

# P51: High Performance Networking Syllabus

Noa Zilberman, Andrew W Moore

**Code:** P51

**Instructors:** Noa Zilberman, Andrew W Moore

**Prerequisites:** Undergraduate courses in digital communication, good working knowledge of C/C++, ECAD, Unix.

**Structure:** Six 1-hour lectures; five 2-hour supervised laboratory sessions

## 1 Aims

This module provides an introduction to High Performance Networking. The module explores both software and hardware aspects in the design for high performance networks and networking devices. The module provides the students an opportunity to experience high performance networking design and usage first hand. This module will:

- Teach basic concepts in high performance networking, with a focus on throughput and latency performance metrics;
- Expose students to real-world architectural constraints in high performance networking;
- Develop design skills through the practical development of a high performance network device;
- Provide research skills for characterization and modelling of high performance network devices.

## 2 Prerequisites

It is required that students have previously (and successfully) completed an undergraduate networking course – or have equivalent experience through project or open-source work. It is strongly recommended to have good working knowledge of C/C++, ECAD, and Unix. Knowledge of Python and P4 is an advantage, but not a prerequisite.

## 3 Lectures and Supervised Lab Sessions

### 3.1 Lectures

**Lecture 1: Introduction and general architecture of high performance network devices.**

This lecture will introduce students to high performance networking, starting from basic aspects in network switching. We will discuss “real world” switches and the drivers of their architecture. Different performance metrics will be introduced and compared. The lecture will also cover basic terms and definitions used in the field.

**Lecture 2: Programmable devices.**

In the second lecture, we will explore recent trends in programmable high performance networking. The lecture will cover programmable data planes, and compare them with “traditional” data planes. The lecture will also explore moving different services to the network.

**Lectures 3: High throughput devices - Part I.**

In this lecture we will start discussing high throughput devices, starting with different types of high performance switches. The lecture will cover cut-through and store-and-forward switches, compare and contrast ToR and Spine

switches, and discuss the differences between general purpose and proprietary ASICs. The lecture will explore different high throughput switch architectures, from the silicon level, through system, to network level.

#### **Lecture 4: High throughput devices - Part II.**

This lecture will continue to explore high throughput devices. The lecture will explore building performance profiles and understanding the bottlenecks of network designs. Amongst others, we will cover bottlenecks in the data path, I/O, memory interfaces, and protocols. The lecture will conclude with a discussion of performance evaluation.

#### **Lecture 5: Low latency devices - Part I.**

This lecture will discuss low latency devices, starting with latency metrics. We will explore basic low latency switch architectures and compare and contrast low latency switches with high throughput switches. The lecture will then focus on bottlenecks in low latency design, such as buffering, header processing, bus conversion and protocols.

#### **Lecture 6: Low latency devices - Part II.**

In this lecture will continue to discuss low latency devices, with a focus on hardware-software integration. The lecture will explore aspects in DMA design, drivers and network stack and their effect on latency and performance. The measurement and evaluation of low latency devices will then conclude the lecture. Finally, we will summarize the course.

### **3.2 Supervised Lab Sessions**

#### **Labs 1-2: NetFPGA hands-on.**

These labs will provide hands-on experience in developing and using the NetFPGA SUME platform and an introduction to P4-NetFPGA.

Part of Lab 2 will be dedicate for a discussion of selected projects.

#### **Lab 3-5: Supervised Lab Sessions.**

These sessions will provide supervision and guidance in the development of the projects. Each session will focus on a different aspect of the project: architecture, performance profile, evaluation plan, debug, etc. In addition, technical assistance will be provided.

A rough skeleton of the goals per session (which may vary between projects) is as follows:

- Lab 2 - project selection and discussion of details.
- Lab 3 - project architecture.
- Lab 4 - performance profile.
- Lab 5 - evaluation plan.

Students will be requested to upload to Moodle the relevant material prior to each lab.

## **4 Objectives**

On completion of this module, students should:

- Describe the role of high performance networking and where it is used;
- Compare and contrast high throughput and low latency networking devices;
- Define the architecture of a high performance networking device;
- Implement a fully functioning high performance networking device;
- Evaluate the performance of a high performance networking device.

## 5 Coursework

Five 2-hour in-classroom supervised lab sessions will ask students to develop and use skills learned in the course and apply them to the design of a high performance device.

The two labs will provide an introduction to a development platform, while the remaining labs will focus on a specific design project (starting from a reference design). Time will be allocated for different design stages (architecture, design and validation). Students may find it useful to work in teams of two within the lab, and the project can be extended in any relevant direction. Instructors will be on-hand throughout labs to provide guidance. Lab participation is not directly included in the final mark, but lab work is a key input to a project that is assessed.

## 6 Assessment

Assessed practical work, based on one practical assignment. The assessment will include both the project submission and related documentation.

## 7 Recommended Reading

The following list provides some background to the course materials, but is not mandatory. A reading list, including research papers, will be provided in the course materials.

- George Varghese. Network algorithmics. Chapman & Hall/CRC, 2010.
- Peterson, L.L. & Davie, B.S. (2019). Computer networks: a systems approach. Morgan Kaufmann (online ed. / 5th ed.).
- Kurose, J.F. & Ross, K.W. (2016). Computer networking: a top-down approach. Addison-Wesley (7th ed.).