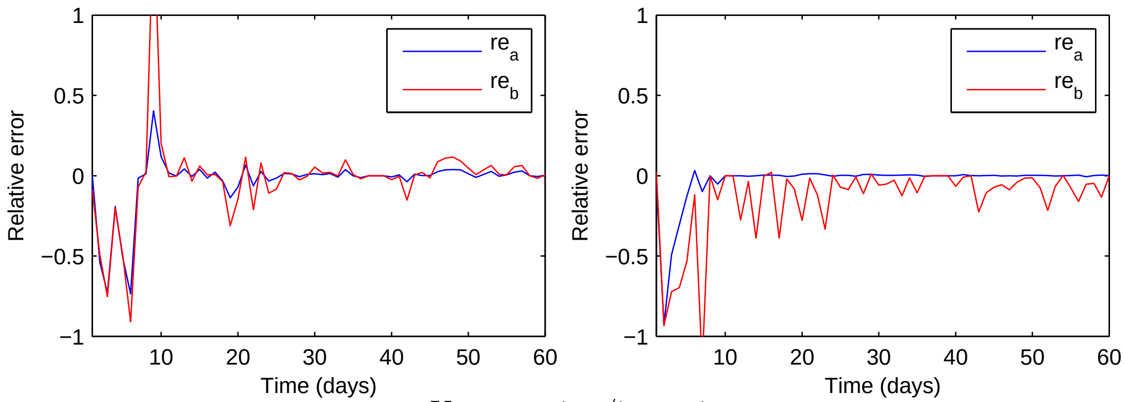

(b) U_{11} outgoing/incoming

Figure 4.5: Relative error for two university-direction pairs for a and b parameters

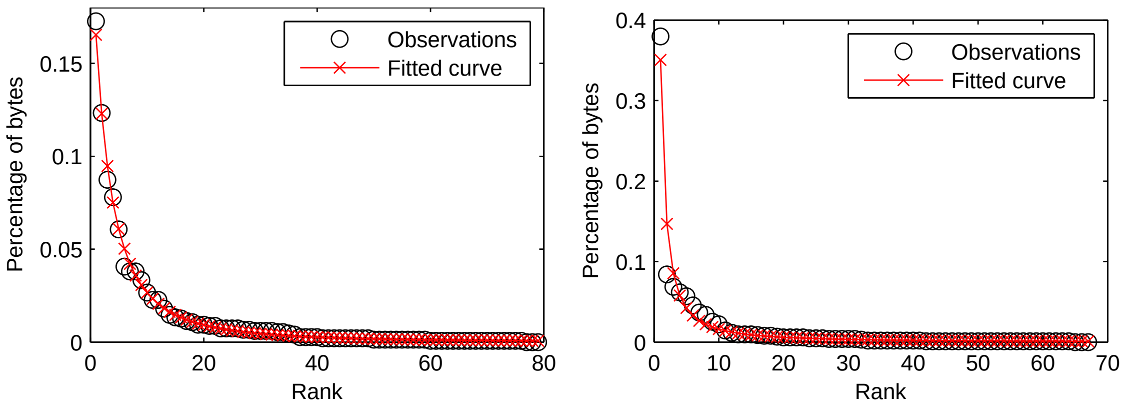

(a) U_{10} outgoing/incoming

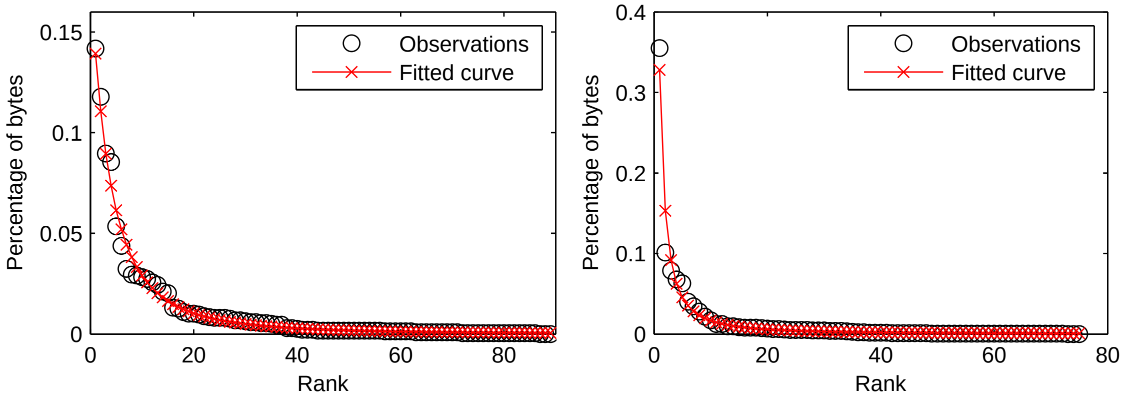

(b) U_{11} outgoing/incoming

Figure 4.6: Examples of empirical versus theoretical distribution for percentage of destinations

University	Direction	Parameter estimate		
		a	b	c
U_1	Outgoing	1.2553	0.0042	0.3254
	Incoming	2.9853	2.6242	19.6172
U_2	Outgoing	1.6252	1.3404	1.0195
	Incoming	1.5122	-0.0038	0.4232
U_4	Outgoing	1.4459	0.7031	0.5940
	Incoming	2.1608	1.1988	2.2645
U_5	Outgoing	1.9891	4.9630	5.6725
	Incoming	1.2407	-0.5353	0.1904
U_6	Outgoing	1.5514	2.1080	1.1068
	Incoming	2.5244	2.4997	8.3640
U_7	Outgoing	1.0050	-0.0036	0.2153
	Incoming	1.2082	-0.1054	0.2949
U_9	Outgoing	2.0804	4.4342	6.3758
	Incoming	1.4100	-0.0118	0.3807
U_{10}	Outgoing	2.1886	5.9179	11.3857
	Incoming	1.4450	0.2046	0.4593
U_{11}	Outgoing	2.4687	9.0764	41.9909
	Incoming	1.5899	0.6253	0.7102
U_{12}	Outgoing	326.71	7.2285	$1.12 \cdot 10^{299}$
	Incoming	1.7949	0.5868	0.9308
U_{13}	Outgoing	2.5076	7.0915	33.1939
	Incoming	1.4256	-0.0003	0.3944
U_{14}	Outgoing	1.7853	3.3648	2.5157
	Incoming	1.7101	0.0004	0.5100

Table 4.4: Results of the maximum likelihood parameter estimation

of traffic.

– Whether or not the set of universities under study connects to the same locations.

– Whether or not the ratio incoming/outgoing traffic is similar from one campus to another and between the countries.

– Whether or not universities connect to the same countries with similar ratio incoming/outgoing traffic.

Consequently, we define three fixed factors and their corresponding interactions (full factorial ANOVA): *Network*, that is the source university network, *Country* that represents the country in which the end-host is placed and *Direction*, either incoming or outgoing traffic. For instance, Figure 4.1 shows the percentages of traffic in bytes (response variable) that U_{10} (factor *Network*) exchanges with *top* 50 contributing countries (factor *Country*) in outgoing direction (factor *Direction*) for a day worth of data. Thus, according to Eq. (4.5) we have the following initial model:

$$
\begin{aligned}
y_{ijku} = \mu &+ Network_i + Country_j + Direction_k \\
&+ (Network\&Country)_{ij} \\
&+ (Network\&Direction)_{ik} \\
&+ (Country\&Direction)_{jk} \\
&+ (Network\&Country\&Direction)_{ijk} \\
&+ \epsilon_{ijku}
\end{aligned}
\tag{4.12}
$$

where y represents any of the measured items.

In the previous section it was shown that at least 30 days worth of data are necessary to obtain stability in the measurements under study. In this light, the ANOVA sample spans the entire month of January 2009. Regarding the number of countries, for the sake of clarity, only the *top* 30 contributing countries in terms of number of bytes were taken into account, although the same analysis can be performed with a larger set. As a result, we have a data base for both directions, involving twelve networks, thirty countries and thirty days, i.e., more than 20,000 samples for each measured item.

4.5.1 ANOVA assumptions

Regarding the ANOVA assumptions introduced in Section 4.2.5, Figure 4.7 shows the autocorrelation function (dots) along with its 95% confidence intervals (solid lines), as described by the Bartlett test [CL66], applied to the averaged number of bytes in both directions for two campus network (U_{10} and U_{11}) and *Spain* as factor. It becomes apparent that the samples are not correlated. It is worth noticing that all the levels showed similar results.

Figure 4.8 shows the Quantile-Quantile plot [DS86, Chapter 2] diagram for the same set of samples. In a QQ-plot, the order statistics of the empirical sample are depicted as a function of the percentiles of the other distribution, in the case the Gaussian distribution. If the data follows such distribution then it nearly fits to a straight line. In general, we have not found evidences of significant deviation from the gaussianity in the measurement set. Conversely, the homoscedasticity hypothesis was rejected by means of the Levene test. However, a non-significant deviation from the homoscedasticity assumption [GPS72] can be accepted in case of Balanced ANOVA with large number of samples, which is the case of our experimental design.

4.5.2 Effect of Network, Country, and Direction factors in the traffic

Table 4.5 shows the results of the ANOVA test with the percentage of bytes per day as the response variable. According to the results, the null hypothesis that supports the homogeneity of means cannot be rejected for the factors *Network*, *Direction* and *Network&Direction*, but it is rejected for *Country*, *Network&Country*, *Country&Direction* and *Network&Country&Direction* at the significance level $\alpha = 0.05$. This gives the following final model according to

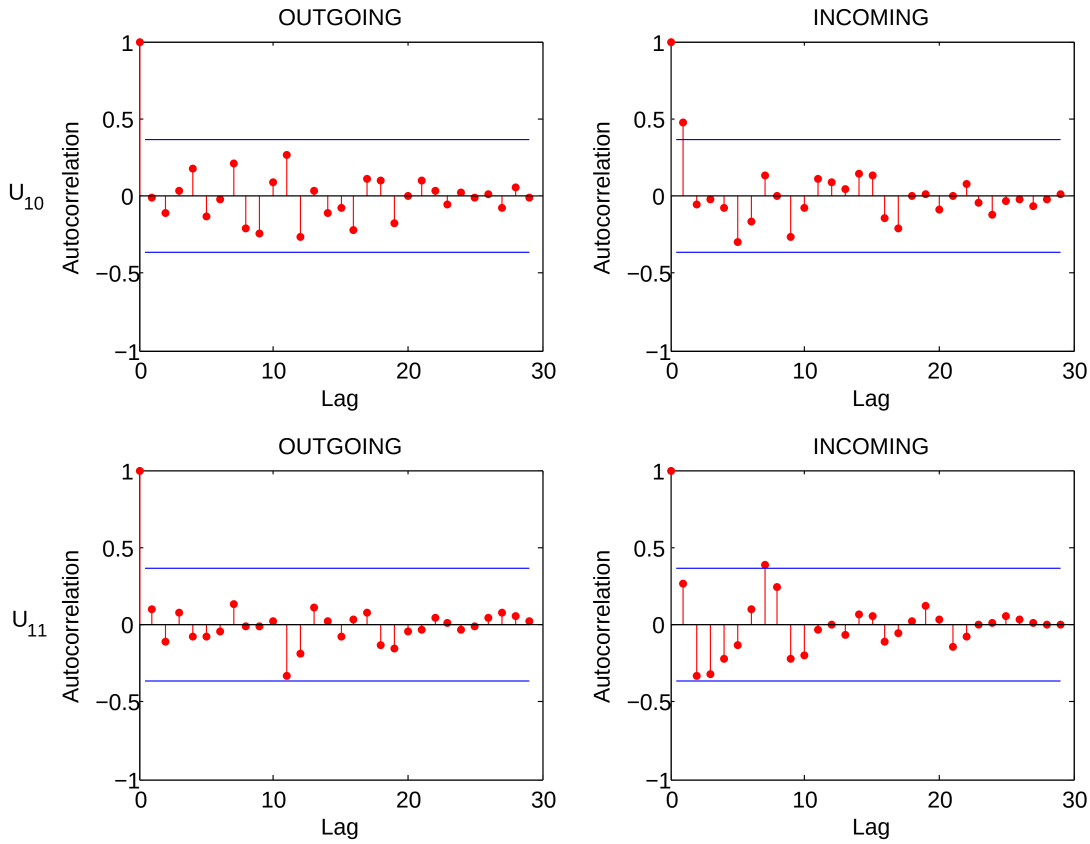

Figure 4.7: Autocorrelation function (dots) and 95%-confidence intervals (solid lines) applied to the averaged number of bytes in both directions with U_{10} (top) and U_{11} (bottom) as factor *Network* and *Spain* as factors *Country*

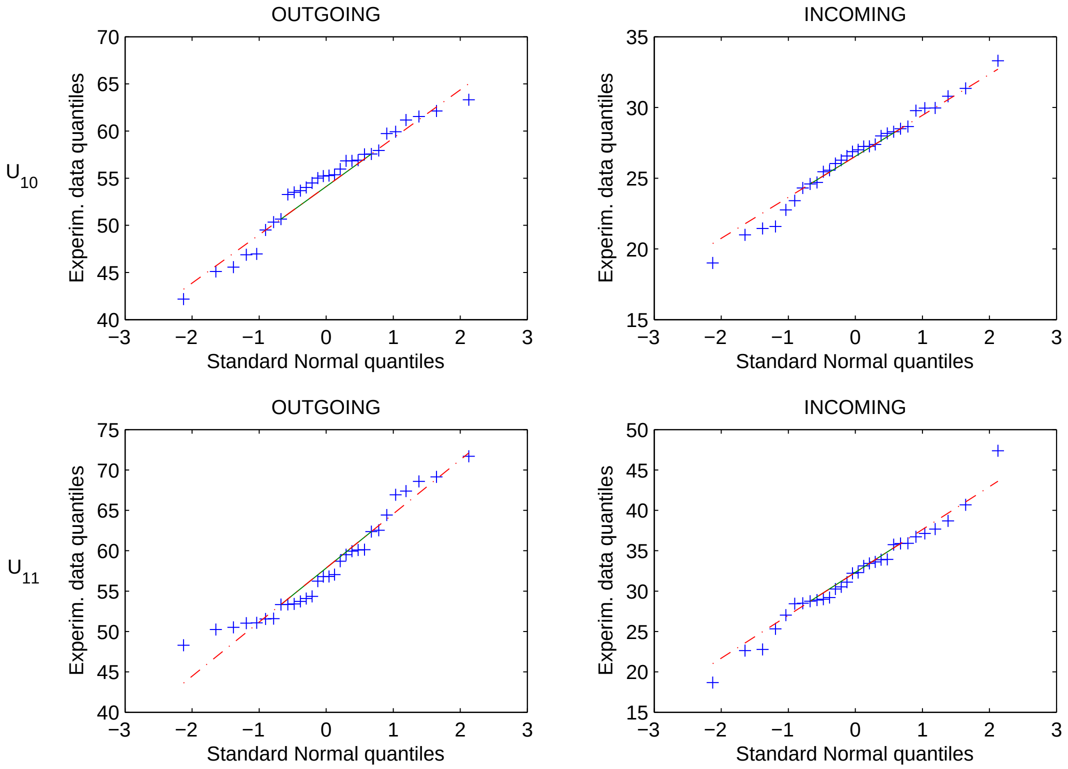

Figure 4.8: QQ-plot diagram of the averaged number of bytes in both directions with U_{10} (top) and U_{11} (bottom) as factor *Network* and *Spain* as factor *Country*

Eq. (4.12):

$$
\begin{aligned}
y_{ijku} = \mu &+ Country_j \\
&+ (Network\&Country)_{ij} \\
&+ (Country\&Direction)_{jk} \\
&+ (Network\&Country\&Direction)_{ijk} \\
&+ \epsilon_{ijku}
\end{aligned}
\tag{4.13}
$$

with $i = 1, 2, 3, \ldots, I = 12$ (number of university networks), $j = 1, 2, 3, \ldots,$ $J = 30$ (number of countries) and $k = 1, 2$.

Several conclusions can be drawn from these results. The homogeneity of means when taking into account factor $Network$ implies that the traffic generated by the campus networks, ignoring destination and direction, has no influence in the measured items. Intuitively, this means that the campus networks has a percentage of traffic per country which is similar, in both directions. However, the countries are not the same per campus network. This is confirmed by the fact that the $Network\&Country$ is clearly significant. Similarly, the fact that factors $Direction$ and $Network\&Direction$ are not significant indicates that the traffic percentages are distributed similarly in both directions, regardless of the network.

The results of the factor $Country$ show strong significance, which ties in with what we expected taking into account the results of the previous section and other works (for instance, [GNS05]). Basically, the location of the end-hosts is far from being homogeneous. In addition, the results of the factor $Country\&Direction$ shows strong significance, which implies that the relation incoming/outgoing traffic depends on the destination country under analysis. This is directly related to the peering agreement decision-making problem introduced in the section 4.1 and 4.2. Actually, one of the most typical peering agreements is the ratio-based paid peering [Nor01b], in which peering is free of charge until traffic asymmetry reaches a certain ratio, commonly 4:1. With the ANOVA results we have shown that such ratio depends heavily on the destination country.

In addition, it is well known that the ratio incoming/outgoing traffic is a good discriminant to differentiate traffic applications. For instance, the ratio of the Hypertext Transfer Protocol (HTTP) is usually low, i.e., more downloaded traffic that uploaded, whereas the ratio of the P2P applications tends to be higher [LFHL07]. Bearing this in mind, ANOVA suggests that users tend to access to certain services in certain countries. Finally, the interaction factor of level 3 $Network\&Country\&Direction$, reveals that not only the countries receive a different sort of traffic but it also depends on the network that generated such traffic. Once more, this is closely related to ratio-based paid peering, since the ratio incoming/outgoing traffic depends both on the country and on the source network. That implies that single-network measurements do not suffice for a meaningful characterization of the distribution of the end-hosts' location, which supports the results presented in Table 4.4.

The results for flows and packets as response variables, tables 4.6 and 4.7, are similar to the previous analysis for bytes. Regarding flows, the factor $Network$ has some significance which indicates that the behavior of the source networks is different even if we ignore the direction or the specific destination country. Regarding the number of packets, the fact that the complete interaction factor is moderate significant (that is, p-value is closer to 0.05) means that the number of packets from any source network per direction to the set of countries is more homogeneous. In any case, the differences are minimal and, regardless of the response variable selected, the conclusions drawn for percentage of bytes remain valid for the other two measured items.

Finally, the tables 4.5, 4.6 and 4.7 also show the (adjusted) coefficient of determination $\bar{R}^2$. It represents a measure of the percentage of variation in the response variable that can be explained by the factors. As $\bar{R}^2$ is close to 1, we can conclude that the factors and their interactions model the measured items distribution accurately.

4.6 Summary and conclusions

In this chapter, we have performed a factor analysis of the Internet end-hosts location, from connections originated in an extensive set of campus networks for a long period of time. The analysis has been performed using Netflow records from the whole Spanish National Research and Education Network.

First of all, we have visualized the results of the geolocation process. This visu-

Dependent variable: Averaged number of bytes

Source	Sum of Squares	df	Mean Square	F	p-value
Network	34.771	11	3.161	0.856	0.584
Country	1454882.713	29	50168.369	13579.707	0.000
Direction	4.909	1	4.909	1.329	0.249
Network & Country	49871.331	319	156.336	42.318	0.000
Network & Direction	11.134	11	1.012	0.274	0.991
Country & Direction	237180.336	29	8178.632	2213.814	0.000
Network & Country & Direction	18286.591	319	57.325	15.517	0.000
Error	77138.303	20880	3.694		
Corrected Total	1837410.089	21599			

Ajusted R^2=0.957

Table 4.5: ANOVA table with *Network*, *Country*, *Direction* and their interactions as fixed factors and average number of bytes as response variable

Dependent variable: Averaged number of flows

Source	Sum of Squares	df	Mean Square	F	p-value
Network	33.089	11	3.008	2.410	0.005
Country	1366731.310	29	47128.666	37755.110	0.000
Direction	0.269	1	0.269	0.215	0.643
Network & Country	61227.702	319	191.936	153.762	0.000
Network & Direction	0.833	11	0.076	0.061	1.000
Country & Direction	1835.568	29	63.295	50.706	0.000
Network & Country & Direction	1049.127	319	3.289	2.635	0.000
Error	26063.930	20880	1.248		
Corrected Total	1456941.828	21599			

Ajusted R^2=0.981

Table 4.6: ANOVA table with *Network*, *Country*, *Direction* and their interactions as fixed factors and average number of flows as response variable

Dependent variable: Averaged number of packets

Source	Sum of Squares	df	Mean Square	F	p-value
Network	40.578	11	3.689	1.286	0.225
Country	1422841.384	29	49063.496	17105.976	0.000
Direction	0.009	1	0.009	0.003	0.956
Network & Country	69959.561	319	219.309	76.462	0.000
Network & Direction	0.856	11	0.078	0.027	1.000
Country & Direction	11580.101	29	399.314	139.221	0.000
Network & Country & Direction	1050.601	319	3.293	1.148	0.036
Error	59888.183	20880	2.868		
Corrected Total	1565361.272	21599			

Ajusted R^2=0.960

Table 4.7: ANOVA table with *Network*, *Country*, *Direction* and their interactions as fixed factors and average number of packets as response variable

alization evidences that the location of end-hosts per country follows a power law distribution. To confirm this hypothesis, we aggregated traffic from several universities for a month and repeated the visualization process. With the aggregated data, we confirmed the power law shape of the measurements but the behavior was not the same per source campus network, even though the aggregation level was very high. Such observations motivated us to perform two different analysis of the data: a Zipf-Mandelbrot characterization of the measurements and a factor analysis to explain the impact of the network, destination country, and traffic direction in our measured items.

In the Zipf-Mandelbrot characterization, we have modeled the traffic volume according to its destination country, concluding that a small set of countries accounts for the most part of the traffic. In addition, we have shown that distribution's parameters a and b need at least one month to be considered stationary, meaning that measurement campaigns should be long enough to be meaningful.

Moreover, the characterization process evidenced that all the networks have different values for the parameters of the distribution. In this light, it becomes necessary to collect the data all across the network and not just from selected measurement points. However, the ZM characterization pays no attention to the countries themselves, but it only takes into account their position in the rank, which calls for subsequent factor analysis.

As factor analysis, we have applied the ANOVA univariate methodology to assess the amount of variance of connection destinations that can be explained in terms of three factors, namely the traffic direction, the source campus network and the destination country. The results show that the factor *Country* is strongly significant, as well as its interactions with *Network* and *Direction* ones. The former issue shows that the distribution of the traffic volume that a network exchanges per country is far from being be homogeneous. Interestingly, the latter result suggests that the sort of traffic, probably at application level, that the campus networks exchange depends on the country in which the end-host is located, i.e., users tend to use certain services in certain countries. Similarly, we conclude that the most popular countries in terms of exchanged traffic differ from one network to another, although all networks were objectively similar and the population size significant. Additionally, we have not found differences in the results when measuring the average number bytes, packets or flows.

Our conclusions have direct application to the ISP capacity planning and routing strategies. It has been shown that there is no canonical model for traffic destinations even though the source user populations are similar. Therefore, the routing policies and peering agreements that may be good for an ISP may not apply to another ISP that serves a similar user population. As an example, let us compare university networks U_5, U_2, and U_9 during January 2009. From Table 4.2 we note that they are very much alike. However, the destination patterns are far from being similar. Just to mention some examples, more than 5% (in bytes) of the U_5 outgoing traffic is destined to Mexico, whereas this country represents less than 0.5% in the two other university networks. Similarly, more than 10% (in bytes) of the U_2 incoming traffic comes from Germany, this amount is three times as less in the other universities. Furthermore, a 30% of the bytes that U_5 and U_2 networks sent were destined to USA. However USA accounts for 50% of the U_9 outgoing traffic. There are a number of similar examples. Thus, our findings show that serving new populations which in principle look similar leads to dramatic changes in the connection destinations and may call for a totally different peering arrangements. Again, we explain this fact by the heavy-hitter phenomenon as explained in Section 3.5.

