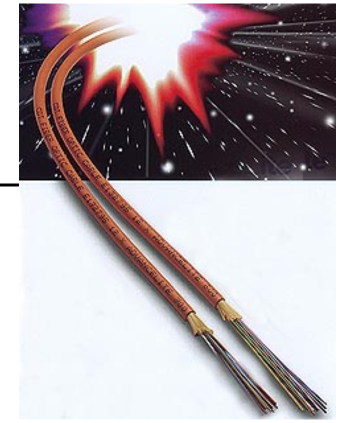

Fundamentals of SONET/SDH

P. Michael Henderson
mike@michael-henderson.us

Agenda

- Why SONET/SDH?
- Historical background
- Basic structure of the SONET/SDH frame
- Transport overhead
- Payload pointer processing
- Payload overhead
- Mapping of SONET payloads
  - Virtual tributaries
  - Handling ATM, POS, and GFP
- Automatic protection switching
- Summary

This presentation is based on a white paper available at http://www.michael-henderson.us/Papers.htm titled "Fundamentals of SONET/SDH."

P. Michael Henderson, mike@michael-henderson.us

What is SONET/SDH

- SONET – Synchronous Optical NETwork (ANSI).
  - ANSI started work on SONET in1985.
- SDH – Synchronous Digital Hierarchy (ITU).
  - The ITU began work in 1986 to achieve the same goal.
- SONET and SDH define a set of physical layer standards for communications over optical fiber.

I will attempt to cover both SONET and SDH in this presentation. However, SONET and SDH use different terminology, which makes it difficult to talk about both at the same time.

I will talk mainly about SONET because you need to understand SONET to understand why some things were done in SDH.

Why SONET/SDH?

- Originally, all communications in the telephone network was analog.
- Analog lines or analog microwave was used to connect to switching offices.

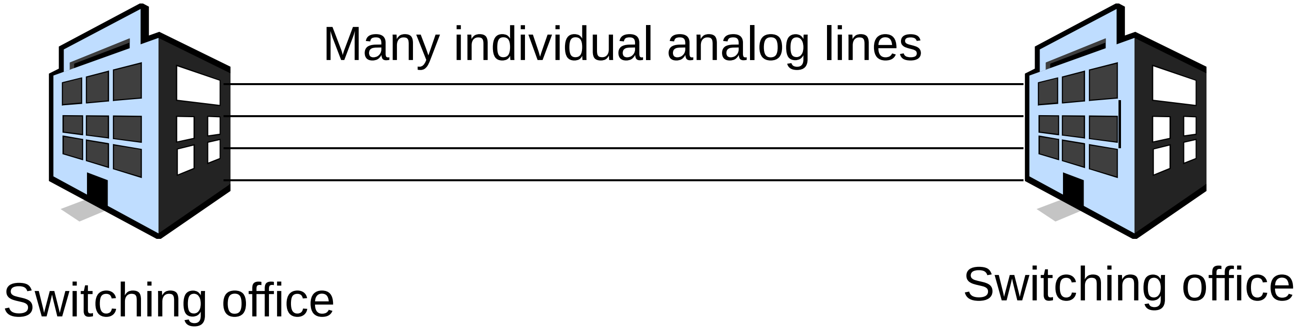

P. Michael Henderson, mike@michael-henderson.us

Why SONET/SDH?

- In about 1962, the network providers began using digital communications between switching centers.
- In the US, this was the D1 channel banks and T-carrier systems.

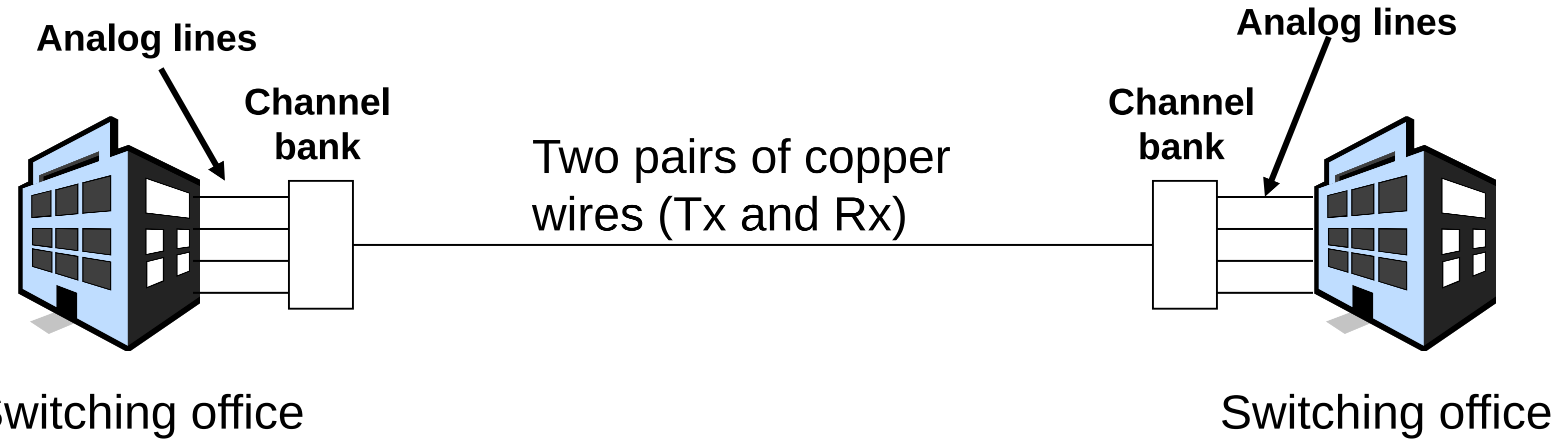

P. Michael Henderson, mike@michael-henderson.us

Why SONET/SDH?

- As communications needs grew, many T- or E-carrier lines were needed between switching centers.
- In the late 1970's optical communications began to be used to interconnect switching offices.

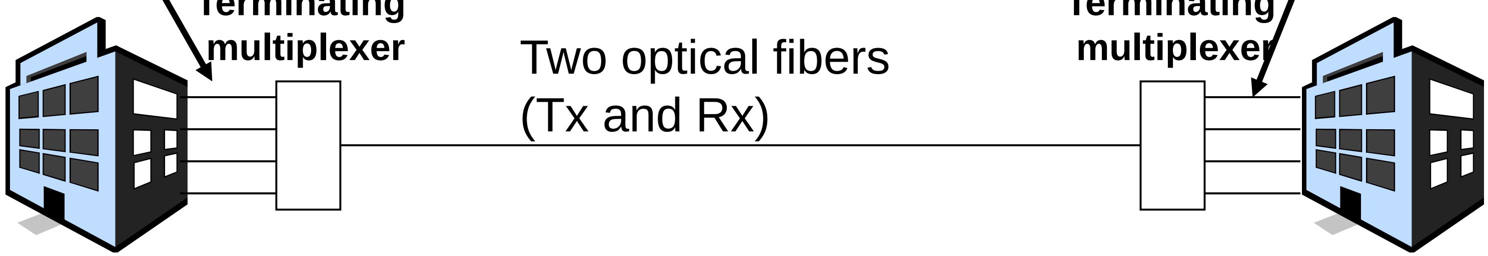

Standardization of SONET/SDH

- Prior to standardization, every manufacturer of optical communications used their own framing.
  - Did not allow multi-vendor networks.
- The ANSI T1X1.5 committee began work in 1985 to define standards for optical communications which would allow “a mid-span meet”.
- The ITU began work in 1986 to achieve the same goal.
- Both bodies finalized their first set of standards in 1988.

