








# Commerce Commission Workshop on UBS Technical Issues – Questions 1-12



Telecom New Zealand  
11 February 2005

# Contention Ratio, jitter, packet delay and latency

*Commission question:*

1. How do the parties define these terms, and how are these parameters appropriately measured?

In order to answer this question, we will define each of the terms, as well as explain how they are inter-related, and affect each other, through some simple examples.

# Telecom manages the DSL infrastructure to achieve the best outcomes for all customers

The DSLAM and associated infrastructure is a very complex shared resource:

- Supporting many services
- For many different types of customer

**TelstraClear is only one of many users of this shared resource**

Hence it is critical that the resource is shared appropriately between the different users and that each user contributes appropriately to the costs of deploying and operating such an infrastructure

- This means that any consideration of service characteristics needs to take into account:
  - the full interaction between all the different types of use
  - and the resulting attribution of costs to achieve the required service characteristics

**Telecom is the only party who is in a position to manage the infrastructure in the best interests of all users**

- It is also natural that TCL will only be concerned about the interests of their customers

# Contention Ratio, Jitter, Packet Delay and Latency

## Definitions:

- Telecom adopts the ITU definitions for these parameters, as first choice
- This includes a preference over IETF definitions where these are either in conflict or offer options

## Measurement:

- Telecom uses a set of active test probes attached to dedicated ports of selected DSLAMs (sample only)
- These probes generate packet streams of the appropriate frequency and size, which are then collected at a specific point within the network wherein they are processed to deliver the desired results as indicated in the following slides
- All asymmetrical, uncommitted bit rate services (retail Jetstream and UBS) are measured by the same probes (distributed evenly across the services)
- Sampling is performed in such a manner as to provide confidence that more than 95% of (Jetstream and UBS) connections are within Telecom's design performance specifications for more than 95% of the time
- Increasing this confidence level beyond 95% would require significant additional investment in probes and allocation of ports (Approx. \$2500 per probe and no revenue per measurement port)

# Contention Ratio: Definition

## Definition:

- The ratio of the number of end users of defined peak bandwidth (PIR) relative to the allocated bandwidth between the DSLAM and the first ATM switch during the peak busy hour of any day
- The peak busy hour is usually defined as the busiest hour of a given year and is typically adjusted annually based on actual measurement
- The above definition is the optimum approach for defining the peak resource required to support a given set of DSL connections
- It relates to the TCL definition of Contention Ratio, which is also in common use, through an assessment of the ratio of active to inactive users during the peak busy hour
  - The TCL definition will result in the allocation of 2-3 times the aggregation bandwidth for a given Contention Ratio and set of end user connections, which creates a cost which will need to be recovered by some means

# Contention Ratio: Measurement

## Measurement

- Contention Ratio is established through a set of dimensioning rules for the associated uncommitted bit rate Virtual Path (VP)
- The measurements used for Packet Delay Variation and Packet Loss (see below) are used to determine that a sample of services using a given VP is performing within its design parameters
- As more customers are connected to the VP and as the measurements indicate variations in the ratio of active to inactive users, the dimensioning rules are adjusted and reconfiguration of the VPs is executed – typically on a monthly cycle basis per VP
- All UBS and retail Jetstream services use the same VP and are treated equivalently by the DSLAM, ATM backhaul and BRAS/L2TP LAC (see later detailed description)
- Contention Ratio variation between DSLAMs will occur over time as customers' churn on and off the DSLAMs between adjustment cycles

# Interaction Between Users and Services


	<b>Single Service per Access</b> (Single user per access - eg. One PC executing web browsing or gaming)	<b>Multiple Services per Access</b> (Single user using multiple services or multiple users using multiple services – eg. Simultaneous web browsing and gaming)
<b>Uncongested Network</b> (each customer's traffic is independent from every other customer's traffic)	Bandwidth will largely determine the end user experience	End user experience is determined by a combination of Bandwidth, Quality and the types of services being used over the access
<b>Congested Network</b> (each customer's traffic is dependent on other customer's traffic using the same common network resources)	End user experience of a given customer is determined by the collective behaviour of all customers for a given available shared bandwidth	End user experience is determined by the interaction of the multiple applications and the collective behaviour of all the customers for a given available shared bandwidth

TCL Preference

# Throughput vs Peak Information Rate

Consider a combined data stream of 3 typical web browsing users over a busy hour:

- Peak Information Rate (PIR) per user = 512kbps
- Simulation of 3.9 MBytes of web browsing per user during the hour period
- Implies an average throughput per user of 8.6kbps
- And a combined peak data rate of 1100kbps




Peak to average ratio =  $1100 / 8.6 = 128 : 1$ .

Data networks need to support high instantaneous peak data rates but per user average throughputs are considerably less than the peak and so contention ratios can be applied enabling efficient statistical multiplexing

# Contention Ratio Example

For these three services with a peak information rate of 512 kbps we define a contention ratio of up to 20:1 in the peak busy hour

Requires a minimum allocated bandwidth of  $(512 \times 3) / 20 = 76.8$  kbps (Red Line)



In practice, for efficient statistical multiplexing we would need a minimum of 20 services at 512 kbps with a 20:1 contention ratio

Then we allocate a minimum bandwidth of 512 kbps, to ensure that the peak information rate (PIR) can be achieved (Green Line)

Even then with only 3 connections active, congestion will occur

# Jitter or Packet Delay Variation (1)

## ITU Definition:

- The difference in packet transfer time between the minimum absolute Packet Transfer Delay and the maximum absolute Packet Transfer Delay of any packet over a 1 minute measurement interval for a stream of packets that are all the same size but in the presence of 1472 Byte packets with randomly varying arrival times
- This is the 2 point definition used in ITU-T Y1541 and Appendix II of ITU-T Y1540
- There appears to be agreement between Telecom and TCL on this basic definition
- We also agree that Jitter and Packet Delay Variation can be used inter-changeably


# Jitter or Packet Delay Variation (2)

## Measurement:

- There may be some disagreement on the measurement of this parameter
- Normally measurement would be achieved by taking 1 minute samples of one way packet transfer delay, using a packet of defined size (eg. TCL request for 64 Bytes) between an end user's modem and a defined point in the network on a continuous basis resulting in a distribution of delay variations from which the maximum and minimum can be calculated over any given 24 hour period
  - During this 24 hour period it is expected that the peak busy hour will be experienced, so providing an approximation to the randomly arriving 1472 Byte packets
  - For best results the Delay Variation should be calculated from the daily distribution averaged over a minimum of 2 weeks and preferably 8 weeks

# Minimising Delay Variation using Policing

Under policing, any data packets above the defined PIR of the aggregated data stream (with defined contention ratio applied) will be dropped – resulting in Packet Loss, but little Delay Variation