While the characterization of the Internet traffic is worthwhile by itself, a number of interesting applications have been pointed out. Our conclusions have direct application to the decision-making process for new transit and peering agreements. ISPs should give priority to those ISPs that maximize the geolocation of their traffic, thus mitigating the path inflation phenomenon and, consequently, producing a latency decrease. This is a fundamental aspect that ISPs should tackle given the advent of new multimedia services that call for more demanding Quality of Service (QoS) requirements. Specially, ISPs should pay attention on the end-hosts location since the Zipf-model characterization has shown that a small set of countries accounts for the most of the traffic.

From a methodology point of view, we have also shown that the length of the measurement campaign needed to obtain a significant characterization of the end-host locations can involve a long period of time. Note that similar conclusion was found in the previous chapter when we analyzed the popularity of the port numbers and IP addresses.

Similarly, from the ANOVA results we have learned that the granularity of such characterization should be very narrow because each network connects to different countries in a different way and the conclusions drawn for one network cannot be extrapolated to the other ones. Note again that this results tie in with the results provided in the previous chapter. Consequently, the ISPs' measurement campaigns should include an extensive set of networks to cope with the space diversity and also encompass a significant period of time due to the large transient time.

Chapter 5

The "queueing equivalent" thresholding method

In the development of accurate capacity planning and network resource dimensioning models, network operators must handle representative information about the traffic volumes traversing its network. However, the amount of traffic measurements available over which to perform such analysis, processing and storage is overwhelming. For this reason, the research community has understood the importance of finding an effective mechanism to reduce (or subsample) such huge amount of data, with minimum loss of information.

In this chapter, we propose a mechanism to downsample traffic time-series using Multiresolution Analysis (MRA) with wavelets, and evaluate the optimal subsampling level based on comparing the queueing behavior of the subsampled and original signals at the output of a router.

The chapter considers the traffic volume traversing a given router per unit of time, which can be obtained by periodically polling the counters of the interfaces table via Simple Network Management Protocol (SNMP). Specifically, we analyze Multi Router Traffic Grapher logs from RedIRIS as traffic time-series. This mechanism is more related to network performance than conventional comparison levels, since queueing delay is a very representative Quality of Service metric. On the one *hand, downsampling allows reducing the storage and processing requirements for the measured time-series. On the other hand, it also allows the reduction of the SNMP polling rate.*

The results show that it is possible to reduce the data to one fourth of its original size for the traffic generated by most analyzed universities, and even to one eighth for data collected from routers with more aggregated traffic, both with a high level of confidence.

This chapter is organized as follows: Section 5.1 presents the introduction and related work of this chapter. Section 5.2 briefly reviews MRA and introduces the queueing equivalent thresholding method. Section 5.3 shows the results of applying the queueing equivalent mechanism to the set of measurements. Finally, Section 5.4 concludes this chapter with the most important findings obtained.

5.1 Introduction and related work

Internet Service Providers (ISP) and network operators have traditionally found challenging the task of network dimensioning and capacity planning for two reasons mainly: the dynamics of network traffic which exhibit high variance and long-term correlation, and secondly the overwhelming amount of traffic measurements over which to perform any model parameters estimation. In this chapter we address this latter problem.

Data collection is often performed by means of polling network devices. For example, Simple Network Management Protocol (SNMP)-enabled devices can be polled in order to obtain the value of a given MIB variable. Precisely, a number of network monitoring tools, such as the Multi Router Traffic Grapher (MRTG) (Section 2.2.2), Cricket [All99] or Cacti [Nag05] among others, perform polling of the *ifInOctets* and *ifOutOctets* counters of the interfaces MIB [MK00].

Essentially, the amount of available traffic measurements is often excessive to be handled efficiently, and the task of reducing its size is becoming more and more important. However, the difficulty lies in the choice of the appropriate timescale in which to represent the data. Clearly, aggregation reduces not only the amount of data but also the available information. In this light, there are few studies concerning this matter: the optimal timescale over which traffic measurements should be taken for analysis, monitoring or storage, among other tasks.

For instance, the Request for Comment (RFC) 1857 [Lam95] recommends to poll every minute on attempts to detect peak behavior. Nevertheless, given that the volume of data generated at such rate might be huge, such RFC proposes to aggregate the data, either using the arithmetic mean or the maximum value of the interval. Finally, this RFC proposes, as a rule of thumb, to use 15-minute aggregates for 24-hour periods, 1-hour aggregates for 1-month periods and 1-day aggregates for periods of years. For instance, the 15-minute aggregates can be justified because the polling intervals should be taken small enough to capture variations in human activity. According to such RFC, 30 minutes is taken as a good estimate of the time at which people remain in one activity. Hence, to track variations in this interval it is necessary to sample twice, which gives us a sample every 15 minutes.

On the other hand, most ISPs and companies collect their router statistics every 5 minutes. This value is compliant with the International Telecommunication Union (ITU)-T recommendation Y.1540 [ITU02], which states that such interval is consistent with practical limits for IP layer operations. These statistics are then stored in a database. In order to limit the database size, only the most recent data is stored in the original timescale, while the older data is aggregated in larger timescales. The aggregation rules are somehow arbitrary and eventually result in the loss of valuable monitoring information.

In this chapter, we have proposed techniques for data aggregation, which are based on a trade-off between storage efficiency and information loss, and avoiding arbitrary rules. It is desired to define a mechanism to subsample the huge amount of traffic measurements available for study, and find a reduced-but-equivalent downsampled signal over which the tasks of analysis, monitoring (using the typical network monitoring tools), storage, and modeling becomes more tractable.

One possible architecture to achieve this goal is shown in Figure 5.1. First the data is captured and stored. Concurrently, a *"thinning" process* is applied to reduce the time-series, and to make it easier for the subsequent analysis process. For example, there is a growing interest in the use of embedding techniques applied to Internet traffic analysis [LZSS06], which could be a possible analysis process in

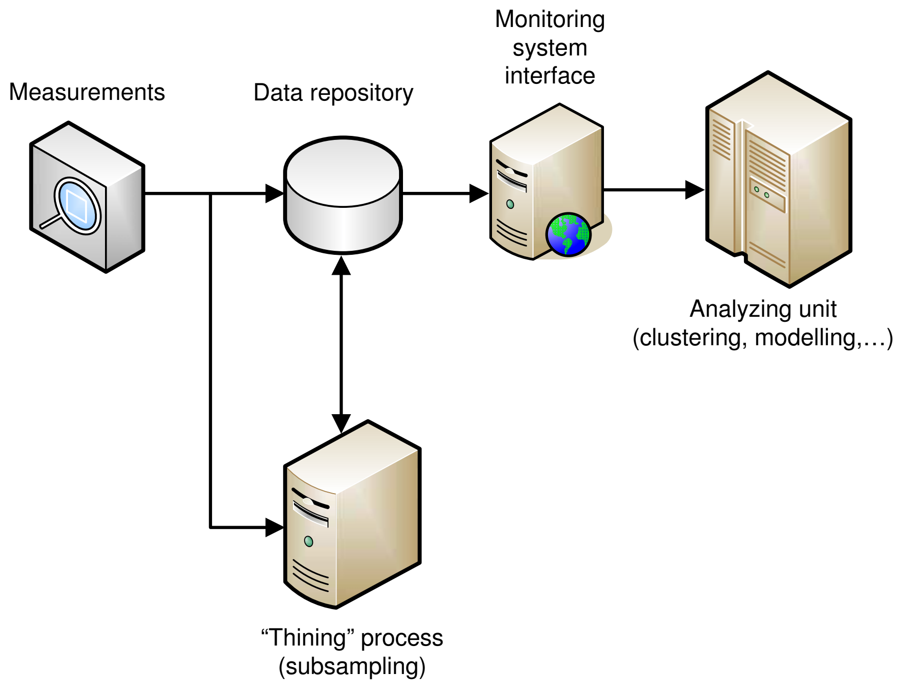

Figure 5.1: Proposed architecture to subsample measurements

the diagram presented in figure 5.1. Such embedding techniques can be used to find similarities between Internet traffic sources or flows with a number of applications in security and capacity planning. In order to perform embedding or clustering it is very convenient to have the time-series size reduced to a minimum, while keeping the original information nearly intact. Clearly, it is easier to perform embedding or clustering if the number of points in the time-series is relatively small. Otherwise, there is a high variability in the data, which makes it more difficult to identify clusters.

Such a *"thinning" process* implies the following two tasks:

1. To propose a mechanism for subsampling the data collected in a way that keeps most of its original information.

2. To define a strategy to decide when to stop the subsampling process based on some sort of equivalence metric.

In the former, Multiresolution Analysis (MRA) with wavelets has been extensively used for time-series analysis, particularly network traffic signals. The main advantage of MRA is its capability to analyze non-stationary time-series, which are common in traffic monitoring comprising day and night periods, but at same time it keeps the information of the series in the frequency domain. Thus, it has a wider applicability than conventional frequency-domain (i.e., Fourier) techniques. Note that if the signal is transformed only into the frequency domain, we lose all information about time, what is not desirable when the final aim is to deal with traffic measurements. Consequently, MRA shows the best results when the signal presents a strong periodicity but the time-resolution information is required. This is the case for most of the internetwork traffic measurements because they are related to the properties of human activity (Section 2.3).

Additionally, MRA has been used to detect traffic anomalies via the isolation of significant increases in local variance [BKPR02]. Also, MRA has proved fundamental in the characterization of network traffic at different aggregation scales, with application in the synthesis of long-range dependent traffic [RRCB00], and further to estimate its characteristic Hurst parameter [AV98].

Concerning the use of wavelets for downsampling traffic series, the work by Papagiannaki et al. [PTZD05] proposes a subsampling strategy based on the 12- and 24-hour periodicity trends found in the measurements. Indeed, the traffic patterns are observed to repeat from day to day and in morning-evening cycles; this phenomenon, known as daily traffic pattern, is close-related to the human behavior patterns (already introduced in Section 2.3). In this light, they choose to take one sample every $\lambda = 90$ minutes, since 12 and 24 hours constitute $2^3 \cdot \lambda$ mins., and $2^4 \cdot \lambda$ mins. respectively.

Once the subsampling strategy is defined, the original and the subsampled signals are compared based on a pre-defined distance metric, typically Euclidean [CFY03] or via the Analysis of Variance (Section 4.2.5). However, with distance metrics, there is no assessment on the quality of the subsampled signal concerning network performance-related information and only second-order moment information is captured but other statistics can be lost, especially in terms of network response.

In this light, since the time-series represents volumes of traffic, it makes sense to determine the difference between the original and subsampled signals based on some sort of network-performance metric, rather than mathematical distance. For this reason, we propose a rather more network-related approach, which consists of feeding a given queueing system with both the original and subsampled signals, and contrast (using a goodness-of-fit test) the two signals based on the queue-occupancy distribution function of the router, which is key in the performance evaluation of networks, as pointed out in [RRCB00].

5.2 "Queueing equivalent" thresholding method

In this section we provide a brief review of MRA. For a thorough explanation and further details the reader is referred to [Chu92]. Then, we introduce our "queueing equivalent" thresholding method.

5.2.1 Multiesolution Analysis review

Let $\{\omega_{j,k} = \sqrt{2^j}\omega(2^j t - k)\}$ for all j, k denote an orthogonal basis for $\mathcal{L}^2$, where $\omega(t)$ spans V_0 (the reference subspace) and $\mathcal{L}^2$ is the space of the functions with finite energy. Let us also consider the subspaces $\{0\} \subset \ldots \subset V_{-2} \subset V_{-1} \subset V_0 \subset V_1 \subset V_2 \subset \ldots \subset \mathcal{L}^2$, and let W_j refer to the orthogonal complement of V_{j+1} with respect to V_j, namely $V_{j+1} = W_j \oplus V_j$. In MRA, a given signal $x(t) \in V_0$ is decomposed into the sum of an approximation signal $A_j(t) \in V_{-j}$ and a set of details $D_j(t) \in W_{-j}$ as follows [Mal89]:

$$\begin{aligned} x(t) &= A_1(t) + D_1(t) \\ &= A_2(t) + D_2(t) + D_1(t) \\ &= A_3(t) + D_3(t) + D_2(t) + D_1(t) \\ &= \ldots \end{aligned} \tag{5.1}$$

where the $A_i(t)$, $D_j(t)$ are the projection of $x(t)$ over the subspaces V_{-i}, W_{-j} respectively. Note that the $A_i(t)$ signals are approximations of $x(t)$ in a larger timescale, i.e., in the timescale 2^i times the timescale of the original subspace V_0.

The aim of MRA is to obtain an adequate approximation for the signal, namely to find the subspace V_i in which the original signal can be projected with minimum information loss. If the original timescale for V_0 is 5 minutes, then the timescale

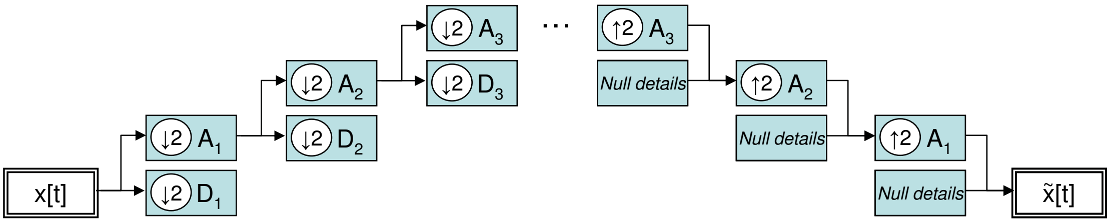

Figure 5.2: Wavelet filter banks. In each step the incoming signal goes through an analysis linear filter and the Approximation signal (A_i) and Detail signal (D_i) are obtained. On the right, the reverse process is shown: $\tilde{x}(t)$ is calculated applying the reconstruction filters to A_i (using null signals as D_i)

for V_i is $2^i \cdot 5$ minutes. The approximations and details are obtained with the analysis filter banks. Basically, the signal goes through an analysis linear filter in order to obtain each successive approximation, as shown in Figure 5.2.

In summary, MRA provides a computationally efficient method for the approximation of a time-series in larger timescales. In what follows, we assume that the original time-series belongs to V_0, namely V_0 is the subspace that groups all the signals in the same timescale of the original time-series. This is only for the sake of notation simplicity. Note that the choice of the V_0 timescale is arbitrary.

5.2.2 “Queueing equivalent” analysis

The goal is to approximate $x(t)$ by the largest timescale approximation $A_i(t)$ such that the information loss, in terms of network behavior, is still acceptable. For example, an approximation of $x(t)$ in the subspace V_1 consists of $\tilde{x}(t) = A_1(t)+0_1(t)$ where $0_1(t)$ is the zero of W_1. In V_2, this is $\tilde{x}(t) = A_2(t)+0_2(t)+0_1(t)$, etc.

As we approximate $x(t)$ by its projection over the subspaces $V_{-1}, V_{-2}, ...$ some information about the signal is lost, since the timescale is larger. In fact, wavelet shrinkage acts as an smoothing operator since it obtains a signal approximation with fewer points. More specifically, if $x(t)$ is a signal of finite length with N points, then $A_1(t)$ has $N/2$ points approximately, and $A_j(t)$ is an approximation of $x(t)$ with length $N/2^j$ points. Our objective is to find the largest timescale approximation which is accurate enough for a given analysis.

Clearly, the appropriate timescale for a given approximation depends on the application of the MRA. For example, if we simply wish to detect an "average" value of the signal, then we may choose to approximate in a very large timescale. The timescale is usually selected by thresholding the energy of the details. However, this is a squared error criterion which is not specifically tailored to any network-related application of the MRA. Moreover, the energy threshold is a heuristic value.

In this thesis we propose and validate an approximation method which relates to queueing performance. Intuitively, a given signal $x(t)$ and approximation $\tilde{x}(t)$ are said to be "queueing equivalent" if an infinite-buffer queue fed with both processes produces the "same" (or very similar) queueing occupancy distribution. If this is the case, then we may take the approximation $\tilde{x}(t)$ instead of $x(t)$ for whatever queueing-related analysis we wish to perform.

Concerning other applications, as mentioned in section 5.1, clustering and embedding applications may benefit from the fact that the time-series length is reduced after applying MRA. If, for instance, we wish to perform a r-dimensional clustering of $x(t)$ with other traffic time-series, then we can take the approximation signal $A_j(t)$ instead of $x(t)$, since $A_j(t)$ has fewer points. This makes the clustering algorithm converge faster.

More formally, let us consider an infinite-buffer single-server system which is governed by the Lindley's equation [Lin52]:

$$Q(t+1) = \max\{0, Q(t) + A(t) - C\}, \quad t = 0, 1, 2, \dots \tag{5.2}$$

where $Q(t)$ is the system occupancy at time epoch t, $A(t)$ are the bytes arriving during such time interval, and C is the router capacity. Let F_A denote the system occupancy distribution under traffic input $A(t)$. The following provides the definition of "queueing equivalent" approximation:

Definition: The signal $x(t)$ and the approximation $A_j(t)$ are equivalent (in the queueing performance sense) for a utilization factor ρ and significance level α if and only if the null hypothesis of goodness-of-fit between F_{A_j} and F_x can be accepted at significance level α. Notation-wise, we say that $x(t)\mathcal{R}_{\rho,\alpha}A_j(t)$.

Remark: Note that $\mathcal{R}_{\rho,\alpha}$ is a binary relationship but not an equivalence relationship in $V_0 \times V_0$. Clearly, $x(t)\mathcal{R}_{\rho,\alpha}x(t)$ and if $x(t)\mathcal{R}_{\rho,\alpha}y(t)$ then $y(t)\mathcal{R}_{\rho,\alpha}x(t)$.

However, the transitive property does not hold. For example, let us consider the Kolmogorov-Smirnov statistic [DS86, Chapter 4] $z(x(t), \tilde{x}(t)) = \max_{\tau \geq 0} |F_x(\tau) - F_{\tilde{x}}(\tau)|$. Then, for any $y(t)$ and $z(t)$ such that $x(t) \neq y(t)$, $y(t) \neq z(t)$ and $x(t) \neq z(t)$,

$$\begin{aligned} z(x(t), z(t)) &= \max_{\tau \geq 0} |F_x(\tau) - F_y(\tau) + F_y(\tau) - F_z(\tau)| \\ &\leq \max_{\tau \geq 0} |F_x(\tau) - F_y(\tau)| \\ &+ \max_{\tau \geq 0} |F_y(\tau) - F_z(\tau)| \\ &= z(x(t), y(t)) + z(y(t), z(t)) \end{aligned} \tag{5.3}$$

and it cannot be assured that if $z(x(t), y(t)) \in \mathcal{S}_\alpha$ and $z(y(t), z(t)) \in \mathcal{S}_\alpha$ then $z(x(t), z(t)) \in \mathcal{S}_\alpha$ for a given significance level α where $\mathcal{S}_\alpha$ is the acceptance region. As a consequence, $x(t)\mathcal{R}_{\rho,\alpha}y(t)$ and $y(t)\mathcal{R}_{\rho,\alpha}z(t)$ do not imply $x(t)\mathcal{R}_{\rho,\alpha}z(t)$. The same result applies to other goodness-of-fit tests such as the χ^2 test ([DS86, Chapter 3] and Section 4.2.5). $\square$.

In conclusion, the "queueing equivalent" thresholding method provides a technique to decide whether a finite-length traffic time series in V_0, say $x(t)$, and its approximation $A_j(t)$ in the 2^j timescale (namely, $A_j(t) \in V_{-j}$) are equivalent in terms of queueing performance. This is the case if and only if $\tilde{x}(t) = A_j(t) + \sum_{i=1}^{j} 0_i$ and $x(t)$ yield queueing occupancy distributions that pass the null hypothesis of a goodness-of-fit test for a given significance level and utilization factor.

Note that, in order to apply the method and obtain the queueing occupancy in the same timescale, one needs to reconstruct the original time-series in V_0 from $A_j(t)$, by means of iterative application of the reconstruction filter j times (upsampling with null details), as depicted in Figure 5.2. However, this is only required to check whether approximation and original time-series are equivalent in the queueing performance sense. Once the original time-series and approximation are considered equivalent by the "queueing equivalent" method both can be used indistinguishable. However, the approximation is smaller in size and it is easier to store and process.

Concerning the computational complexity of this technique, we found that most computational cost is carried by the MRA decomposition function. Therefore, the use of the “Queueing equivalent” method does not involve any substantial overload increase, with respect to other techniques such as the squared error computation.

5.3 Results and discussion

In this section we first describe the analysis scenario and data sample used to experiment and assess the *“thinning” process*. Then, we provide the results for a number of traffic time-series. The aim is to determine to which extent the queueing equivalent approximation presented in the previous section can effectively perform thinning of the original time-series, i.e., reduction in the time-series length.

5.3.1 Analysis Scenario

The “queueing equivalent” thresholding method has been applied to a data collection which spans one month (February to March 2007) of MRTG logs from the output routers of eight universities in Spain (labeled as U_a, U_b, …, U_h), four regional networks (RN_a, …, RN_d) which aggregate traffic from several universities, hospitals, computing and research centers, etc. and five external links (EL_a, …, EL_e) that RedIRIS has provided us (see Section 3.2).

This data comprises the average incoming and outgoing traffic volumes collected in 5 min time intervals. Thus, each monitored variable contains 288 samples per day, 2016 per week and 8640 measurements per month. A sample plot is shown in Figure 5.3.

In this figure the overall trend over a week becomes apparent: traffic volumes at weekends are smaller than during weekdays, which exhibit similar patterns among them. Within a day, the traffic signal grows in the morning, it shows a significant decrease at lunchtime, and then it increases until the users go home. These patterns are well known by the researchers dedicated to networks studies, actually diurnal patterns of Internet activity is considered one of the best-known invariants that can be found in the Internet (Section 2.3).

Note that the traffic sample is diverse in terms measurement points (nationwide). Furthermore, the campuses from which the data was collected have a large number of Internet users (both students and staff members), which favors that the

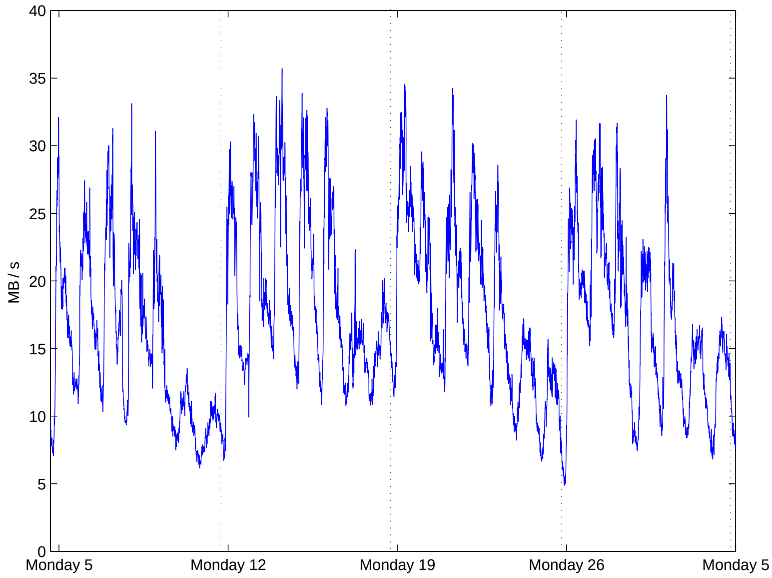

Figure 5.3: Incoming traffic for university U_h for one month, a sample per 5 minutes

traffic sample is statistically representative. Table 5.1 shows the total number of Internet users per university campus. Concerning the regional networks and the external links, the data provided is even more representative since these networks group up the traffic generated by several universities, hospitals, computing and research centers, thus aggregating a very large number of Internet users.

