46PaQ: IPv4 and IPv6 Performance and QoS

http://www.cs.ucl.ac.uk/research/46paq/

Saleem Bhatti *S.Bhatti@cs.ucl.ac.uk* *Computer Science*
Dimitris Miras *D.Miras@cs.ucl.ac.uk* *Computer Science*
Nicola Pezzi *np@hep.ucl.ac.uk* *High Energy Physics*
Centre of Excellence in Networked Systems
University College London, Gower Street, London WC1E 6BT

1 Introduction

There is much ongoing work to define new protocols and mechanisms for high-speed, QoS-controlled networking within Internet Protocol (IP) based environments. This includes work on QoS mechanisms and services, congestion control mechanisms (e.g. ECN and TCP-like adaptation for UDP applications) and work on new transport protocols for specific purposes (e.g. DCCP). However, as these new mechanisms have been developing, there has not been much activity in trying to observe and analyse the behaviour of these systems working together, within a real, integrated networked environment. Additionally, the way in which applications make use of the new services, mechanisms and protocols have also not been observed and analysed within a real network environment using these mechanisms. With the rapidly increasing deployment of IPv6, it is also vital to consider carefully the differences in behaviour in the use of these mechanisms compared to IPv4. In this project, we will have two broad areas of work. Firstly we will, through experimentation and analysis, examine how QoS mechanisms and congestion control can operate together in a very high-speed (multi-Gb/s) IPv4 and IPv6 environment to support TCP- and UDP-based applications. Secondly, we will propose ways in which such networks can be instrumented in order to provide performance and operational data to network operators as well as users and applications through appropriate APIs and using network monitoring equipment. We will examine the performance of the system in a real networked environment operating at very high speeds. To stretch the network services, we intend to test with selected applications from the e-Science/Grid community with very demanding needs.

2 Rationale

The networking landscape is changing. Network speeds of 100Mb/s on the desktop are common, and 1Gb/s interfaces are becoming more widely used (available even on laptops now), whilst 10Gb/s Ethernet has now been ratified. More users from a diverse range of disciplines are using the network for their data-to-day activities and this has been encouraged by the worldwide Grid funding. Whereas in the past, the core network has relied on over-provisioning to satisfy user needs, this will not be possible for the future, as desktop access speeds approach the same order of magnitude as the core network speeds. It is recognised that the future demands on network connectivity must be investigated now so that research communities can cope with the needs of the users as they arise. This is evidenced by various initiatives and projects around the world to investigate very-high speed networking with QoS-control, e.g. QBone, CANARIE (CA*net), 6NET, Internet2, EU-DataGrid, DataTAG, NetherLight and StarLight. Also, some UK projects in these areas have been funded - MB-NG, GridProbe and GRS. These projects investigated, respectively; QoS provisioning in the core at high speed (multi-Gb/s);

network monitoring and measurement; end-to-end QoS control.

So, there have been distinct and separate activities in the areas of Quality of Service (QoS) mechanisms, high-speed networking modifications to network- and transport-level protocols, congestion control mechanisms, and network performance measurement and monitoring. However, a real network may need to run all of these currently disparate systems together to provide services for a wide variety of users. We would like to use QoS services, congestion control and network performance measurement and monitoring systems ***together*** in a single network environment that also offers very high-speed connectivity (1Gb/s and beyond). The following questions arise:

1. Is it indeed possible to use these mechanisms together in a true end-to-end controlled service?
2. How do these different mechanisms interact with each other?
3. How do we monitor and measure performance of such networks?
4. How do such mechanisms function and perform in very-high speed network environment?
5. Is it possible for the application to exploit some knowledge of the operation of the network, with such functions in place, to adapt its behaviour and so make better use of the resources available?

Our aim in this project is to seek answers to the questions posed above. We believe that to enable the kind of multi-disciplinary research that is required, for example, by the e-Science and Grid community, requires that the data-intensive, distributed applications that need to make use of the more sophisticated network services should be able to do so easily. We also believe that, in order to allow true end-to-end connectivity in a scaleable fashion for the ever-increasing applications and users that inhabit the IP network space, the development and use of IPv6 infrastructure and IPv6 applications is vital. So, 46PaQ intends to look at the issues in mechanisms that would allow QoS-capable, congestion-controlled, high-speed network protocols that can operate at very high speeds (multi-Gb/s) with IPv6 as well as IPv4, and which can provide performance monitoring information, via tools and APIs, so that applications can adapt to the available network conditions.

3 Current status

46PaQ started in December 2004 and will last 24 months. The partners are: UCL, University of Cambridge, UKERNA, Cisco and Sun Microsystems.. Currently, UCL are testing multi-Gb/s protocol operation in the laboratory, using a mix of transport protocols, including UDP, TCP and DCCP. The experiments and tests that are currently being carried out in the lab will then be moved onto UKLIGHT, first at 1Gb/s and then at higher speeds.

Our investigations will include:

1. Examining the raw performance of the protocols.
2. Analysing the interaction of mechanisms such as congestion control and protocol data transfer behaviour.
3. Looking at the performance and tuning of the end-system.
4. Examining what kind of feedback is possible and how it would help the performance of the application.

We have direct links with two other UKLIGHT projects.

The protocol-performance work will also interact with the project ESLEA (Exploiting Switched Lightpaths in E-Science Applications).

The work related to monitoring and measurement will also interact with the projects MASTS (Measurement at All Scales in Time and Space).

ESLEA, as well as other collaborative e-Science/Grid projects at UCL, have a range of applications that could also be used within out test environment.