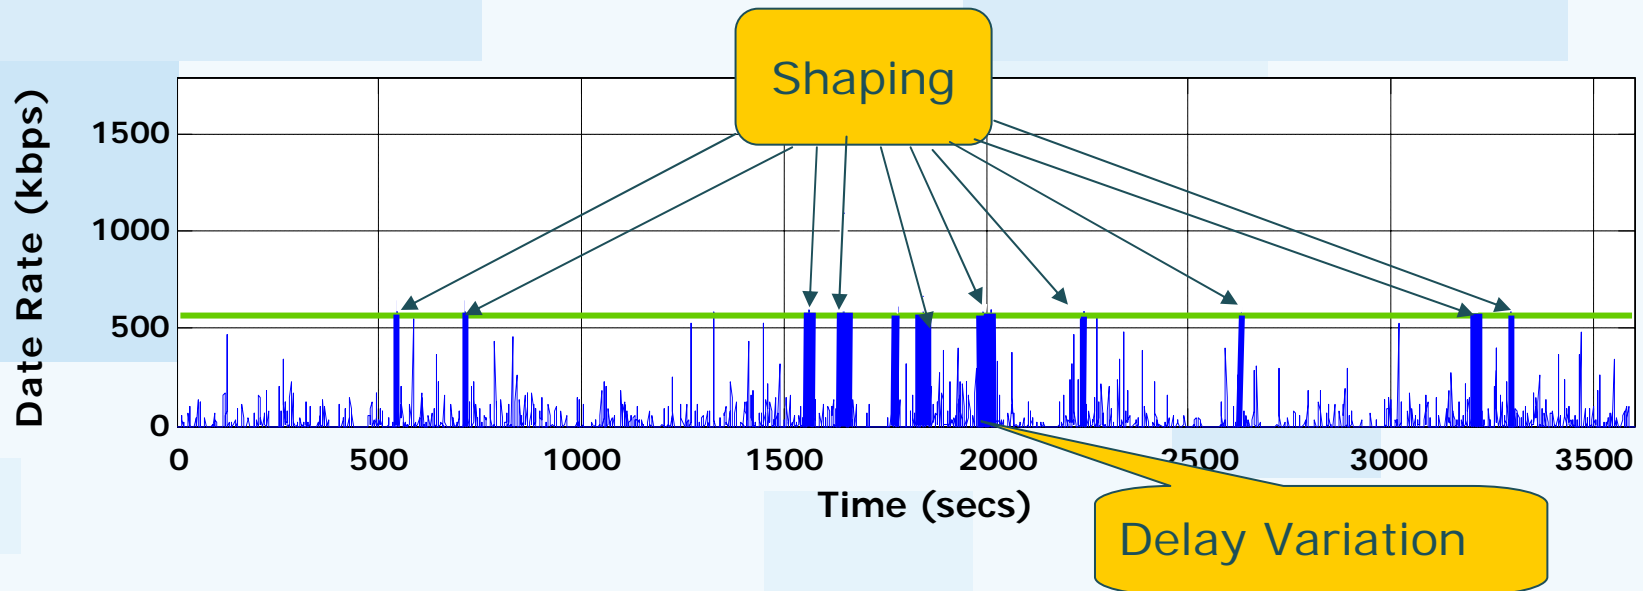


Under this regime we rely on the higher layer protocols, such as TCP, to resend the discarded packets and hence to manage the integrity of the end to end communication. The impact of this on end user experience will be to transfer the delay associated with resending the discarded data to the application layer

# Delay Variation Example

Under shaping, data packets above the defined PIR of the combined aggregated data stream (with defined contention ratio applied) are delayed until there is a lull in the data stream and then re-inserted:

- Most of the time this will result in zero Packet Loss, but will increase Delay Variation
- Under severe congestion, with a defined buffer size, the buffer can over-flow, resulting in packet loss



Hence as can be seen, there is a trade-off between buffer size, which determines Delay Variation and Packet Loss Ratio

# Latency

## ITU Definition:

- The absolute time measured from the time the first bit of an IP packet is passed to the source (ingress) User Network Interface to the time the last bit of the packet is transmitted by the destination (egress) User Network Interface towards the user's equipment or the host (Ref: ITU-T-Y1540 and Y1541)
- My understanding is that Telecom and TCL are in agreement over this definition

## Measurement:

- The measurement procedure is usually the same as that for the measurement of Delay Variation, with the minimum of the Delay Variation measurements for a given packet size (eg. 64 Bytes) being defined as the Latency

# Packet Loss

## ITU Definition:

- Packet Loss Ratio is the ratio of the total lost packet outcomes to the total transmitted IP packets in a population of interest (Ref: ITU Rec. Y1540)
- Lost packets are defined as those discarded or those severely delayed in relation to the service contract
- The population of interest is defined as the sample of packets measured over a 15 minute interval for all flow bandwidths above 64 Kbps
- The populations are sampled regularly and averaged over a period of one month
- My understanding is that TCL and Telecom are aligned on this definition

## Measurement:

- The measurement of packet loss is typically combined with that of Delay Variation by detecting those packets from the source that do not reach the egress destination by identifying missing packet numbers
- The number of missing packets per 15 minute sample period are collected and then these samples are averaged on a 4 weekly basis

# The Trade-off between Contention Ratio, Delay Variation and Packet Loss (1)

From the previous discussion, it can be seen that there are trade-offs between Contention Ratio, Delay Variation and Packet Loss

- For example, for a given Contention Ratio, it is possible to define a delay variation and then this choice will determine the Packet Loss Ratio – all within defined statistical limits
- Other combinations of these parameters are possible, but in all cases, it is difficult to determine all three parameters independently, unless the Contention Ratio approaches 1
- Furthermore, defining the bounds for any one or a combination of these parameters, to a high degree of confidence (such as the 99.9% required by TCL for Delay Variation and Packet Loss Ratio) it is even more difficult to achieve, unless the Contention Ratio approaches 1
  - Under these conditions, the Delay Variation and Packet Loss Ratios must be achieved for all but 1.44 minutes per day, including during the busy hour

# The Trade-off between Contention Ratio, Delay Variation and Packet Loss (2)

- In order to achieve a Contention Ratio approaching 1, and meet the defined Delay Variation and Packet Loss Ratio for more than 99.9% of the time, it is necessary to allocate resources so that the Committed Information Rate (CIR) is close to that of the PIR for each connection
  - This defeats the ability to have efficient statistical multiplexing
- In addition, any requirement to establish confidence bounds at the 99.9% level for the Delay Variation and Packet Loss Ratio would require a much higher sampling rate for the measurement of actual performance than is currently the case
  - This would require many more probes and the resulting loss of port revenue, both of which increase the service costs significantly

All of the above needs to be put in the context of the Telecom retail JetStream services, where from a design perspective:

- The nominal Contention Ratio is  $<50:1$ , for approximately 95% of the time;
- The one way Delay Variation is  $<500\text{ms}$  for more than 95% of the time; and
- The one way Packet Loss Ratio is  $<3\%$  for more than 95% of the time.