5.3.2 Goodness-of-fit test

As stated previously, MRA can be used for subsampling a time-series, since we can reduce the original N-size traffic signal into a smaller signal of length $N/2^j$ just by its projection onto the subspace V_{-j}. For instance, Figure 5.4 (top) shows the original monthly signal (from a given university), and its projections onto subspaces V_{-1} to V_{-10} using the Daubechies wavelet family [Dau90] and upsampling with null details.

It can be seen that subsequent projections cause the original signal to lose resolution gradually, but at the same time the signal length is divided by two. Additionally, Figure 5.4 (bottom) shows the details (projections onto W_{-j}) which, added to their associated approximation signals through the appropriate recon-

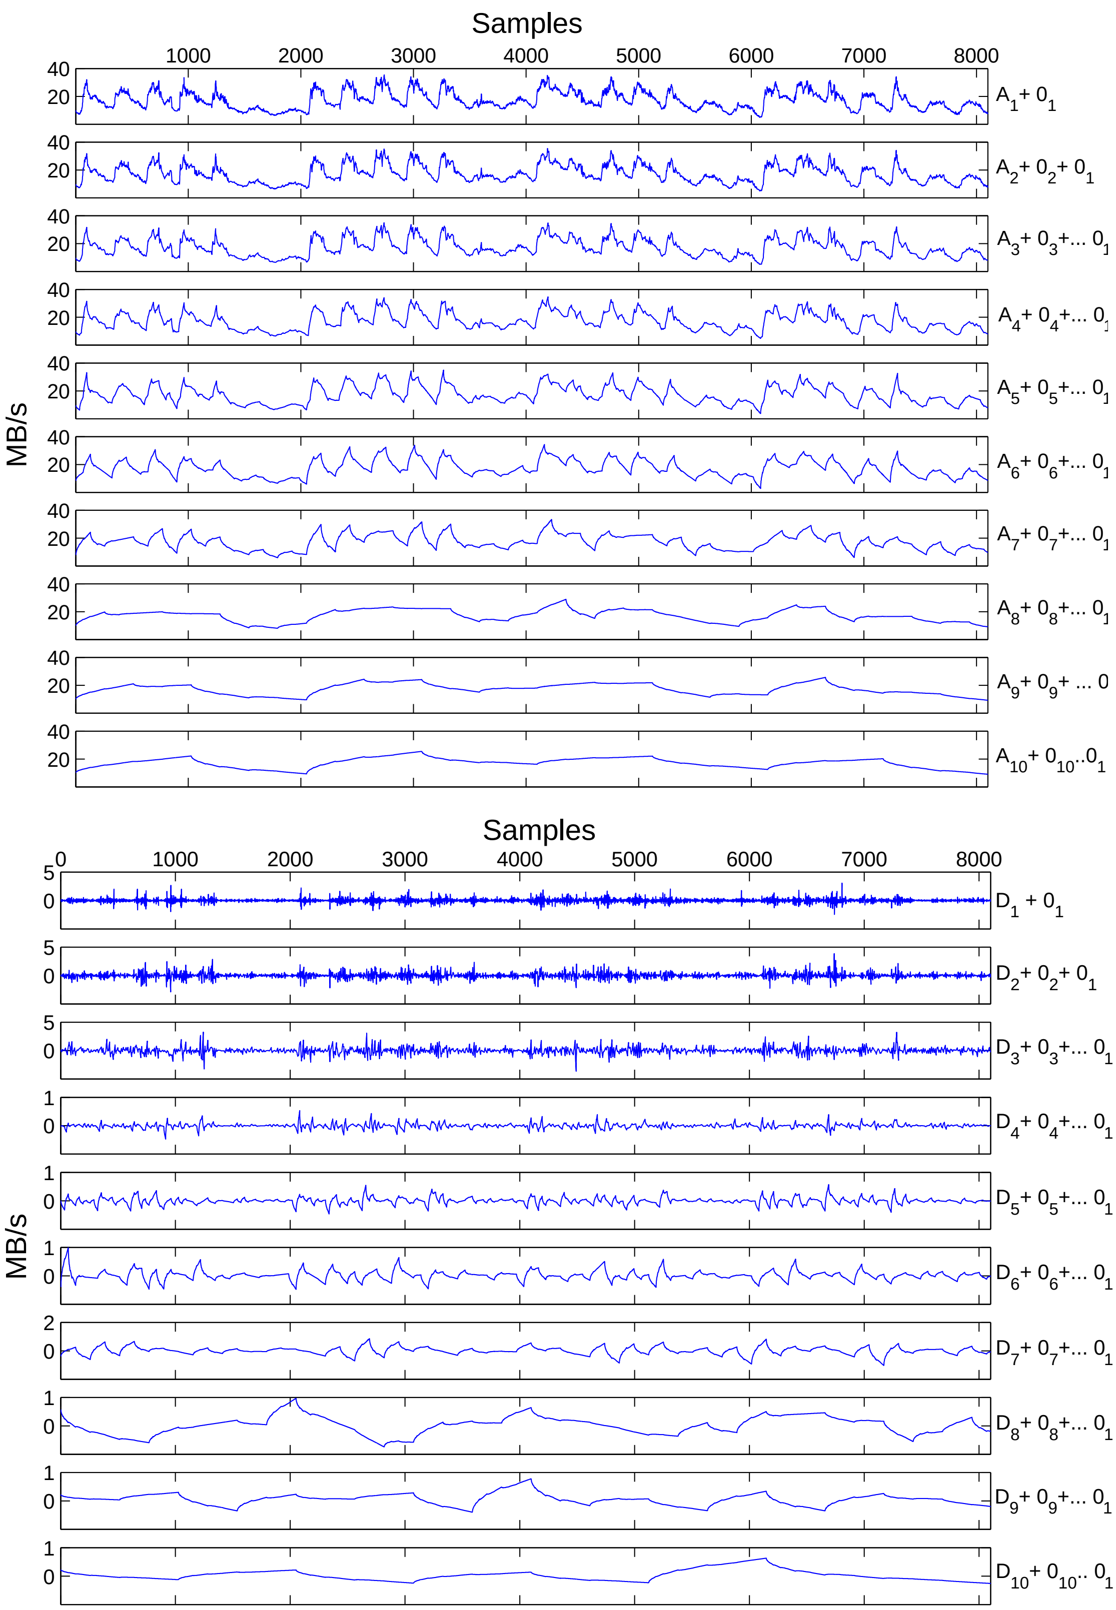

Figure 5.4: First ten Approximation and Detail signals for university U_h

University	Population
U_a	1,500
U_b	8,000
U_c	8,500
U_d	12,000
U_e	22,000
U_f	30,000
U_g	16,000
U_h	40,000

Table 5.1: Population size per university

struction filters, give the original time-series.

Figure 5.5 shows three different measures of similarity for this example: the energy of the details D_i, the Mean Squared Error (MSE), defined as:

$$\frac{1}{N}\sum_{t=1}^{N}(x(t)-\hat{x}(t))^2 \tag{5.4}$$

and the Euclidean distance [CFY03]:

$$\left(\sum_{t=1}^{N}(x(t)-\hat{x}(t))^2\right)^{\frac{1}{2}} \tag{5.5}$$

As shown, it is not easy to determine a relevant cut-off approximation level, since no knee point becomes evident. Nevertheless, even if a knee point is evident, there is no possible way to know whether the information lost has a clear impact on network performance or not.

On the other hand, Figure 5.6 shows the queueing occupancy Complementary Cumulative Distribution Function (CCDF) at different approximation levels, for the same signal. As shown, the queueing behavior significantly degrades for coarser approximations of the original signal. A goodness-of-fit test can be used to check whether such queueing occupancy distributions are "equal" or not.

For instance, Figure 5.7 shows the lowest significance level α such that the subsampled signal passes the χ^2 test for several approximations A_j. As shown, a clear step concerning significance level occurs from approximation A_3 to A_4. For

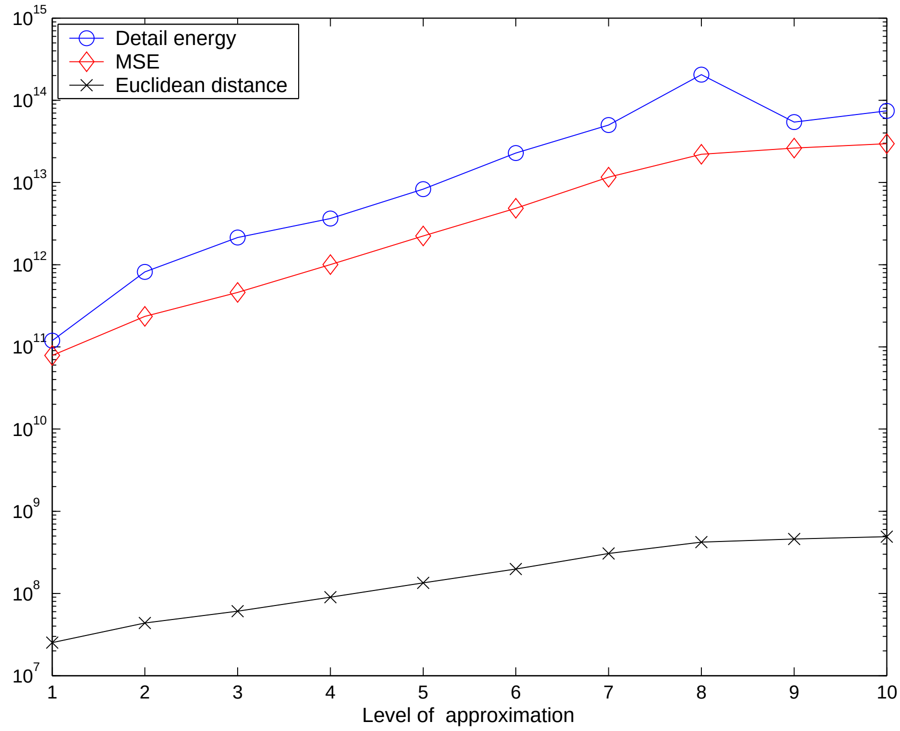

Figure 5.5: Detail energy, MSE and Euclidean distance for several approximation for university U_h

A_3, $\alpha \approx 10^{-6}$ ($1 - \alpha \approx 100\%$ confidence level) and it passes the χ^2 test, however in the fourth approximation the null hypothesis is absolutely rejected. This step-change behavior in terms of significance level has been found in the vast majority of the measurement sets studied, which gives a clear cut-off level.

This is a distinguishing feature of the proposed technique, because the goodness-of-fit test provides a clear *pass* or *fail* outcome. Thus, it clearly allows determining which is the best approximation for a given time-series, in contrast to previous methods based on continuous distances.

Hereafter, this test will be evaluated and compared with the Euclidean distance for all the available data. The number of histogram bins will be approximately $\sqrt{N_s}$, where N_s is the number of samples. All bins are required to contain at least one sample to be taken into account. Table 5.2 shows the results of applying the χ^2 test with significance level $\alpha = 0.05$ to the full measurement set. The results

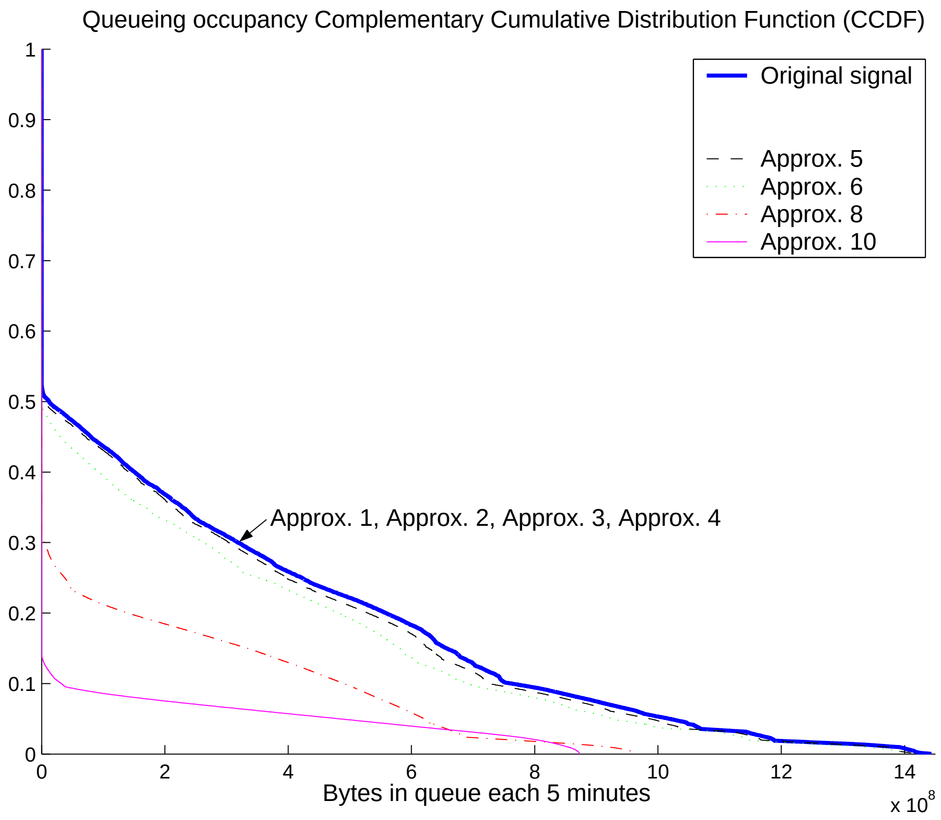

Figure 5.6: Distributions of delay suffered by a queue fed with the original and sampled signals, university U_h ($\rho = 0.8$)

show that it is possible to reduce the number of samples at least by a factor of 2^2 (1 sample per 20 mins) in all the cases studied. Moreover, 70% of the time-series can be further reduced to the 2^3 timescale (1 sample per 40 mins). For the external links, the data can be subsampled by a factor of 2^4 (1 sample per 80 mins).

Interestingly, the universities show the worst results, and the external links, with highest aggregation level, the best ones in terms of subsampling. A possible explanation is that the larger the aggregation level the smoother the traffic in terms of marginal distribution variance, by the Central Limit Theorem (CLT). Thus, it is expected that external links, with highly aggregated traffic, provide the best results in terms of approximation in a larger timescale.

Table 5.3 gives the Euclidean distance for the same original traffic signals and their approximations. The comparison of the tables shows that there is no clear

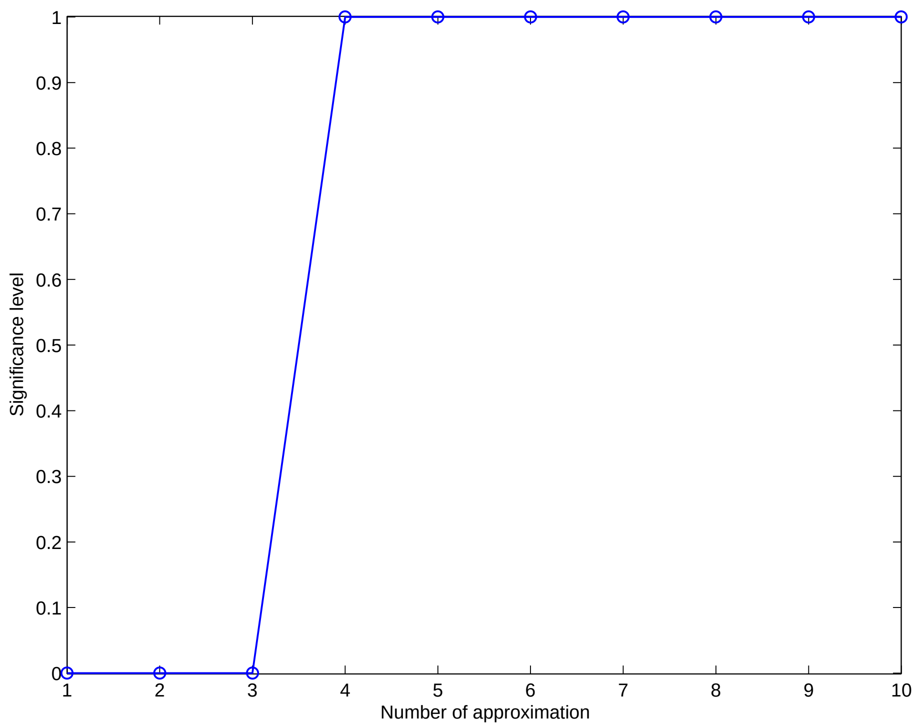

Figure 5.7: Lowest significance level (α) such passes χ^2 test between distributions delay suffered by a queue fed with the original and first ten approximation signals for university U_h ($\rho = 0.8$)

relationship between the Euclidean distance metric and the “queueing equivalent” thresholding method. For instance, university U_c shows larger Euclidean distance than university U_b, however, the former can be subsampled by 2^4 according to χ^2 test and the latter only by 2^3. In the same manner, EL_a can be subsampled more than RN_b, but the latter shows smaller Euclidean distance. This is not surprising because the Euclidean distance is only affected by the amount of data being removed and not by the real impact in queueing performance, as the “queueing equivalent” thresholding method does.

There is no a evident way to contrast if the “queueing equivalent” criteria permits to subsample the data more than the Euclidean distance does, since there is no method to determinate the optimal subsampling level with this last method. Nonetheless, in order to show a comparison, a reasonable error of 5% per sample, in average, is accepted as the threshold for the Euclidean distance. In this way if the distance exceeds this threshold the subsampling level is rejected, otherwise, it is accepted. Table 5.4 shows the results after applying this test. Obviously this threshold has been fixed arbitrarily but actually, there is no other way to set it. This is one of the major advantages of the "queueing equivalent" criteria as previously stated. Assuming these premises, overall, we find that, with the "queueing equivalent" criteria, the subsampling ratio achieved is larger than with other metrics such as Euclidean distance. For the given error of 5% per sample, with the latter metric only 50% of the analyzed networks can be subsampled by 2^2 whereas, with the former metric, all networks pass the χ^2 test. The same results arise for the next approximation level A_3, that is, all the regional networks and external links can be subsampled by 2^3 according to the "queueing equivalent" metric, however two of these networks would not admit approximation level A_3 following the other criteria. Similar results are obtained from the fourth approximation. Moreover Table 5.4 shows, as Table 5.3, that the traffic time-series from the external links (highly aggregated traffic) can be subsampled up to larger levels than universities or regional networks.

5.4 Conclusions

In this chapter we have proposed a new approach for downsampling a traffic time-series in a large timescale –in the MRA sense– where the subsampled threshold is determined based on queueing performance impact. Specifically, both the original time-series and approximation signals are fed to a single-server infinite queue system, and a goodness-of-fit test between the resulting occupancy distributions is performed for a given significance level and utilization factor.

In addition, we take advantage of the fact that the internetwork traffic measurements show a strong periodicity, due to their relation to the human activity patterns. Thus, by applying MRA with wavelets we take into account information not only in the time domain but also in the frequency domain, which in turn gives better results than those based on well-known time-based methodologies.

The results of applying our approach to real traffic time-series obtained from RedIRIS show that all the time-series can be approximated in the 2^2 timescale (20 mins), and 70% of them in the 2^3 timescale (40 mins). In the case of highly aggre-

Type	Center	Average Traffic (Mb/s)	Approx. 1 (10 mins.)	Approx. 2 (20 mins.)	Approx. 3 (40 mins.)	Approx. 4 (80 mins.)	Approx. 5 (160 mins.)
Universities	U_a	0.46	✓	✓	×	×	×
	U_b	1.87	✓	✓	✓	×	×
	U_c	4.90	✓	✓	✓	✓	×
	U_d	5.36	✓	✓	×	×	×
	U_e	10.17	✓	✓	×	×	×
	U_f	13.95	✓	✓	×	×	×
	U_g	15.30	✓	✓	✓	×	×
	U_h	138.24	✓	✓	✓	×	×
Regional Networks	RN_a	41.25	✓	✓	✓	✓	×
	RN_b	42.89	✓	✓	✓	×	×
	RN_c	650.38	✓	✓	✓	✓	×
	RN_d	803.64	✓	✓	✓	×	×
External Links & IXPs	EL_a	20.46	✓	✓	✓	✓	×
	EL_b	240.98	✓	✓	✓	✓	×
	EL_c	338.66	✓	✓	✓	✓	×
	EL_d	383.72	✓	✓	✓	✓	×
	EL_e	516.21	✓	✓	✓	✓	×

Table 5.2: Test χ^2 for first five approximations $(\alpha = 0.05, \rho = 0.8)$

Type	Center	Approx. 1 (10 mins.)	Approx. 2 (20 mins.)	Approx. 3 (40 mins.)	Approx. 4 (80 mins.)	Approx. 5 (160 mins.)
Universities	U_a	0.110	0.192	0.240	0.301	0.349
	U_b	0.244	0.360	0.491	0.594	0.701
	U_c	0.544	0.817	1.053	1.287	1.574
	U_d	0.656	1.231	1.455	1.871	2.135
	U_e	0.785	1.401	1.862	2.182	2.622
	U_f	0.623	1.190	1.564	2.483	2.971
	U_g	0.558	0.917	1.261	1.628	2.068
	U_h	2.526	4.367	6.100	9.022	13.449
Regional Networks	RN_a	1.276	2.212	3.117	3.873	4.813
	RN_b	1.320	2.421	3.699	4.808	6.565
	RN_c	9.004	15.228	23.317	30.286	41.183
	RN_d	20.723	38.961	56.127	78.406	102.296
External Links & IXPs	EL_a	1.098	2.254	3.696	4.930	6.570
	EL_b	4.865	8.492	15.205	21.021	39.893
	EL_c	22.969	41.697	63.026	80.201	98.845
	EL_d	13.097	19.604	33.871	42.902	84.455
	EL_e	8.931	15.291	29.383	40.785	81.034

Table 5.3: Euclidean distance between original signal and first five approximations. Values should be multiplied by 10^7

Type	Center	Approx. 1 (10 mins.)	Approx. 2 (20 mins.)	Approx. 3 (40 mins.)	Approx. 4 (80 mins.)	Approx. 5 (160 mins.)
Universities	U_a	✓	✓	×	×	×
	U_b	✓	✓	×	×	×
	U_c	✓	×	×	×	×
	U_d	✓	×	×	×	×
	U_d	✓	×	×	×	×
	U_e	✓	×	×	×	×
	U_f	✓	✓	×	×	×
	U_g	✓	✓	✓	×	×
Regional Networks	RN_a	✓	✓	✓	×	×
	RN_b	✓	✓	×	×	×
	RN_c	✓	✓	✓	×	×
	RN_d	✓	✓	✓	×	×
External Links & IXPs	EL_a	✓	✓	×	×	×
	EL_b	✓	✓	✓	×	×
	EL_c	✓	✓	✓	✓	×
	EL_d	✓	✓	✓	✓	×
	EL_d	✓	✓	✓	×	×

Table 5.4: Test based on Euclidean distance, error must be lower than 5%

gated traffic, most time-series were successfully approximated in the 2^4 timescale (80 mins). These results overcome the ones obtained applying the typical distance metrics.

Chapter 6

Characterization of the busy-hour traffic based on IP networks' intrinsic features

Internet Traffic measurements collected in the busy hour constitute a key tool to evaluate the operation of networks under the heaviest-load case scenarios, and further provide a means to network dimensioning and planning future capacity upgrades. In this light, this chapter provides an in-deepth analysis of the busy-hour traffic measurements from an extensive set of universities, regional networks and external links collected from the Spanish Research and Education Network (RedIRIS).