Why do we need SONET/SDH

- All data requires framing.
- Since optical networks are complex, provisions are made for management information.
- Many other things are provided.
  - Multiplexing.
  - Error checking.
  - Handling variations in clocks.
  - Mapping of plesiochronous voice and data traffic.
  - Signaling for automatic switching in case of a fiber or node failure.

SONET/SDH Terminology

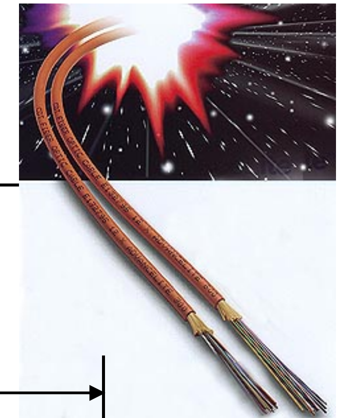

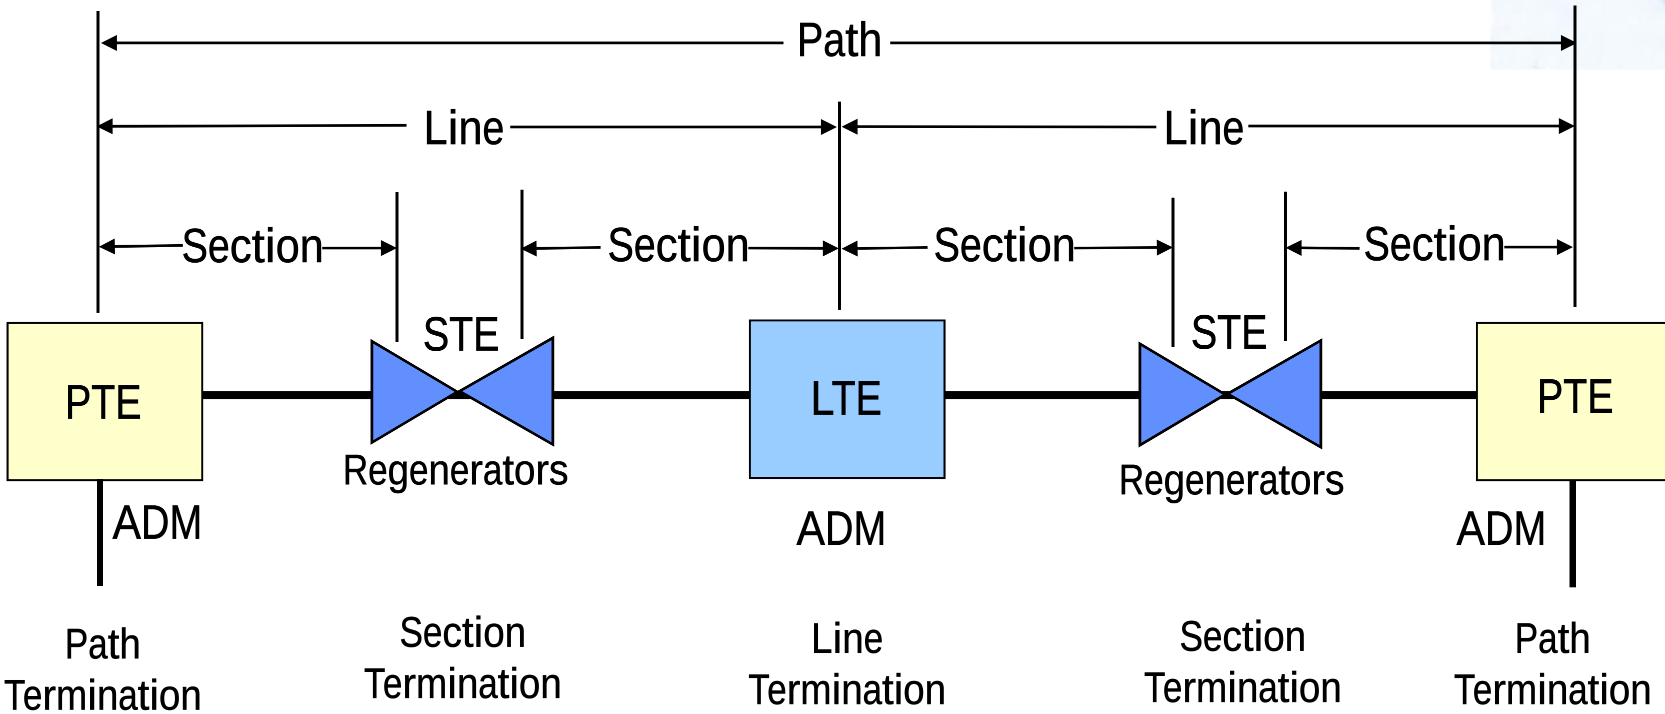

PTE = Path Terminating Equipment
LTE = Line terminating Equipment
STE = Section Terminating Equipment
ADM = Add/Drop Multiplexer

