INTERMON – Advanced Architecture for Inter-Domain Quality of Service Monitoring, Modelling and Visualisation

by Ulrich Hofmann and Christof Brandauer

Salzburg Research Advanced Networking Centre (ANC) is working jointly with European partners from industry and research institutions on Internet inter-domain performance modelling and visualisation for network planning and inter-domain routing.

Over a number of years, the integration of QoS (Quality of Service) provisioning into Internet technologies was the motivation for several European 5th Framework Program projects (eg AQUILA and TEQUILA). Internet-Telephony, video conferencing and content distribution via the Internet were the driving forces. The requirements are well known (eg telephony requires a transmission delay less than 100ms), but the current QoS provisioning standard ‘Integrated Services’ architecture has failed because of the large overhead generated by the per-flow resource reservations. The recently developed ‘Differentiated Services’ provisioning schemes avoid this overhead. These standards and algorithms must be integrated into routers, and must be managed by the network providers per network domain (‘intra-domain’). With these mechanisms, the QoS provisioning task can be technically solved per network domain.

However, a large portion of Internet traffic carried by a provider comes from and goes to other domains of the provider or other providers’ networks. In such a scenario the provider acts as a transit carrier. Many network providers want to increase their competitiveness in this business area by offering global QoS provisioning. They therefore need a global view of the transmission qualities of other possible transit providers on the path to specified destinations, to find the optimal (eg by cost/performance) inter-domain path.

In Figure 1 for example, network provider A receives a transmission request from the video conferencing application server: {destination = application client, QoS_delay < 100 ms, QoS_rate = 1 Mbit/s }. Provider A knows from the inter-domain routing protocol BGP (Border Gateway Protocol) the two possible paths to the destination provider D: path_1=ABCD, path_2=ACD. Because each of the providers B and C wants to get the contract for the transit transmission from A to D, both offer information on their transmission quality and costs to provider A. The information from provider B to provider A is aggregated for the path BC and path CD (which provider B got from provider C). Now provider A is able to select between the two paths and informs the application server of the delay and the cost of a transmission via domain A to the application client. In order to support such business scenarios, the INTERMON project is focusing on the inter-domain traffic measurement and modelling aspects.

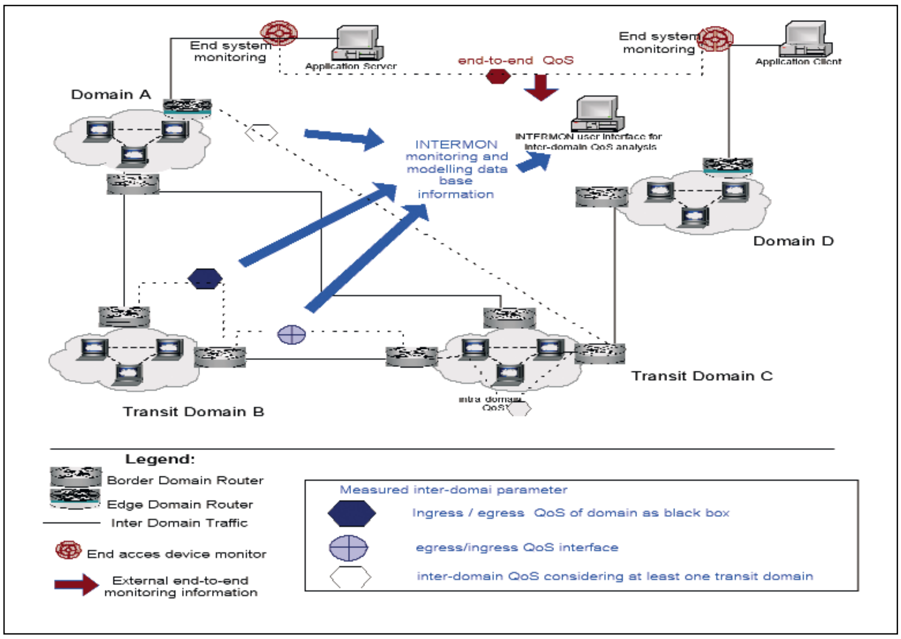

Figure 1: Inter-domain QoS analysis with policy controlled data collection.

Links between network providers generally have high capacities (ranging from 100 Mbit/s to 10 Gbit/s). However for transit traffic, this bandwidth may be reduced by the peering contracts between the providers. To monitor the link load, the data packets must firstly be monitored with the parameters {arrival time [μs or ns], length [bytes], flow_id}. The flow_id allows per-flow traffic monitoring, eg for incoming transit flows with different destinations. So the domain B provider in Figure 1 is able to monitor the traffic between the application servers at the entry and exit routers. In the INTERMON project this is called ‘policy-driven monitoring and measurement’, because different monitoring application requirements need different monitoring methods (eg for delay measurements, two monitoring points are necessary).