After proving that the traffic volumes observed in the busy hour over time can be modeled by a white Gaussian process, this work takes one step further and inspects the influence of the networks' intrinsic feature, primarily population size and access link capacity, on the characterization of busy-hour traffic as a Gaussian model. Analysis of Variance and Covariance methodologies are applied to the data, and the results show that the network size in terms of number of users accounts for the most of the variance of busy-hour traffic information. We further provide a linear-regression model that adjusts the amount of traffic that each network user contributes to the busy-hour traffic mean and variance values, with direct application to the link-capacity-planning *problem of IP networks.*

The remainder of this chapter is organized as follows: Section 6.2 reviews previous work in the field, presents the goals of this study and provides some extra details of the measurement set under study. The sections 6.3 and 6.4 comprise the core results of this chapter, which are finally summarized in Section 6.5.

6.1 Introduction

The characterization of Internet Traffic has received much attention from both network operators and the research community over years (Section 2.3). Indeed, there has been an intensive research effort in the characterization of the packet and byte counting process at small time-scales (say seconds and smaller), giving raise to a number of long-range dependence models [AFT98, HPA07]. Also, the estimation of Internet bandwidth demands has been a subject of study, either from a long-term [PTZD05] or a short-term [vdBMvdM+06, FTD03, MACL06] point of view. While these analyses and models serve to better understand the dynamics of Internet traffic, it turns out that network operators often use a different metric for capacity planning purposes: the total traffic volume observed in a given link during its busiest hour [vdMMP07, MACL06]. Obviously, network operators base their capacity planning strategies on worst-case scenarios, that is, on measurements collected when the network is most heavily loaded. This justifies the interest by the research community on studying and characterizing the busy-hour traffic, and its evolution over time. Typically, as noted by the authors in [vdMMP07], network operators use the following rule of thumb:

$$C = d \cdot M \tag{6.1}$$

where C is the target link capacity, M represents the bandwidth demand over the link under study, and d is some constant. Clearly $d \geq 1$, and is often much greater than one to provide sufficient capacity C to satisfy the burstiness of the bandwidth demand M.

The goal of this chapter is to characterize such bandwidth demand M for university access links during the busy hour, and further study the impact of intrinsic features [GDHA+08] of the universities on such bandwidth demands. Examples of intrinsic features of networks comprise their population (i.e., number of users) and access link capacity, among others. More specifically, this work studies the busy-hour traffic observed in the access links of a numerous set of large-size networks, focusing on its statistical properties and applicability to network dimensioning tasks. To this end, RedIRIS has kindly donated the traffic measurements of the access links to a large number of universities, regional networks and external links over a six-month period.

More specifically, we show that the busy-hour traffic can be accurately characterized by a pure Gaussian process, i.e., the average traffic volume during the most loaded hour is independent from one day to another, and it is further Gaussian distributed over time. Then, we go one step further and examine the influence of intrinsic network features, such as its population size (number of users) and its actual access link capacity, on the mean and variance of such a Gaussian model. After performing an Analysis of Variance (ANOVA) test, it is found that both access link capacity and population size are significant factors that greatly influence the traffic demands during the busy hour. However, after this, an Analysis of Covariance (ANCOVA) test is applied to check whether or not the access link capacity factor has a significative effect after removing the variance for which the population size accounts. Interestingly, in the set of networks under study (which show high capacity over-provisioning), it is only the population size that matters in the characterization of the busy-hour Gaussian process.

6.2 Preliminaries

6.2.1 Related work and contributions

In spite of its paramount importance for capacity planning and network design purposes, the network research community has surprisingly paid little attention to the study of the busy-hour traffic observed in network links, on the contrary to Plain Old Telephone System (POTS) designers. In fact, the network research community has addressed the problem of capacity planning by modeling the whole traffic process at different aggregation scales: packet, flow, application and aggregated traffic volumes.

At the packet level, the classical queueing theory has provided a framework for capacity planning, considering Markovian arrivals and service times. However, such assumptions no longer apply in light of the observed self-similar features of Internet traffic [CB97, PF95, LTWW94].

A flow-based approach is proposed in [BK00] whereby the authors base their capacity planning models on flow metrics. In such work, the authors end up with a model that considers the bandwidth mean and distribution tail of simulated TCP flows. On the downside, such flow-based dimensioning models are hardly feasible in practice and very sensitive to changes in the profiles of flows. The use of aggregated busy-hour traffic values provides a more robust approach to the process of traffic characterization.

As introduced in Section 2.4, the authors in [vdBMvdM+06] take one step further and propose a hybrid model $\rho + \alpha\sqrt{\rho}$ which considers both aggregated (the network load ρ) and per-flow (by means of α) metrics. Such parameter α is related to some characteristics of individual measured flows, for instance, their size and peak rate. A further refinement to this approach is proposed in [vdMMP07] where the burstiness of traffic is modeled from the variance of aggregated traffic, rather than following a flow-based approach (parameter α). In both approaches [vdBMvdM+06, vdMMP07], the estimation of average demands is kept constant throughout the entire analysis, and the authors mostly focus on the burstiness of traffic. However, given the assumption of traffic stationarity (at small time-scales) these approaches are only valid for capacity planning over short periods of time (in the order of few hours or so), which makes them impractical for long-mid term planning purposes. In fact, such a stationarity assumption breaks with the well-known fact that traffic patterns follow human behavior [FP01].

At the application level, the authors in [MNAR+04] characterize the traffic demand of individual users as a combination of the typical application sessions started by them: web browsing, Peer-to-Peer (P2P), Instant Messaging (IM) and email. However, such an application-based model requires network operators to correctly identify each application (which is not straightforward [MP05]). Additionally, this model is extremely sensitive to changes in user request patterns and hardly viable for forecasting purposes.

Finally, the authors in [PTZD05] have addressed the problem of capacity planning and network dimensioning by modeling the whole traffic measurement plot, as already introduced in Chapter 2.4.2. In [PTZD05], a long-term forecast of network traffic load based on three years of aggregated traffic measurements is presented. Essentially, the authors apply Auto Regressive Integrated Moving Average (ARIMA) models to the measurements collected on attempts to infer future network load values. Such a model is further applied in [FML$^+$03] to characterize the end-to-end traffic demands between each pair of Points of Presence (POPs) in a backbone network. Interestingly, such work shows that the bandwidth over-provisioning could be lower than usually assumed for a given Quality of Service (QoS) requirements. For instance, only about 15% of extra bandwidth (that is, $d = 1.15$ parameter of Eq. (6.1)) is required to ensure less than 3 ms of queuing delay.

Given the large size of measurements involved in such studies (one measurement every five minutes), the authors in [PTZD05] firstly aggregate the data to 90-minute intervals and then, they apply wavelets and ANOVA to further reduce the data volume. In contrast, using the busy-hour traffic approach only requires one measurement per day (the throughput value during the most loaded hour) and data preprocessing is barely required, which first simplifies the process of data collection, storage and management, and secondly considers the worst case scenario for capacity planning purposes.

The model proposed in this work tries to overcome the above limitations found in the literature. More precisely, our model studies only the traffic volumes during the busy hour over a relatively long period of time (in the range of several months), and finds that such busy-hour traffic characterization is accurately modeled with a Gaussian process. Additionally, the model only requires one aggregated traffic-related value to be collected every day: the busy-hour traffic mean, which makes it more practical and robust (less sensitive to fine-grane measurements). Moreover, the model relates the bandwidth demand results to the intrinsic characteristics of networks, such as its population size and access link capacity, which is novel.

Concerning traffic Gaussianity, the authors in [KN02, vdMMP06] test whether or not aggregated traffic follows a Gaussian distribution at different aggregation levels in terms of number of users and time-scales. Both studies find Gaussian behavior from 5-ms to 5 seconds of time-granularity. It is worth noticing that we are facing a different problem: The Gaussian modeling of the busy-hour traffic over a number of consecutive days, rather than characterizing the aggregated traffic sample itself.

6.2.2 Definition of busy-hour traffic

Let $A(t)$ be the instantaneous network throughput measured (in units of Mbits/s for instance) on a given access link. Here, t spans a day of throughput measurements, that is, $t \in [0, 24)$ hours. Also, let $H_T(t)$ denote an average throughput metric computed over a given range $[t - \frac{T}{2}, t + \frac{T}{2}]$ of length T, typically one hour:

$$H_T(t) = \frac{1}{T} \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} A(\tau) d\tau \tag{6.2}$$

According to this, the busy-hour traffic X is the value that maximizes the above equation, i.e.:

$$\begin{aligned} X \quad = \quad & \max_t H_T(t), \quad t \in [0, 24) \text{ hours} \\ & T = 1 \text{ hour} \end{aligned} \tag{6.3}$$

which gives the average throughput (in Mbits/sec) during the busiest hour of a given day, and such value occurs at time t that maximizes $H_T(t)$.

Additionally, let V be the variance of $A(t)$ during the busy hour $[t - \frac{T}{2}, t + \frac{T}{2}]$, that is

$$V = \frac{1}{T} \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} (A(\tau) - X)^2 d\tau \tag{6.4}$$

Since the traffic is collected in intervals of five minutes the above equations are discretized accordingly.

Thus, for each data unit (either university, regional network or external links), the above equations define the time-series $\{X_i, i = 1, \ldots, N\}$ and $\{V_i, i = 1, \ldots, N\}$ which comprise the average traffic volume observed during the busy hour and its variance for different days $i = 1, \dots, N$. In addition, it is also interesting to study the time of day t_i at which the traffic busiest hour occurs. This is given by the time-series $\{t_i, i = 1, \dots, N\}$. Figure 6.1 illustrates how these metrics are computed for a given network over three days.

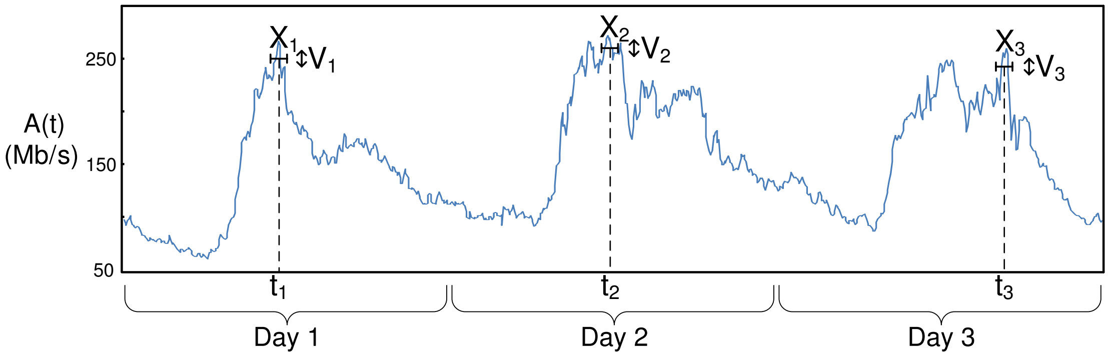

Figure 6.1: A three-day example of traffic measurements to illustrate X_i, V_i, and t_i

Finally, with daily values of X_i and V_i, it is also possible to compute the Coefficient of Variation (CV) as defined by:

$$CV_i = \sqrt{\frac{V_i}{X_i^2}}, \quad i = 1, \dots, N \tag{6.5}$$

which gives a measure of the variability of the busy-hour traffic volume with respect to its mean.

6.2.3 Measurement set description

The Spanish National Research and Education Network, RedIRIS, kindly provided the measurements to carry out this study. RedIRIS infrastructure and available data was presented in Section 3.2. Specifically, we use traffic capture measurements ranging from April to October 2007, in which traffic was monitored in both directions of the access link, say incoming (from the Internet to the network under study) and outgoing (sourced on the network under study and destined to the Internet), as was shown in Figure 3.2.

The traffic trace collected is based on Multi Router Traffic Grapher (MRTG) logs and Cisco's Netflow data. As explained in Section 3.2.3, such Netflow data provides another means to obtain the inbound and outbound traffic volumes, which were shown to validate the traffic values given by the MRTG data. We preformed this experiment with the RedIRIS' data and the results were satisfactory.

Both these data provide an approximation to the instantaneous network throughput $A(t)$ defined previously, and consequently have been used to calculate the daily busy-hour traffic volume mean X and variance V as stated in Eqs. (6.3) and (6.4) for each access link.

After the time-series $\{X_i, i = 1, \ldots, N\}$ for each access link was calculated over different days, it is worth mentioning that the values obtained on both local and bank holidays as well as other non-teaching periods were removed from all time-series. The reason is that only the working-day values are of interest since the network operators are interested in worst-case scenarios for capacity planning purposes. In addition to this, note that both network upgrades and configuration changes, for instance infrastructure improvements, new filtering policies, new killer applications appraisal, etc. may also have a negative effect on the long-term characterization of the busy-hour throughput, yielding to a non-stationary process. Hence, in order to avoid these problems and to maximize the number of available days in the time-series, we have carefully selected a set of network access links that minimizes all these restrictions, thus removing all those which have shown a clear non-stationary behavior over year 2007. Such a refined measurement set comprises data from four regional networks (in what follows $RN_a \ldots RN_d$) which aggregate traffic from several universities, hospitals, computing and research centers; five external links/Internet Exchange Point (IXP) (namely $EL_a \ldots EL_e$); and 22 university networks generically labeled as U_j.

As ANOVA requires to group the networks by factors and the group sizes must be similar, we have extended the analysis to a larger set of university networks and we have carefully selected them to obtain homogeneous sets of universities by factor. Thus, following the labeling carried out in the previous chapters, the set of 22 university networks that are studied in this chapter are labeled as U_2, U_7, U_9, U_{10}, U_{13}, U_{15}, $\ldots$, U_{30}. This set comprises some campus networks earlier analyzed and other that included specifically to this study ($U_{15}, \ldots, U_{30}$).

The set of university networks was completed with the so-called network intrinsic features, that is, the values of their population size and access link capacity. There exist well-documented central repositories which describe the university networks' user population, Internet access capacity and organization [Con]. This information has allowed us to select a representative set of universities regarding such intrinsic features (see Table 6.1), and rank them by means of both population size and access link capacity. However, we note that the same does not apply to the EL and RN networks as the population size of each of them is unknown.

Finally, the *capping effect*, as introduced in [VK00], states that the traffic demands may be affected (bounded) by a limiting bandwidth capacity. In this light, Table 6.1 summarizes the access link capacity C_{U_j} for each university network U_j (third column), along with the maximum average busy-hour traffic over the N days of measurement (last column). The reader should note that all links show low-levels of utilization even at highly-loaded days, typically below 40%. Indeed, such over-provisioning of access links is a common practice by network operators, as noted in Section 2.4. Actually, the average utilization during the busy hour in our set of measurements was under 25% in all cases, and the most loaded network, U_{20} showed a maximum utilization lower than 70%. Such low levels of utilization make the capping effect negligible in this data [NP08, Chapter 4], since the maximum busy hour traffic volumes found in the measurements are far from reaching the maximum capacity of access links. Obviously, in other under-provisioned scenarios with higher levels of utilization, the capping effect cannot be ignored.

6.3 Characterization and dynamics of the busy-hour traffic process

The following experiments firstly study the marginal distribution of the busy-hour traffic volume or throughput, and then focus on its correlation structure. The results obtained for all ELs, RNs and Us are summarized in Table 6.2.

Networks	Capacity (Mb/s)	Population size (thousands)	$max\ (X_i)$ (Mb/s)	Max. utilization (Mb/s)
U_{15}	1000	58	328 / 190	0.33 / 0.19
U_{10}	1000	55	150 / 115	0.15 / 0.12
U_{13}	1000	46	328 / 207	0.33 / 0.21
U_7	1000	34	107 / 126	0.11 / 0.13
U_9	1000	28	158 / 126	0.16 / 0.13
U_2	1000	28	83 / 111	0.08 / 0.11
U_8	1000	26	115 / 90	0.12 / 0.09
U_{16}	1000	14	23 / 73	0.02 / 0.07
U_{17}	1000	11	17 / 54	0.02 / 0.05
U_{18}	1000	8	17 / 41	0.02 / 0.04
U_{19}	1000	7	8 / 21	0.01 / 0.02
U_{20}	200	54	133 / 137	0.67 / 0.69
U_{21}	100	37	35 / 65	0.35 / 0.65
U_{22}	200	35	120 / 111	0.60 / 0.56
U_{23}	100	20	46 / 38	0.46 / 0.38
U_{24}	200	18	40 / 57	0.20 / 0.29
U_{25}	200	15	13 / 29	0.07 / 0.15
U_{26}	100	15	4 / 23	0.04 / 0.02
U_{27}	100	13	30 / 57	0.30 / 0.57
U_{28}	100	9	20 / 29	0.20 / 0.29
U_{29}	100	7	7 / 17	0.07 / 0.17
U_{30}	100	4	17 / 11	0.17 / 0.11

Table 6.1: Description of universities, their intrinsic features, and maximum utilization in outgoing/incoming direction ranked by population size

6.3.1 Gaussian marginal distribution

The first two columns of Table 6.2 show the estimated mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$ of the busy-hour throughput distribution over time, measured at each monitored point in both outgoing and incoming directions. Essentially, for a given university U_j, whose busy-hour traffic over N days is defined in the set $\{X_1^{U_j}, \ldots, X_N^{U_j}\}$, such mean and standard deviation values are computed as:

$$\hat{\mu}_{U_j} = \frac{1}{N}\sum_{i=1}^{N} X_i^{U_j} \tag{6.6}$$

$$\hat{\sigma}_{U_j} = \frac{1}{N}\sqrt{(\sum_{i=1}^{N}(X_i^{U_j} - \hat{\mu}_{U_j})^2)^{\frac{1}{2}}} \tag{6.7}$$

The fourth column in the table shows the maximum coefficient of variation (CV^{max}) during each busy hour, calculated following Eq. (6.8). Essentially, for each university U_j, we compute its CV for each day i and take its maximum value over all N days:

$$CV_{U_j}^{max} = \max_i \sqrt{\frac{V_i}{X_i^2}}, \quad i = 1, \ldots, N \tag{6.8}$$

This value represents the ratio of the variance V to the mean X, and it is very useful for comparing the degree of variation of the busy hour traffic over different days. This maximum considers the worst possible case: The day which shows highest variability ratio (highest bursty behavior). This result is discussed at the end of this section.

Finally, Table 6.3 provide the results of different Gaussian goodness-of-fit tests applied to the measurements. Essentially, the null hypothesis assumes that the busy hour traffic follows a Gaussian distribution with parameters $\hat{\mu}_{U_j}$ and $\hat{\sigma}_{U_j}$ for university network U_j. The easiest way to visually assess on the validity of the null hypothesis is via the Quantile-Quantile plot [DS86, Chapter 2], which plots the pairs $x_{(i)}$ versus $Q(i/n))$, whereby $x_{(i)}$ is the order statistics of the empirical sample and $Q(\cdot)$ is the Quantile function (inverse of the cumulative distribution function) of the reference distribution (in this case, Gaussian distribution). If indeed the measured data follows the Gaussian distribution $N(\hat{\mu}_{U_j}, \hat{\sigma}_{U_j})$, the points depicted overlap the angle bisector (line $y = x$). This is the case of Figure 6.2, where the QQ-plot technique is applied to the busy-hour measurements of EL_a, RN_a and U_{15} respectively in both incoming and outgoing directions. The same experiment has been applied to all other measurement sets, showing linear QQ-plots in all cases.

Besides visual matches, it is desirable to assess Gaussianity objectively following the most common goodness-of-fit tests found in the literature, say: Kolmogorov-Smirnov [DS86, Chapter 4], Shapiro-Wilks [DS86, Chapter 9], Anderson-Darling [DS86, Chapter 9] and correlation-based [KN02]. Basically, the correlation test consists in checking whether or not the linear correlation coefficient R computed between the pairs $x_{(i)}$ and $Q(i/n))$ in the QQ-Plot gives a relatively high value, say 0.9 [vdMMP06].

Table 6.3 gives the results obtained after applying such tests. As shown, all empirical distributions pass the correlation test. Also, we observe that the goodness-of-fit tests support the null hypothesis (Gaussianity) for most of the cases and further show visual Gaussianity too. However, as noted in [vdMMP06], conventional goodness-of-fit tests are usually excessively demanding with traffic measurements. Note that certain outliers may arise from events such as network misuse, power cuts, temporal malfunctioning, etc. instead of typical network behavior, hence making the tests fail. For this reason, we conclude that the busy-hour traffic measurements for the access links of all universities, regional networks and external links can be considered “fairly Gaussian”, borrowing the term from [vdMMP06].