STS-1 SONET Frame

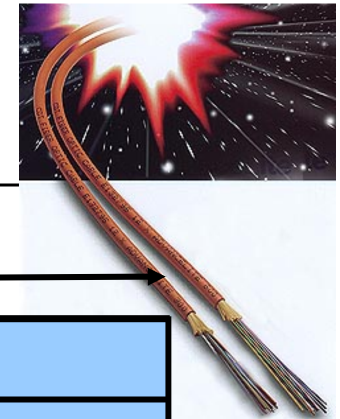

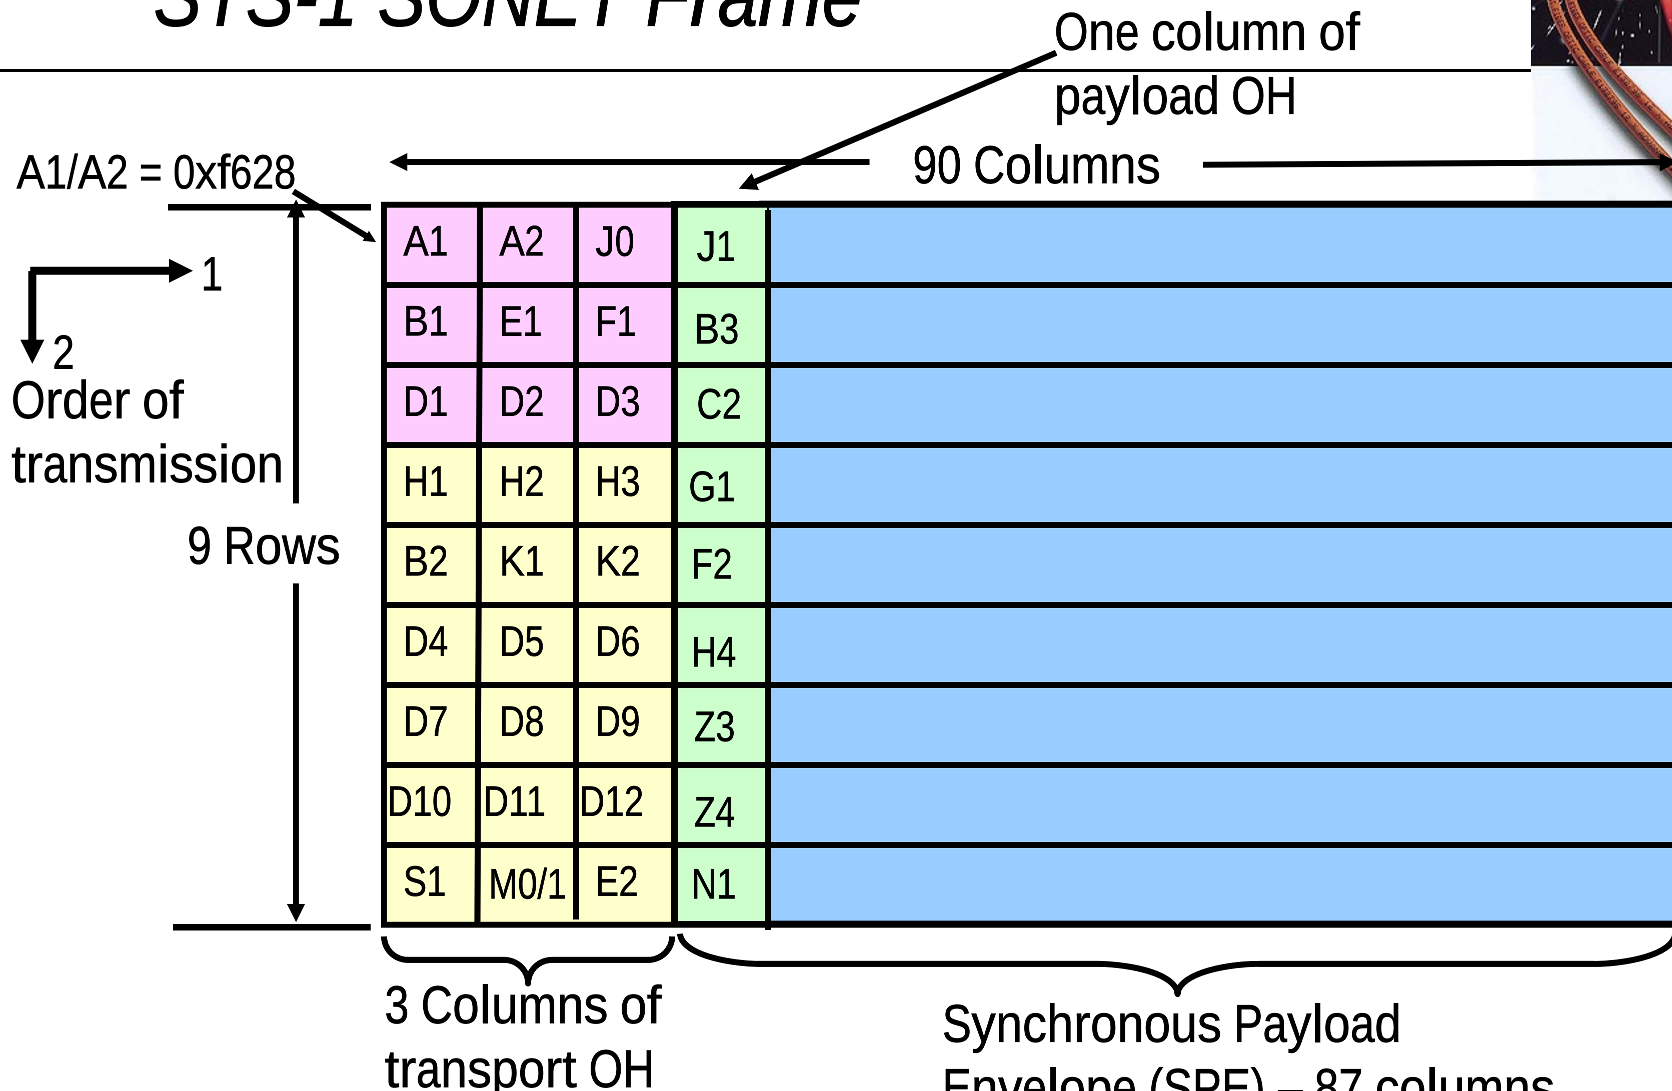

Pink	Section overhead	Green	Payload overhead
Yellow	Line overhead	Blue	Data

STS-1 Frame Format

- 9 rows by 90 columns – 810 octets in the frame.
- Frame is transmitted from left to right, by row.
- Frames are transmitted 8,000 times per second, every 125 μseconds.
- STS-1 bit rate is therefore 51.84 Mbps (810 octets x 8,000 times per second x 8 bits per octet).
- This lowest level SONET signal is called a Synchronous Transport Signal, level 1 (STS-1). Once the scrambler is applied, it is known as an Optical Channel, level 1 (OC-1).
- The lowest level SDH signal is know as a Synchronous Transport Module, level 1 (STM-1).

SONET/SDH data rates

SONET name	SDH name	Line rate (Mbps)	SPE rate (Mbps)	Overhead rate (Mbps)
OC-1	STM-0	51.84	50.112	1.728
OC-3	STM-1	155.52	150.336	5.184
OC-12	STM-4	622.08	601.344	20.736
OC-48	STM-16	2,488.32	2,405.376	82.944
OC-192	STM-64	9,953.28	9,621.504	331.776
OC-768	STM-256	39,813.12	38,486.016	1,327.104