# Summary: Meeting TCL's Requirements

Specification	TCL Requirement	Telecom Position
Latency (See Note 1)	< 50ms using a 64 Byte packet for > 99.9% of the time with interleaving off	< 50ms can be achieved using a 64 Byte packet with interleaving on for > 95% of the time
Delay Variation (See Note 2)	< mean packet delay for 64 Byte packets averaged over a 1 minute interval for > 99.9% of the time	< 200ms for 64 Byte packets averaged over a 1 minute interval for > 95% of the time at a cost
Packet Loss Ratio (See Note 3)	< 0.1% for 64 Byte packets averaged over a 1 minute interval for > 99.9% of the time	< 0.1% for 64 Byte packets averaged over a 1 minute interval for > 95% of the time at a cost
Contention Ratio (See Note 4)	< 20:1 for allocated connections for > 99.9% of time	< 20:1 as measured in the peak busy hour for more than 95% of the time at a cost

# Summary Notes

**Note 1:** The one way latency of  $< 50\text{ms}$  can be achieved with a 64 Byte packet with interleaving turned on or off

- Telecom has strong reasons for requiring Interleaving to remain on at this time – see later slides

**Note 2:** If the Contention Ratio of 20:1 is to be implemented, then Telecom is of the view that a Delay Variation of  $< 200\text{ms}$  for more than 95% of the time is inadvisable (and even this would have cost implications)

- Especially if a Packet Loss Ratio of less than 1% is to be achieved simultaneously for  $> 95\%$  of the time

**Note 3:** A Packet Loss Ratio of  $< 0.1\%$  cannot be achieved for typical use together with a Contention Ratio of 20:1 and a short Delay Variation

- With 500ms Delay Variation and 5:1 Contention Ratio, it might be achieved for  $> 95\%$  of the time (with cost implications)

**Note 4:** Telecom always deploys Contention Ratios relative to peak busy hour usage to ensure efficient use of resource (see later)

- A 20:1 Contention Ratio will have a cost impact relative to 50:1

**Note 5:** Any requirement to achieve a confidence level of  $> 95\%$  on any of the above parameters will have a direct cost impact for Probes and will require further optimisation of the relative values of the parameters, particularly reduction in Contention Ratio (with resulting cost impacts)

# Provision of non-shaped downstream speed

## *Commission's questions:*

2. Telecom state "there are certain technical and operational practicability issues associated with providing multiple downstream speed configurations" – What are these limitations and if so, can they feasibly be overcome?
3. What are the cost implications for additional network maintenance and resources on Telecom?
4. Can you rate shape the upstream speed without limiting downstream UBS speed?

In order to fully answer these questions, we need to traverse the range of services that are provided through Telecom's DSLAMs, how they are configured, how they are jointly managed, and the resulting cost impacts of the resource demands made by the various services.

# Provision of non-shaped downstream speed (1)

It is clear from the submissions, that TCL and their Consultants assume that the Telecom DSLAMs and the associated aggregation infrastructure are only used to support the retail JetStream type services and the UBS type services

- This is definitely not the case as we will illustrate
- This infrastructure supports a range of Telecom business and consumer, retail and wholesale services today and will support more in the future
- Furthermore, these services share the available resource on a fairness basis and with value based pricing applied in the market

Hence in order to address the issue of Non-shaped downstream speed we need to consider the following factors:

- Technical and operational complexity
- Fairness in the allocation of shared resources
- Direct infrastructure costs

Each of these factors will be discussed in turn within the slides which follow

It should also be noted that although Telecom has unconstrained bit rate services in the market today (related to those being requested by TCL), these are being capped in the market for precisely the reasons described in the following slides

# Provision of non-shaped downstream speed (2)

- As the market demand has grown, it has become increasingly difficult to sustain these unconstrained bit stream products in the market

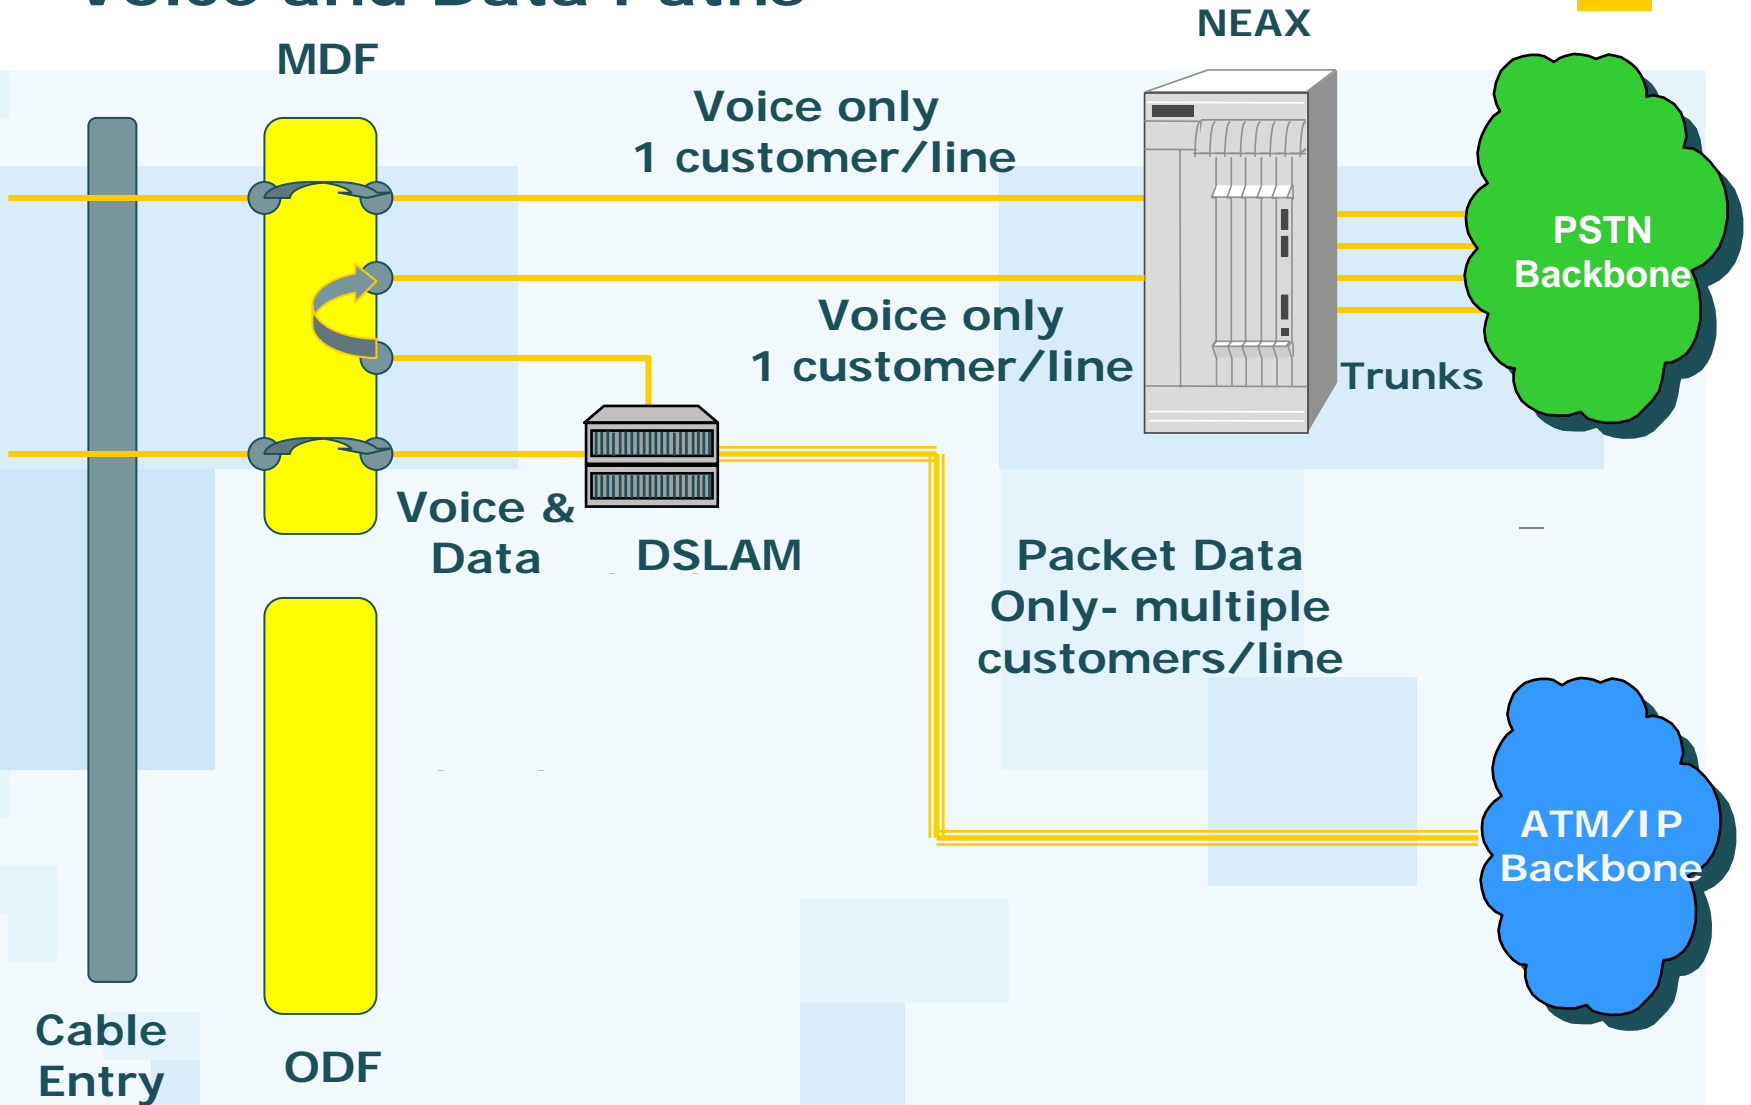
In order to consider the above factors it is essential to understand how shaping and policing is applied to the asymmetrical bit stream services – This occurs in three places:

