TECHNICAL TUTORIAL

Subject: **SDH**

Date: October, 2003

Prepared by: John Rumsey

SDH

Synchronous Digital Hierarchy

1. Introduction

This document provides a technical discussion of SDH.

It explains the differences between the Plesiochronous Digital Hierarchy and the Synchronous Digital Hierarchy.

Before we start it is important that you understand the difference between Synchronous, Plesiochronous, Isochronous and Asynchronous.

These are all different ways of synchronising a data stream between the transmitter of the data stream and the receiver.

They all refer to how a data stream is clocked.

- **Synchronous (Synchronised)**

All of the clocks are synchronised to a master reference clock. They may be out of phase with each other but they will run at exactly the same frequency.

- **Plesiochronous (Almost Synchronised)**

All of the clocks run at the same frequency to a defined precision. These clocks are not synchronised to each other so the data streams will run at slightly different rates.

- **Isochronous (Synchronised)**

An Isochronous data stream has the timing information embedded in it (eg. a G.704 stream). These data streams can be carried over Synchronous or Plesiochronous networks.

- **Asynchronous (Not Synchronised)**

The clocks are not synchronised. The transmitter and the receiver have independent clocks that have no relationship with each other.

2. The Plesiochronous Digital Hierarchy (PDH)

In a PDH network you have different levels of Multiplexers.

Figure 1 shows three levels of multiplexing:-

- 2Mbit/s to 8Mbit/s
- 8Mbit/s to 34Mbit/s
- 34Mbit/s to 140Mbit/s

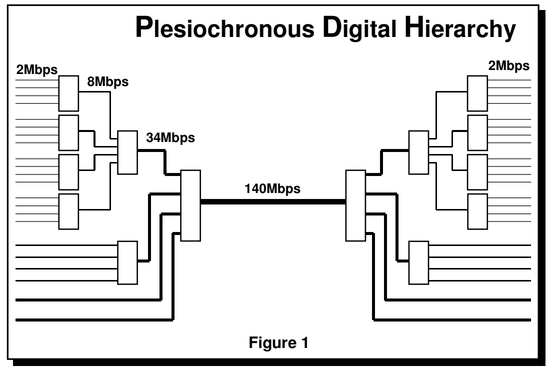

Figure 1

So to carry a 2Mbit/s data stream across the 140Mbit/s trunk requires it to be multiplexed up through the higher order multiplexers into the 140Mbit/s trunk and then to be multiplexed down through the lower order multiplexers.

Because Plesiochronous is not quite Synchronous each of the multiplexers need a little bit of overhead on their high speed trunks to cater for the slight differences in data rates of the streams on the low speed ports.

Some of the data from low speed ports (that are running too fast) can be carried in the trunk overhead, and this can happen at all multiplexing levels. This is known as Justification or Bit Stuffing.

2.1 PDH Multiplexing Hierarchy

Figure 2 shows that there are two totally different hierarchies, one for the US and Japan and another for the rest of the world.

The other thing to notice is that the different multiplexing levels are not multiples of each other.

For example CEPT2 supports 120 Calls but it requires more than 4 times the bandwidth of CEPT1 to achieve this.

This is because PDH is not exactly synchronous and each multiplexing level requires extra bandwidth to perform Bit Stuffing.

So the Plesiochronous Hierarchy requires "Bit Stuffing", at all levels, to cater for the differences in clocks.

PDH Multiplexing Levels

	United States & Japan			Europe & Australia		
Multiplexing Level	*Name*	*# calls*	*Rate (Mbps)*	*Name*	*# calls*	*Rate (Mbps)*
1	DS1	24	1.544	CEPT1	30	2.048
2	DS2	96	6.312	CEPT2	120	8.448
3	DS3	672	44.736	CEPT3	480	34.368
4	DS4	4032	274.176	CEPT4	1920	139.264

Figure 2

This makes it particularly difficult to locate a particular 2Mbit/s stream in the 140Mbit/s trunk unless you fully de-multiplex the 140Mbit/s stream all the way down to 2Mbit/s.