Interleaving of SONET/SDH signals

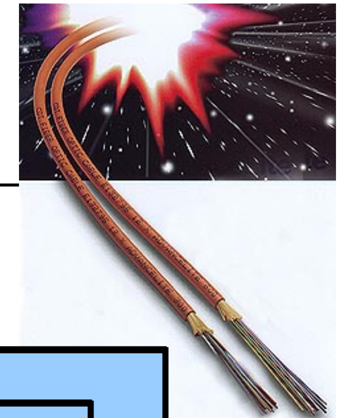

Order of transmission

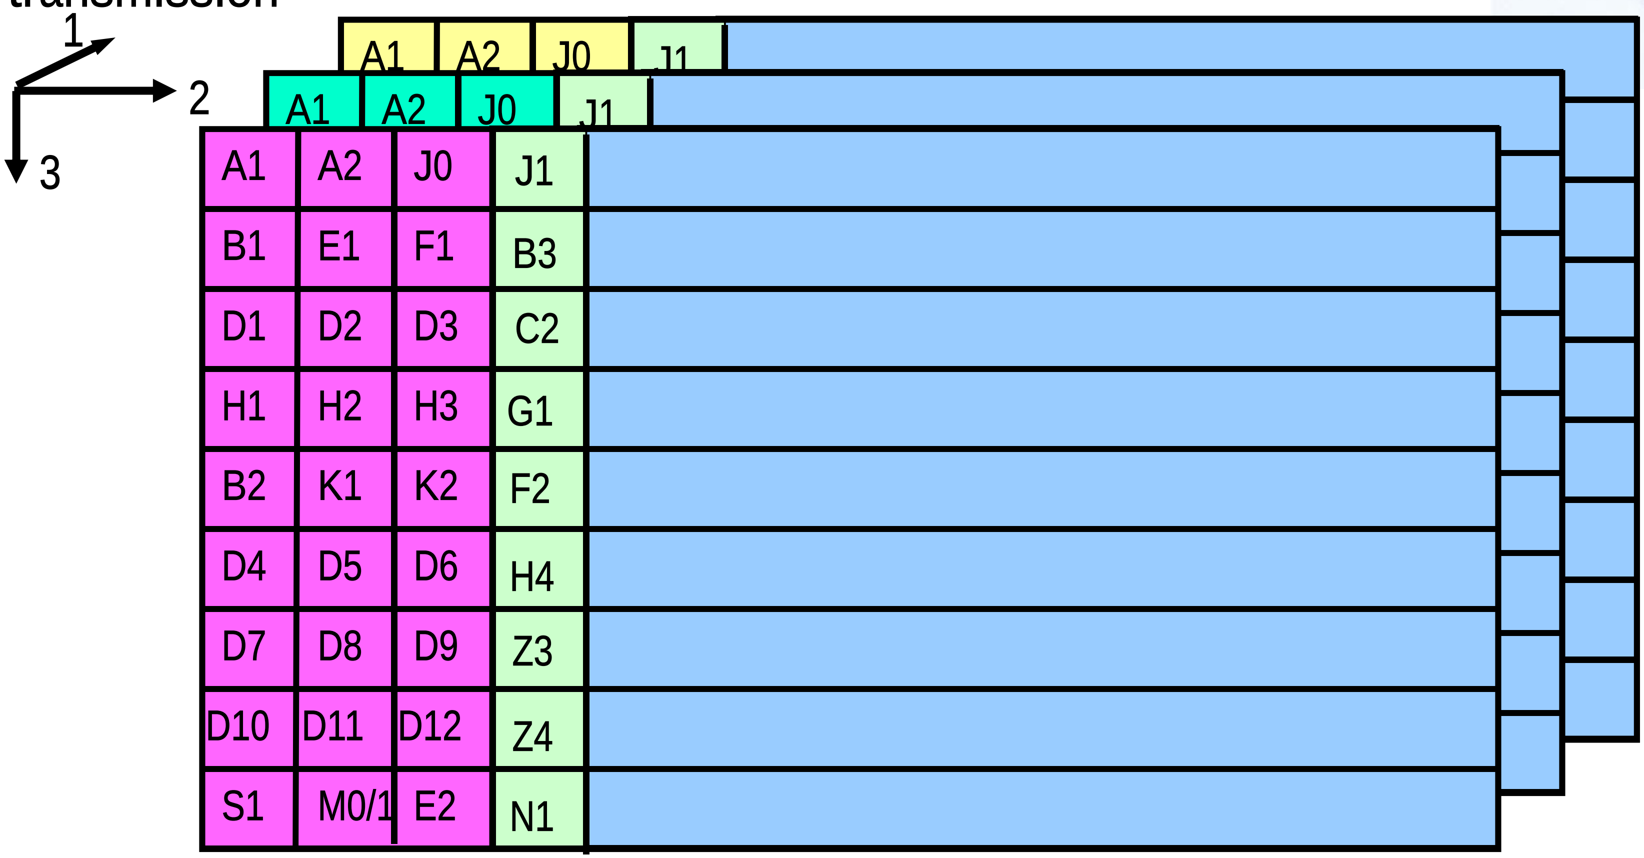

P. Michael Henderson, mike@michael-henderson.us

Interleaving of SONET/SDH signals

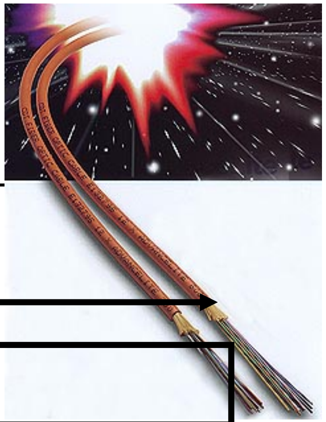

270 Columns

A1	A1	A1	A2	A2	A2	J0	Z0	Z0	
B1	X	X	E1	X	X	F1	X	X	
D1	X	X	D2	X	X	D3	X	X	
H1	H1	H1	H2	H2	H2	H3	H3	H3	
B2	B2	B2	K1	X	X	K2	X	X	
D4	X	X	D5	X	X	D6	X	X	
D7	X	X	D8	X	X	D9	X	X	
D10	X	X	D11	X	X	D12	X	X	
S1	Z1	Z1	M0/1	Z2	M2	E2	X	X	

First STS-1 (magenta) — Second STS-1 (cyan) — Third STS-1 (yellow)

