ERCIM “Alain Bensoussan”
Fellowship Scientific Report

Fellow: Henrik Lundqvist
Visited Location : Q2S, NTNU, Trondheim
Duration of Visit: 1/12-2005, 31/8-2006

I - Scientific activity)

The common thread of the research has been to investigate how to use explicit feedback from the network to provide quality of service for real-time traffic. The first part considered how to use early congestion notification (ECN) in a single class network where TCP and real-time traffic use the same congestion feedback and have to coexist in the same class. In the second part, the scenario is that the real-time traffic is separated from the best-effort TCP traffic by means of scheduling in the routers.

The goal of the research is to provide real-time multimedia applications with a quality that is consistent once a session has started. Instead of continuously adapting the rate, new sessions have to be blocked in case the network does not have sufficient free resources to provide good quality. From the users point of view this is more attractive than an adaptive application where the quality varies too much during the life-time of the session. If the application is adaptive the changes have to be rather small and infrequent in order not to be annoying.

The admission control that determines whether to admit new sessions into the network can either be centralized or decentralized. A decentralized approach, as studied here, avoids problems with scalability and can be implemented at the end-systems or edge routers, ideally without requiring all network operators to cooperate. When the admission control is performed in the end-system the condition of the network path has to be estimated in order to determine whether it can support a new flow. The estimation can either be based on implicit congestion information, such as the loss rate or delay, or on explicit congestion information, such as marking of packets by the routers. The advantage of explicit marking is that the information can be fed back to the end systems before the quality experienced by the user is degraded.

The feedback information can be used both for proactive admission control, reactive rate control by rate-adaptive applications and pre-emption of low priority flows when high priority traffic needs the resources. The first part of this project has been an extension of a host based differentiated services architecture previously developed at KTH. The architecture is based on a best-effort network where all the differentiation is made in the end-systems, which should allow a simple deployment. The improvement studied is to add explicit congestion feedback rather than solely using loss estimates to determine whether to admit a new session or not. Since there is only one class, the same feedback information is used by TCP and for admission control of real-time traffic, therefore the feedback information and the admission policies have to be adapted to coexist with TCP. The study is based on simulations and the results indicate both that it is possible to make more accurate admission decisions and thereby differentiate between users of different priority, and that the architecture can be extended to scenarios with lossy links, for example wireless access networks, since the losses are not mistaken as signs of congestion.

To provide real-time applications with better quality it is an advantage to separate it into a separate class so that it does not have to share its capacity with TCP traffic. There is ongoing work in the IETF transport working group to provide congestion feedback for real-time traffic in such scenarios, which uses the ECN bits in the IP header to signal three different levels: no-congestion, low congestion and high congestion. These levels can then be used to determine whether to admit or pre-empt sessions.

As an alternative we investigate the possibility to use a stochastic marking where the packet marking probability is based on the arrival rate at the queue in a router. If the marking probability as a function of the load is a simple invertible function it would be possible to calculate the load of the router. However, there are typically many routers on an Internet path, and a path passing many heavily loaded routers should be more likely to be blocked than a session over a single router. Hence, it is difficult to determine a marking function that works well regardless of the number of routers, the traffic patterns and the requirements of different applications. We therefore choose a simple quadratic function, and investigate how this behaves compared to a deterministic multi-level marking function. Preliminary simulation results indicate that the probabilistic marking gives almost the same behavior as deterministic marking in simple scenarios, while more complex scenarios are still under investigation. The possible benefits of the probabilistic marking are to provide more flexibility to extensions to more classes and a stochastic behavior that may limit the need for extra functionality for determining which sessions to block or pre-empt.

II- Publication(s) during your fellowship

Please insert the title(s), author(s) and abstract(s) of the published paper(s). You may also mention the paper(s) which were prepared during your fellowship period and are under reviewing.

Henrik Lundqvist, Gunnar Karlsson, “Host-Based Service Differentiation with Congestion Feedback”, International Workshop on Quality of Service, IWQoS 2006, June 19-22, 2006.

***Abstract*—This paper investigates the possibility to differentiate services by using endpoint traffic controls, while sharing a single best-effort network. The starting point is a scheme based on a combination of probe-based admission control in the end-systems for streaming traffic and error correction to isolate different flows. There are inherent limitations in pure host-based control mechanisms and we therefore investigate the improvement that can be achieved when explicit congestion notification is used in routers. In particular we investigate the sensitivity to different active queue management mechanisms and parameter settings. Simulations show that a high ECN marking rate helps to differentiate between flows, but it is important to make sure that the chosen AQM works well for the chosen parameters. This study also shows that a main benefit of ECN is to enable the extension of the scheme to wireless environments.**

III -Attended Seminars, Workshops, and Conferences

Please identify the name(s), date(s) and place(s) of the events in which you participated during your fellowship period.

International Workshop on Quality of Service, IWQoS 2006, June 19-22, 2006, New Haven, Connecticut, USA.