Internet Traffic & the Behavior of Processing Workloads

UNIVERSITAT POLITÈCNICA DE CATALUNYA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Ruken Zilan†*, Javier Verdú†*, Jorge Garcia Vidal*, Mario Nemirovsky†, and Mateo Valero‡,*

* DAC. UPC. Spain.   †Barcelona Supercomputing Center (BSC), Spain.

Introduction

- Nowadays, the evolution of network services provided at the edge of Internet increases **the number** and **complexity of network applications.**

- Such applications result in **complexities** thus, the processors need to execute more complex workloads that can deal not only with the packet header, but also with the **packet payload** (e.g. Deep Packet Inspection).

- Unlike common routing applications that show similar processing among packets, next-generation of network applications present **variations in the processing procedure among packets.**

- Therefore, **different traffic behaviors** can produce **different process patterns** and present different memory and processing requirements.

Related Work

- **Ramaswamy et al. [RWW05]** analyze the processing characteristics of individual packets and the variations between them to understand the network processing workloads. They present low network layer applications and thus, they cannot include any result about flow related impact, due to traffic features.

- **Verdú et al. [Ver05]** analyze architectural bottlenecks running stateful network applications. However, this work is done from the processor performance and requirements' point of view, instead of the processing workload.

- There are many studies which show that Internet traffic continues to grow and provided services become more complex, over relatively short time scales. This is not only simply due to **the traffic volume**, but also due to **heterogeneity** which is caused by the mixed-traffic, protocols, applications, and users **[ZBPS02]**.

- **Holanda et al. [HVGV05]** show that, using clustering techniques, there is no much variety in high speed link flows so that flows can be grouped in few clusters.

Workload Features

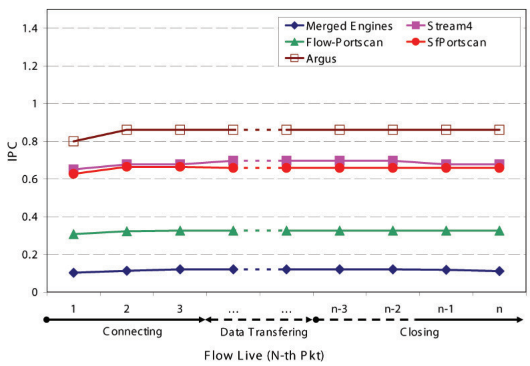

Figure 1. IPC vs Flow Lifetime

- Next-generation of network applications especially deal with high network layers of the OSI Model (i.e. from Transport to Application layer). **[VGNV05]** focuses on applications that process up to Transport layer (e.g. TCP, UDP).

- The workload generated through the flow lifetime may be quite different depending on the tasks of the application. **Figure 1** shows the average of instructions executed by the n-th packet processing through the flow lifetime. The life of a connection presents mainly three stages; connecting, data transferring, and closing.

- While **Figure 1** shows workloads generated by the flows through their lifetime, **Figure 2** depicts the IPC for different TCP flow lifetime durations.

The main goal is to identify the network traffic effects on the processing workload in terms of identified workload parameters. Considered features are:

- Number of instructions; to see the processing requirements throughout the packet processing considering different traffic features.

- The instructions footprint; to study the variations of instructions executed throughout the processing of network traffic.

-Data footprint and the Data working set; to address the effects of memory demand.

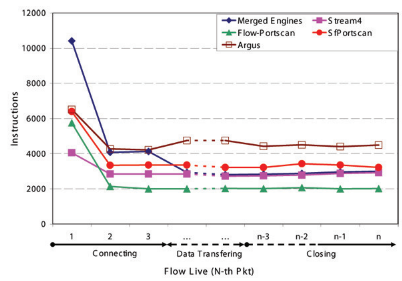

Figure 2. Instructions generated through the flow lifetime

Traffic Features

- New generation applications bring complex requirements, so that the **traffic features affect processed instruction workloads** due to the variable payload inspections.

- Therefore it is important to **characterize the main traffic behaviors** in order to see different effects on the processing workload.

- The main traffic features that we will use to analyze the traffic behaviors are: **Traffic aggregation**, **Burstiness, Packet size, Flow temporal distributions** and **Flow size distributions.**

- **Some traffic characteristics**, such as burstiness and distributions, will be analyzed. Special emphasis will be on **the upper layer characteristics of the traffic.**

Modeling

These features will be used together in order to model the whole system that is shown in **Figure 3.**

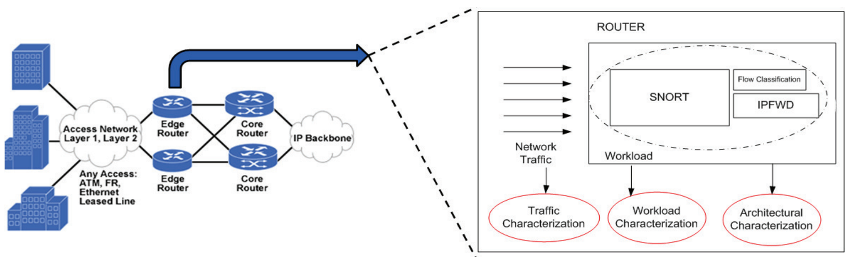

Figure 3. The General Scene of the Study

Future Work

- A study in progress which focuses on correlating Internet traffic features with variations of processing workloads on the next-generation of edge routers is presented.

- We will analyze the impacts of upper network layers' behavior on multithreaded processing workloads.

- We will extend our analysis to multithreaded processors, such as thread contention in critical sections and bottlenecks on shared resources.

- We will perform the modeling considering the workload, traffic and CPU requirements for network processors using these identified critical features.

References

[HVGV05] R. Holanda, J. Verdú, J. García, and M. Valero. Performance analysis of a new packet trace compressor based on TCP flow clustering. In Proc. of ISPASS'05, pages 219–225, Austin, TX, March 2005.

[RWW05] R. Ramaswamy, N. Weng, and T. Wolf. Analysis of network processing workloads. In Proc. of ISPASS'05, pages 226–235, Austin, TX, March 2005.

[VGNV05] J. Verdú, J. García, M. Nemirovsky, and M. Valero. Architectural impact of stateful networking applications. In Proc. of ANCS'05, NY, USA, 2005. ACM.

[ZBPS02] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. The characteristics and origins of Internet flow rates. In Proc. of ACM SIGCOMM, 2002.

This work has been supported by the Ministry of Science and Technology of Spain under contract TIN-2007-60625, Ministry of Science and Education of Spain under contract TSI2007-66869-C02, the HiPEAC European Network of Excellence and a Collaboration Agreement between Cisco Systems and BSC. The authors wish to thank Rodolfo Milito for his support.