Payloads in Interleaved Signals

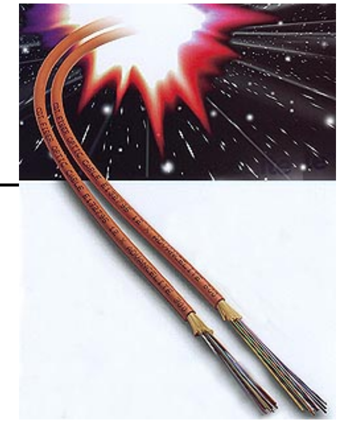

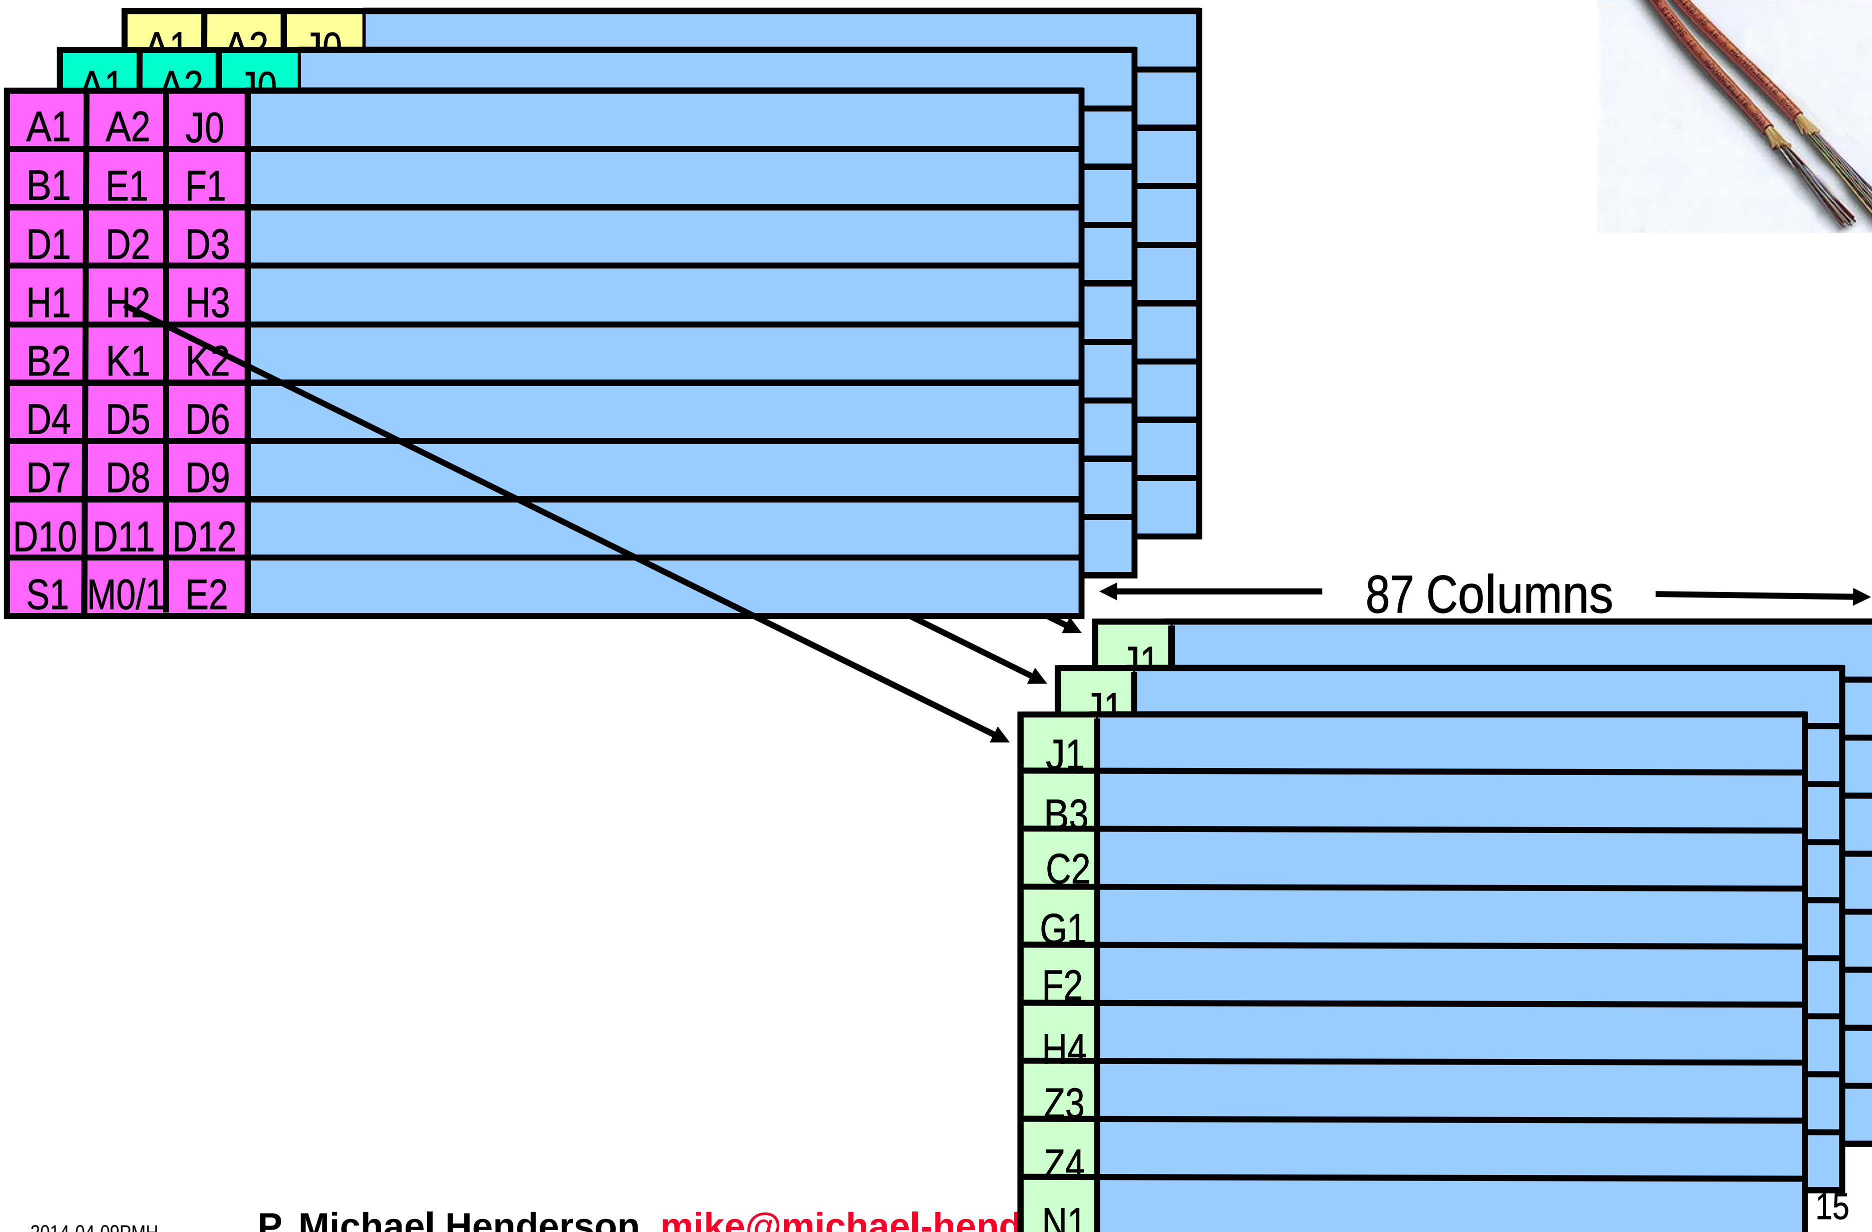

Concatenated Payloads

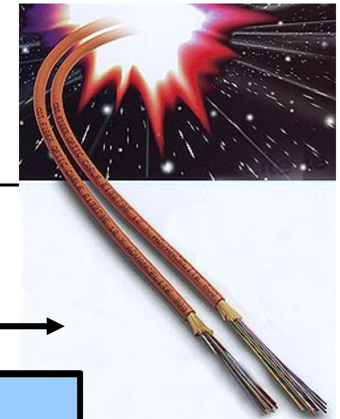

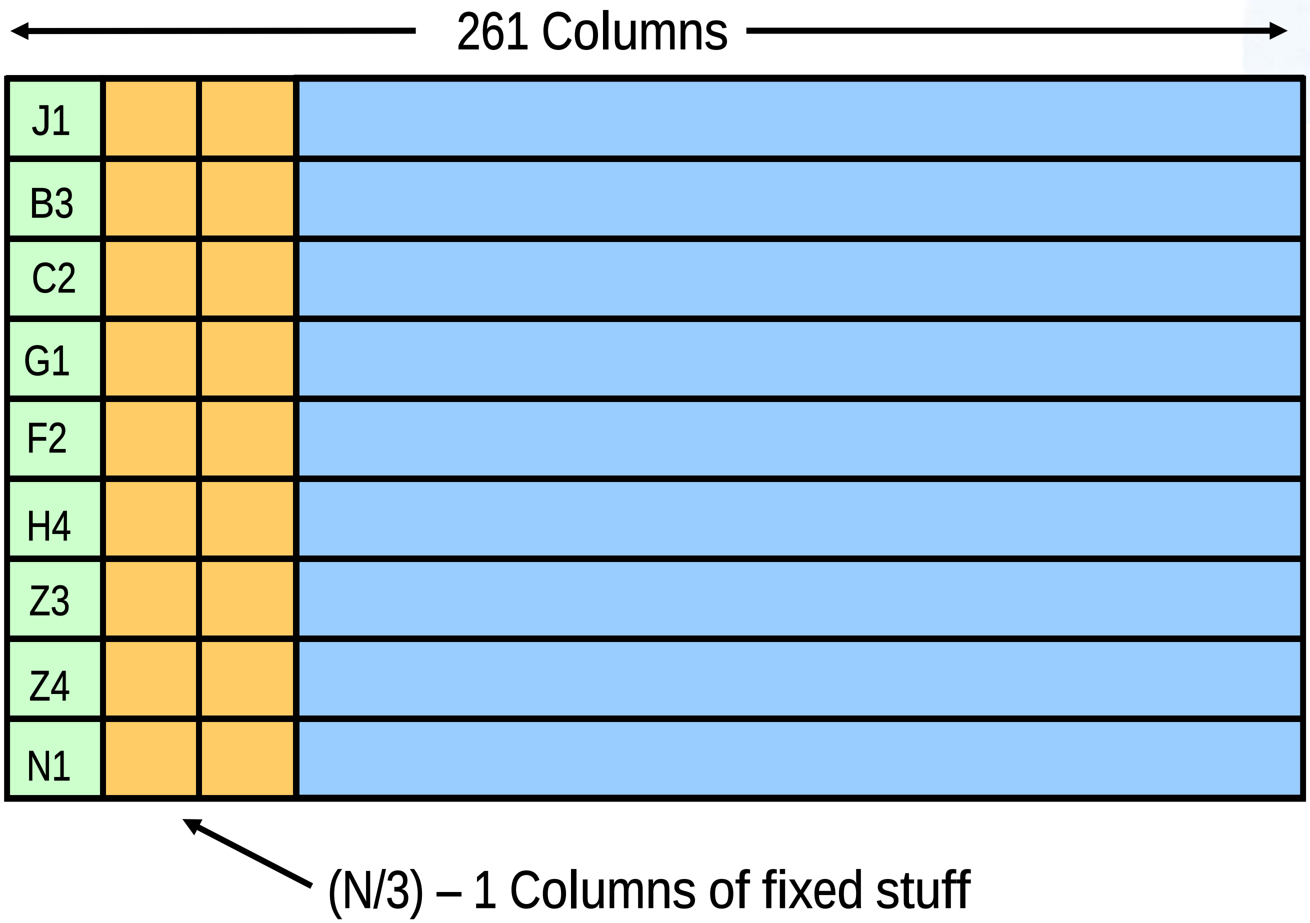

P. Michael Henderson, mike@michael-henderson.us

Concatenated Payloads

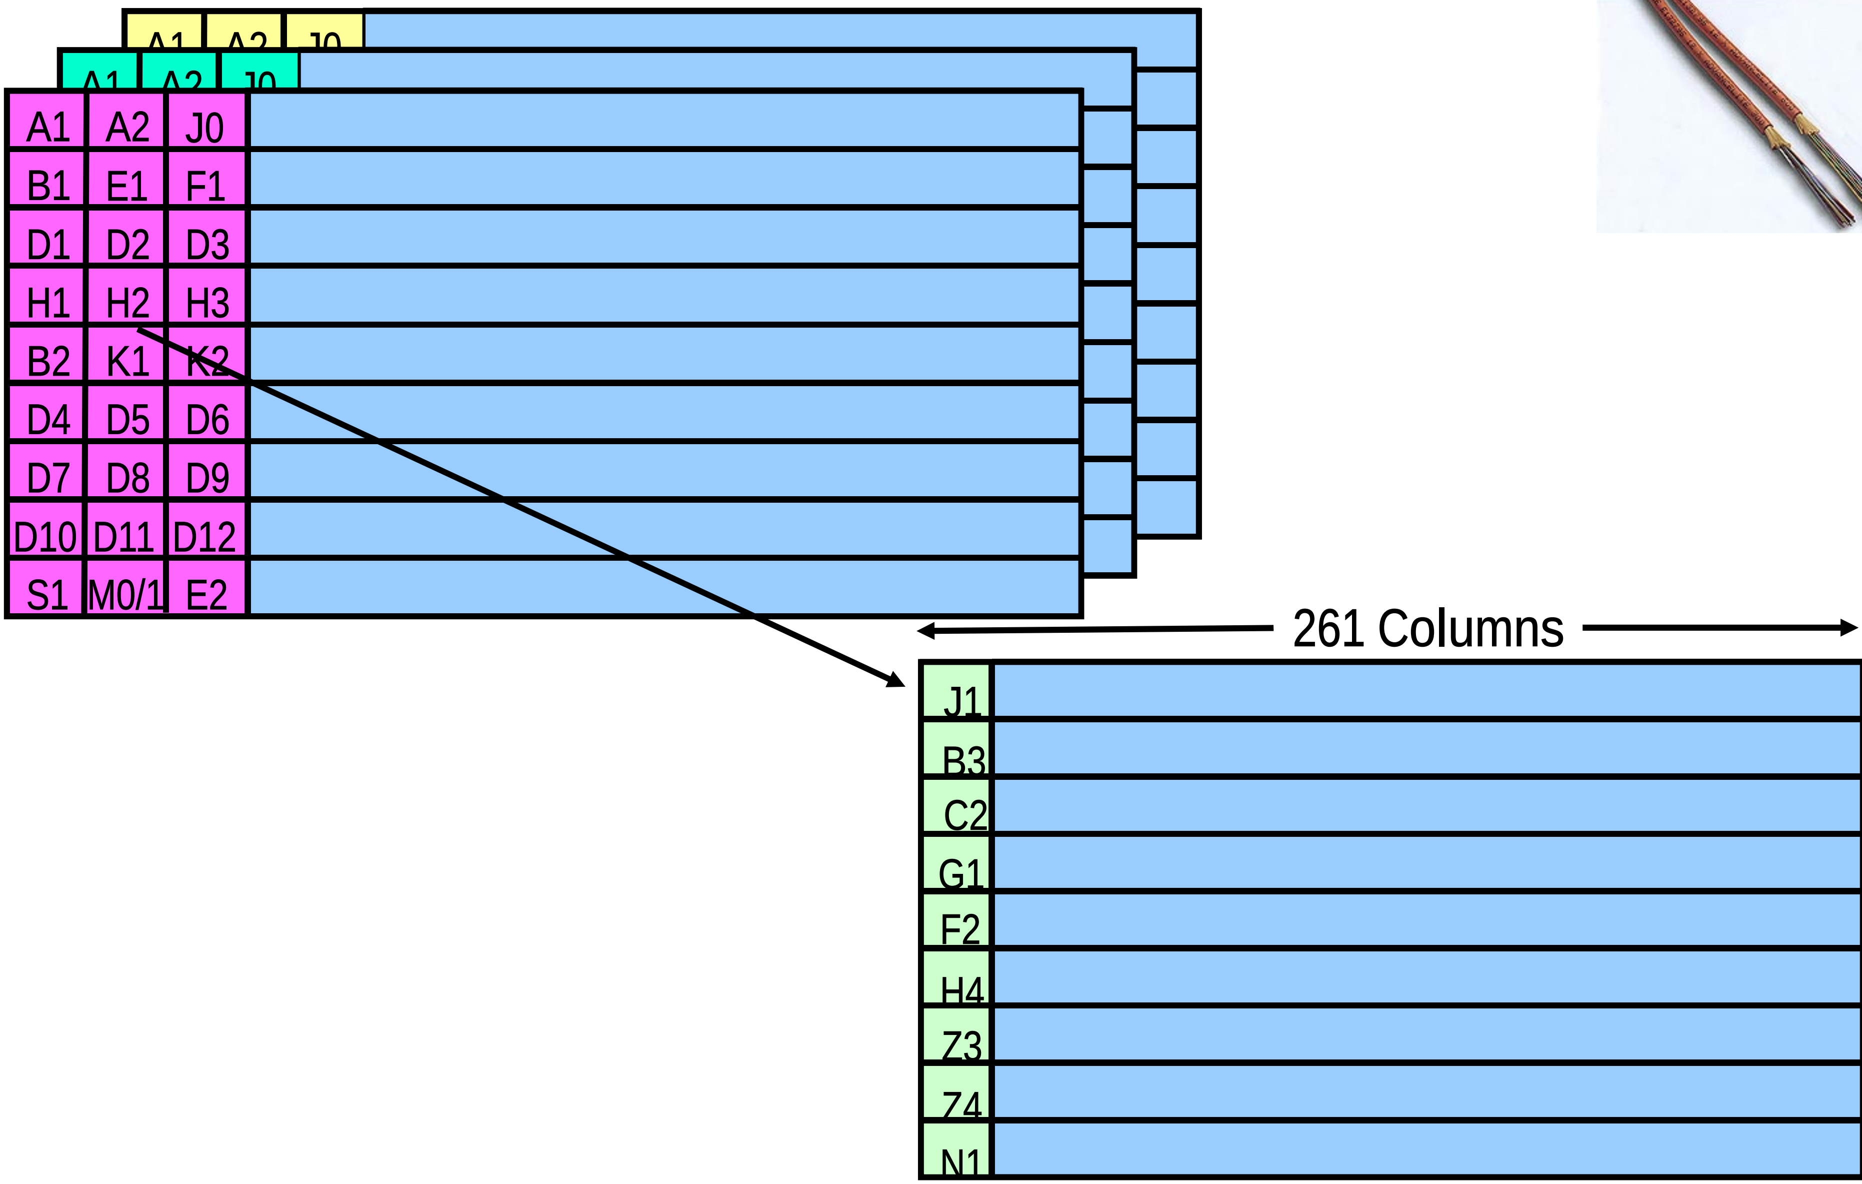

P. Michael Henderson, mike@michael-henderson.us

Framing and Section Trace

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- Two octets used for framing: A1, A2.
- Bit pattern is A1 = 1111 0110, A2 = 0010 1000
- Higher levels of STS-N have N A1 octets and N A2 octets.
- Section Trace (J0) is used to verify continued connection between section entities.

Monitoring for Bit Errors

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- Does a bit interleaved parity (BIP-8) over the octets in the (previous) frame.
- Even parity.
- Parity check applies to *previous* frame
- Separate check for section and line.
- Only one B1 octet, no matter what the STS-N.
