

Performance Analysis of Peer-to-Peer Traffic

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

Peer-to-peer (P2P) networks have recently emerged as an attractive solution to enable large-scale content distribution without requiring major infrastructure investments. Recent developments have led to a significant maturity increase of peer-to-peer technologies, which are currently available as tools for performing core tasks in virtual and networked organizations.

Peer-to-peer technologies can be employed, among other tasks, to implement service-oriented architectures as well as in the business-to-business (B2B) area, specially for highly dynamic virtual organizations, where they can provide flexible fail-over scenarios and an increase in service availability (Fattah, 2002). P2P systems are also useful in the implementation of request-based virtual organizations (RBVO). Thus, expanding collaboration possibilities and availability for business partners (Aberer & Hauswirth, 2005; Subramanian & Goodman, 2005).

Because of these properties, P2P architectures are specially suited for virtual organizations and virtual enterprises in particular. For instance, workflow technology seeks to support or automate the performance of organizational processes. When the organization is a virtual enterprise, the traditional client-server workflow architecture is not suitable: virtual enterprises lack centralized resources and administration to support server installations. Further, virtual enterprise participants can be highly autonomous and heterogeneous.

A workflow system to support virtual enterprises must account for autonomy and heterogeneity, as well as wide geographic distribution. Peer-to-peer architectures are being proposed for deploying workflow

in virtual enterprises engaged in knowledge-intensive activities. In this context, a virtual enterprise is seen as a loosely coupled organization of autonomous, distributed actors supported by individual workflow nodes that enact local processes to guide the actor, and share events to coordinate activities with other members of the organization (Fattah, 2002; Noll, 2005; Yan, Yang, Kowalczyk, & Nguyen, 2005).

This article discusses specific models for peer-to-peer networks traffic, including models for content delivery systems, file-sharing systems, networked games and other overlay networks. Characteristics of typical peer-to-peer traffic dynamics are outlined as well as traffic analysis methodologies, summarizing recent results in measurement and evaluation of peer-to-peer architectures and protocols.

BACKGROUND

Peer-to-Peer is an approach to distributed computing based on sharing computing services and resources among systems without intermediaries, that is, through symmetric communication. In a peer-to-peer architecture, computers communicate directly among themselves and can act as both clients and servers, as opposed to client-server architectures such as the World Wide Web. The implementation of P2P systems usually involves the creation of overlay networks (Foster & Iamnitchi, 2003).

The definition of P2P is sometimes relaxed so that systems that include centralized elements in their overlay networks are also considered P2P systems. In this case, P2P is broadly defined as a class of applications

that takes advantage of resources available at the edges of the Internet (Foster & Iamnitchi, 2003). Among other characteristics, traffic patterns of P2P systems are fundamentally different to that of client-server and grid systems. Particularly, the deployment of P2P systems usually leads to an increase of overall network traffic volume while at the same time distributing the traffic among different sub-networks.

P2P systems such as DNS and Usenet have existed for many years but may represent just the beginning of what peer-to-peer has to offer. P2P proponents often cite music and other file sharing systems, or distributed processing intensive applications like SETI@home, as evidence that this technology will someday lead to popular and profitable business models. Although they have actually existed for many years, P2P technologies provide a network computing model that can radically change the future of networking (Aberer & Hauswirth, 2005).

Although peer-to-peer systems are rapidly evolving and challenging to analyze, latest measurement studies provide an insight in the dynamics of peer-to-peer traffic and how it relates different architectures and protocols. These results can in turn be used to classify and select peer-to-peer architectures as well as optimizing peer-to-peer systems (Theotokis & Spinellis, 2004).

PEER-TO-PEER TRAFFIC MODELS

Most recent P2P traffic measurement studies have been performed on popular unstructured P2P networks for file sharing, in which there are no restrictions on data placement in the overlay network. These studies can be classified into application level data collection and lower layers network monitoring and sniffing. Measurements of P2P systems are useful in evaluating different aspects of system performance and the effect of P2P traffic on the network in terms of bandwidth usage as well as spatial and temporal patterns. System performance of P2P systems is closely related to the total number of peers, total number of resources shared and the total network capacity provided by the peers (Guo, Chen, Xiao, Tan, Ding, & Zhang, 2005).