- On the copper access through rate limiting (more correctly policing)
- In the ATM aggregation
- And in the L2TP origination to facilitate congestion management and fairness


The impact of all of these components in terms of both cost and performance will be discussed in the slides which follow.

# Introduction: The Local Exchange: DSL

## Voice and Data Paths



# Fast Internet only DSLAM DSL Logical Function



# DSLAM Overview (1)

Telecom has the following types of DSLAMs deployed:

- Nokia DSLAMs in exchange sites
- Nokia remote DSLAM shelves in cabinets
- Alcatel DSLAMs in exchange sites (some co-located with the Nokia DSLAMs)
- Alcatel sub-tended DSLAMs in smaller exchange sites
- Conklin DSLAMs sub-tended off exchange based Alcatel DSLAMs

All of these DSLAMs support a range of Telecom data services

- Typically 2Mbps and below today, with some asymmetrical services beyond this
- All future symmetrical services below 10Mbps and asymmetrical services up to about 20Mbps will be supported off this infrastructure
- This includes data services to meet the needs of business, consumer, wholesale and access seeker customers

## DSLAM Overview (2)

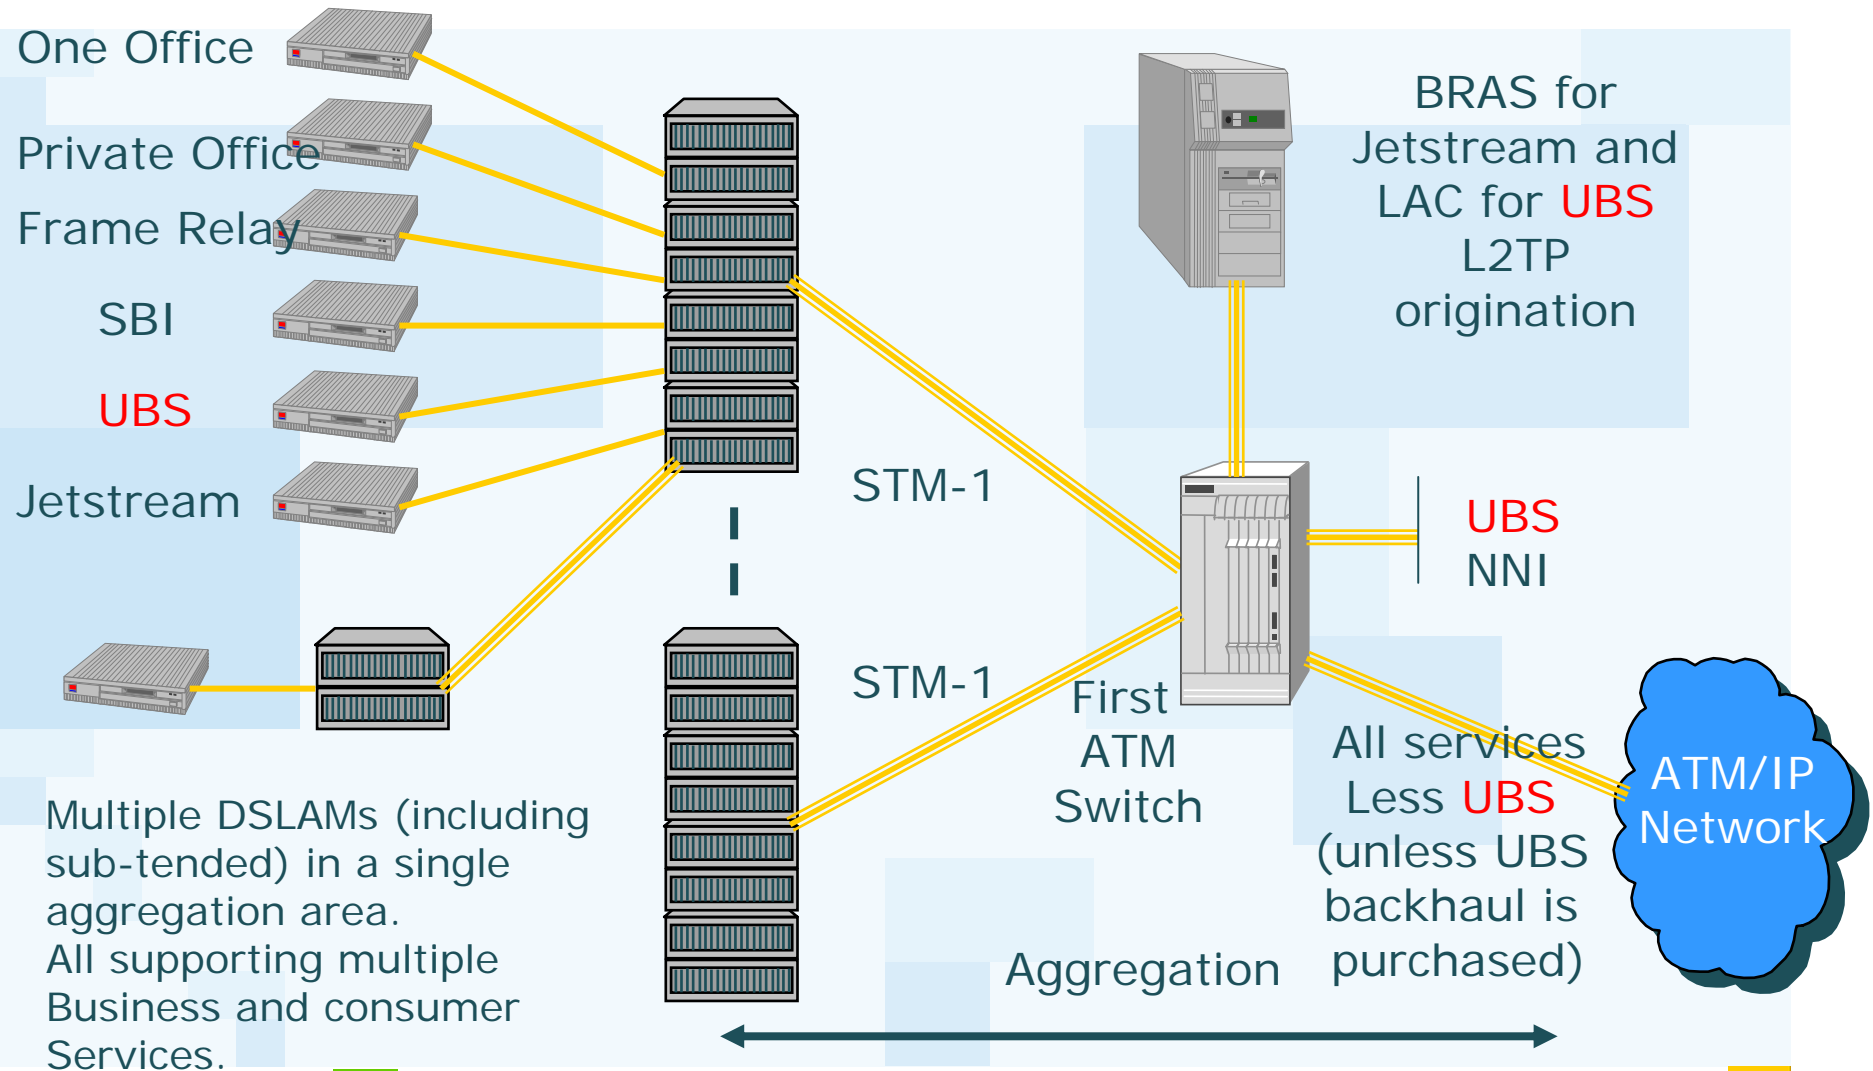
All of the main exchange based DSLAMs have a single STM-1 (155Mbps) backhaul transport link