2.2 Drop & Insert a 2Mbit/s stream

To drop & insert a 2Mbit/s stream from a 140Mbit/s trunk you need to break the 140Mbit/s trunk and insert a couple of "34Mbit/s to 140Mbit/s" multiplexers. You can then isolate the appropriate 34Mbit/s stream and multiplex the other 34Mbit/s streams back into the 140Mbit/s trunk.

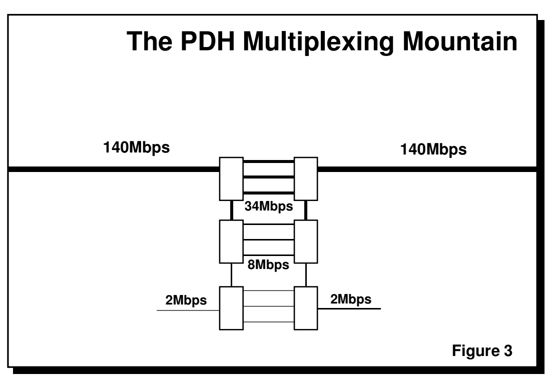

Figure 3

Then you de-multiplex the 34Mbit/s stream, isolate the appropriate 8Mbit/s Stream and multiplex the other 8Mbit/s streams through the higher layer multiplexer, into the 140Mbit/s trunk.

Then you de-multiplex the 8Mbit/s stream, isolate the 2Mbit/s Stream that you have been looking for and multiplex the other 2Mbit/s streams up through the higher layer multiplexers into the 140Mbit/s trunk.

Figure 3 shows the "multiplexer mountain" required to drop & insert the 2Mbit/s stream.

2.3 The limitations of PDH:-

- ***PDH is not very flexible***

As previously explained, it is not easy to identify individual channels in a higher order bit stream.
You must multiplex the high rate channel down through all multiplexing levels to find a particular lower speed channel. This requires an expensive and complex "multiplexer mountain".

- ***Lack of Performance***

It is not easy to provide good performance if you can't monitor the performance in the first place.
For PDH there is no international standard for performance monitoring and no agreed management channels. There are some spare overhead bits that are being used for management but they have limited bandwidth and are hard to locate in a 140 Meg stream without de-multiplexing.

- ***Lack of standards***

Not only does PDH have two totally different multiplexing hierarchies but it is quite weak on standards. For example there are no standards for data rates above 140Mbit/s and no standards for the line side of a "Line Transmission Terminal".

3. The Synchronous Digital Hierarchy (SDH)

SDH, like PDH is based on a hierarchy of continuously repeating, fixed length frames designed to carry isochronous traffic channels.

SDH was specifically designed in such a way that it would preserve a smooth interworking with existing PDH networks.

The developers of SDH also addressed the weaknesses of PDH. They recognised that it was necessary to adopt not only a Synchronous frame structure but one that also preserves the byte boundaries in the various traffic bit streams.

Because SDH is synchronous it allows single stage multiplexing and de-multiplexing. This eliminates hardware complexity. You don't need multiplexer mountains.

3.1 SDH multiplexing levels

Figure 4 shows the SDH multiplexing levels.
The US and Japan use SONET while most of the rest of the world use SDH.

Apart from using some different terminology, there is very little difference between SONET and SDH.

SDH Multiplexing Hierarchy

Data Rate (Mbps)	SONET (USA)	SDH (Europe)
51.84	STS-1, OC-1	*(not defined)*
155.52	STS-3, OC-3	STM-1
466.56	STS-9, OC-9	STM-3
622.08	STS-12, OC-12	STM-4
933.12	STS-18, OC-18	STM-6
1244.16	STS-24, OC-24	STM-8
1866.24	STS-36, OC-36	STM-12
2488.32	STS-48, OC-48	STM-16
9953.28	STS-192, OC-192	STM-64

STS = Synchronous Transport Signal

OC = Optical carrier

STM = Synchronous Transport Module

Figure 4

You can see that the data rates are the same except SDH doesn't specify a 51 Meg rate.