6.3.2 Autocorrelation experiments

This section studies the correlation between consecutive busy-hour traffic measurements, that is, whether or not the busy-hour traffic experienced on one day depends on the values measured the previous days. To this end, the autocorrelation function was calculated for all data items (Us, RNs and ELs) together with their confidence intervals (with significance level $\alpha = 0.05$) as described by the Bartlett test [CL66] for the autocorrelation of a pure white Gaussian process, i.e., a Dirac delta at the origin of the autocorrelation function. Figure 6.3 shows the autocorrelation (dots) and the confidence intervals (solid lines) applied again to

Networks (outgoing / incoming)	$\hat{\mu}_{U_j}$ (Mb/s)	$\hat{\sigma}_{U_j}$ (Mb/s)	$CV_{U_j}^{max}$
EL_a	1460 / 570	18 / 8.0	0.05 / 0.10
EL_b	1111 / 806	18 / 19	0.59 / 0.55
EL_c	1510 / 2223	11 / 9.1	0.20 / 0.03
EL_d	52 / 33	3.2 / 2.3	0.31 / 0.40
EL_e	683 / 895	8.2 / 12	0.17 / 0.22
RN_a	288 / 558	6.6 / 9.2	0.22 / 0.17
RN_b	80 / 136	4.2 / 5.0	0.40 / 0.22
RN_c	71 / 213	4.1 / 6.3	0.42 / 0.70
RN_d	920 / 661	11 / 10	0.13 / 0.18
U_{15}	255 / 199	6.1 / 4.1	0.20 / 0.42
U_{10}	109 / 84	21 / 13	0.10 / 0.17
U_{13}	206 / 113	6.1 / 4.9	0.32 / 0.27
U_7	63 / 102	18 / 9.0	0.21 / 0.14
U_9	55 / 117	4.1 / 4.0	0.32 / 0.26
U_2	51 / 84	14 / 13	0.20 / 0.42
U_8	94 / 73	15 / 11	0.10 / 0.20
U_{16}	32 / 90	2.4 / 4.6	0.20 / 0.18
U_{17}	10 / 41	3.2 / 7.0	0.70 / 0.25
U_{18}	8.1 / 30	2.6 / 6.6	0.35 / 0.47
U_{19}	4.6 / 16	1.4 / 2.7	0.57 / 0.61
U_{20}	97 / 118	17 / 10	0.10 / 0.10
U_{21}	15 / 40	2.6 / 4.1	0.50 / 0.34
U_{22}	102 / 88	22 / 11	0.33 / 0.41
U_{23}	26 / 26	6.0 / 5.6	0.20 / 0.27
U_{24}	31 / 46	3.2 / 5.4	0.28 / 0.16
U_{25}	8.7 / 21	1.3 / 3.5	0.32 / 0.19
U_{26}	3.1 / 17	0.5 / 3.1	0.41 / 0.50
U_{27}	21 / 38	4.1 / 7.1	0.49 / 0.41
U_{28}	11 / 22	2.7 / 3.8	0.35 / 0.47
U_{29}	4.6 / 10	1.4 / 3.0	0.80 / 0.52
U_{30}	6.6 / 7.7	3.2 / 1.6	0.34 / 0.31

Table 6.2: Gaussian characterization of busy-hour traffic $N(\hat{\mu}, \hat{\sigma})$ in both incoming/outgoing directions of traffic for the set of network under study

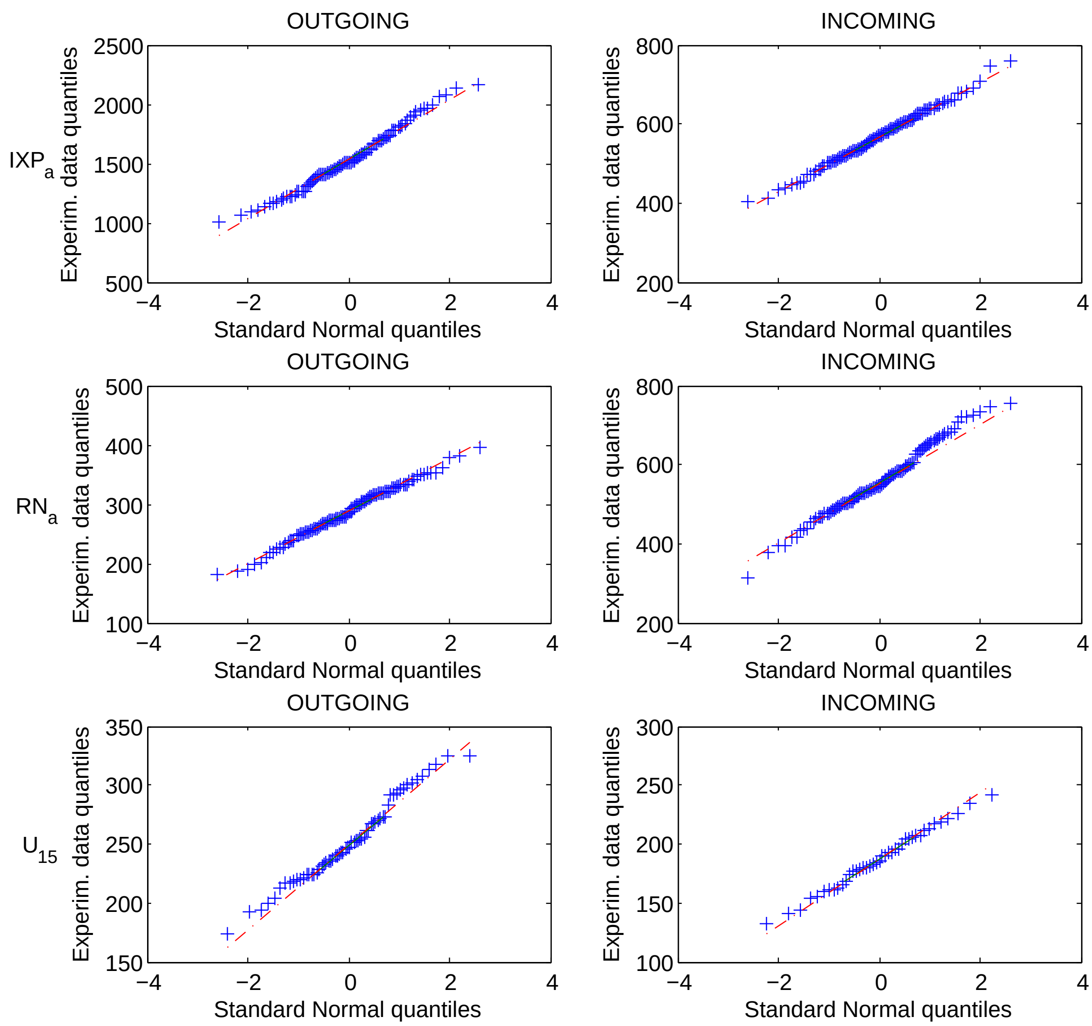

Figure 6.2: QQ-Plot for EL_a, RN_a and U_{15}

Networks (outgoing / incoming)	Kolmogorov-Smirnov ($\alpha = 0.05$)	Shapiro-Wilks ($\alpha = 0.05$)	Anderson-Darling ($\alpha = 0.05$)	Correlation Test ($R > 0.9$)
EL_a	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
EL_b	✓ / ✓	✓ / ✓	✓ / ×	✓ / ✓
EL_c	✓ / ✓	✓ /✓	✓ / ✓	✓ / ✓
EL_d	× / ✓	✓ / ✓	× / ✓	✓ / ✓
EL_e	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
RN_a	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
RN_b	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
RN_c	✓ / ✓	✓ / ✓	✓ / ×	✓ / ✓
RN_d	✓ / ✓	✓ / ✓	✓ / ×	✓ / ✓
U_{15}	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_{10}	✓ / ✓	✓ / ✓	✓ / ×	✓ / ✓
U_{13}	✓ / ✓	✓ / ✓	× / ✓	✓ / ✓
U_7	× / ✓	✓ / ✓	× / ✓	✓ / ✓
U_9	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_2	✓ / ×	✓ / ✓	✓ / ×	✓ / ✓
U_8	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_{16}	× / ✓	✓ / ✓	× / ✓	✓ / ✓
U_{17}	× / ✓	× / ✓	× / ✓	✓ / ✓
U_{18}	× / ✓	× / ✓	× / ✓	✓ / ✓
U_{19}	✓ / ✓	✓ / ✓	✓ / ✓	× / ✓
U_{20}	✓ / ✓	✓ / ✓	✓ / ×	✓ / ✓
U_{21}	✓ / ×	✓ / ✓	✓ / ×	✓ / ✓
U_{22}	× / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_{23}	✓ / ✓	× / ✓	× / ✓	✓ / ✓
U_{24}	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_{25}	× / ✓	× / ✓	× / ✓	✓ / ✓
U_{26}	× / ✓	× / ✓	✓ / ✓	✓ / ✓
U_{27}	✓ / ✓	✓ / ✓	✓ / ✓	✓ / ✓
U_{28}	× / ✓	× / ✓	× / ✓	✓ / ✓
U_{29}	✓ / ✓	× / ✓	× / ✓	✓ / ✓
U_{30}	✓ / ✓	× / ✓	× / ✓	✓ / ✓

Table 6.3: Goodness-of-fit test results for Gaussian characterization in both incoming/outgoing directions of traffic

EL_a, RN_a and U_{15} respectively. Interestingly, all networks pass this test, which proves the independence of the busy-hour traffic values from one day to another after removing both weekends and holidays.

6.3.3 Distribution of the Busy-hour times

Figure 6.4 shows the Cumulative Distribution Function (CDF) of the time instants when the daily busy hour occurs, that is, the value of t in Eq. (6.3). For the sake of clarity only the results for six universities during two months are shown, although similar behaviors were observed for the rest of the networks under study. As shown, the CDFs for all six universities are almost overlapped in both directions of traffic. In the incoming direction of traffic, the busiest hour typically occurs in the range from 10:00 a.m. to 1:00 p.m. However, the outgoing direction shows a bimodal behavior with its busiest hour typically found either around 11:00 a.m. or around 5:00 p.m.

These results are consistent with the "Daily traffic pattern" invariant defined in [TMW97, FP01] and presented in Section 2.3.

Essentially, the authors in [FP01] expose that some traffic patterns follow strictly the human behavior which, in the case of an academic network, this seems to show two peaks of traffic: one in the morning and another one in the afternoon.

As shown, the busiest hour t never occurs at night, which gives at least 12 hours between any two consecutive busy-hour traffic measurements X. Intuitively, this can be the reason that explains why the correlation structure in the busy-hour traffic time-series $\{X_i, i = 1, \ldots, N\}$ vanishes since there is a gap of at least 12 hours between any two consecutive busy-hour traffic measurements X (see Figure 6.4).

6.3.4 Applicability to capacity planning problem

On one hand, the above results show that the busy-hour traffic samples $\{X_1, \ldots, X_N\}$ are both uncorrelated and Gaussian distributed. Hence, the busy-hour traffic process can be modeled by a white Gaussian process.

Additionally, the maximum coefficient of variation, which gives the maximum ratio of the variance V to the mean X over different days i is always smaller than

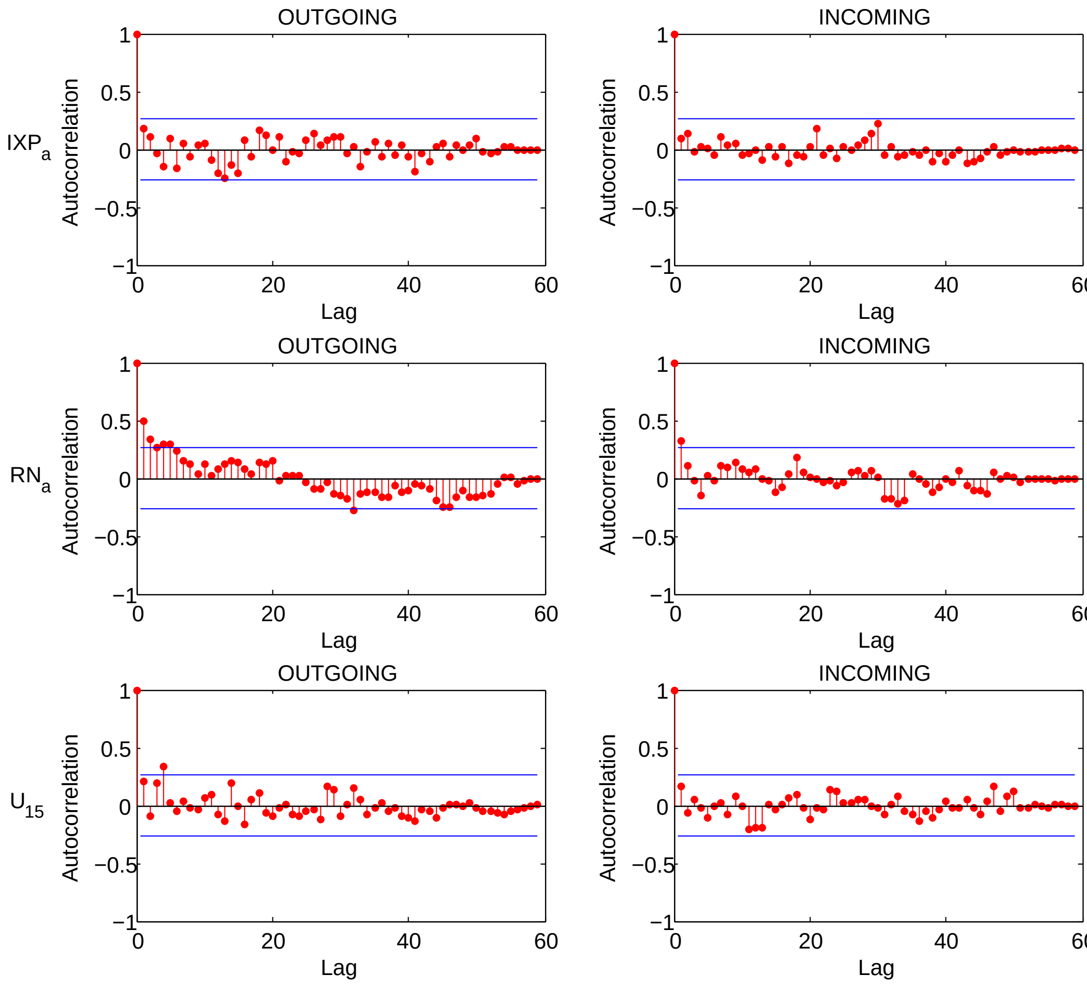

Figure 6.3: Autocorrelation function and Bartlett Test (solid lines) for EL_a, RN_a, and U_{15}

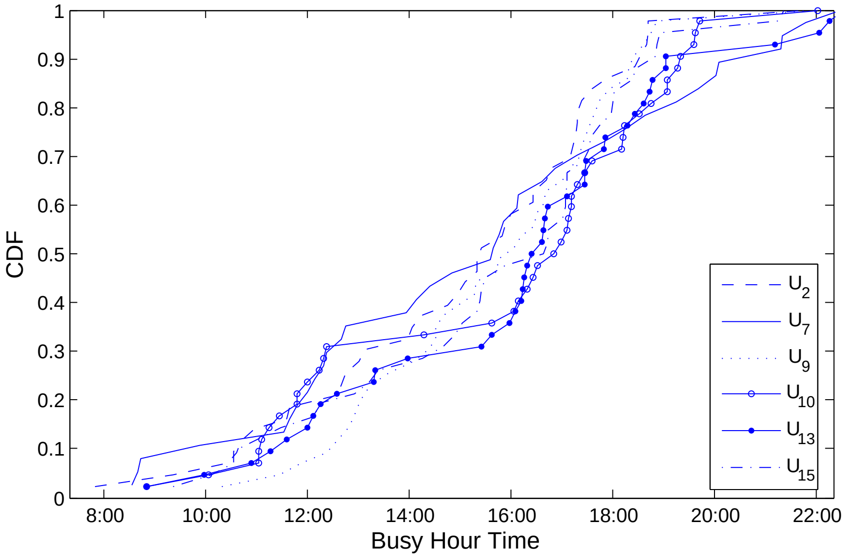

(a) Outgoing

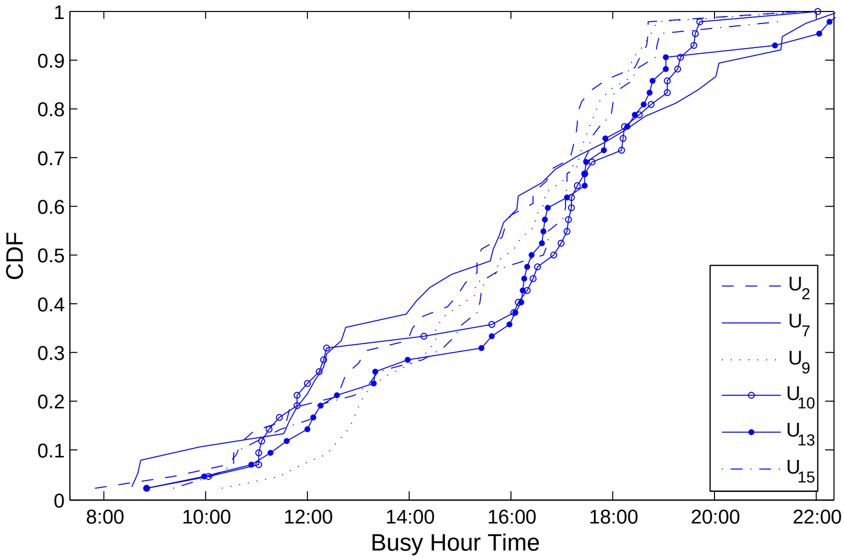

(b) Incoming

Figure 6.4: Busy-hour time CDF in both outgoing (a) and incoming (b) directions of traffic

one, hence showing sub-exponential behavior in all cases. In other words, the traffic during the busy hour is close to the average value (small variation with respect to the mean), which is of clear importance for capacity planning purposes.

Having found that the process is white and Gaussian, network operators can apply the conventional sample mean and variance estimator to a measurement set. With such parameters at hand, operators can use the following formula to derive the access link capacity C required such that the busy-hour traffic volume is met with probability $1 - \varepsilon$ (typically $\varepsilon \leq 0.1$):

$$\begin{array}{lc} C_{U_j} \text{ such that} & \text{Prob}(d \cdot X^{U_j} < C_{U_j}) \geq 1 - \varepsilon, \\ \text{with} & X^{U_j} \sim N(\hat{\mu}_{U_j}, \hat{\sigma}_{U_j}) \end{array} \tag{6.9}$$

Finally, it is worth noticing that this section's conclusions are derived based on measured busy-hour traffic volumes only. This capacity planning application is not valid for designing new networks over which no measurements have already been taken. For this reason, the next section is devoted to extracting how much information of the busy-hour traffic is directly related to the number of users (population size) for a given network, on attempts to refine the dimensioning rule above (Eq. (6.9)).

6.4 Factor analysis of access link capacity and population size

The previous tests have shown that the busy-hour traffic distribution of university networks can be accurately characterized by a Gaussian distribution $N(\mu, \sigma)$, whereby its characteristic parameters μ and σ can be estimated from measurements. The next set of experiments aim to study whether or not the intrinsic features of the networks (population size and access link capacity), which are denoted as explanatory variables in what follows, have any influence on such parameters μ and σ, which are denoted as response variables. To do so, ANOVA and ANCOVA methodologies are first overviewed, and then applied to the measurement set. Before that, we remark that this section only takes into account the measurements collected at the 22 university access links U_j. The ELs and RNs measurements are not considered in these experiments since their population size is unknown.

Both the ANOVA and ANCOVA methodologies require the data to meet a few requirements, as introduced in Chapter 4: First, the explanatory variables must be independent and Gaussian distributed; and second, all of them must share the same intra-group variance (exhibit homoscedasticity). However, the results of ANOVA and ANCOVA are generally accepted provided that the number of elements in each group are similar and there is a non-excessive deviation from the homoscedasticity assumption. This is the case for our measurements. For further explanation see Section 4.2.5. Additionally, the ANCOVA model assumes a linear relationship between the response and the explanatory variables (as will be shown in Section 6.4.2).

Table 6.4 summarizes the factors (access link capacity C_{U_j} and population size P_{U_j}) for each university under study. As noted from the table, the universities have been split into two groups depending on the capacity of their access links: The universities with 1 Gb/s of capacity belong to group G_{high} (which stands for high-speed access link), thus leaving G_{low} to the universities with low access capacity (ranging from 100 to 200 Mb/s). This classification is important to apply ANOVA, as shown in the following.

6.4.1 Effect of access link capacity: ANOVA

This section studies the effect of the access link capacity only on the busy-hour traffic volumes for each university characterized by $N(\mu_{U_j}, \sigma_{U_j})$. Remark that, for each university U_j, its access link capacity C_{U_j} and population size P_{U_j} are known but, for this former experiment, P_{U_j} is ignored.

ANOVA is a statistical methodology whereby the observed variance of a given response variable is split into explanatory factors or categories and provides a means to determine if the factors have any importance in explaining such a response variable, and how much this accounts for.

In our example, ANOVA proceeds as follows: It first splits the response variable μ_{U_j} into two categories: G_{high} and G_{low}. Then, it computes the adjusted mean squares for each category and for the total. The difference between both is due to the experimental error.

Networks	CapacityAccess		PopulationSize (users)	
	Group	(Mb/s)	Group	(Thousands)
U_{15}	$Ghigh$	1000	Large	58
U_{10}		1000		55
U_{13}		1000		46
U_7		1000	Medium	34
U_9		1000		28
U_2		1000		28
U_8		1000		26
U_{16}		1000	Small	14
U_{17}		1000		11
U_{18}		1000		8
U_{19}		1000		7
U_{20}	G_{low}	200	Large	54
U_{21}		100		37
U_{22}		200		35
U_{23}		100	Medium	20
U_{24}		200		18
U_{25}		200		15
U_{26}		100		15
U_{27}		100	Small	13
U_{28}		100		9
U_{29}		100		7
U_{30}		100		4

Table 6.4: Set of universities grouped by access link capacity and population size

Finally, ANOVA performs a contrast test using the ratio between the adjusted mean square of each factor and the total, which follows a Snedecor-F distribution. The null hypothesis considers that the total adjusted mean square is due to the experimental error, and not to differences in the population when grouped by categories. However, if the null hypothesis is not accepted, this means that the factor used to build the groups (access link capacity) is statistically significant according to the F-test. Moreover, the ANOVA test provides a p-value which determines if the null hypothesis should be accepted or not, according to a given pre-defined significance level α (typically $\alpha = 0.05$). Basically, if $p > \alpha$, then the null hypothesis is accepted (non-significative factor), and rejected otherwise. Furthermore, ANOVA also quantifies the amount of variance explained by the factors (explained variance) and the amount of variance that remains unexplained (non-explained or residual variance).

It is worth noticing that this test will be applied to both μ and σ in both outgoing and incoming directions of traffic. For now, let us refer to them as a generic response variable y.

The ANOVA model for response variable y with the access link capacity as its only factor is given by:

$$y_{U_j}^{group} = k_y + \alpha^{group} + \epsilon_{U_j}^{group} \tag{6.10}$$

Here, k_y is the overall means response for the response variable under study (y), typically named as μ but, in this case, to avoid confusion with the response variable μ_{U_j} we have renamed this term as k. $y_{U_j}^{group}$ refers to the mean or variance (in incoming or outgoing direction) value of university U_j which belongs to a given *group* (either G_{high} or G_{low}). The value of α^{group} represents the effect because of a given network U_j belongs to a given *group*.

Finally, the value of $\epsilon_{U_j}^{group}$ refers to the experimental error introduced above. Clearly, large values of $\epsilon_{U_j}^{group}$ means that the link access capacity factor explains little variance and, perhaps, other factors that explain more variance must be included in the model given by Eq. (6.10).

Table 6.5 shows the results after applying the ANOVA test to the busy hour traffic mean μ and standard deviation σ in both incoming and outgoing directions of the university access routers under study. The third column gives the sum of

Response Variable (direction)		Source of Variation	Sum of Squares	df	Adj. Mean Square	F	p-value
outgoing	μ	Capacity	13425	1	13425	2.67	0.118
		Error	100386 (88%)	20	5019		
		Total	113811	21			
	σ	Capacity	95.5	1	95.5	1.22	0.282
		Error	1561 (94%)	20	78.1		
		Total	1657	21			
incoming	μ	Capacity	8158	1	8158	4.44	0.048
		Error	36760 (82%)	20	1838		
		Total	44918	21			
	σ	Capacity	168	1	168	4.77	0.041
		Error	704 (82%)	20	35.2		
		Total	872	21			

Table 6.5: ANOVA table with access link capacity as factor and μ and σ parameters as response variables (in both directions)

squares for each source of variation: Capacity and Error. According to the results, the access link capacity factor shows moderate significance (that is, $p \approx \alpha = 0.05$) only in the test for the measurements in the incoming direction (both mean and standard deviation). On the other hand, the access link capacity factor has no influence for the case of mean and standard deviation in the outgoing direction of traffic.

Finally, the third column also shows the percentage that the error represents of the total variance (inside brackets). Although the ANOVA test shows that there exists some degree of influence between the access link capacity and the busy hour traffic, it can be noted that the amount of unexplained variance remains high after the test is applied. More specifically, these values are 88%, 94%, 82% and 82%, of the total variance for μ and σ in the outgoing and incoming directions respectively.

Indeed, such large values of error reinforce the conclusion that the access capacity barely influences the measurements, namely the measurements are not distorted by capping effects. Following this, the next section checks whether or not the other intrinsic network parameter, population size, explains more variance than that of the access link capacity.

6.4.2 Combined effect of the access link capacity and population size: ANCOVA

This section aims to repeat the previous experiment but taking into account both intrinsic network factors: the access link capacity and population size. In this case, the model of Eq. (6.10) is extended to:

$$y_{U_j}^{group} = k_y + \alpha^{group} + \beta^{group} P_{U_j} + \epsilon_{U_j}^{group} \tag{6.11}$$

where the term $\beta^{group} P_{U_j}$ has been included with respect to Eq. (6.10) to deal with the population size intrinsic feature of networks.

In this case, the population size factor appears as a quantitative variable rather than a categorical group as it is the case for the access link capacity. When this occurs, it is recommended to use the ANCOVA methodology instead of ANOVA. Additionally, ANCOVA is recommended when the two factors are strongly correlated since it helps to separate the amount of variance explained by each factor.