- This transport link carries ATM over SDH
- This transport link terminates on an ATM aggregation node (the first ATM switch)
- This ATM aggregation node will typically support ATM aggregation from several DSLAMs

Smaller exchange sites and cabinets are served (today) using:

- Sub-tended remote line shelves, with 45 Mbps backhaul (Nokia)
- Sub-tended remote DSLAMs, with 45Mbps backhaul (Alcatel)
- Sub-tended mini DSLAMs, with up to 4xE1 backhaul (Conklin)
- All of the above share the main DSLAM backhaul capacity

# Multi-service Physical Implementation



# Copper Rate Limiting (1)

Copper rate limiting is a form of policing which is applied on each copper access line connected to the DSLAM, in a manner appropriate to the service deployed over that copper line:

- The copper rate limiting is applied at the DSLAM through a service profile
- The rate limiting is independently applied at the PIR of the upstream and downstream bit rates (plus about 20% overhead to accommodate the average packet to cell conversion overheads)

The copper rate limiting achieves the following outcomes:

- It stops customers using more resource from the DSL modem to the first ATM switch, than their access provider has purchased from Telecom
- It reduces the power spectral density for a given service per copper cable pair, within a given cable sheath and hence reduces the impacts of cross-talk in the copper cable network
- It reduces the number of re-synchronisation events by improving the noise margins for any given service

## Copper Rate Limiting (2)

Copper rate limiting can typically be set in 32 kbps increments from 32 kbps up to about 8Mbps for the downstream and 32kbps to 640kbps for the upstream:

- The upstream and downstream values can be set independently
- Some DSLAMs can only be configured in 64kbps increments

Hence it is possible to configure the copper rate limiting to support a service with 8Mbps downstream and 128kbps upstream as requested by TCL

- On the other hand, none of Telecom's lines will synchronise to a downstream speed of more than 7.6Mbps and then only about 20% will achieve this rate
- Hence the actual downstream rate achieved for any given access line will vary considerably, with the majority synchronising between 2 and 4 Mbps (with a good modem and clean in-premise wiring)
- Unfortunately, it is essential to allocate resource relative to the maximum PIR expected for all lines, as there is no way to tell how the variable PIRs will be distributed across a given DSLAM (see later)
  - Especially when we use TCL's definition of Contention Ratio

# Fairness

Under congested conditions, it is essential that all the different services using the DSLAM receive an appropriate amount of resource relative to their defined service characteristics:

- Services with a defined Committed Information Rate (CIR) will receive sufficient resource to deliver on this guarantee for better than 99.9% of the time
  - These are typically high value business grade services with commensurate prices in the market
- Services with a defined PIR, no defined CIR and a defined Contention Ratio (referred to as uncommitted bit rate (UBR) services) need to receive sufficient resource to deliver the defined Contention Ratio during periods of congestion and the full PIR during periods of no congestion (as illustrated previously)
- Furthermore the service performance should degrade in proportion to the PIR under congestion conditions
  - The different PIR services all experience bit rate degradation in proportion to their PIR and Delay Variation and Packet Loss performance is distributed in an equitable manner for all services

# Fairness Implementation

The Committed Information Rate services receive a Virtual Circuit which has characteristics defined to ensure the defined CIR is achieved for better than 99.9% of the time:

- This typically means using CBR, VBR-rt or VBR-nrt ATM shaping in conjunction with the allocation of bandwidth per service approximately equal to the defined CIR

On the other hand, the UBR services are given a Virtual Circuit which has a defined PIR and Uncommitted Bit Rate (UBR+) ATM shaping

- All the UBR VCs are combined in a single Virtual Path from the DSLAM to the BRAS
- The Fairness for these services is achieved in the Broadband Remote Access Server (BRAS/L2TP LAC)
- At the BRAS both the aggregated VP and each of the individual VCs is shaped by a Hierarchical Scheduler
- All retail Jetstream and UBS services in the market today are treated in exactly the same manner from the customer modem through to the output of the BRAS/L2TP LAC and so they are all subject to the same shaping and Fairness criteria

# DSLAM Service Types

## Current services:

- One Office: CIR = PIR, VBR-nrt or VBR-rt, one VC per connection
- UPC: CIR = PIR, CBR or VBR-nrt, one VC per connection
- Private Office: CIR = PIR/2, VBR-nrt, one VC per connection
- Frame Relay: CIR = PIR/x, where x can be 1, 2 or 4, VBR-nrt, up to 16 VCs per connection
- Secure Business Internet: PIR defined, Contention Ratio <50:1, UBR+, single shared VP
- Jet Services: PIR defined, Contention Ratio <50:1, UBR+, single shared VP
- UBS (commercial): PIR defined, Contention Ratio <50:1, UBR+, single shared VP

