RTT TECHNOLOGY TOPIC
January 2003

IP network processor performance limitations

In our last **HOT TOPIC (December)**, we reviewed the impact of handset power budgets on session amplitude (session complexity) and session duration (session persistency) and by implication, the related impact on session value.

In this month's **HOT TOPIC**, we move into the network and look at router packet processing. In particular, we want to look at some of the performance metrics needed in order to … yes, you've guessed already, preserve session value.

Definition of a Real Time Network

The IEEE defines a real time operating system as 'A system that responds to external asynchronous events in a predictable amount of time'. Handset hardware and software engineers, for example, need to make sure that application layer and physical layer functions are executed within known and predictable time scales in response to known and predictable external events and requests - this is often described as a deterministic process.

We can apply a similar definition to a real time network as 'A network that can provide throughput to asynchronous traffic in a predictable amount of time', ie network hardware and software engineers need to make sure that router processors can process offered traffic within known and predictable time scales. This in turn requires a reasonably intimate understanding of the offered traffic mix and offered traffic properties, particularly any time inter-dependencies present between multiple coded channel streams.

Consider for example a user with several OVSF code streams simultaneously multiplexed on to a 3G radio channel. The code streams could be carrying separate voice, image, video and application data streams which may or may not be time inter-dependent. The properties of this relatively complex single user multiplex will need to be preserved as the traffic streams are moved into and through the network. A certain amount of delay might be tolerable but differential delay (some channel streams delayed longer than others) could potentially destroy the integrity (and value) of the complex session exchange.

Switching or Routing

The general assumption is that there will be more routing and less switching in a next generation network - less use of hardware, greater use of software to reduce cost and improve flexibility.

However, software based routing introduces delay and delay variability - the time taken to capture a packet, the time taken to check the header and routing table, the time packets spend in buffers waiting for the router to deal with other packets and the time packets spend in buffers waiting for egress bandwidth to become available (queuing delay).

Traffic shaping protocols such as Diffserv and MPLS can help manage queuing delay (ie reduce delay for some packets by increasing delay for other packets) but queuing delay can still be 20 or 30 milliseconds and varies depending on traffic load. Transmission re-tries (using TCP/IP) introduce additional variable delay.

IP Switching as an Option

Rather than treating packets individually, a sequence of packets can be processed as a complete entity - the first packet header contains the properties and bandwidth requirements of the whole packet stream - this is described as flow switching or IP switching - 'a sequence of packets treated identically by a possibly complex routing function'. This would seem like a good idea, particularly as session persistency increases over time. However, session properties and session characteristics tend to change as a session progresses so more flexibility is needed. It is difficult to deliver flexibility and deterministic routing or switch performance.

Significance of the IPv6 Transition

An additional challenge for an IPv4 router is that it never quite knows what type of packet it will have to deal with next. IPv6 tries to simplify things by using a fixed length rather than variable length header and by reducing the 14 fields used in IPv4 to 8 fields - the protocol version number, traffic class (similar to type of service in IPv4), a flow label to manage special priority requests, payload length, next header, hop limit, source address and destination address. This is one of the reasons Japanese vendors are keen on mandatory IPv6 in routers - it makes deterministic performance easier to achieve.

Router Hardware/Software Partitioning

Performance can also be made more predictable by adding in a hardware co-processor (or parallel co-processors) to the router.

Figure 1 shows the packet flow sequence in a packet processor.

To minimise delay, the router divides packet processing into a number of tasks.

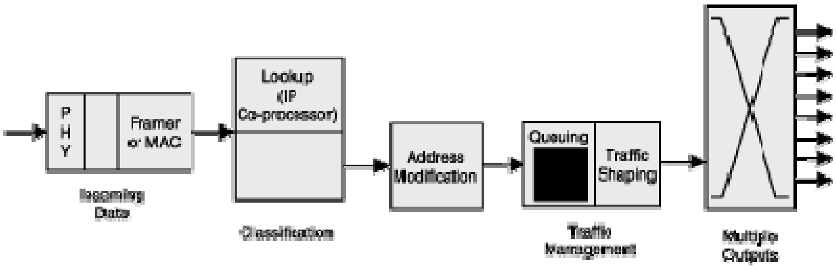

Figure 1: Router Packet Processing

The data comes in from the physical layer and is demultiplexed in accordance with the MAC (medium access control) layer rule set, ie packets within frames within multi-frames. The packet is then sent for Classification. At this point a hardware co-processor may be used to improve look up performance. Performance is defined by the number of searches per second, the number of entries in the table, and whether or not multi-protocol tables are used. Multi-protocol tables are needed if differentiated classes of service are supported. A software based standard processor based solution can take several hundred instruction cycles to classify a packet with QoS and/or security attributes. A co-processor can perform the task in a single clock cycle but lacks flexibility, ie it can only be used when the decisions to be taken are largely pre-defined and repetitive, ie this is hardware based switching rather than software based routing.

The Impact of Increasing Policy Complexity

This looks like a good solution but there's a drawback. We have described in previous HOT TOPICS how session complexity increases over time - for example we might need to support an increasing number of multi-user to multi-user multimedia exchanges. In these exchanges, policy management can become really quite complex. Policy rights might be ascribed to the user such as a priority right of access to delivery or storage bandwidth. The user, the user's device or the user's application may be (likely will be) authenticated and the user's traffic may have end to end encryption added for additional security. In a multi-user session, the security context may change/will change as users join, leave or rejoin the session. This implies substantial flexibility in the way in which packets are handled, hard to realise in a hardware co-processor.

Summary

In an ideal world we would like to combine consistent deterministic throughput and flexibility. In practice this is hard to achieve. One (simple) way to achieve consistent performance and flexibility is to over-dimension transmission bandwidth in the IP network but this has a cost implication. It is difficult to escape the reality that consistent network performance comes with a price tag attached. Having to handle highly asynchronous traffic increases that price tag. Bursty bandwidth is expensive bandwidth.

About RTT Technology Topics

RTT Technology Topics reflect areas of research that we are presently working on.

We aim to introduce new terminology and new ideas to clarify present and future technology and business issues.

Do pass these Technology Topics on to your colleagues, encourage them to join our Push List and respond with comments.

Contact RTT

RTT, the Shosteck Group and The Mobile World are presently working on a number of research and forecasting projects in the cellular, two way radio, satellite and broadcasting industry.

If you would like more information on this work then please contact

geoff@rttonline.com

00 44 208 744 3163