The next step is for this information to become the input for a ‘what if’ analysis,

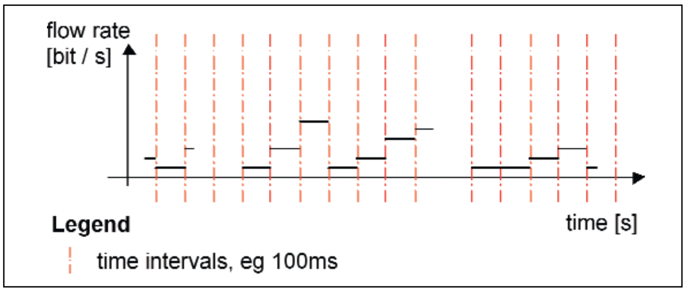

Figure 2 : The simulation process is driven by traffic streams that are modelled as chunks of fluid flows.

eg: “How much increase of delay or loss is caused by an increase of the transit traffic by x%?” To find the answer, the INTERMON project derives inter-domain QoS simulation models. One of these is a so-called fluid approach. Obviously the simulation of complex inter-domain scenarios cannot be done by a per-packet simulation, due to the explosion in the number of simulation events that occur in such large-scale scenarios. For each event the simulator has to update the system state.

The first step towards a scalable simulation model is an approach in which traffic streams are modelled as chunks of fluid flows. The monitored individual packet arrivals are aggregated to traffic load events, eg per 100 ms, and the simulation process is now triggered by these traffic-load events (see Figure 2).

A rigorous next step in this traffic modelling is the transformation of this discrete load process into a continuous ‘fluid’ process. To retain the most important process characteristics - mean, variance, and autocorrelation - the continuous fluid process is derived from the discrete fluid process by a newly developed iterative algorithm. With this abstraction, inter-domain links become continuous queuing systems and the dynamic relations between incoming fluid-traffic, service link rate, buffer occupation, loss rate etc, are described by differential equations.

On the basis of these models, the ‘what-if’ scenarios can be executed with the help of powerful existing continuous simulation modelling tools like SIMULINK. To optimally support the network operator/planner, a high degree of automation is employed to create the simulation. Topology information is imported into a graphical user interface (GUI) where the scenarios of interest (‘what-if’ changes) are configured, and the current load situation is fed from the monitoring database (IPFIX standard) into the simulation.

In general, the fluid simulation algorithms run on a digital computer and the time progress of the differential equations will be approximated by discrete ‘sufficiently small’ time steps. The dependencies between simulation performance (accuracy, simulation time) and the size of the simulation time step will be investigated in the next project period. Running the fluid simulation with sufficiently small time intervals may be interpreted as a step backward to the granularity of the criticised per-packet simulation, but today’s powerful numerical processors are able to solve the approximations for differential equations very efficiently. Finally, a very high degree of accuracy may not always be the main interest in large inter-domain scenarios and the ability to arbitrarily choose the trade-off between simulation accuracy and performance is a powerful feature.

Link:
http://www.ist-intermon.org/

Please contact:
Christof Brandauer
Salzburg Research Forschungsgesellschaft
E-mail:
christof.brandauer@salzburgresearch.at

Ulrich Hofmann,
Salzburg Research Forschungsgesellschaft and Salzburg University of Applied Sciences and Technology
E-mail: ulrich.hofmann@salzburgresearch

Inference for Random Sets

by Marie-Colette van Lieshout

Image analysis and spatial statistics are widely used in medical diagnostics (eg body scans), satellite technology (eg cartography) and analysis of spatial correlation (eg in forestry and epidemiology). At CWI, scientists in the group Signals and Images have studied the problem of extracting linear features such as road networks from remotely sensed images. A new set of methods has been constructed based on Monte Carlo Markov Chain simulation. The corresponding new simulation procedures are faster and more precise than earlier methods.

Many images found in microscopy, material science and biology can be described as a set of independent, randomly placed particles. Think for example of the distribution of pine trees in a forest or, more aggressively, the distribution of fallen bombs. A formal description of such a random set is called a Boolean model.

Notwithstanding the strong independence assumptions, inference is not trivial because of the occlusion effect: Only an image of all the particles – rather than individual particles – is observable.

However, due to interaction, not all images can be described as completely random spatial patterns. Examples are the distribution of cells in the cat retina (eyes), metal particles in an alloy, and road networks. Cats’ eye cells arrange themselves in two lattices, hard particles cannot penetrate each other and roads tend to be long and straight and have perpendicular crossings. Therefore, a