## TCL proposed new service:

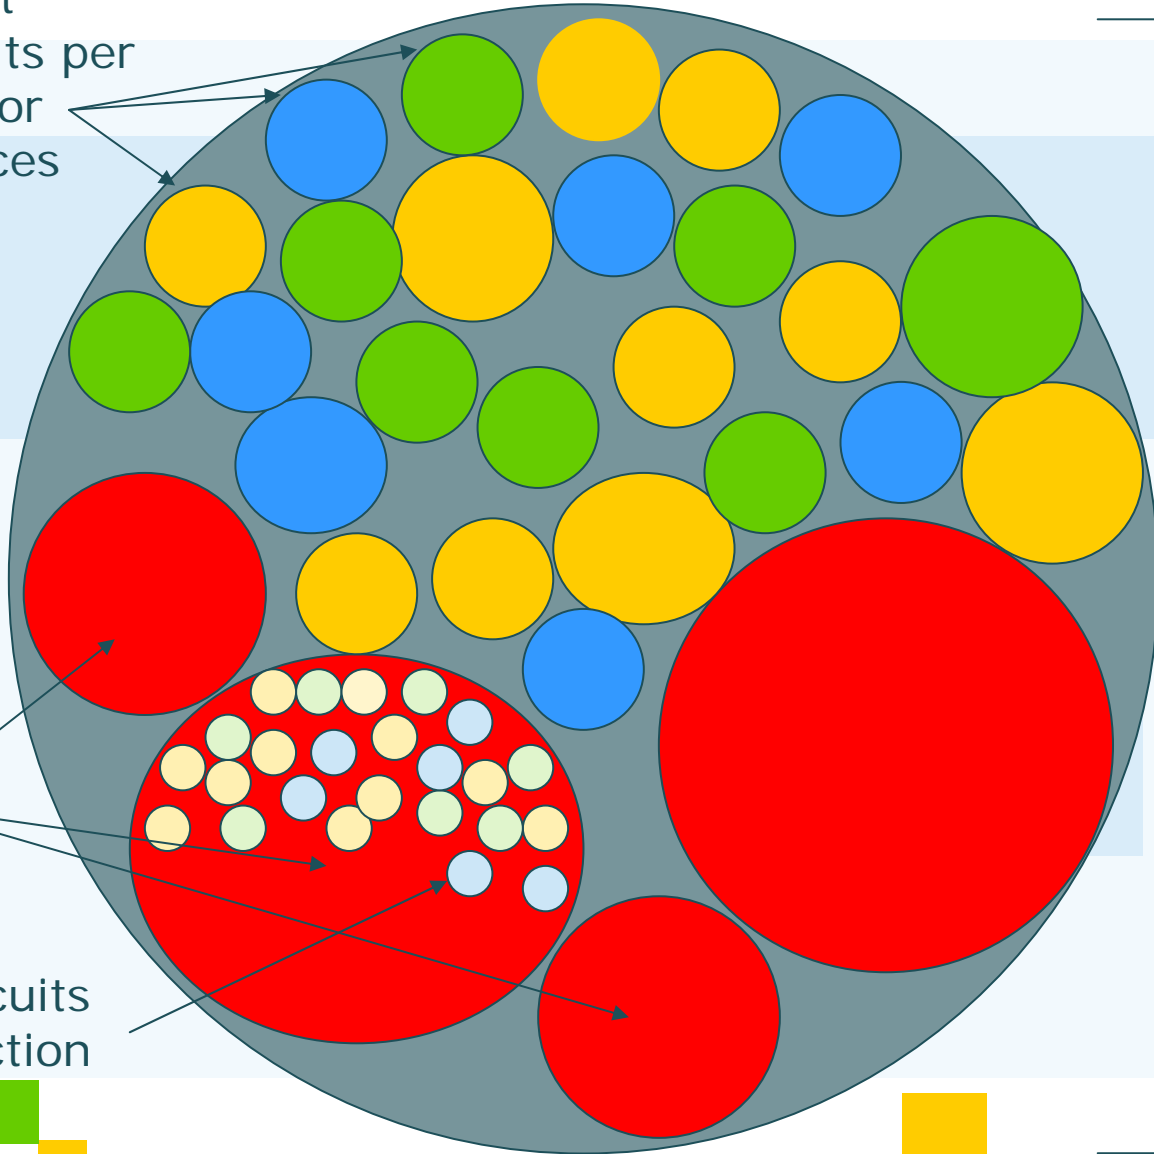
- UBS (TCL): PIR defined and unconstrained, Contention Ratio <20:1, UBR+ (?), separate VP (?)

# DSLAM Backhaul Logical Configuration

Point to Point  
Virtual Circuits per  
Connection for  
all CIR services

Aggregated  
Virtual Paths  
for all UBR  
Services per  
DSLAM

Virtual Circuits  
Per connection



STM-1  
155Mbps

# Fairness and the BRAS (1)

Fairness is best achieved at the BRAS/L2TP LAC as the ATM infrastructure alone is poor at achieving reliable fairness outcomes in a highly contended environment

- Even then there is considerable operational complexity to achieve a satisfactory outcome

At the BRAS/L2TP LAC, the Hierarchical Scheduler uses a number of mechanisms to ensure that under congestion conditions:

- Each PIR service is speed degraded in proportion to its PIR
- All the different PIR services experience a similar delay variation probability distribution and packet loss probability distribution

It uses weighting coefficients associated with each PIR service to achieve this outcome:

- These weighting coefficients need to be configured for each new service deployed at both the VC and VP levels

## Fairness and the BRAS (2)

- Then the interactions between the various services needs to be tested to ensure that the “Fairness” criterion is actually being achieved in practice across all parameters:
  - Speed degradation, delay variation and packet loss
- Even then “Fairness” will vary considerably across different DSLAMs and over time, as customers are connected and disconnected from the DSLAMs
  - Dimensioning rules are used to keep the Fairness outcomes reasonably consistent, but at best, it can only be assumed to apply correctly for 95% of all connections, for 95% of the time

# Fairness and Unconstrained Bit Rate Services

Given the “Fairness” objectives outlined above, it is obviously more difficult to deal with Unconstrained Bit Rate services than those that are constrained

- In order to deal with unconstrained bit rate services, they all need to effectively be defined at the maximum PIR possible
- This is of course wasteful of resource relative to the constrained bit rate services, unless the full PIR is actually used on all connections – which is not possible (as described previously)
- It also leads to greater variation in fairness across DSLAMs over time, as it is difficult to determine the actual PIRs that the unconstrained services actually synchronise to on a particular line:
  - How a given DSLAM will be loaded with low or high PIR traffic depending on what access lines it is connected to
  - And how churn in connections over time will change this mix
- It is these difficulties which have encouraged Telecom to cap its unconstrained bit rate services

# DSL based service relative costs

Capital Cost =  $Y + X$ , where:

- $X$  = the variable cost per connection (eg. Per port used)
- $Y$  = the proportion of the fixed cost of the DSLAM and associated infrastructure (to the first ATM switch) which is attributable to a given service

A typical exchange based DSLAM has the following characteristics:

- Fixed costs with backhaul and BRAS costs included of \$ $Z$  (where  $Z$  is typically 100s of thousands of dollars)
- Supports around 1500 ports, with each port costing \$ $X$  for an asymmetrical port and \$ $2X$  for a symmetrical port (where  $X$  is typically in the low 100s of dollars)
- A usable throughput and backhaul capacity of about 150 Mbps

# DSL Costs: High CIR Services Only Example

One extreme is to load a DSLAM up with 2Mbps symmetrical services that have CIR = PIR

- Then the DSLAM can only support  $150/2 = 75$  such services with its total resource fully utilised
- The cost of each such service is then:
  - Capital Cost =  $Z/75 + 2X$
  - As  $Z$  is in the 100s of thousands of dollars and  $X$  is in the low 100s of dollars, then it is easily seen that in this situation the fixed costs dominate and the total cost of providing such services is high
  - All the resource associated with a 1500 port DSLAM is fully used up by 75 ports

