Performance Measurement of Computer Networks

Federico Montesino Pouzols
University of Seville, Spain

Angel Barriga Barros
University of Seville, Spain

Diego R. Lopez
RedIRIS, Spain

Santiago Sánchez-Solano
CSIC - Scientific Research Council, Spain

INTRODUCTION

In this article, general findings about Internet traffic models are first reviewed, with emphasis on two important invariants or characteristics that are observed with some reproducibility and independently of the precise settings of the network under consideration: self-similarity and heavy-tail marginal distributions. Then metrics and measurement techniques and tools will be discussed.

This article deals with generic network performance measurement systems and outlines models, measurement techniques and tools that measure performance at the network and transport layers and can thus be applied regardless of the application layer protocols being employed. These systems are useful for analyzing performance of any network application and are an important foundational tool for enabling advanced virtual organizations (Foster, Kesselman, & Tuecke, 2001). Note, however, that application-level (or specific application details aware) measurements are commonly needed to complement generic tools so as to achieve a clear understanding of overall applications performance, which cannot be synthesized from lower level data with ease (Andrews, Cao, & McGowan, 2006).

BACKGROUND

The quality of network connectivity among the agents that interact within networked and virtual organizations has a deep impact on the overall performance that users perceive from the organizations (Foster, Kesselman, & Tuecke, 2001). Thus, tasks such as dimensioning and performance analysis of organizations do not only require the analysis of computational resources but also measuring and analyzing network connections (DeSanctis & Jackson, 1994). To that end, a number of measurement techniques and tools can be employed.

Throughout the last years, a number of performance metrics, measurement practices and tools have been developed and standardized. Performance measurement techniques are often used by network engineers in diagnosing network problems; however, most recently their application has been by network users, service providers and researchers in analyzing traffic behaviour across specific paths or the performance associated with individual Internet Service Providers (ISPs) (Nassar, 2000). A recent development in the industry is the offering of service level agreements (SLAs), contracts to guarantee a specified level of service, subject to cost rebates or other consumer remuneration should measurements suggest that the ISP did not adhere to the SLA. SLAs are rather controversial in the community since there is no standard metric or even measurement methodology for calibrating them (Hardy, 2001).

TRAFFIC MODELS

Network traffic is in general a complex, nonlinear, non-stationary process. Since self-similarity and burstiness at multiple time scales was discovered in network traffic (Leland, Taqqu, Willinger, & Wilson, 1994), a number of studies have shown the failure of Poisson modelling for Internet traffic at various levels of aggregation (from local area networks to backbones) (Paxson & Floyd, 1995). A large number of alternative models have been proposed to capture the properties of packet and byte arrival processes (Cappé, Moulines, Pesquet, Petropulu, & Yang, 2002).

The big challenge in developing Internet traffic models is that everything keeps changing. For example, WWW traffic grew from zero in 1995 to more than 80% at most sites by the early 2000s. However, WWW proportion of total traffic is dropping on most networks while peer-to-peer traffic is steadily rising as peer-to-peer-related technologies are getting more mature (Karagiannis, Molle, & Faloutsos, 2004).

A number of empirical studies (Park & Willinger, 2000) of high-resolution traffic measurements from a variety of communications networks have provided clear evidence that actual network traffic is self-similar or fractal in nature, that is, bursty over a wide range of time scales.

Both self-similarity and heavy-tail marginal distributions arise as characteristics of most scalar signals that can be extracted from network traffic traces. Typical scalar signals include arrival times of successive packets at some point of the network or time series obtained by counting the size of the data transferred within a time interval. This observation is in contrast to commonly made modelling choices in classical engineering theory and practice, where exponential assumptions still dominate and are only able to reproduce the bursty behaviour of measured traffic either on a pre-specified time scale or over a very limited range of time scales. It has been shown that the observed self-similarity characteristic of measured aggregate network traffic has a direct link to heavy-tailed, infinite variance phenomena at the level of individual network connections (Willinger, Paxson, & Taqqu, 1998).

With the discovery of self-similarity and long-range dependence (LRD) in network traffic, the research community has undergone a mental shift from poisson and memory-less processes to LRD and bursty processes (Karagiannis, Molle, & Faloutsos, 2004). Despite its widespread use, LRD analysis is hindered by the difficulty in actually identifying dependence and estimating its parameters unambiguously. Currently, there is a lack of accuracy and robustness in LRD estimation (Veitch, Hohn, & Abry, 2005).

The dependence structure of Internet traffic is, however, such a complex phenomenon that no model is entirely satisfactory (Gibbens, 1996), and recent studies have even proposed a return to poisson models for the smallest time scales in high-aggregation environments (Cao, Cleveland, Lin, & Sun, 2003; Karagiannis, Molle, Faloutsos, & Broido, 2004). Queueing analysis (Daigle, 2004) is a well- established field and complementary engineering tool to study network traffic (Rolls, Michailidis, & Hernández-Campos, 2005). Additionally, self-similarity and long-range dependencies can be analyzed by means of wavelets (Abry & Veitch, 1998).

PERFORMANCE METRICS

Measurement and performance evaluation studies usually consider concepts such as connectivity, availability, accessibility, routing speed, connection reliability, routing reliability and connections quality (for different classes of traffic, such as data and voice) as well as continuity (Hardy, 2001). We note that evaluation of security and privacy issues, which have an effect on user perceived performance is outside the scope of this article.

Reproducible and systematic evaluation of these concepts requires the adoption of metrics and definitions that are also the basis for QoS procedures (Firoiu, Le Boudec, Towsley, & Zhang, 2002). These metrics can be classified into end-to-end (or path level) and per link (or per hop).

A number of metrics are being standardized by the IP Performance Metrics (IPPM) Working Group of the Transport Area of the Internet Engineering Task Force (IETF) (see http://www.ietf.org/html.charters/ippm-charter.html). The goal of the IP Performance Metrics effort is to achieve a situation in which users and providers of Internet transport services have an accurate, common understanding of the performance and reliability of the Internet component that they use and provide (Paxson, Almes, Mahdavi, & Mattis, 1998).

The IPPM working group develops a set of standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services. These metrics are designed such that they can be performed by network operators, end users or independent testing groups. The following main criteria are considered for metrics: metrics are concrete and well-defined, a repeatable measurement methodology is provided and they are useful to users and providers in understanding the network performance.

Using a Design Science Research Approach in Human-Computer Interaction (HCI) Project: Experiences, Lessons and Future Directions
Muhammad Nazrul Islam (2017). *International Journal of Virtual and Augmented Reality (pp. 42-59).*
www.irma-international.org/article/using-a-design-science-research-approach-in-human-computer-interaction-hci-project/188480