- One B2 octet for each STS-N.

Communication Channels

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- E1, E2 (orderwire) was intended to be used by craftspersons while installing a line. Craftspeople use cellular phones so this octet is not used very much.
- F1 is available for the network provider to use as they wish.

P. Michael Henderson, mike@michael-henderson.us

Communication Channels

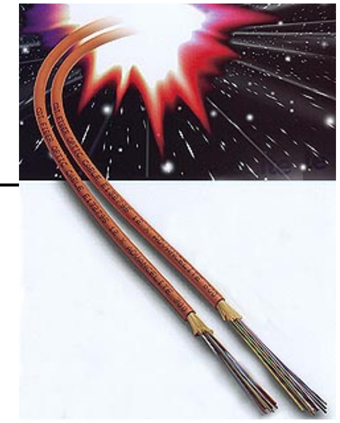

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- These channels are used for network management.
- Technically, D1, D2, D3 is intended for section messages but this is not always adhered to.

Automatic Protection Switching

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- These octets (K1, K2) are used to send messages between two nodes when a failure is detected between them.
- Messages are sent both ways around the ring (if possible).
- Last four bits of K1 specifies the address of the addressed node.
- First four bits of K2 specifies the address of the sending node.
- Means that there can only be 16 nodes on a ring.

Sync and Error Indication

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- S1 is used for synchronization status.
- M0/M1 is used to send back to the sender, the error status of the received signal (determined by the Bx octets).

Payload Pointers

A1	A2	J0
B1	E1	F1
D1	D2	D3
H1	H2	H3
B2	K1	K2
D4	D5	D6
D7	D8	D9
D10	D11	D12
S1	M0/1	E2

- The H1, H2 octets are joined to form a pointer which points to the first octet of the SPE.
- The first 4 bits of the 16 bits indicates if the pointer is changing.
- The next two bits are not used.
- The last 10 bits are the pointer and can have a value from zero to 782.
- H3 is used to carry a payload octet when a negative pointer adjustment is done.
- These octets are covered in additional detail in the next slides.

Why do we need Payload Pointers?

Network box which takes traffic from one side and passes it along to the other side