# DSL Costs: UBR Services Only Example

Another extreme is to load up a DSLAM with UBR services only, with a contention ratio of less than 50:1 (all ports active) and for illustrative purposes, a PIR of 2 Mbps:

- The DSLAM can then support services on all 1500 ports and only uses some  $2/50 \times 1500 = 60$  Mbps of capacity
- If the total cost of the DSLAM is allocated to these services, then the cost of each such service is:
  - Capital Cost =  $Z/1500 + X$
  - As  $Z$  is in the 100s of thousands of dollars and  $X$  is in the low 100s of dollars, then it is easily seen that in this situation the proportion of fixed costs becomes comparable to the variable costs and the cost of providing each service is relatively low.
  - In this case 40% of the resource associated with a 1500 port DSLAM is used up by all 1500 ports
  - In practice a DSLAM used like this, would have remote DSLAMs sub-tended to it (up to an additional 2000 ports), or it would include a mix of CIR and UBR services to ensure that all the resource is used efficiently

# Active vs Inactive ports

In today's environment, even with "always on" connectivity offered by broadband service providers, few customers have their broadband connection running 24 hours of the day, 365 days of the year

- Today, we find that there is a ratio of active to inactive users of about 1:3
- We expect this ratio to reduce over time, but it is unlikely to ever get much below 1:1

Given this situation, we can further refine the definition of Contention Ratio to relate to active users only

- This ensures even better utilisation of the DSL resource and enables further reduction in cost per connection
- In the previous example, if this definition of contention ratio is used relative to the peak busy hour of usage, then with a ratio of active to inactive users of 1:4, the resource used by 1500 ports could be as little as 10% of the total, allowing even lower cost per connection (over 10,000 connections per primary DSLAM)

# DSL Costs: Mixed Service Example

Now consider the optimum loading of the DSLAM with a mix of services supported:

- What is the optimum loading of a 1500 port DSLAM with a mix of 2Mbps symmetrical CIR=PIR services and Uncommitted bit rate services with PIR = 2Mbps, with contention ratio of >50:1 and 1 active user for each inactive user during the peak busy hour?
  - This is a resource optimisation problem with the following outcome:
    - No. of PIR=2Mbps UBR users = 1440
    - No. of CIR=PIR=2Mbps users = 60
  - Under these conditions each of the services would use the optimum amount of DSL resource and hence be optimal from a capital cost perspective

Of course, the real situation encountered in the field is much more complicated than that illustrated in these few idealised examples:

- There are many more PIR and CIR variants for both the CIR services and the UBR services and the DSLAM may have other DSLAMs subtended from it, which in turn may have a complex mix of services associated with them.

# Provision of non-shaped downstream speed: Summary (1)

With respect to technical and operational complexity:

- Copper rate limiting of the downstream can be performed independently of the upstream copper rate limiting and so unconstrained bit rate services can be provided - however doing so adds the following complexity
  - It requires the implementation and management of a new service profile on the DSLAM
  - It increases the power spectral density per copper pair within a given cable sheath and hence increases the impacts of cross-talk in the copper cable network
  - It increases the probability of re-synchronisation events by reducing the noise margins for any given connection
- Establishing fairness for all services in the presence of non-shaped downstream speed services is complicated
  - The setting of weighting coefficients in the hierarchical scheduler must be established at the maximum PIR, but in practice, the actual PIR per line can vary considerably from this maximum
  - This leads to the interactions between the various services being difficult to predict, even with normal test practices, as the interactions will change dynamically over time on a given DSLAM, with customer churn
    - Resulting in customer service issues associated with both the unconstrained bit stream services and the other services with which they interact

# Provision of non-shaped downstream speed: Summary (2)

With respect to fairness in the allocation of shared resources:

- The non-shaped downstream speed leads to greater variation in fairness degradation of performance under conditions of congestion
- Operationally, it is difficult to determine the actual PIRs that the unconstrained services actually synchronise to on a particular line and hence to determine how a given DSLAM will be loaded with low or high PIR traffic depending on what access lines it is connected to and how churn in connections over time will change this mix

With respect to direct infrastructure costs:

- The setting of the PIR to 8Mbps for the downstream for all unconstrained bit rate connections will result in the allocation of resources relative to the PIR at 8Mbps (especially under the TCL definition of Contention Ratio)
- With a contention ratio of 20:1 (with all ports active) as defined by TCL, only 375 such connections could be configured per exchange based DSLAM
  - As compared to 1500 connections at 2Mbps PIR with the same contention ratio
  - **Hence the costs of these two options will vary by a factor of around 4:1 in favour of the constrained downstream speed service**

# Difficulties in providing multiple downstream/upstream speed configurations

*Commission's question:*

5. What are these technical and operational practicability issues associated with providing multiple downstream/upstream speed considerations?

# Difficulties in providing multiple downstream/upstream speed configurations

Telecom supports a wide range of downstream/upstream speed configurations today

- Supporting more speed configurations only, can be arranged, but there will be a cost
  - Setting up new DSLAM profiles
  - Setting up new Hierarchical Scheduler profiles
  - Testing to ensure Fairness is being achieved
  - Possible requirement for additional active test probes
- Supporting more downstream/upstream speed configurations with significantly different quality characteristics adds additional complexity
  - The configuration of new Virtual Paths per DSLAM, in addition to the above

# Difficulties in providing multiple downstream/upstream speed configurations

Furthermore, any access seeker can create many downstream/upstream speed configurations from the existing Telecom offered service set by

- Selecting a higher speed service
- Shaping it in the desired upstream and downstream speed combination themselves
  - The configurations are of course constrained within the bounds of the access limits (eg. Upstream limited to 128 Kbps maximum)
- In this way, any access seeker can differentiate their service from those offered by Telecom
  - And can control their costs associated with doing this within their own business

# Policing

*Commission's questions:*

6. Is this necessary?
7. What is the effect on latency?

Question 7 has been substantially answered already.

# Policing

Policing is only applied in the form of copper rate limiting on a line by line basis at the DSLAM to:

- Stop customers using more resource (from the customer modem right through to the first ATM switch) than they have purchased from their Service Provider
- Reduce the power spectral density per copper pair within a given cable sheath and hence reduce the impacts of cross-talk in the copper cable network
- Reduce the number of re-synchronisation events by improving the noise margins for any given service

Congestion management between services sharing the same resource is achieved through the use of shaping in the Hierarchical Scheduler within the BRAS/L2TP LAC, which ensures:

- Optimum performance of TCP sessions
- Optimal management of Contention Ratio, Delay Variation and Packet Loss Ratio
- Optimum allocation of resources between the different services to ensure fairness is achieved under congestion conditions

