# Topics in Computing Science: Distributed Systems and Networks

Dr. Colin Perkins
<csp@dcs.gla.ac.uk>

9 January 2005

## 1 Introduction

The aims of this module are to immerse students in some seminal topics in distributed systems and networks via papers from the literature, reinforce critical reading and reviewing skills, and enable students to develop their critical thinking and group discussion capabilities.

The intended learning outcomes for the module are that students should: 1) exhibit understanding of a broad range of seminal research in the area of distributed systems and networking by summarising relevant papers, and by answering examination questions that require an overall understanding of distributed systems and networks as a whole; and 2) demonstrate the ability to present a critical analysis of relevant literature by presenting verbal summaries of selected research papers to the class, and leading class discussions.

## 2 Module Description

The module comprises 12 tutorial sessions held from 2-3pm on Tuesdays in F121. The first session will be on 11 January 2005 (week 13), continuing weekly to 15 March, then resuming after the Easter holiday with two final sessions on 12 and 19 April. The module is structured around discussion of one research paper each week. Students *must* read the papers in advance of each session, and come prepared to discuss the subject. All students will be expected to participate in the discussion.

The module will focus on providing students with a deep understanding of the behaviour of the Internet and how this effects the design of network protocols and distributed systems. Accordingly, you are expected to read additional material relating to the topics under discussion, rather than limiting yourself to a narrow reading of the papers we will consider in the class.

### Week 13: Introduction

The module begins with a study of the design of the Internet protocol suite. Three papers will be discussed this week: one ("The Design Philosophy of the DARPA Internet Protocols") is new to you, we have previously discussed the other two papers in the Advanced Topics in Computing Science module last semester. These papers set the scene for the remainder of the module, describing the structure of the Internet protocol and the design philosophy underpinning that structure. You should read these papers to understand why the designers made particular choices, how those choices impact the structure of the network, and what could be done differently.

- D. D. Clark "The Design Philosophy of the DARPA Internet Protocols", Proceedings of ACM SIGCOMM '88, Stanford, CA, USA, August 1988.

- J. H. Saltzer, D. P. Reed and D. D. Clark, “End-to-End Arguments in System Design”, ACM Transactions on Computer Systems”, volume 2, number 4, November 1984.
- V. G. Cerf and R. E. Kahn, “A Protocol for Packet Network Intercommunication”, IEEE Transactions on Communications, volume 22, number 5, May 1974.

## Week 14: Measuring Network Performance – The Core Network

We begin our survey of network performance by studying some recent measurements of the performance of the core network. You should read this paper to understand how the measurement infrastructure has been developed, how the network behaves, and how the performance of the network measured in this paper might differ from other networks.

- C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, C. Diot, “Packet-level Traffic Measurement from the Sprint IP Backbone”, IEEE Network Magazine. November 2003.

## Week 15: Measuring Network Performance – Measurement Infrastructure

This week we consider the difficulties in building and operating a large-scale distributed computing infrastructure to collect network performance measurements. You should read this paper with a view to how the infrastructure could be improved, considering which issues are inherent to the problem domain and which result from the setup used in these particular experiments. Consider how other large-scale distributed applications (e.g. Grid computing, peer-to-peer systems) might be impacted by the same issues. How does the system described here differ from that discussed last week? Are these differences important?

- V. Paxson, A. K. Adams and M. Mathis, “Experiences with NIMI”, Proceedings of the Passive and Active Measurement Workshop, Hamilton, New Zealand, April 2000.

## Week 16: Measuring Network Performance – Congestion Control

This week we will study how network performance affects the dynamics of transport protocols. The paper describes how the ARPANET suffered congestion collapse in the late 1980s, and how TCP was modified to prevent that collapse. In addition to understanding the problem of congestion control for TCP in the ARPANET, you should consider how the design choices made affect today’s networks. With the vast increases in network performance, and the growing prevalence of wireless networks and non-TCP traffic, are the design choices expressed in this paper still relevant today? How does the presence of network features such as RED, ECN, and QoS affect the congestion response of the network and the design of transport protocols?

- V. Jacobson and M. J. Karels, Congestion Avoidance and Control, Proceedings of ACM SIGCOMM ’88, Palo Alto, CA, USA, August 1988.

## Week 17: Measuring Network Performance – Traffic Dynamics

We continue our discussion of how properties of the network can affect the performance of transport layer protocols, using the example of TCP. You should read the paper to understand some of the ways in which TCP connections are influenced by changes in network performance, to gain an understanding of how measurements can be used to improve protocol performance, and to understand the difficulties inherent in inferring protocol performance data from measurement studies.

- M. Allman and V. Paxson, “On Estimating End-to-End Network Path Properties”, Proceedings of ACM SIGCOMM’99, Cambridge, MA, USA, August 1999.