A number of studies of P2P traffic performed by measuring flow-level information at border routers (Sen & Wang, 2002) has shown that, as opposed to WWW or overall Internet traffic, because of the stability properties of P2P traffic it can be managed

by ISPs via application specific transport layer traffic engineering. The distribution of traffic among content units has been found to obey Zipf's law in most P2P analyzed to date. The same applies to the distribution of traffic among nodes in the overlay networks (Tian, Wu, & Ng, 2006).

Self-similarity patterns (Crovella & Bestavros, 1997) as well as long-range dependencies (Cappé, Moulines, Pesquet, Petropulu, & Yang, 2002) common in overall network traffic have been also identified in P2P traffic to a varying extent (Gummadi, Dunn, Saroiu, Gribble, Levy, & Zahorjan, 2003; Guo, Chen, Xiao, Tan, Ding, & Zhang, 2005).

Recent measurement results provide an insight into overlay topologies of P2P systems, characterizing routing and path properties (Lua, Crowcroft, Pias, Sharma & Lim, 2005; Stutzbach, Rejaie & Sen, 2005; Touch, 2005). It has been demonstrated that the employment of P2P content delivery architectures can lead to a decrease of the computing resources required by content providers and consumers (by replacing traditional content distribution networks (CDNs), server farms and other infrastructures requiring large investments) at the expense of a higher usage of communications resources provided by ISPs (Karagiannis, Rodriguez, & Papagiannaki, 2005). Query patterns of popular P2P systems have been characterized as well (Klemm, Lindemann, Vernon, & Waldhorst, 2004).

PERFORMANCE ANALYSIS OF PEER-TO-PEER SYSTEMS

The characteristic feature of a pure P2P network is that it is a distributed autonomous system which does not rely on a specific server for communications. As a result, such systems are expected to exhibit scalability in processing power and load balancing at the end computers (Aberer & Hauswirth, 2005). However, the P2P traffic volume is becoming much larger than that of the previous Internet applications and the bottlenecks in processing power are shifting from the end computers to the network. In addition, traffic control is very difficult because there is no administrator in the overlay networks and on account of the anonymous nature of the traffic (Taylor, 2004; Theotokis & Spinellis, 2004).

The domain name system (DNS) is an example of a system that combines peer-to-peer networking with a hierarchical model of information ownership. The

4 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/performance-analysis-peer-peer-traffic/17745

Related Content

Advanced Visual SLAM and Image Segmentation Techniques for Augmented Reality

Yirui Jiang, Trung Hieu Tran and Leon Williams (2022). *International Journal of Virtual and Augmented Reality* (pp. 1-28).

www.irma-international.org/article/advanced-visual-slam-and-image-segmentation-techniques-for-augmented-reality/307063

Fast Single Image Haze Removal Scheme Using Self-Adjusting: Haziness Factor Evaluation

Sangita Roy and Sheli Sinha Chaudhuri (2019). *International Journal of Virtual and Augmented Reality* (pp. 42-57).

www.irma-international.org/article/fast-single-image-haze-removal-scheme-using-self-adjusting/228945

Supporting Facilitators in Communities of Practice via Design and Technology

Halbana Tarmizi, Gert-Jan de Vreede and Ilze Zigurs (2009). *Virtual Team Leadership and Collaborative Engineering Advancements: Contemporary Issues and Implications* (pp. 65-83).

www.irma-international.org/chapter/supporting-facilitators-communities-practice-via/30876

Private Label Sales through Catalogs With Augmented Reality

Gerardo Reyes Ruiz, Samuel Olmos Peña and Marisol Hernández Hernández (2018). *Virtual and Augmented Reality: Concepts, Methodologies, Tools, and Applications* (pp. 1500-1531).

www.irma-international.org/chapter/private-label-sales-through-catalogs-with-augmented-reality/199751

Tourism Networks and Clusters

Patrice Braun (2008). *Encyclopedia of Networked and Virtual Organizations* (pp. 1653-1658).

www.irma-international.org/chapter/tourism-networks-clusters/17804