# Interleaving

## *Commission's questions:*

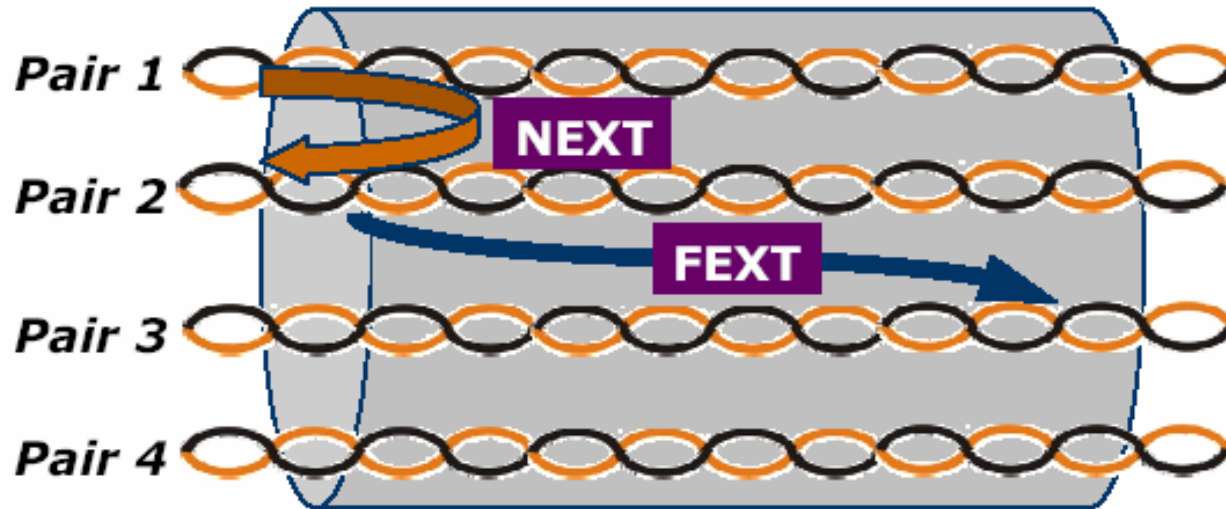
8. What are the technical reasons why interleaving is unable to be turned on or off by Telecom at the DSLAM?
9. What impact would interleaving off have on service levels provided by Telecom?
10. Can interleaving be turned off solely for TelstraClear at the wholesale level without affecting the other wholesale/retail customers?

# Interleaving is present for a purpose

The DSL specifications include interleaving to enable carriers to better manage the noise and interference present on copper cables today

- This includes:
  - Noise
  - Crosstalk from within the cable sheath
  - Interference from external sources
- All these sources of interference and noise degrade the performance of DSL when operating on a given cable pair within a defined cable sheath
- The noise and interference manifest themselves as packet loss as the signal traverses the copper cable between the customer and the DSLAM
  - Furthermore high packet loss can force the DSL modem to resynchronise with the DSLAM, which can cause traffic outages of 10-20 seconds
- Our measurements show that on many copper cable pairs within the Telecom network, the positive impact of interleaving on packet loss ratio and resynchronisation events is considerable (increases reach by up to 50% and reduces resynchronisation events by around 20%)

# The Pain Called Crosstalk



- > Past technological advances
  - Higher bandwidth, better spectral efficiency
- > Future technological advances
  - Mitigation of NEXT and FEXT crosstalk

# Crosstalk in Cables

Crosstalk is always present in multi-pair copper cables, and increases as the spectral power density of the signals carried on individual pairs in that cable increase in frequency

- Hence for voice signals only, which were constrained to around 3Khz the problem was minimal
- Over the last 20 years we have deployed a wide range of line codes in our copper cables, all of which have substantial power densities up into the MHz range
  - ADSL is one of the more recent of these and has some of the greatest interference protection mechanisms, of which interleaving is one
  - However, ADSL is still susceptible to interference from other line codes operating in the same cable sheath, such as HDB3 (E1 MUX), 2B1Q (ISDN BRA), SHDSL, etc and in fact with itself
  - Telecom is in the process of removing as many sources of interference from the copper cable network as possible, but this will take many years to complete and even then we will still have to deal with SHDSL and its own self interference
- In order to properly manage noise and interference to the benefit of all users of the copper cable network, Telecom has implemented a policy of having interleaving always turned on

# Can Interleaving be turned off?

Interleaving can be turned off at the DSLAM on a line by line basis

- Today this facility is not included within our automated management capability as Telecom policy is to have interleaving on for all asymmetrical bit rate services
- Our objective is to provide as much performance as possible to as many users of the copper cable network as possible
- At some time in the future, this policy may change as newer versions of ADSL (eg. ADSL2+) which have improved management features included within the technology are introduced
  - To better manage the impact of either turning interleaving off or changing its characteristics
- Turning interleaving off or changing its characteristics with existing ADSL will result in highly unpredictable performance, especially for longer cable reach situations
- Hence even if Telecom could provide interleaving as optional for individual lines, Telecom has no means of predicting the resulting performance, and so it would be impossible to offer any assurance around packet loss ratio or resynchronisation events under these conditions (Packet Loss Ratio and Availability being defined TCL requirements)

# Experience with Interleaving Elsewhere

The provision of interleaving on or off is variable across carriers worldwide

- The general rule is that it is either applied universally on or universally off, options to choose by line are rare
- Those carriers that choose the interleaving off option, generally ensure that the services offered under these conditions have absolutely no service commitments associated with them – the services are purely “best effort”
- Those carriers that retain the interleaving on option, generally have at least some ADSL based services that do have some form of service commitments and/or are striving to achieve the best possible reach from their copper network
  - Both are true of Telecom in New Zealand
    - We are delivering services with long reach
    - Some services using ADSL, such as our Private Office services are sold with some service commitments
- Some carriers such as BT and Belgacom started out with interleaving off and are now in the process of transitioning to interleaving on
  - To achieve more predictable performance, particularly on longer cable runs

# The Bottom Line on Interleaving

Telecom policy today is to implement all asymmetrical bit stream services with interleaving on

- Interleaving can be turned on or off on a per line basis
  - Telecom has no automated capability for provisioning services with interleaving as an option so there would be considerable cost to implement this option specifically for TCL
- The performance achieved on a per line basis is currently highly unpredictable with interleaving switched off
  - It is known that many lines would not meet TCL's defined Packet Loss and availability specifications
  - Telecom could not offer any performance commitments under these conditions
  - Telecom would have increased difficulty in managing the performance of all of its DSL based services due to the increased impact of crosstalk
- Telecom can meet the latency specification requested by TCL with interleaving turned on
- Telecom may change its position on interleaving at some time in the future (due to changes in technology capability) and this would be an appropriate time to have further discussion on this topic

# Internal Service Levels

## *Commission Questions:*

11. TelstraClear's application requests certain service parameters for the UBS service, for example latency, jitter, contention ratio and their effect on the overall service level independently or as a whole.

What levels of service are offered internally to Telecom for Jetstream and UBS on these parameters, and are these translatable to wholesale services?

12. How are these measured?

These questions have been substantially answered earlier, however some additional points need to be made.

# Internal Service levels

- Any discussion of the levels of service offered internally to Telecom is artificial. Telecom operates as an integrated business. It does not have a retail division that is supplied by a wholesale division. Telecom's wholesale division is a channel to market that delivers services to Telecom's wholesale customers; it does not deliver services to Telecom's retail division
- Telecom does not currently offer service levels around latency, delay variation, packet loss and the contention ratio in relation to its JetStream services. This is because JetStream services have been pitched to provide a functional best efforts, internet grade service at an acceptable price for the mass market

## Internal Service levels *cont'd*

- For engineering purposes, Telecom does have a set of design parameters around latency, delay variation, packet loss and the contention ratio for its JetStream services. [Table 2 of Telecom's cross-submission]
- The design parameters for Telecom's current JetStream services have already been "translated" into Telecom's commercial UBS service
- Telecom's commercial UBS service is configured and managed so that to the greatest extent possible, it is equivalent to Telecom's JetStream services. There are minor technical differences between Telecom's current UBS and JetStream services. However, these differences are virtually immaterial to the end user experience
- Any requirements for additional reporting above that generated for Jetstream and commercial UBS would be at an additional cost