STM-1 forms the basis of the SDH frame structure. For example an STM-4 is a frame consisting of 4 x STM-1s.

In Sonet, the STS levels refer to the speed of the bit stream. When these bits are converted to a train of optical pulses in a fibre, they are called an Optical Carrier (OC).

You may also see "OC-3c" referred to. This is simply the same bit rate as OC-3, but interpreted as one channel instead of 3 multiplexed OC-1s. The "c" stands for "Concatenated".

3.2 SDH Network Elements

Figure 5 shows the elements that make up an SDH network.

- ***Path Terminating Element***

These are the end point devices where the lower speed channels enter and leave the SDH Network. These are known as “Path Level” devices.

- ***Digital Cross Connect***

These devices can x-connect at the STM level down to individual E1 streams. So an E1 stream on one STM trunk could be x-connected to another STM trunk.

- ***Regenerator***

Is a device that regenerates the signal.

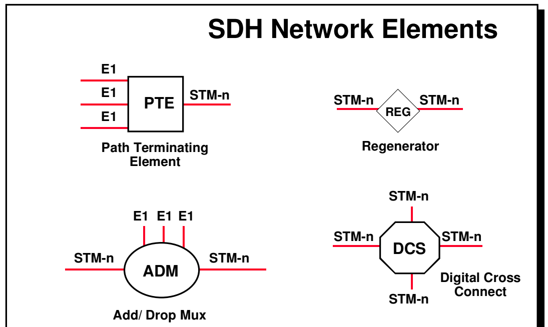

Figure 5

- ***Add/Drop Multiplexer***

The Add/Drop mux has the ability to breakout and insert low speed channels into an STM stream.

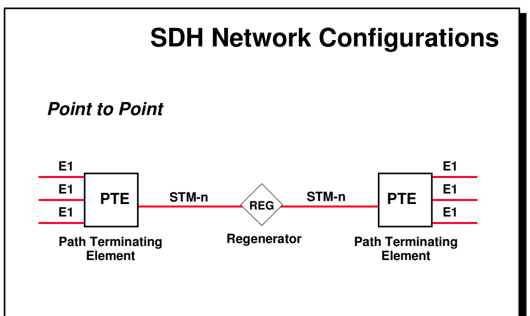

Figure 6

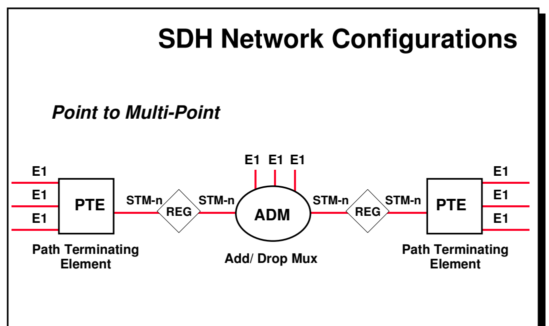

Figure 7

3.3 SDH Network Configurations

The simplest network configuration is a Point to Point network as shown in figure 6. This involves two terminal multiplexers linked by fibre with of without a regenerator in the link.

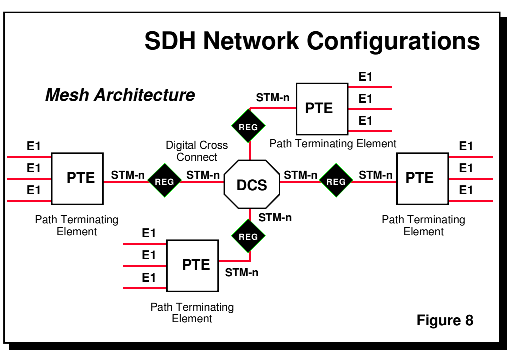

Figure 8

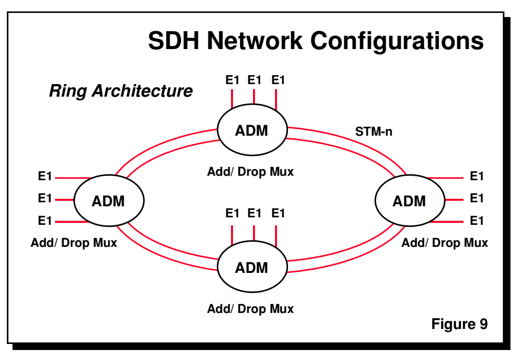

Figure 9

If we include an Add/Drop mux we can now have a point to multi-point configuration as shown in figure 7.