## Week 18: Simulation and Modelling

The first paper to be discussed this week ("Why We Don't Know How To Simulate The Internet") was previously read as part of the Advanced Topics in Computing Science module. It outlines challenges encountered when developing realistic simulator models of the Internet. The discussion will cover the difficulties of capturing scale, topology and heterogeneity of protocols and traffic, along with strategies to cope with these difficulties.

The second paper ("Performance Evaluation with Heavy Tailed Distributions") discusses the evidence for heavy tailed distributions as a realistic model for network traffic as well as computation, and discusses implications on the design of future networks and servers. The following points will be covered during the discussion: characteristics of heavy tailed distributions, characteristics of self similar traffic, evidence of this traffic model, and performance implications network and server design.

- S. Floyd and V. Paxson, "Why We Don't Know How to Simulate the Internet", Proceedings of the 1997 Winter Simulation Conference, December 1997.
- M. E. Crovella, "Performance Evaluation with Heavy Tailed Distributions (Extended Abstract)", Lecture Notes in Computer Science 1786, 2000. Slightly revised version appeared in Proceedings of the 7th International Workshop on Job Scheduling Strategies for Parallel Processing, June 2001.

This session will be led by Dr. Mohamed Ould-Khaoua.

## Week 19: Simulation and Modelling

This paper compares four task scheduling algorithms in a distributed server. Two task execution times (or task sizes) are examined: exponential and Pareto distributions. The discussion will cover the differences between the four task scheduling algorithms, the characteristics of the Pareto distribution, performance metrics and analysis of the simulation results.

- M. Harchol-Balter, M. E. Crovella and C. D. Murta, "On Choosing a Task Assignment Policy for a Distributed Server System", Journal of Parallel and Distributed Computing, Volume 59, Number 2, 1999.

This session will be led by Dr. Mohamed Ould-Khaoua.

## Week 20: Simulation and Modelling

This paper assesses the performance of the ServerNet network architecture when subjected to self-similar traffic, and contrasts it to that under Poisson traffic. The following points will be covered during the discussion: Modelling of self-similar traffic, Poisson traffic, ServerNet architecture, switching technique used, performance metrics, and analysis of the simulation results.

- D. R. Avresky, V. Shurbanov, R. Horst and P. Mehra, "Performance Evaluation of the ServerNet SAN under Self-Similar Traffic", Proceedings of the 10th Symposium on Parallel and Distributed Processing, San Juan, Puerto Rico, April 1999.

This session will be led by Dr. Mohamed Ould-Khaoua.

### Week 21: Naming and Addressing – Performance

The paper for this week is a study of the performance of the root DNS servers, discussing why there is so much traffic and how the performance of the DNS system can be improved. You should read this paper to understand not just the issues exposed with DNS implementations, but if/how other protocols and their implementations might be affected in similar ways. What are the implications of the data presented here for denial-of-service protection? What are the implications for the design of alternative name resolution systems?

- D. Wessels and M. Fomenkov, "Wow, That's a Lot of Packets", Proceedings of the Passive and Active Measurement Workshop, La Jolla, CA, USA, April 2003.

### Week 23: Naming and Addressing – Alternative Architectures

After our discussion of the measured performance of the DNS last week, we now discuss an alternative, peer-to-peer, name resolution system for the Internet. You should consider how the performance of this system would compare with the DNS, and if you believe this system to be deployable?

- V. Ramasubramanian and E. G. Sirer, "The Design and Implementation of a Next Generation Name Service for the Internet", Proceedings of ACM SIGCOMM 2004, Portland, OR, USA, August 2004.

### Week 23: Naming and Addressing – Alternative Architectures

The paper this week discusses how the naming architecture of the Internet might be re-designed and why confusing host names and routing identifiers was an unfortunate design choice. You should read the paper to understand how peer-to-peer algorithms might be used for naming and addressing, why linking routing and application-level addressing is problematic, and how the ideas discussed impact mobility, security, etc.

- H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica and M. Walfish, "A Layered Naming Architecture for the Internet", Proceedings of ACM SIGCOMM 2004, Portland, OR, USA, August 2004.

### Week 24: Future Directions

We conclude the module with a discussion of the future of the network, looking at how the design of networked systems may evolve, and how we can attempt to direct that evolution through careful system design.

- D. D. Clark, K. R. Sollins, J. Wroclawski and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", Proceedings ACM SIGCOMM '02, Pittsburgh, PA, USA, August 2002.

## 3 Assessment

Assessment is my means of a presentation given to the Networked Systems, Measurement and Control (NSMC) research group (30%) and by an open book examination (70%). This is a level M module, worth 5 credits.

Presentations should be given during weeks 19 to 24. Students must contact Dr. Olufemi Komolafe to arrange the date of their presentation. Each student is responsible for choosing the subject for his own presentation, related to the topics under discussion in the module. Subjects must be approved by the module co-ordinator by the end of week 18 (18th February); students may not choose a topic already chosen by another student.