Basically, ANCOVA is the result of removing the variance for which some covariates or quantitative variables (in this case, the population size) account by means of a linear regression and, after this, applying a regular ANOVA with the access link capacity as unique factor. Note that such a linear regression does not assume that the value of the slopes β^{group} in Eq. (6.11) for groups G_{high} and G_{low} are equal.

Following this, Table 6.6 shows the results obtained after applying ANCOVA to the whole set of universities. The table shows a new row that quantifies the adjusted sum of squares of the explained variance by the population size as covariate. As shown, including the population size in the analysis brings two important conclusions: (i) the amount of total unexplained variance reduces very significantly with respect to the previous experiment; and (ii) the amount of variance explained by the access link capacity factor becomes negligible. Concerning the former conclusion, note that the amount of variance that remains unexplained for the mean busy-hour traffic μ has been reduced to 24% and 38% for the incoming and outgoing directions respectively. Remark that, in the previous model which only took into account the access link capacity, the unexplained variance was much higher, up to 82% and 88% (see Table 6.5) for the incoming and outgoing

Response Variable (direction)		Source of Variation	Sum of Squares	df	Adj. Mean Square	F	p-value
outgoing	μ	Popula.	68491	1	57569	25.6	0.000
		Capacity	2504	1	2504	1.11	0.305
		Error	42817 (38%)	19	2254		
		Total	113811	21			
	σ	Popula.	943	1	851	22.8	0.000
		Capacity	3.66	1	3.66	0.10	0.758
		Error	710 (58%)	19	37.4		
		Total	1657	21			
incoming	μ	Popula.	32014	1	25944	45.6	0.000
		Capacity	2089	1	2089	3.67	0.071
		Error	10815 (24%)	19	569		
		Total	44918	21			
	σ	Popula.	349	1	252	10.6	0.003
		Capacity	71.1	1	71.1	2.99	0.100
		Error	452 (56%)	19	23.8		
		Total	872	21			

Table 6.6: ANCOVA table with access link capacity as factor, population size as covariate, and μ and σ parameters as response variables (in both directions)

directions respectively.

Indeed, the amount of variance explained by the access link capacity was due to the correlation between the population size and the access link capacity of universities, rather than on the latter factor only. This null effect of access link capacity is consistent with the premise of negligible capping effect (Section 6.2.3).

We believe that the population size has a more clear impact in the busy hour traffic in the incoming direction than in the outgoing direction. Apparently, it seems to be the university which is responsible for the amount of traffic observed in the incoming direction (the university users request such information), whereas in the outgoing direction, it is the Internet's population (unknown) which request such traffic. Probably, the amount of traffic in the outgoing direction is proportional to the network services (web services, Virtual Private Networks (VPNs), databases, etc) that universities offer to the outside world, which also have a relationship to the population size of universities, but not as strong as in the incoming direction of traffic.

Given that the access link capacity is not significant, the following section is focused on a simplified model that only takes into account the population size via linear regression.

6.4.3 Focusing on the population size: Linear Regression

As stated before, ANCOVA performs a linear regression in order to remove the variance explained by the covariates. Next, we assess whether such a linear regression can be useful to estimate the busy hour traffic distribution $N(\mu, \sigma)$ based on the university's population size only.

Note that the previous model (Eq. (6.11)) assumes a different β^{group} for each group of universities G_{high} and G_{low}. The next model simplifies this issue assuming a common slope β for all universities. Such assumption is known as the homogeneity of regression coefficients [All97]. We did not find any evidence that such assumption is violated, consequently we can use the same β parameter for all groups (G_{high} and G_{low}), given by the following simplified model:

$$y_{U_j} = k_y + \beta P_{U_j} + \epsilon_{U_j} \tag{6.12}$$

which only takes into account the university's population size P_{U_j} as the only source of influence in the busy-hour traffic distribution $N(\mu, \sigma)$. We remark that the value of β represents the slope in the linear regression model, and can be viewed as the amount of traffic that each network user contributes to the average busy-hour traffic value μ_{U_j}. This is a parameter of key importance in the capacity planning of university access links based on their population size.

After applying the linear regression, Table 6.7 shows the regression coefficients for each response variable, together with their 95% confidence intervals. The fourth column in the table provides the coefficient of determination, $\bar{R}^2$, which gives the amount of variance explained by the linear regression model. The results show that the population size explains 62% and 71% of the variance of the busy-hour process mean μ in the outgoing and incoming directions of traffic respectively. For σ, the experiment results give 56% and 40% of explained variance, again in the outgoing and incoming directions respectively.

Response Variable (direction)		Coefficients (Mb/s)	95%-confidence intervals	$\bar{R}^2$
outgoing	μ	k = -24.960 $\beta = 0.0034$	-63.024 / -13.104 0.0021 / 0.0047	62.18%
	σ	$k = 0.44$ $\beta = 0.0004$	-4.339 / 4.251 0.0002 / 0.0006	56.09%
incoming	μ	$k = 5.10$ $\beta = 0.0024$	-15.211 / 25.411 0.0017 / 0.0031	71.27%
	σ	$k = 3.117$ $\beta = 0.00025$	-0.969 / 7.206 0.00011 / 0.00039	40.10%

Table 6.7: Regression coefficients for μ and σ in both directions

Furthermore, Figure 6.5 shows the regression lines estimated by ANCOVA for each parameter along with the data, on attempts to provide a visual contrast of the results.

Concerning capacity planning, a given university U_j with population size P_{U_j} requires a capacity in the incoming direction of $k = 5.1$ Mb/s constant plus 2.4 Kb/s (and depending on the selected significance level an extra addend that takes into account the standard deviation (240 b/s)) per user, as given by Table 6.7, since the error ϵ is zero-mean Gaussian distributed (modeled with $N(0, \sigma_\epsilon)$ and the representativeness of $\bar{R}^2$. This provides a simple rule for dimensioning the access link capacity of a new university based on its expective population (number of users), as:

$$
\begin{array}{rcl}
C_{U_j} & \text{such that} & \text{Prob } (d \cdot X^{U_j} > C_{U_j}) \leq \varepsilon, \\
 & \text{with} & X^{U_j} \sim N(\mu_{U_j}, \sigma_{U_j}) \\
 & \text{where} & \mu_{U_j} = k_\mu + \beta_\mu P_{U_j}, \\
 & \text{and} & \sigma_{U_j} = k_\sigma + \beta_\sigma P_{U_j} \\
 & & \text{for each of the directions}
\end{array}
\tag{6.13}
$$

This constitutes a further refinement of Eqs. (6.1) and (6.9).

This methodology makes it possible to estimate the demand for bandwidth

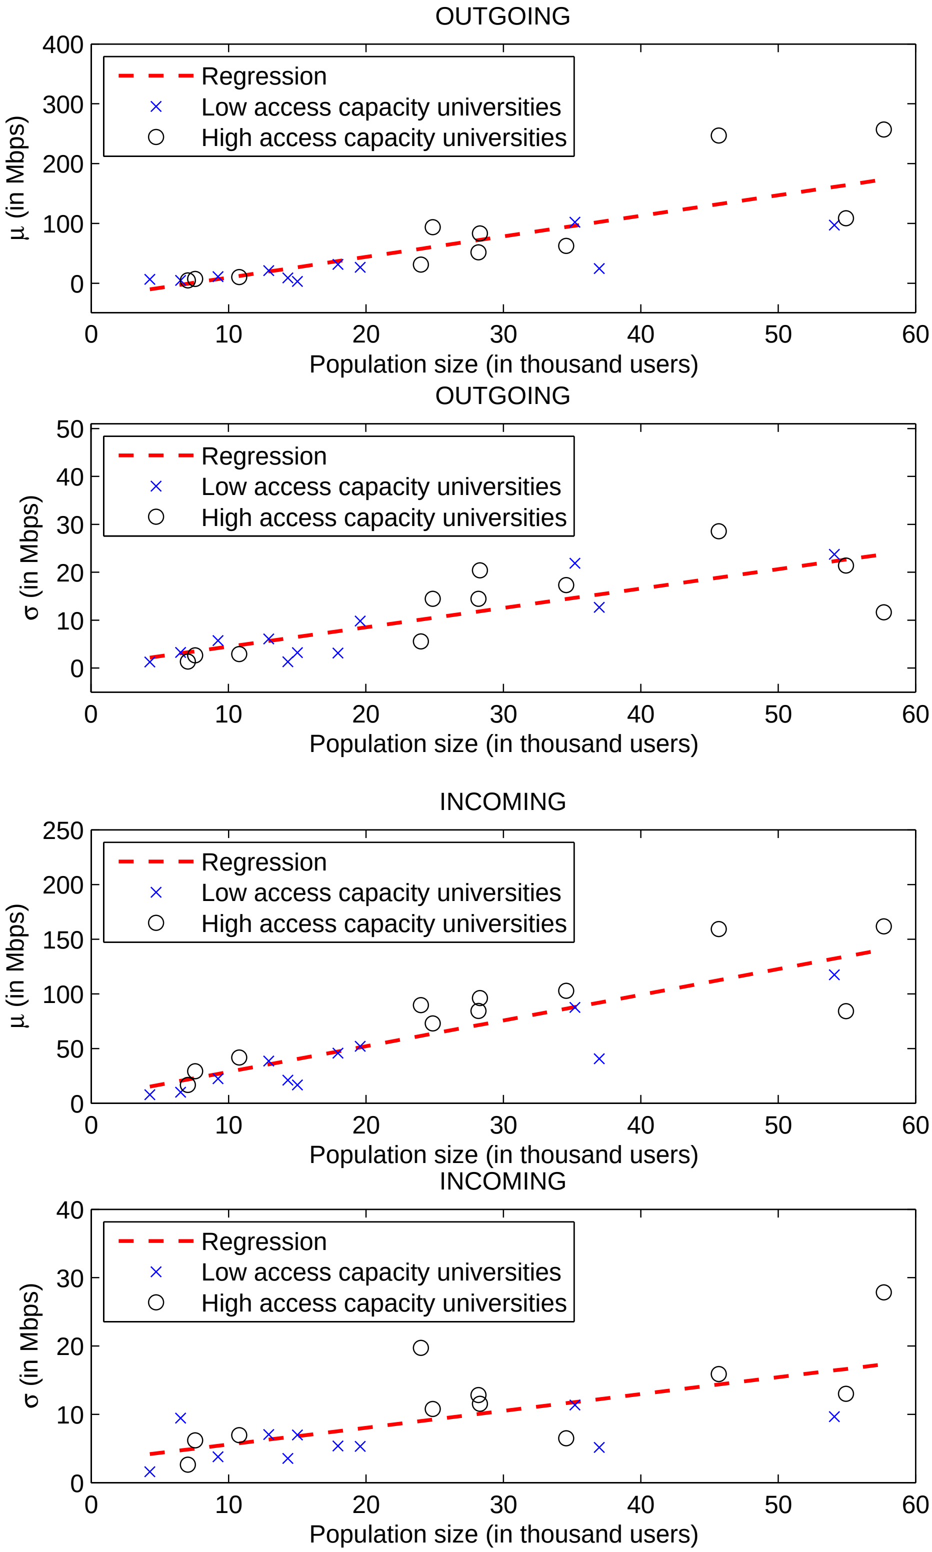

Figure 6.5: ANCOVA linear regression for the μ and σ parameters in both directions of traffic

for new universities, over which no previous measurement experiments have been carried out, in contrast with Eq. (6.9) which requires a set of busy-hour daily traffic measurements. Furthermore, Eq. (6.13) can be used to estimate the bandwidth demands for a university network whose population changes with time, that is, whose student body either increases or decreases every academic year.

6.4.4 On the relationship between heavy-hitters and population size

Previous studies have pointed out that most of the Internet traffic is generated by a small fraction of network users [Bro02, PTB+02], often referred to as *heavy-hitters* [FGL+01], this phenomenon was introduced in Section 2.3.2. According to the definitions presented in such chapter, Figure 6.6 shows that there is a clear correlation between the population size of a given university and the number of heavy-hitters observed during its busiest hour. Figure 6.6(a) considers as heavy-hitters all those IP addresses which account for 90% of the total traffic. Figure 6.6(b) defines heavy hitters as those users who exchange more than 1 Gbit (about 100 times the average) of traffic, and Figure 6.6(c) marks as heavy hitters to the users than account for more than 0.5% of the total traffic, in all cases measured in the busy hour. The points depicted are computed as the average number of heavy-hitters per day found during the busy-hour over the six-month experiment. In the plots, we have removed those university networks in which the use of Network Address Translation (NAT) is a common practice, resulting in a set of eleven networks under study.

These results give support to the idea that heavy-hitters are homogeneously distributed with respect to the population across the universities. In other words, the larger a university is in terms of population size, the more number of heavy-hitters is expected to be found in its busy-hour traffic measurements. Nevertheless, it is more interesting to define link dimensioning rules based on well-documented intrinsic characteristics such as the population size of universities rather than on measurement-based metrics such as the number of heavy-hitters since the latter requires extensive measurement experiments and computational analysis.

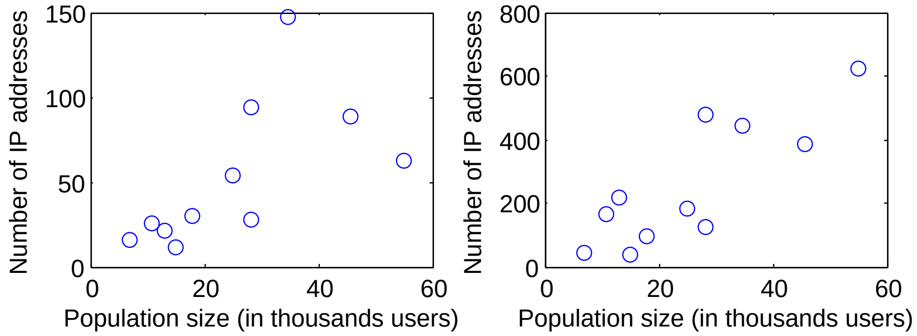

(a) IP addresses that account for 90% of the total traffic volume (outgoing/incoming) during the busy hour per day

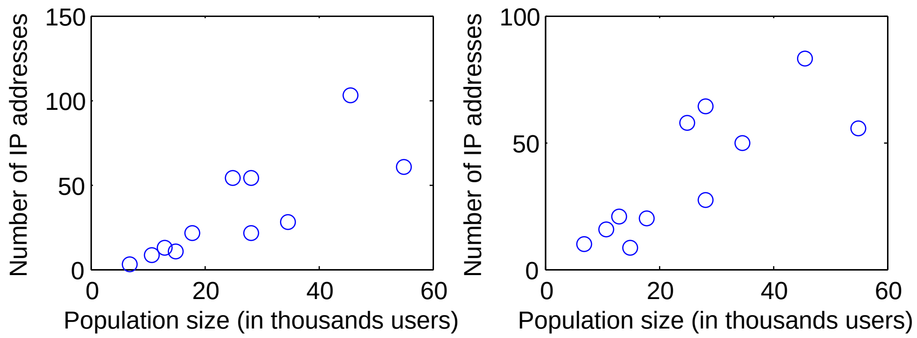

(b) IP addresses that send/receive more than 1 Gbit (outgoing/incoming) during the busy hour per day

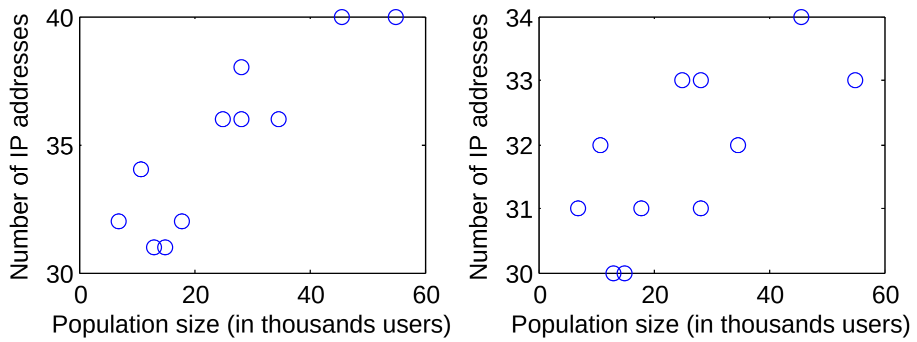

(c) IP addresses that account for more than 0.5% of the total traffic volume (outgoing/incoming) during the busy hour per day

Figure 6.6: Average number of different IP addresses per day according to several heavy-hitter definitions

6.5 Conclusions

This chapter provides an in-depth analysis on the traffic volumes observed during the busy hour of the access links of universities, regional networks, and external links in the Spanish Research and Education Network from a mid-term point of view. First, it is shown that such busy-hour traffic is Gaussian distributed, and shows no correlation between measurements over different days, hence accurately characterized by a white Gaussian process. Therefore, the collection of measurements over different days can be used to estimate the mean and standard deviation of such a white Gaussian model for capacity planning purposes. Nevertheless, the network operator must use this methodology only while the stationarity of the measurements remains, that is, no infrastructure upgrades, new killer applications appraisal, P2P filtering policy changes, etc. In such cases, the operator should restart the traffic measurement campaign.

In this regard, we proposed in [MAGD09] an algorithm to detect stationary changes of demand for bandwidth in links whose load is increasing over time. The result showed that the variation of the load over time (not necessarily during the busy hour) can be characterized by a Gaussian process in periods of some weeks or months in several links of the RedIRIS' network. That is, the increase/decrease of the demand for bandwidth are not changing at a constant rate. But, it changes periodically from a stationary process, in this case, a Gaussian process with certain mean and variance, to another stationary process with different parameters. This supports the hypothesis that the demand for bandwidth during the busy hour can be modeled as a stationary process at the mid-term, as we have already shown in this chapter for some months in the RedIRIS' network.

Additionally, this chapter goes one step further and aims to characterize the mean and variance of such process based on the population of the network for which the measured access link gives service. The ANOVA and ANCOVA methodology are applied over the Gaussian models that characterize the busy-hour traffic volumes measured for different universities on attempts to check whether or not the universities' intrinsic features (population size and access link capacity) explains part of the busy-hour traffic volume generated. The experiments show that the access link capacity feature shows little influence on the busy-hour traffic for networks whose maximum utilization are far from reaching the maximum available capacity. On the other hand, the population size accounts for the majority of explained variance in the ANCOVA test. Furthermore, the test provides a linear regression model and estimates its parameters, making possible the task of dimensioning access links for university networks based on their population size. However, after applying ANCOVA the unexplained variance still accounted for some percentage. The use of more features may improve the results that we have shown obtaining more accurate estimations. To this end, we plan to find and analyze these other features, such as, for instance, the sort of university under study (technical versus non-technical) or the ratio between staff members and student body.

Chapter 7

Conclusions and future work

This chapter first summarizes the main contributions and conclusions for the research described in this thesis (Section 7.1). Then, in Section 7.2, the hypotheses and objectives presented in the first chapter are reviewed . Finally, Section 7.3 identifies possible directions for future work.

7.1 Contributions and conclusions

The results of this thesis can be summarized in four items. First, we have analyzed the representativeness and generality of the Internet traffic measurements. Second, we have addressed the problem of subsampling the traffic captures. Third, we have characterized the Internet traffic busy hour and we have shown how it can be used to tackle the network capacity planning problem. Finally, we have explained the relation between the traffic in the busy hour and the intrinsic features of IP networks.

1. *We have analyzed the representativeness and generality of the Internet traffic measurements*, which were kindly donated by RedIRIS. We have shown that although the frequency statistics of IP addresses and port numbers of an extensive set of university networks follow a Zipf distribution, the Zipf's parameter is very different from one network to another, even for networks with very similar features. We have further analyzed the distribution of the geolocation of the Internet connections obtained similar conclusions although, in

this case, we have modeled the measurements by means of a Zipf-Mandelbrot distribution and the Analysis of Variance (ANOVA). Furthermore, it has been shown that more than one month worth of measurements are necessary to obtain a sampled distribution that is stable in time. In other words, a large time horizon is required in order to capture stationarity. Hence, a very large measurement experiments must be performed, both in duration and spatial diversity, in order to generalize the derived conclusions. Finally, we have shown that the erratic and non-homogeneous traffic patterns of a small set of users, known as heavy-hitters, causes this behavior.

Chapters 3 and 4 include this contribution, which has led the following publications:

– José Luis García-Dorado, José Alberto Hernández, Javier Aracil, Jorge E. López de Vergara, Francisco Montserrat, Esther Robles, and Tomás P. de Miguel, *On the duration and spatial characteristics of Internet traffic measurement experiments,* IEEE Communications Magazine **46** (2008), no. 11, 148-155.

– F. Mata, J. L. García-Dorado, J. Aracil, and J. E. López de Vergara, *Factor analysis of Internet traffic destinations from similar source IP subnetworks,* submitted to Elsevier Computer Networks.

– J. L. García-Dorado, J. E. López de Vergara, J. Aracil, V. López, J. A. Hernández, S. Lopez-Buedo, and L. de Pedro, *Utilidad de los flujos Netflow de RedIRIS para análisis de una red académica (On the use of RedIRIS' Netflow information for academic networks),* Boletín de RedIRIS, no. 82-83, Jornadas Técnicas RedIRIS 2007 (Mieres, Spain), November 2007.

2. *We have shown a new method to subsample optimally network traffic measurements.* This method is based on the Multiresolution Analysis (MRA) with wavelets, thus subsampling the measurements in both time and frequency planes. Moreover, it is based on metrics that are related to the traffic queueing behavior instead of the typical Euclidean distance and Mean Squared Error (MSE) methodologies. This method has been applied to an

extensive set of real measurements collected from RedIRIS. The results show that it is possible to reduce the data to one-fourth of its original size for the traffic generated by most universities, and even to one-eighth for data collected from routers with more aggregated traffic, both with a high level of confidence.

Chapter 5 includes this contribution, which has led the following publication:

– José Luis García-Dorado, Javier Aracil, José Alberto Hernández, and Jorge E. López de Vergara, *A queueing equivalent thresholding method for thinning traffic captures,* in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (Salvador, Brazil), April 2008.

3. *We have analyzed the dynamics of Internet traffic busy hour.* We have found that the average traffic during the busy hour through several months can be modeled by a Gaussian distribution with no correlation between measurements over different days (white noise) and subexponential behavior. This finding provides a refinement for capacity planning purposes, whereby it is possible to design the capacity needed such that the traffic volume does not exceed the network capacity with a given probability.

   Section 6.3 of this thesis includes this contribution.

4. Finally, *we have found relations between the intrinsic features of a network,* basically the population and the network infrastructure, *and the traffic that goes through it.* Specifically, we have used them to predict the Gaussian distribution parameters that model the traffic of the busy hour. Knowing that the busy hour follows a normal distribution, we have used ANOVA and the Analysis of Covariance (ANCOVA) to determine in what extent the observed variance is explained by each intrinsic network feature. Consequently, given the characterization of a network, infrastructure and population, an operator can estimate the traffic that a network is expected to receive/generate. Thus, Internet operators can dimension their links with a formal and objective methodology instead of the current arbitrary rules based on previous experiences or those methods that estimate the traffic in a short-term fash-

ion. In addition, an operator can foresee the impact of changes on the intrinsic features (network upgrades and changes on the population size, for instance) will have on the demand for bandwidth or on the demand for any other network resource.

This contribution explained in Chapter 6 together with the previous one make up the following article:

– J. L. García-Dorado, J. A. Hernández, J. Aracil, J. E. López de Vergara, and S. Lopez-Buedo, *Characterization of the busy-hour traffic of IP networks based on their inherent features,* submitted to IEEE/ACM Transactions on Networking.

7.2 Assessment of the objectives and hypotheses

This section reviews the hypotheses and objectives posed in the first chapter.