A meshed network architecture (as shown in figure 8) uses a “Digital Cross Connect” to concentrate traffic at a central site and allow easy re-provisioning of the circuits.

The most popular network configuration is the Ring Architecture shown in figure 9. Here we have four Add/Drop muxes interconnected by 2 fibre rings. The main advantage of this architecture is its survivability. If a fibre is cut or an Add/Drop mux dies the multiplexers have the intelligence to heal the network.

3.4 The SDH Frame

The basis of SDH is the STM-1 Frame as shown in figure 10.

The STM-1 frame runs at 155.52Mbit/s, and is 125uS long. This means that you get 8,000 STM-1 frames per second.

8,000 frames a second is a very common rate in telecommunications networks for example G.704 operates at 8,000 frames a second.

This means that each Byte in the frame is equal to a 64kbit/s channel.

The Frame is made up of a "Section Overhead" field and a "Payload" field.

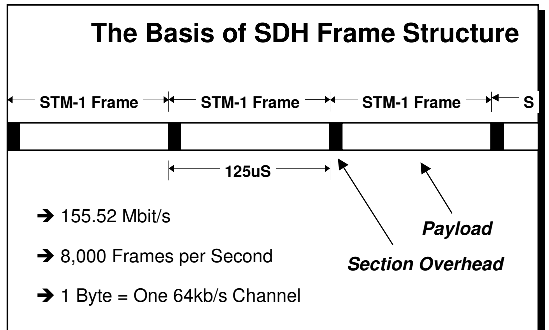

Figure 10

STM-1 Frames are usually represented as 9 Rows by 270 Columns for a total of 2430 Bytes as shown in figure 11. The bytes are transmitted from Left to Right, Top to Bottom.

The first 9 Columns are the section overhead and the other 261 columns are used to carry the payload.

The Section Overhead has three parts:-

* Regenerator Section Overhead
* Pointers
* Multiplex Section Overhead

In SDH the actual user data is carried in "Virtual Containers". The Virtual Containers have a Path Overhead field and they come in a number of different sizes.

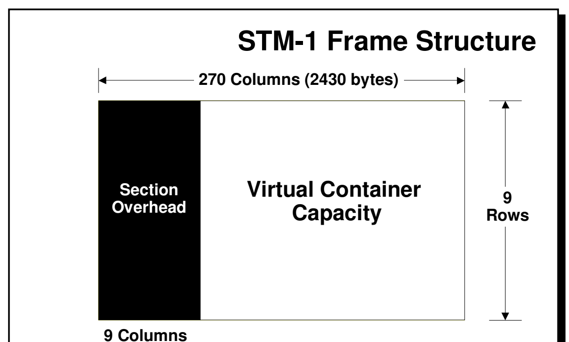

Figure 11

We will look in detail at Virtual Containers later in the tutorial.

But first we will have a look at the SDH Overhead.

3.5 SDH Overhead

Figure 12 shows how the SDH overhead and transport functions are divided into the following layers:-

- ***Regenerator Section***

The Regenerator Section Overhead contains information required for the elements located at both ends of a section. This might be between two Regenerators or Line Termination Equipment.

- ***Multiplex Section***

The multiplex section overhead contains information required between the multiplex section termination equipment at each end of the Multiplex Section. Basically that means between consecutive network elements excluding Regenerators.

- ***Path***

The Path overhead is created and terminated by the Path Terminating Equipment at either end of the link. It is transported in the Virtual Container with the user data.

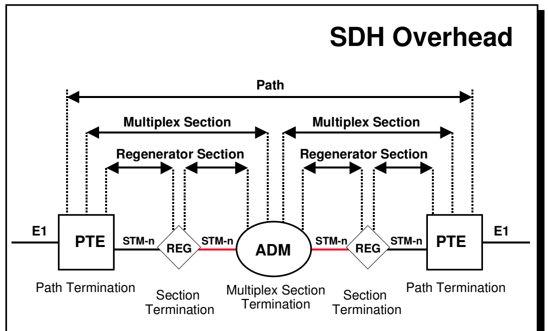

Figure 12

In SONET terms these are known as "Path", "Line" and "Section". And sometimes these terms are incorrectly used for SDH.

Figure 13 shows the structure of the overhead bytes.

The first 3 rows are Regenerator Section Overhead:-

Section Overhead

Regenerator Section Overhead	A1	A1	A1	A2	A2	A2	J0		
	B1			E1			F1		
	D1			D2			D3		
Pointers	H1	H1	H1	H2	H2	H2	H3	H3	H3
Multiplex Section Overhead	B2	B2	B2	K1			K2		
	D4			D5			D6		
	D7			D8			D9		
	D10			D11			D12		
	S1						E2		

Figure 13

A1 and A2 are framing bytes and indicate the start of the STM-1 frame.

J0 contains a trace message that is continually transmitted between Regenerator Sections so that they know they are still connected.

B1 provides parity checking. Calculated over all bytes of the previous STM-1 frame.

E1 can be used for voice communications over the "regenerator Section"

F1 is set aside for the users purposes.

D1, D2 and D3 form a 192kbit/s message channel for Operations, Administration and Maintenance. Eg. Alarms, control, monitoring.

The last five rows of the Section Overhead are used for the Multiplex Section:-

B2 24 bit Parity Check

K1 and K2 Automatic Protection Switching. This is used to provide automatic switching to a backup facility in the event of a failure.

D4 to D12 form a 576kbit/s message channel for Operations, Administration and Maintenance. Eg. Remote provisioning, monitoring, alarms etc.

S1 is the synchronous status message byte used to carry synchronisation messages ie: "I am primary reference clock" or "Do not use me as clock reference".

M1 Remote error indication

E2 A 64kbit/s voice channel between "Multiplex Sections".

Path Overhead
The Path overhead is part of the Virtual Containers. Figure 14 shows the overhead associated with high order and low order Virtual Containers.

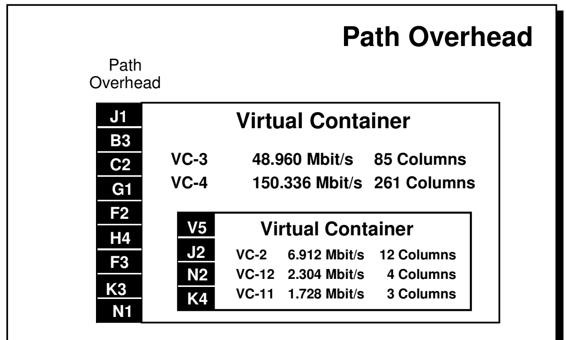

Figure 14

This overhead is transported between Path Terminating Devices:-

High Order Virtual Containers
J1 This byte repetitively transmits a user programmable data string so the receiving path device can see that it is still connected to the intended transmitting path device.

B3 Parity bits

C2 Specifies the type of mapping in the Virtual Container. For example it tells you if the High Order Virtual Container contains Lower Order Virtual Containers or ATM CELLs etc.

G1 Conveys performance of the path

F2 Communications between path elements

H4 Multiframe indicator

Low Order Virtual Containers
There is only a single byte of overhead in a low order Virtual Container. This byte is used for the following purposes over a 4 byte multiframe.

V5 Error Checking, Path Status and Signal label (async, byte sync, bit sync)

J2 Access path identifier (so the receiver knows he is still connected to the intended transmitter)

N2 Connection Monitoring

K4 Automatic Protection Switching. This is used to provide automatic switching to a backup facility in the event of a failure.