– *Hypothesis: Traffic measurements gathered from a limited number of networks and limited duration cannot be considered to be sufficiently representative of the Internet.*

  In the chapters 3 and 4 we have shown that although some internetwork statistics (specifically, the popularity of IP addresses and port numbers and the geolocation of the Internet connections) can be modeled by well-known probability distributions, the distributions' parameters are very different from one network to another, and more than one month worth of data are necessary to obtain a sampled distribution that is stable in time. Thus, a limited number of networks and duration of the measurement campaigns cannot be sufficient to obtain representative conclusions.

– *Hypothesis: If the Internet traffic measurement campaigns must last for long periods of time, the volume of data that such campaigns entails can result by itself difficult to analyze, monitor, and store.*

  Chapter 5 has shown a novel method to subsample traffic measurements over time in order to reduce the computational load of analyzing and monitoring such measurements as well as their storage capacity requirements. This method subsamples the measurements in both time and frequency domains with subsampling-level thresholds based on close-related metrics to traffic queueing behavior. The results have shown how the proposed method outperforms previous well-known techniques.

– *Hypothesis: The demand for bandwidth in the busy hour over mid-length periods can be characterized by a stochastic process.*

  This issue has been tackled in Section 6.3, we have found that a white Gaussian process models the average traffic during the busy hour through several months in an extensive set of networks.

– *Hypothesis: The demand for bandwidth during the busy hour over long periods calls for a non-stationary process model.*

  In Section 6.5 it was pointed out that the RedIRIS institutions' demands for bandwidth (not necessarily during the busy hour) increase/decrease as a staircase function with time intervals between two consecutive changes in the range of several weeks or months. This suggests that, within the time interval between two change occurrences, the busy-hour traffic demand can be modeled by a stationary process. Therefore, in the long-term case, the busy hour process' parameters call for a new estimation every time a change occurs.

– *Hypothesis: The demand for bandwidth in low-utilized networks are not polluted by their access capacities. As RedIRIS networks' utilizations are typically low, we support the hypothesis that access capacities are not "capping" the demand of the users.*

  Section 6.4.1 has shown how the factor *Capacity* is not a significant factor by applying ANCOVA methodology. This implies that, in the case of the networks under study with low utilization, the demand for bandwidth per user does not depend on the access capacity of the network.

– *Hypothesis: The process that models the demand for bandwidth over time can be estimated by means of the networks' intrinsic features. Consequently, the demand for bandwidth in a network can be estimate in an objective and* *fairly fashion, and, even, avoiding to carry out a dedicated measurement campaign.*

The mean and standard deviation of the process that models the busy hour can be fairly estimated by means of the population size, as shown in Section 6.4.2. It allows the operators to base their estimations of the demand for bandwidth on intrinsic features with the mentioned advantages.

7.3 Future work

We suggest some future research lines to continue the work presented in this thesis:

– *Characterization of Internet behavior.* In this thesis we have focused on the characterization of several facets of the Internet behavior. To this end, we have used actual measurements but only from campus networks. Thus, we would like to apply the proposed methodologies to non-academic networks and compare the results. The final aim would be determinate if the conclusions obtained in some RedIRIS' institutions also hold in other kinds of networks.

– *Geolocation of the Internet connections.* We plan to extend the geolocation analysis studying not only the distribution of the connections by country but also by region and even by Internet Service Provider (ISP) and Autonomous System (AS).

– *Internetwork measurement subsampling.* In Chapter 5 we have proposed a novel method to subsample internetwork measurement, the "queueing equivalent thresholding method". As further work, we will apply some typical data mining algorithms (for instance, Principal Component Analysis and clustering algorithms) to perform network capacity planning tasks over the measurements previously subsampled applying the proposed method.

– *Stationarity of the busy hour process.* We plan to pay spatial attention on the stationarity of the demand for bandwidth in the busy hour. We have shown that its stationary holds in RedIRIS networks for some months, however, it is well-known that the demand for bandwidth varies over time.

This implies the traffic mean and variance in the busy hour may change over time. Consequently, we plan to study the duration and occurrence of the temporal frames in which stationarity can be assumed.

- *Overdimensioning factor.* In the state-of-the-art chapter we have presented the overdimensioning factor (named as d parameter). That is, the extra bandwidth capacity over the traffic mean that a link requires to meet the Service Level Agreements (SLA). We think that this parameter can be also estimated using the intrinsic features of the networks.

- *Busy-hour model.* We are interested in including more features to ANOVA and ANCOVA methodologies. Actually, campus networks can be characterized by more features than the population size and the access capacity, for instance, ratio between student and teaching bodies, accessibility to the Internet, filtering policies, typology of the campus network (private/public, technical university/social science university, etc.) among others. All these features can be known in advance for a campus network and be useful to improve the accuracy of the model.

Conclusiones y trabajo futuro

Este capítulo, primero, resume las principales contribuciones y conclusiones obtenidas en esta tesis. A continuación, las hipótesis y objetivos presentados en el primer capítulo son revisados. Finalmente, se determinan posibles continuaciones al trabajo aquí desarrollado.

Contribuciones y conclusiones

Los resultados de esta tesis pueden resumirse en cuatro puntos. Primero, hemos analizado la representatividad y generalidad de las medidas de tráfico de Internet. Segundo, hemos tratado el problema del submuestreo de capturas de red en el tiempo. Tercero, hemos caracterizado la hora cargada de Internet y hemos mostrado como puede resultar útil en el problema de asignación de capacidad a la redes. Finalmente, hemos explicado la relación entre el tráfico durante la hora cargada y las características inherentes de las redes IP.

1. *Análisis de la representatividad y generalidad de las medidas de tráfico de Internet,* las cuales han sido amablemente cedidas por RedIRIS. En esta tesis se ha comprobado que, aunque la frecuencia estadística de la popularidad de las direcciones IP y puertos de un numeroso conjunto de redes académicas siguen una distribución Zipf, los parámetros de esta distribución son muy diferentes en un conjunto extenso de redes universitarias, incluso cuando las redes presentan características muy similares. Al analizar la distribución de la geolocalización de las conexiones de Internet se han obtenido resultados equivalentes, aunque en este caso las medidas han sido modeladas con una distribución Zipf-Mandelbrot y Analisis de Varianza (ADEVA). Además, se ha concluido que es necesario más de un mes de medidas para obtener una

distribución muestral que permanezca estable en el tiempo. En otras palabras, un extenso horizonte temporal es requerido para que cierta medidas muestren estacionariedad. Por lo tanto, las campañas de captura de medidas deben afrontar la diversidad espacial y temporal para obtener conclusiones realmente representativas del tráfico de Internet. Finalmente, hemos comprabado que el comportamiento errático y no homogéneo de un pequeño conjunto de usuarios causa este fenómeno.

Las capítulos 3 y 4 incluyen esta contribución, que ha dado lugar a los siguientes artículos:

- José Luis García-Dorado, José Alberto Hernández, Javier Aracil, Jorge E. López de Vergara, Francisco Montserrat, Esther Robles, and Tomás P. de Miguel, *On the duration and spatial characteristics of Internet traffic measurement experiments,* IEEE Communications Magazine **46** (2008), no. 11, 148-155.
- F. Mata, J. L. García-Dorado, J. Aracil, and J. E. López de Vergara, *Factor analysis of Internet traffic destinations from similar source IP subnetworks,* enviado a Elsevier Computer Networks.
- J. L. García-Dorado, J. E. López de Vergara, J. Aracil, V. López, J. A. Hernández, S. Lopez-Buedo, and L. de Pedro, *Utilidad de los flujos Netflow de RedIRIS para análisis de una red académica,* Boletín de RedIRIS, no. 82-83, Jornadas Técnicas RedIRIS 2007 (Mieres, Spain), November 2007.

2. *Esta tesis propone un nuevo método para submuestrear medidas de tráfico de red.* Este método esta basado en el Análisis Multiresolución (MRA) con wavelets, de modo que el submuestreo de las medidas se hace tanto en el dominio del tiempo como en el de la frecuencia. Además, se basa en métricas estrechamente relacionadas con la teoría de colas, en vez de en distancias euclideas o mínimos cuadrados. Este método ha sido aplicado a un conjunto extenso de medidas reales de red capturadas en RedIRIS. Los resultados muestran que es posible reducir a un cuarto de su tamaño original el tráfico generado por la mayoría de las universidades, e incluso, a un octavo para

medidas tomadas en routers con mayor agregación de tráfico, ambas con altos niveles de confianza.

El capítulo 5 incluye esta contribución, que ha dado lugar al siguiente artículo:

– José Luis García-Dorado, Javier Aracil, José Alberto Hernández, and Jorge E. López de Vergara, *A queueing equivalent thresholding method for thinning traffic captures,* in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (Salvador, Brazil), April 2008.

3. *Análisis de la dinámica de la hora cargada de Internet.* En esta tesis hemos comprobado que el volumen de tráfico intercambiado durante la hora cargada durante varios meses puede ser modelado mediante una distribución normal, sin correlación entre las medidas a lo largo de los días (ruido blanco) y comportamiento subexponencial. Este resultado representa un avance en el problema de asignación de capacidad a la redes, puesto que hace posible determinar la capacidad necesaria tal que el volumen de tráfico intercambiado no exceda con cierta probabilidad dicha capacidad.

   La sección 6.3 de esta tesis incluye esta contribución.

4. Finalmente, *hemos relacionado las características inherentes de una red,* esencialmente la población y la infraestructura de red, *y el tráfico que la atraviesa.* En concreto, hemos usado estas características para predecir los parámetros de la distribución normal que modela el tráfico de la hora cargada. Sabiendo que la hora cargada sigue la distribución normal, hemos usado ADEVA y Analisis de Covarianza (ADECOV) para determinar en que medida la varianza observada es explicada por cada característica inherente. Consecuentemente, dada la caracterización de una red, infraestructura y población, un operador puede estimar el tráfico que se espera que la red genere y reciva. De este modo, los operadores de Internet pueden dimensionar sus enlaces con una metodología formal y objetiva en lugar de las actuales reglas arbitrarias basadas en experiencias previas o en métodos válidos a corto plazo. Además, un operador puede preveer el impacto que

los cambios en las características inherentes (actualización de la red y variaciones en el número de usuarios, por ejemplo) conllevarán en la demanda de ancho de banda o en la demanda de cualquier otro recurso de red.

Esta contribución expuesta en el capítulo 6 junto con la contribución anterior forman el siguiente artículo:

– J. L. García-Dorado, J. A. Hernández, J. Aracil, J. E. López de Vergara, and S. Lopez-Buedo, *Characterization of the busy-hour traffic of IP networks based on their inherent features,* submitted to IEEE/ACM Transactions on Networking.

Evaluación de las hipótesis y los objetivos

Esta sección examina las hipótesis y objetivos planteados en el primer capítulo.

– *Hipótesis: Las campañas de captura de medidas de tráfico sobre un número limitado de redes y de duración acotada no son suficientemente representativas de Internet.*

  En los capítulos 3 y 4 hemos mostrado que aunque algunas estadísticas de Internet (en concreto, la popularidad de las direcciones IP y puertos, y la geolocalización de las conexiones de Internet) pueden ser modeladas usuando distribuciones de probabilidad bien conocidas, los parámetros de estas distribuciones son muy diferentes de una red a otra, y más de un mes de medidas es necesario para obtener una distribución muestral estable en el tiempo. De este modo, un pequeño número limitado de redes y un tiempo acotado de captura de medias puede no ser suficiente para obtener conclusiones representativas.

– *Hipótesis: Si las campañas de captura de medidas de Internet deben durar peridos largos de tiempo, el volumen de datos que estas campañas conllevan puede resultar, por si mismo, difícil de analizar, monitorizar y almacenar.*

  El capítulo 5 ha mostrado un nuevo método que submuestrea medidas de red para reducir el coste computacional de analizar y monitorizar tales medidas así como los requisitos de almacenamiento. Este método submuestrea las medidas en el dominio del tiempo y de la frecuencia con umbrales de submuestreo basados en métricas relacionadas con la teoría de colas. Los resultados han demostrado como el método propuesto mejorar técnicas previas bien conocidas.

– *Hipótesis: La demanda de ancho de banda durante la hora cargada a medio plazo puede ser caracterizada por un proceso estocástico.*

  Esta cuestión fue analizada en la sección 6.3, se comprobó que un proceso blanco gaussiano modela el tráfico medio durante la hora cargada durante varios meses en un conjunto extenso de redes.

– *Hipótesis: La demanda de ancho de banda durante la hora cargada a largo plazo requiere un modelo no estacionario.*

  En la sección 6.5 se señaló que la demanda de ancho de banda (no necesariamente durante la hora cargada) crece/decrece como una función a escalones con intervalos entre cambios consecutivos en el rango de semanas o meses. Esto sugiere que en los periodos entre cambios, la demanda de tráfico durante la hora cargada puede ser modelada como un proceso estacionario. De este modo, a largo plazo, los parámetros del proceso que modela la hora cargada pueden requerir una nueva estimación cada vez que un cambio ocurre.

– *Hipótesis: La demanda de ancho de banda en redes con baja utilización no está limitada por la capacidad de acceso de las redes. Como la utilización de las redes de RedIRIS es típicamente baja, pensamos que las capacidades de acceso no están limitando la demanda de los usuarios.*

  La sección 6.4.1 mostró como el factor *Capacity* era un factor no significante al aplicar ADECOV. Esto implica que, en el caso de las redes en estudio con baja utilización, la demanda de ancho de banda por usuario no depende de la capacidad de acceso de la red.

– *Hipótesis: Los parámetros del proceso que modela la demanda de ancho de banda en el tiempo pueden ser estimados usando las características inherentes de la redes. Consecuentemente, la demanda de ancho de banda en una red puede ser estimada de una manera objetiva, evitando incluso la captura de medidas de red.*

La media y desviación estándard del proceso que modela la hora cargada pueden ser ajustadamente estimados por medio del tamaño de la población, como fue mostrado en la sección 6.4.2. Esto permite a los operadores basar sus estimaciones de la demanda de ancho de banda en características inherentes con las ventajas ya mencionadas.

Trabajo Futuro

En esta sección se presentan posible lineas de investigación que continuan el trabajo presentado en esta tesis:

- *Caracterización del comportamiento de Internet.* En esta tesis nos hemos centrado en la caracterización de varias facetas del comportamiento de Internet. Para ello, hemos usado medidas reales de tráfico, pero sólo de redes universitarias. De este modo, queremos aplicar las metodologías propuestas a redes no académicas y comparar los resultados. El objetivo final sería determinar si las conclusiones obtenidas en una serie de instituciones académicas también son validas en otros escenarios.

- *Geolocalización de las conexiones de Internet.* El estudio de la geolocalización de las conexiones de Internet se puede extender para incluir el análisis no sólo a nivel de paises, sino también a nivel regional o incluso por proveedores de servicios de Internet y sistemas autónomos.

- *Submuestreo de medidas de red.* En el capítulo 5 hemos propuesto un nuevo método para submuestrear medidas de red. Como trabajo futuro, aplicaremos varios algorítmos del area de la minería de datos (por ejemplo, Análisis de Componentes Principales y algorítmos de clustering) sobre las medidas de red previamente submuestreadas con este método, y así, reevaluar el rendimiento.

- *Estacionariedad del proceso hora cargada.* En esta tesis se ha mostrado que la estacionariedad puede ser asumida en muchas de las instituciones de RedIRIS durante varios meses, sin embargo, es bien conocido que la demanda de ancho de banda cambia con el tiempo. Esto implica que la media y varianza del tráfico en la hora cargada puede cambiar con el tiempo. Consecuentemente, hemos planificado estudiar la duración y ocurrencia de las ventanas temporales en las cuales la estacionariedad puede ser asumida.

– *Factor de sobredimensionado.* En el capítulo de estado del arte se introdujó el factor de sobredimensionado (denominado como parámetro d). Esto es, la capacidad extra que se requiere sobre la media del tráfico para cumplir los acuerdos de calidad de servicio. Pensamos que este parámetro puede ser también estimado usando las características inherentes de la redes.

– *Modelo hora cargada.* Estamos interesados en incluir más características a los modelos ADEVA y ADECOV. De hecho, las redes de la universidades pueden ser caracterizadas con más características que el número de usuarios y la capacidad de acceso, por ejemplo, el ratio entre el número de estudiantes y profesores, accesibilidad a Internet, políticas de filtrado, clases de universidades (privada/pública, técnicas/humanísticas, etc.) entre otras. Todas estas características pueden resultar útiles para mejorar la precisión del modelo propuesto.

References

[AB97] R. E. Ahmed and S. H. Bakry, *New topology designs for the future expansion of the academic network of the gulf countries*, International Journal of Network Management **7** (1997), 18–32.

[AFT98] R. Adler, R. Feldman, and M. S. Taqqu (eds.), *A practical guide to heavy tails: Statistical techniques for analyzing heavy tailed distributions*, Birkhäuser Verlag, Boston, USA, 1998.

[AH02] L. A. Adamic and B. A. Huberman, *Zipf's law in the Internet*, Glottometrics **3** (2002), 143–150.

[All97] M. P. Allen, *Understanding regression analysis*, Springer-Verlag, New York, USA, August 1997.

[All99] J. R. Allen, *Driving by the rear-view mirror: Managing a network with Cricket*, In Proceedings of the USENIX Conference on Network Administration (Santa Clara, USA), April 1999, pp. 1–10.

[All08] J. Allspaw, *The art of capacity planning*, O'Reilly Media, September 2008.

[ALML06] J. Almhanaa, Z. Liua, R. McGorman, and B. Lanteignea, *Bandwidth provisioning and capacity planning tools*, In Proceedings of Annual Communication Networks and Services Research Conference (Moncton, Canada), May 2006, pp. 328–330.

[AV98] P. Abry and D. Veitch, *Wavelet analysis of long-range-dependent traffic*, IEEE Transactions on Information Theory **44** (1998), no. 1, 2–15.

[AW97] M. F. Arlitt and C. L. Williamson, *Internet Web servers: workload characterization and performance implications*, IEEE/ACM Transactions on Networking **5** (1997), no. 5, 631–645.

[BAD99] S. H. Bakry and A. Al-Dhelaan, *Plannig information networks: The scope and the customer issues*, International Journal of Network Management **9** (1999), 28–37.

[BBD05] R. Bolla, R. Bruschi, and F. Davoli, *Capacity planning in IP virtual private networks under mixed traffic*, Computer Networks **50** (2005), no. 8, 1069–1085.

[BC02] N. Brownlee and K. Claffy, *Understanding Internet traffic streams: Dragonflies and tortoises*, IEEE Communications Magazine **40** (2002), no. 10, 110–117.

[Ber06] D. López Berzosa, *Contribución a la caracterización de los requisitos de calidad de los servicios de telecomunicación*, Ph.D. thesis, Universidad Politécnica de Madrid, November 2006.

[BK00] A. W. Berger and Y. Kogan, *Dimensioning bandwidth for elastic traffic in high-speed data networks*, IEEE/ACM Transactions on Networking **8** (2000), no. 5, 643–654.

[BKPR02] P. Barford, J. Kline, D. Plonka, and A. Ron, *A signal analysis of network traffic anomalies*, In Proceedings of ACM SIGCOMM Workshop on Internet Measurement (Marseille, France), November 2002, pp. 71–82.

[BM95] R. H. Brown and I. Matin, *Development of artificial neural network models to predict daily gas consumption*, In Proceedings of the IEEE International Conference on Industrial Electronics, Control, and Instrumentation, vol. 2, November 1995, pp. 1389–1394.

[BM03] A. Z. Broder and M. Mitzenmacher, *Survey: Network applications of bloom filters: A survey*, Internet Mathematics **1** (2003), no. 4, 636–646.

[BMM+07] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, and P. Tofanelli, *Revealing Skype traffic: When randomness plays with you*, ACM SIGCOMM Computer Communication Review **37** (2007), no. 4, 37–48.

[BMM+09] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, and P. Tofanelli, *Detailed analysis of Skype traffic*, IEEE Transactions on Multimedia **11** (2009), no. 1, 117–127.

[Bor00] M. S. Borella, *Source models of network game traffic*, Computer Communications **23** (2000), no. 4, 403–410.

[BR02] N. Benameur and J. W. Roberts, *Traffic Matrix Inference in IP Networks*, In Proceedings of International Telecommunication Network Strategy and Planning Symposium (München, Germany), June 2002, pp. 103–114.

[Bro02] N. Brownlee, *Understanding Internet traffic streams: Dragonflies and tortoises*, IEEE Communications Magazine **40** (2002), 110–117.

[BYCE07] R. Baeza-Yates, C. Castillo, and E. N. Efthimiadis, *Characterization of national Web domains*, ACM Transactions on Internet Technology **7** (2007), no. 2, 9.

[CB97] M. E. Crovella and A. Bestavros, *Self-similarity in World Wide Web traffic: Evidence and possible causes*, IEEE/ACM Transactions on Networking **5** (1997), 835–846.

[CB05] B. Choi and S. Bhattacharyya, *Observations on Cisco sampled Netflow*, ACM SIGMETRICS Performance Evaluation Review **33** (2005), no. 3, 18–23.

[CEBRSP] V. Carela-Español, P. Barlet-Ros, and J. Sole-Pareta, *Traffic classification with sampled netflow*, Tech. report, `http://personals.ac.upc.edu/pbarlet/reports/netflow_classification-techrep.pdf`.

[CFSD90] J. Case, M. Fedor, M. Schoffstall, and J. Davin, *Simple Network Management Protcol*, RFC 1157, May 1990.

[CFY03] F. K.-P. Chan, A. W.-C. Fu, and C. Yu, *Haar wavelets for efficient similarity search of time-series: with and without time warping*, IEEE Transactions on Knowledge and Data Engineering **15** (2003), no. 3, 686–705.

[Chu92] C. Chui, *Wavelets: A tutorial in theory and applications*, Academic Press, Boston, USA, 1992.

[Cis] Cisco, *Introduction to Cisco IOS NetFlow - A Technical Overview*, `http://www.ciscosystems.com/en/US/prod/collateral/iosswrel/ps6537/ps6555/ps6601/prod_white_paper0900aecd80406232.html`.

[CK06] M. Crovella and B. Krishnamurthy, *Internet measurement: Infrastructure, traffic and applications*, John Wiley and Sons Inc., New York, USA, 2006.

[CL66] D. R. Cox and P. A. W. Lewis, *The statistical analysis of series of events*, Chapman and Hall, London, United Kingdom, 1966.

[Con] Consejo Coordinación Universitaria, *Estadística Universitaria del alumnado, profesorado y PAS universitario*, `http://www.educacion.es/educacion/universidades/estadisticas-informes/estadisticas/alumnado.html`.

[CPZ04] B.-Y. Choi, J. Park, and Z.-L. Zhang, *Adaptive packet sampling for accurate and scalable flow measurement*, In Proceedings of the IEEE Global Telecommunications Conference (Dallas, USA), vol. 3, November-December 2004, pp. 1448–1452.

[Dau90] I. Daubechies, *The wavelet transform, time-frequency localization and signal analysis*, IEEE Transactions on Information Theory **36** (1990), no. 5, 961–1005.

[DC74] O. J. Dunn and V. A. Clark, *Applied satistics: Analysis of variance and regression*, John Wiley and Sons Inc., New York, USA, 1974.

[dFV02] Y. d'Halluin, P. A. Forsyth, and K. R. Vetzal, *Managing capacity for telecommunications networks under uncertainty*, IEEE/ACM Transactions on Networking **10** (2002), 579–587.

[DLT05] N. Duffield, C. Lund, and M. Thorup, *Estimating flow distributions from sampled flow statistics*, IEEE/ACM Transactions on Networking **13** (2005), no. 5, 933–946.