3.6 Virtual Containers
Figure 15 shows the 5 Virtual Container sizes and the services they are intended to carry.

It is a very rigid structure and is not very bandwidth efficient.

For example to carry an E3 (CEPT3) service requires more than 30% overhead.
Even the best case an E1 requires more than 10% overhead.
And there is no provision to carry N x 64kbit/s, E2 (8Mbit/s) or 10 meg or 100 Meg.

Virtual Containers

Virtual Container	Bandwidth	For carriage of	
VC-4	150.336 Mbit/s	CEPT4	(139.264 Mbit/s)
VC-3	48.960 Mbit/s	DS3	(44.736 Mbit/s)
		CEPT3	(34.368 Mbit/s)
VC-2	6.912 Mbit/s	DS2	(6.312 Mbit/s)
VC-12	2.304 Mbit/s	CEPT1	(2.048 Mbit/s)
VC-11	1.728 Mbit/s	DS1	(1.544 Mbit/s)

Figure 15

The Virtual Containers provide a permanent nailed up path across the network and there is no Dynamic Bandwidth Allocation.

3.7 SDH Multiplexing Structure

Figure 16 shows how to combine multiple services onto the STM trunk.

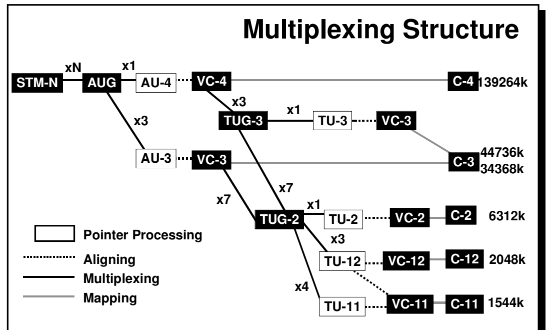

Figure 16

If we look at the example of a E1 (2Mbit/s) service.

1. The E1 frame is placed into a C-12 Container
2. A Path overhead is added and it becomes a VC-2 Virtual Container.
3. Multiple VC-12s are assigned Pointers and become a TUG-2 Tributary Unit Group. The pointers indicate the location of the first byte of each of the Virtual Containers.
4. Seven of these TUGs can be Mapped into a VC-3 Virtual Container
5. Multiple VC-3 Virtual Containers will be assigned Pointers and placed into an AUG Administration User Group
6. And the AUG will be placed in the STM Frame.

The Pointers are used to locate individual 2 meg streams in the STM Frame.

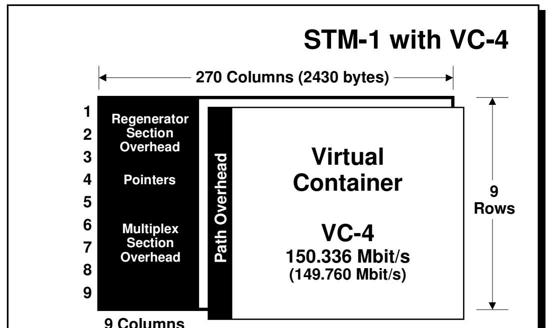

Figure 17

3.8 Virtual Container Examples

Figure 17 shows how a VC-4 Virtual Container fits into the STM-1 Frame.

The VC-4 fits perfectly in the STM-1 Frame and a Pointer indicates the location of the first byte of the VC-4.

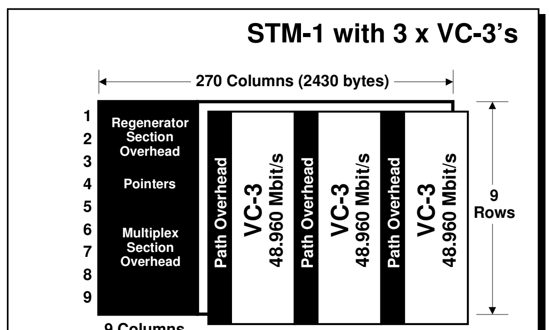

Figure 18

If the VC-4 is not properly synchronised with the STM-1 frame it can slip position in the Frame.

There is actually a Byte of bandwidth in the Pointer Section that can be used by the VC-4 if it is running faster than the STM-1 Frame.

Figure 18 shows how 3 x VC-3s fit into the STM-1 Frame.

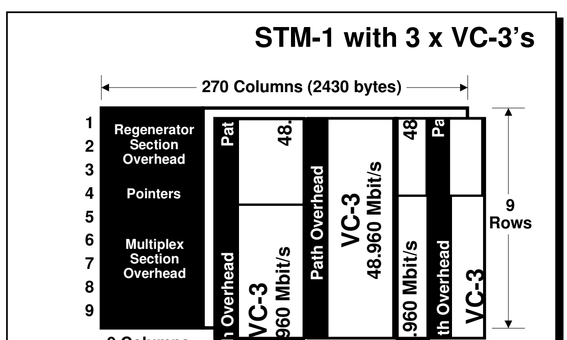

Figure 19

Although it is not shown in this diagram the VC-3 Virtual Containers do not take up all the bandwidth of the STM-1. They are allocated a certain fixed space and they can move around in it. There is extra bandwidth if the VC-3s are running from faster clocks than the STM-1 Frame.

The pointer indicates the first byte of the Virtual Containers.

Figure 19 shows how the VC-3 Virtual Containers may start in different places within their assigned area in the STM-1 Frame.

When a device places a Virtual Circuit into the STM-1 Frame it also sets the Pointer Value so the receiving device can locate the beginning of the Virtual Circuit.

Figure 20 shows how Lower Order Virtual Containers are Carried in the STM-1 Frame:-

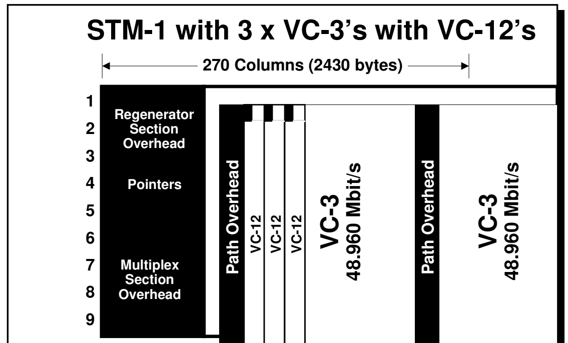

Figure 20

To hold "Lower Order" Virtual Containers we need to use a TUG (Tributary Unit Group). Each TUG has a fixed place in the Virtual Circuit and it has a number of Pointers in fixed positions in the TUG.

The TUG shown in figure 20 has 3 pointers for VC-12 Virtual Containers. The VC-12 is allocated a fixed space that is a little larger that it requires so it has room to move around.

The Pointer shows the location of the first byte of the VC-12.

So when the VC-12 arrives at it destination the receiving device will locate the VC-12 by looking at the pointers.

3.9 Bandwidth Efficiency

Figure 21 shows a number of common data rates and the corresponding Virtual Containers that would need to be used to carry these data streams.

Bandwidth Efficiency

N x 64 Kbit/s	→	VC-12 (2.3 Mbit/s)	
2.048 Mbit/s	→	VC-12 (2.3 Mbit/s)	☺
8.192 Mbit/s	→	VC-3 (48.9 Mbit/s)	
10 Mbit/s	→	VC-3 (48.9 Mbit/s)	
34 Mbit/s	→	VC-3 (48.9 Mbit/s)	
45 Mbit/s	→	VC-3 (48.9 Mbit/s)	☺
100 Mbit/s	→	VC-4 (150 Mbit/s)	

Figure 21

4. Conclusion

SDH has addressed the weaknesses of PDH. It transmits data in Virtual Containers and uses pointers to locate a low speed channel in a high speed trunk.

Carriers like SDH because it provides:-

- A Robust Ring Architecture with Self Healing Capabilities
- Good Provisioning and Management Attributes
- Strong International Standards