[dOSVP06] M. R. de Oliveira, L. Sertorio, R. Valadas, and A. Pacheco, *Decomposition of network flows with applications to cluster analysis of Internet access traffic*, In Proceedings of Conference on Communications Environments Next Generation Internet Design and Engineering (Valencia, Spain), April 2006.

[DS86] R. B. D'Agostino and M. A. Stephens (eds.), *Goodness-of-fit techniques*, Marcel Dekker, Inc., New York, USA, 1986.

[EC96] A. Espasa and J. R. Cancelo, *Modelling and forecasting daily series of electricity demand*, Investigaciones Económicas **20** (1996), no. 3, 359–376.

[EV03] C. Estan and G. Varghese, *New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice*, ACM Transactions on Computer Systems **21** (2003), no. 3, 270–313.

[FCFW05] W.-C. Feng, F. Chang, W.-C. Feng, and J. Walpole, *A traffic characterization of popular on-line games*, IEEE/ACM Transactions on Networking **13** (2005), no. 3, 488–500.

[FGL$^+$01] A. Feldmann, A. G. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, *Deriving traffic demands for operational IP networks: methodology and experience*, IEEE/ACM Transactions on Networking **9** (2001), no. 3, 265–279.

[FK03] S. Floyd and E. Kohler, *Internet research needs better models*, ACM SIGCOMM Computer Communication Review **33** (2003), no. 1, 29–34.

[FML+03] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, *Packet-level traffic measurements from the Sprint IP backbone*, IEEE Network **17** (2003), no. 6, 6–16.

[FP01] S. Floyd and V. Paxson, *Difficulties in simulating the Internet*, IEEE/ACM Transactions on Networking **9** (2001), no. 4, 392–403.

[FTD03] C. Fraleigh, F. Tobagi, and C. Diot, *Provisioning IP backbone networks to support latency sensitive traffic*, In Proceedings of IEEE INFOCOM (San Francisco, USA), March-April 2003, pp. 375–385.

[GDHA+08] J. L. García-Dorado, J. A. Hernández, J. Aracil, J. E. López de Vergara, F. J. Montserrat, E. Robles, and T. P. de Miguel, *On the duration and spatial characteristics of Internet traffic measurement experiments*, IEEE Commununications Magazine **46** (2008), no. 11, 148–155.

[GMS+07] I. Gutiérrez, J. Martínez, P. M. Santiago, J. L. García Dorado, J. E. López de Vergara, J. Aracil, F. J. Montserrat, E. Robles, and T. P. de Miguel, *Análisis de la relación entre la intensidad del tráfico de datos y el número de alumnos en universidades españolas*, Actas de las Jornadas de Ingeniería Telemática (Málaga, Spain), 2007, pp. 353–359.

[GNS05] E. Giovannetti, K. Neuhoff, and G. Spagnolo, *Agglomeration in Internet co-operation peering agreements*, Cambridge Working Papers in Economics 0505, Faculty of Economics, University of Cambridge, January 2005.

[GPS72] G. V. Glass, P. D. Peckham, and J. R. Sanders, *Consequences of failure to meet assumptions underlying the fixed effects analysis of variance and covariance*, Review of Educational Research **42** (1972), no. 3, 237–288.

[GUF07] B. Gueye, S. Uhlig, and S. Fdida, *Investigating the imprecision of IP block-based geolocation*, In Proceedings of Passive and Active Measurement Conference (Louvain-la-neuve, Belgium), April 2007, pp. 237–240.

[GZCF06] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, *Constraint-based geolocation of Internet hosts*, IEEE/ACM Transactions on Networking **14** (2006), no. 6, 1219–1232.

[HPA07] J. A. Hernández, I. W. Phillips, and J. Aracil, *Discrete-time heavy-tailed chains, and their properties in modeling network traffic*, ACM Transactions on Modeling and Computer Simulation **17** (2007), no. 4.

[ITU02] ITU-T, *Internet protocol data communication service - IP packet transfer and availability performance parameters*, ITU-T Recommendation Y.1540, December 2002.

[IUKB+04] M. Izal, G. Urvoy-Keller, E. W. Biersack, P. A. Felber, A. Al Hamra, and L. Garcés-Erice, *Dissecting Bittorrent: Five months in a Torrent's lifetime*, In Proceedings of Passive and Active Network Measurement Workshop (Juan Les Pins, France), April 2004, pp. 1–11.

[Izs06] F. Izsak, *Maximum likelihood estimation for constrained parameters of multinomial distributions: Application to Zipf-Mandelbrot models*, Computational Statistics & Data Analysis **51** (2006), no. 3, 1575–1583.

[Jai91] R. Jain, *The art of computer system performance analysis*, John Wiley and Sons Inc., New York, USA, 1991.

[JBOK03] A. Jauk, M. Bonač, I. Ozimek, and G. Kandu, *Academic and research network of Slovenia*, Computer Communications **26** (2003), no. 5, 444–450.

[JKK05] N. L. Johnson, S. Kotz, and A. W. Kemp, *Univariate discrete distributions*, Wiley-Interscience, New York, USA, 2005.

[JTO09] W. John, S. Tafvelin, and T. Olovsson, *Passive Internet measurement: Overview and guidelines based on experiences*, Computer Communications (2009), `http://dx.doi.org/10.1016/j.comcom.2009.10.021`.

[KN02] J. Kilpi and I. Norros, *Testing the Gaussian approximation of aggregate traffic*, In Proceedings of ACM SIGCOMM Workshop on Internet Measurement (Marseille, France), November 2002, pp. 49–61.

[KV07] P. H. Kvam and B. Vidakovic, *Nonparametric statistics with applications to science and engineering*, Wiley-Interscience, New York, USA, 2007.

[Lam95] M. Lambert, *A Model for Common Operational Statistics*, RFC 1857 (Informational), October 1995.

[LASP+99] P. J. Lizcano, A. Azcorra, J. Solé-Pareta, J. Domingo-Pascual, and M. Álvarez Campana, *MEHARI: A system for analyzing the use of the Internet services*, Computer Networks **31** (1999), no. 21, 2293–2307.

[LBCM03] A. Lakhina, J. W. Byers, M. Crovella, and I. Matta, *On the geographic location of Internet resources*, IEEE Journal on Selected Areas in Communications **21** (2003), no. 6, 934–948.

[LdVBF04] D. López, J. E. López de Vergara, L. Bellido, and D. Fernández, *Monitoring an academic network with netflow*, In Proceedings of the EUNICE Open European Summer School (Tampere ,Finland), June 2004.

[Lei04] S. Leinen, *Evaluation of Candidate Protocols for IP Flow Information Export (IPFIX)*, RFC 3955, October 2004.

[LFHL07] H. Liu, W. Feng, Y. Huang, and X. Li, *A peer-to-peer traffic identification method using machine learning*, In Proceedings of International Conference on Networking, Architecture, and Storage (Guilin, China), July 2007, pp. 155–160.

[Lia93] X. J. Liang, *Network planning methodology and practice in China*, IEEE Communications Magazine **31** (1993), 34–37.

[Lin52] V. Lindley, *On the theory of queues with a single server*, In Proceedings of the Cambridge Philosophical Society **48** (1952), 277–289.

[LMRT01] J.-J. Laffont, S. Marcus, P. Rey, and J. Tirole, *Internet peering*, The American Economic Review **91** (2001), no. 2, 287–291.

[LPC+04] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. Kolaczyk, and N. Taft, *Structural analysis of network traffic flows*, ACM SIGMETRICS Performance Evaluation Review **32** (2004), no. 1, 61–72.

[LS08] S. Lippert and G. Spagnolo, *Internet peering as a network of relations*, Telecommunications Policy **32** (2008), no. 1, 33–49.

[LSLY05] Y. Li, A. Slagell, K. Luo, and W. Yurcik, *Canine: A Netflow converter/anonymizer tool for format interoperability and secure sharing*, In Proceedings of Network Flow Analysis Workshop (Pittsburg, USA), 2005.

[LTWW94] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, *On the self-similar nature of Ethernet traffic*, IEEE/ACM Transactions on Networking **2** (1994), no. 1, 1–15.

[LZSS06] S. Lee, Z. Zhang, S. Sahu, and D. Saha, *On suitability of Euclidean embedding of Internet hosts*, SIGMETRICS Performance Evalualtion Review **34** (2006), no. 1, 157–168.

[MA98] D. A. Menascé and V. A. F. Almeida, *Capacity planning for Web performance: metrics, models, and methods*, Prentice Hall, Upper Saddle River, USA, 1998.

[MACL06] R. McGorman, J. Almhanaa, V. Choulakian, and Z. Liua, *Empirical bandwidth provisioning models for high speed Internet traffic*, In Proceedings of Annual Communication Networks and Services Research Conference (Moncton, Canada), May 2006.

[MAGD09] F. Mata, J. Aracil, and J. L. García-Dorado, *Automated detection of load changes in large-scale networks*, In Proceedings of the IFIP International Workshop on Traffic Monitoring and Analysis (Aachen, Germany), May 2009.

[Mal89] S. G. Mallat, *A theory for multiresolution signal decomposition: A wavelet representation*, IEEE Transactions on Pattern Analysis and Machine Intelligence **11** (1989), no. 7, 674–693.

[Mat87] J. Matsukata, *Computer networking for large computers in universities*, ACM SIGCOMM Computer Communication Review, October 1987, pp. 78–87.

[Max] L. L. C. MaxMind, *GeoLite free country database*.

[MCS+06] J. Mai, C. N. Chuah, A. Sridharan, T. Ye, and H. Zang, *Is sampled data sufficient for anomaly detection?*, In Proceedings of the ACM SIGCOMM Conference on Internet Measurement (Rio de Janeiro, Brazil), 2006, pp. 165–176.

[MK00] K. McCloghrie and F. Kastenholz, *The Interfaces Group MIB*, RFC 2863, June 2000.

[MKYK89] J. Murai, H. Kusumoto, S. Yamaguchi, and A. Kato, *Construction of Internet for Japanese academic communities*, In Proceedings of ACM/IEEE Conference on Supercomputing (Heraklion, Greece), 1989, pp. 737–746.

[MM09] M. Mellia and M. Meo, *Measurement of IPTV traffic from an operative network*, European Transactions on Telecommunications (2009), `http://dx.doi.org/10.1002/ett.1401`.

[MNAR+04] H. T. Marques-Neto, J. M. Almeida, L. C. D. Rocha, W. Meira, P. H. C. Guerra, and V. A. F. Almeida, *A characterization of broadband user behavior and their e-business activities*, SIGMETRICS Performance Evaluation Review **32** (2004), no. 3, 3–13.

[MP05] A. W. Moore and K. Papagiannaki, *Toward the accurate identification of network applications*, In Proceedings of the Passive and Active Measurement Conference (Boston, USA), 2005, pp. 41–54.

[MRT] *MRTG log-files*, `http://oss.oetiker.ch/mrtg/doc/mrtg-logfile.en.html`.

[Nag05] A. Nag, *Network monitoring with Cacti*, Linux.com (2005).

[Nic87] P. T. Nicholls, *Estimation of Zipf parameters*, Journal of the American Society for Information Science **38** (1987), no. 6, 443–445.

[Nor01a] W. B. Norton, *A business case for ISP peering*, 2001, `http://www.equinix.com`.

[Nor01b] ———, *Internet service providers and peering*, 2001, `http://www.equinix.com`.

[NP08] A. Nucci and K. Papagiannaki, *Design, measurement and management of large-scale IP networks*, Cambridge University Press., Cambridge, United Kingdom, 2008.

[OA84] S. F. Olejnik and J. Algina, *Parametric ANCOVA and the rank transform ANCOVA when the data are conditionally non-normal and heteroscedastic*, Journal of Educational Statistics and Behavioral Statistics **9** (1984), no. 2, 129–149.

[Odl03] A. Odlyzko, *Data networks are lightly utilized, and will stay that way*, Review of Network Economics **2** (2003), no. 3, 210–237.

[Oet98] T. Oetiker, *MRTG - the Multi Router Traffic Grapher*, In Proceedings of USENIX Conference on System Administration (Boston, USA), December 1998, pp. 141–148.

[OGK06] J. Oberheide, M. Goff, and M. Karir, *Flamingo: Visualizing Internet traffic*, In Proceedings of IEEE/IFIP Network Operations and Management Symposium (Vancouver, Canada), April 2006, pp. 150–161.

[Pax94] V. Paxson, *Growth trends in wide-area TCP connections*, IEEE Network **8** (1994), no. 4, 8–17.

[PF95] V. Paxson and S. Floyd, *Wide area traffic: the failure of Poisson modeling*, IEEE/ACM Transactions on Networking **3** (1995), no. 3, 226–244.

[PMDC03] R. Prasad, M. Murray, C. Dovrolis, and K. Claffy, *Bandwidth estimation: Metrics, measurement techniques, and tools*, IEEE Network **17** (2003), no. 6, 27–35.

[PNvdMM09] A. Pras, L. Nieuwenhuis, R. van de Meent, and M. Mandjes, *Dimensioning network links: A new look at equivalent bandwidth*, IEEE Network **23** (2009), no. 2, 5–10.

[PP00] L. Pedraja Ponce and E. Rodríguez Ponce, *Modelo estocástico Wiener Gauss: Una aplicación la economía financiera en el mercado de captitales de España*, Revista Facultad de Ingeniería **8** (2000), 37–47.

[PS01] V. N. Padmanabhan and L. Subramanian, *An investigation of geographic mapping techniques for Internet hosts*, In Proceedings of ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (San Diego, USA), August 2001, pp. 173–185.

[PTB+02] K. Papagiannaki, N. Taft, S. Bhattacharyya, P. Thiran, K. Salamatian, and C. Diot, *A pragmatic definition of elephants in Internet backbone traffic*, In Proceedings of ACM SIGCOMM Workshop on Internet Measurement (Marseille, France), November 2002, pp. 175–176.

[PTZD05] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot, *Long-term forecasting of Internet backbone traffic*, IEEE Transactions on Neural Networks **16** (2005), no. 5, 1110–1124.

[RB89] J. C. W. Rayner and D. J. Best, *Smooth tests of goodness of fit*, Oxford University Press, New York, USA, and Oxford, United Kingdom, 1989.

[Red] RedIRIS, *What is RedIRIS*, `http://www.rediris.es/rediris/index.html.en`.

[RFL00] S. Romig, M. Fullmer, and R. Luman, *The OSU Flow-tools package and Cisco Netflow logs*, In Proceedings of Systems Administration Conference (New Orleans, USA), December 2000.

[RGMG+02] R. Romeral, A. García-Martínez, A. B. García, A. Azcorra, and M. Álvarez Campana, *MIRA: A distributed and scalable WAN/LAN real-time measurement platform*, Lecture Notes in Computer Science **2511/2008** (2002), 263–272.

[RRCB00] V. J. Ribeiro, R. H. Riedi, M. S. Crouse, and R. G. Baraniuk, *Multiscale queuing analysis of long-range-dependent network traffic*, In Proceedings of INFOCOM (Tel-Aviv, Israel), vol. 2, March 2000, pp. 1026–1035.

[SF02] R. Sommer and A. Feldmann, *Netflow: Information loss or win?*, In Proceedings of the ACM SIGCOMM Workshop on Internet Measurement (Marseille, France), November 2002, pp. 173–174.

[SGU08] S. S. Siwpersad, B. Gueye, and S. Uhlig, *Assessing the geographic resolution of exhaustive tabulation for geolocating Internet hosts*, In Proceedings of Passive and Active Measurement Conference (Cleveland, USA), April 2008, pp. 11–20.

[Shi08] S. Shipway, *Using MRTG with RRDtool and Routers2*, Cheshire Cat Press, November 2008.

[SMA03] N. Spring, R. Mahajan, and T. Anderson, *Quantifying the causes of path inflation*, In Proceedings of ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (Karlsruhe, Germany), August 2003, pp. 113–124.

[SPK02] L. Subramanian, V. N. Padmanabhan, and R. H. Katz, *Geographic properties of Internet routing: Analysis and implications*, In Proceedings of USENIX Annual Technical Conference (Monterey, USA), June 2002.

[SS06] S. Shakkottai and R. Srikant, *Economics of network pricing with multiple ISPs*, IEEE/ACM Transactions on Networking **14** (2006), no. 6, 1233–1245.

[Tera] Terena, *Terena compendium of national research and educational networks in Europe*, `http://www.terena.org/activities/compendium/ToC.html`.

[Terb] ———, *Trans-european research and education networking association*, `http://www.terena.org/`.

[TMW97] K. Thompson, G. J. Miller, and R. Wilder, *Wide-area Internet traffic patterns and characteristics*, IEEE Network **11** (1997), no. 6, 10–23.

[vdBMvdM+06] H. van den Berg, M. Mandjes, R. van de Meent, A. Pras, F. Roijers, and P. Venemans, *QoS-aware bandwidth provisioning for IP network links*, Computer Networks **50** (2006), no. 5, 631–647.

[vdM06] R. van de Meent, *Network link dimensioning: A mearurement and modeling based approach*, Ph.D. thesis, University of Twente, March 2006.

[vdMMP06] R. van de Meent, M. R. H. Mandjes, and A. Pras, *Gaussian traffic everywhere?*, In Proceedings of IEEE International Conference on Communications (Istanbul, Turkey), vol. 2, June 2006, pp. 573–578.

[vdMMP07] ______, *Smart dimensioning of IP network links*, In Proceedings of IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (San José, USA), October 2007, pp. 86–97.

[VK00] N. Vicari and S. Köhler, *Measuring Internet user traffic behavior dependent on access speed*, ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management (Monterey, USA), September 2000.

[WF06] J. Wallerich and A. Feldmann, *Capturing the variability of Internet flows across time*, In Proceedings of IEEE INFOCOM (Barcelona, Spain), April 2006, pp. 1–6.

[WGT95] O. J. Wasem, A. M. Gross, and G. A. Tlapa, *Forecasting broadband demand between geographic areas*, IEEE Communications Magazine **33** (1995), no. 2, 50–57.

[WP05] W. V. Wanrooij and A. Pras, *Data on retention*, In Proceedings of IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (Barcelona, Spain), October 2005, pp. 60–71.

[Yur06] W. Yurcik, *Visflowconnect-IP: A link-based visualization of Netflows for security monitoring*, In Proceedings of the Annual FIRST Conference on Computer Security Incident Handling (Baltimore, USA), June 2006.

List of publications

Publications related to thesis

1. José Luis García-Dorado, José Alberto Hernández, Javier Aracil, Jorge E. López de Vergara, Francisco Montserrat, Esther Robles, and Tomás P. de Miguel, *On the duration and spatial characteristics of Internet traffic measurement experiments,* IEEE Communications Magazine **46** (2008), no. 11, 148-155.
Chapter 3 in this thesis.

2. J. L. García-Dorado, J. E. López de Vergara, J. Aracil, V. López, J. A. Hernández, S. Lopez-Buedo, and L. de Pedro, *Utilidad de los flujos Netflow de RedIRIS para análisis de una red académica (On the use of RedIRIS' Netflow information for academic networks),* Boletín de RedIRIS, no. 82-83, Jornadas Técnicas RedIRIS 2007 (Mieres, Spain), November 2007.
Section 3.2 in this thesis.

3. F. Mata, J. L. García-Dorado, J. Aracil, and J. E. López de Vergara, *Factor analysis of Internet traffic destinations from similar source IP subnetworks,* submitted to Elsevier Computer Networks.
Chapter 4 in this thesis.

4. José Luis García-Dorado, Javier Aracil, José Alberto Hernández, and Jorge E. López de Vergara, *A queueing equivalent thresholding method for thinning traffic captures,* in Proceedings of the IEEE/IFIP Network Operations and Management Symposium (Salvador, Brazil), April 2008.
Chapter 5 in this thesis.

5. J. L. García-Dorado, J.A. Hernández, J. Aracil, J. E. López de Vergara, and S. López, *Characterization of the busy-hour traffic of IP networks based on their inherent features,* submitted to IEEE/ACM Transactions on Networking.
Chapter 6 in this thesis.

Other publications ranked chronologically

6. J. L. García-Dorado, J. E. López de Vergara, and J. Aracil, *Analysis of OBS burst scheduling algorithms with stochastic burst size and offset time,* in Proceedings of the Conference on Optical Network Design and Modelling (Copenhagen, Denmark), May 2006.

7. I. Gutiérrez, J. Martínez, P. M. Santiago, J. L. García Dorado, J. E. López de Vergara, J. Aracil, F. J. Montserrat, E. Robles, and T. P. de Miguel, *Análisis de la relación entre la intensidad del tráfico de datos y el número de alumnos en universidades españolas (Analyzing the relation between the demand for bandwidth and the number of users in the Spanish Academic Network),* in Actas de las Jornadas de Ingeniería Telemática (Málaga, Spain), September 2007.

8. J. L. García-Dorado, J. Aracil, J. A. Hernández, S. Lopez-Buedo, J. E. López de Vergara, P. Reviriego, G. Huecas, S. Pavón, and J. Quemada, *A quality of service assessment technique for large-scale management of multimedia flows,* in Proceedings of the IFIP/IEEE International Conference on Management of Multimedia and Mobile Networks and Services (San José, USA), October 2007.

9. V. López, J. L. García-Dorado, J. A. Hernández, J. Aracil, *Performance comparison of scheduling algorithms for IPTV traffic over Polymorphous OBS routers,* in Proceedings of the International Conference on Transparent Optical Networks - Mediterranean Winter (Sousse, Tunisia), December 2007.

10. L. de Pedro, J. Aracil, J. A. Hernández, J. L. García-Dorado , *Analysis of the processing and sojourn times of burst control packets in optical burst*

switches times, in Proceedings of the International Conference on Optical Networking Design and Modeling (Vilanova i la Geltrú, Spain), March 2008.

11. F. Mata, J. Aracil, J. L. García-Dorado, *Automated detection of load changes in large-scale networks,* in Proceedings of the IFIP International Workshop on Traffic Monitoring and Analysis (Aachen, Germany), May 2009.

12. J. L. García-Dorado, V. Lucas, J. Aracil, J. E. López de Vergara, A. Cuadra, M. Cutanda, and F. Garces *Caracterización avanzada de uso de servicios y aplicaciones en redes de datos mediante el análisis detallado de tráfico (Characterization of the use of services and applications with deep packet inspection),* en Actas de Telecom I+D (Madrid, Spain), November 2009.

13. V. Lucas Segura, J. L. García Dorado, I. González Martínez, F. J. Gómez Arribas, and J. Aracil Rico, *Paralelización del análisis de tráfico de red mediante computación con GPU (Parallelizing the inspection of Internet traffic traces using GPUs),* en Actas del Workshop en Aplicaciones de Nuevas Arquitecturas de Consumo y Altas Prestaciones (Madrid, Spain), November 2009.

14. Book Chapter Contribution, Editors: J. Aracil and F. Callegatti, *Optical Burst Switching, enabling optical Internet with advanced network technologies,* Ed. Springer, Series "Computer Communications and Networks", 2009.

15. J. A. Hernández, J. Aracil, V. López, J. L. García-Dorado, and L. De Pedro, *Performance analysis of asynchronous best-effort traffic coexisting with TDM reservations in polymorphous OBS networks,* Photonic Network Communications **17** (2009), no. 2, 93-103.

16. D. Morató, J. Aracil, J.A. Hernández, and J. L. García-Dorado, *On the blocking time distribution of core OBS switches,* Photonic Network Communications **18** (2009), no. 3, 314-322.

17. J. A. Hernández, P. Reviriego, J. L. García-Dorado, V. López, D. Larrabeiti, and J. Aracil, *Performance evaluation and design of Polymorphous OBS*

networks with guaranteed TDM services, IEEE/OSA Journal of Lightwave Technology **27** (2009), no. 13, 2495-2505.