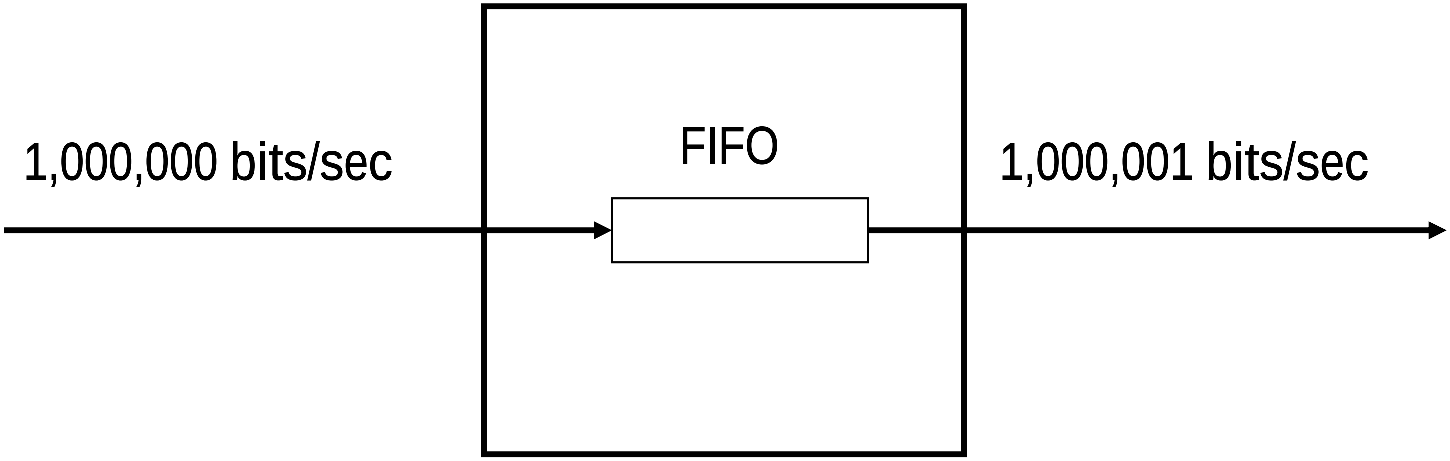

Every eight seconds, the FIFO will run dry. If an extra, "meaningless" octet can be sent at that time, it will give the FIFO time to fill up again. Problem: How to tell the receiver that this one octet is meaningless?

Format of the H1, H2, H3 Octets

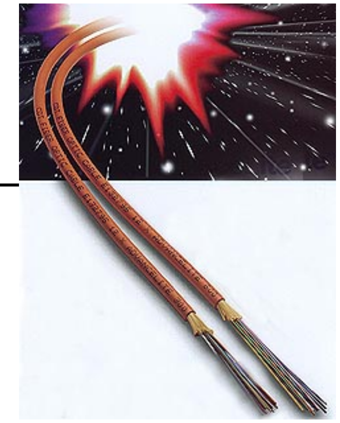

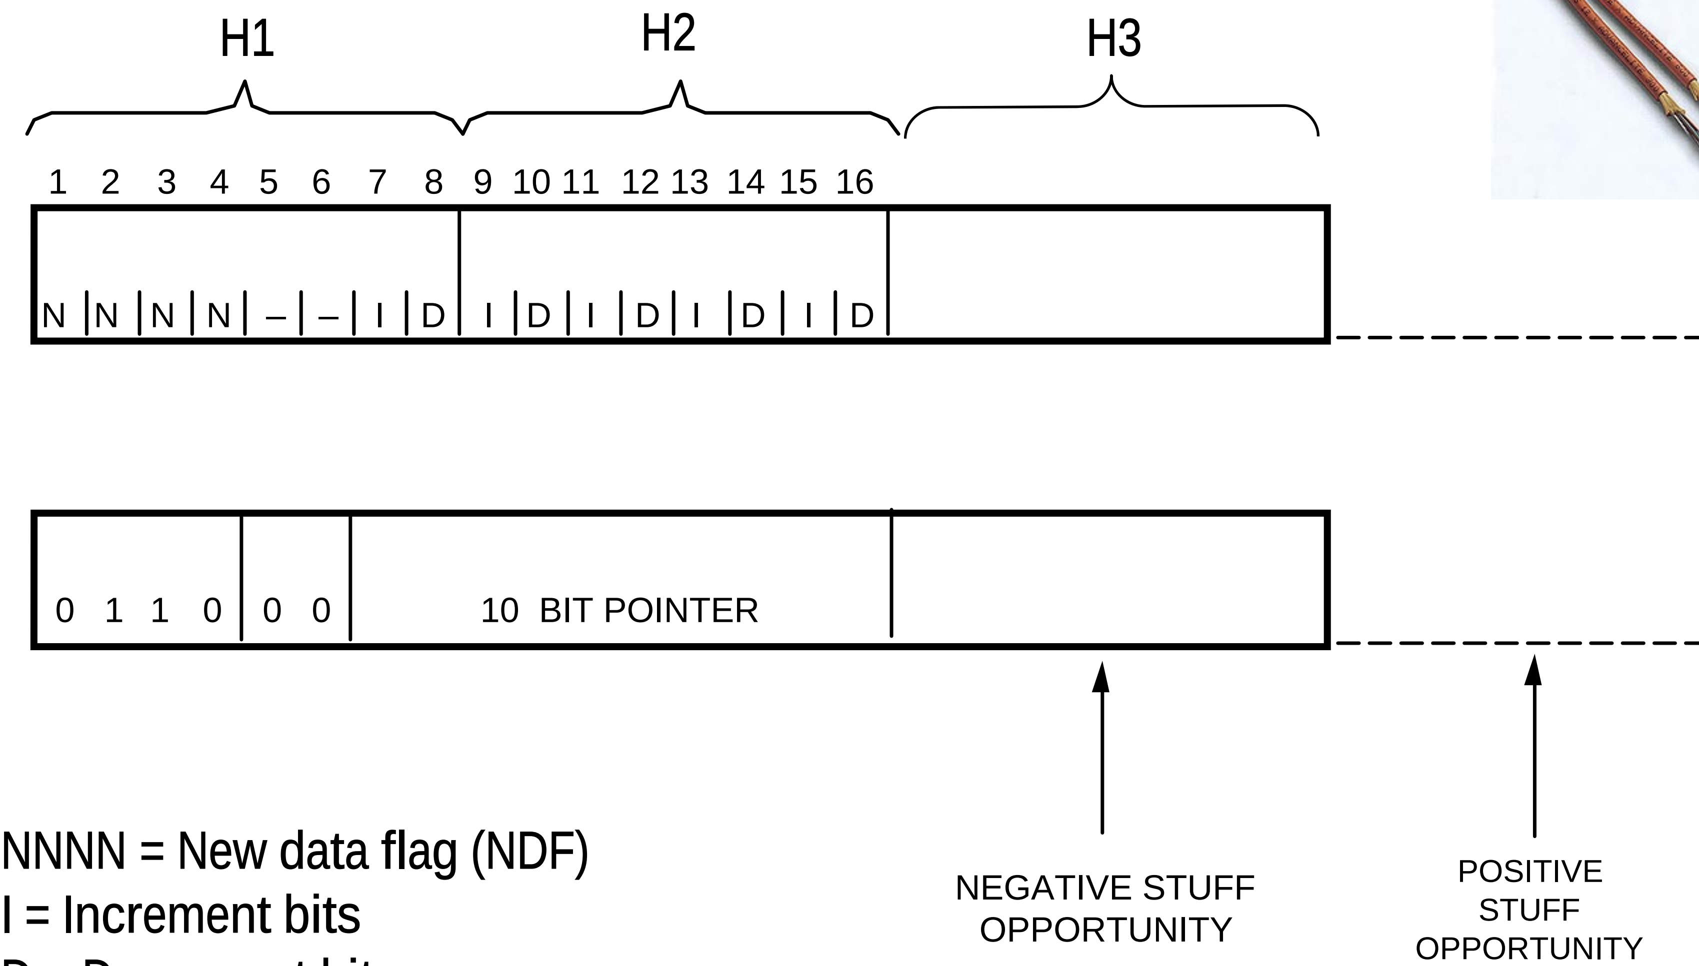

NNNN = New data flag (NDF)
I = Increment bits
D = Decrement bits

P. Michael Henderson, mike@michael-henderson.us

A Negative Pointer Adjustment

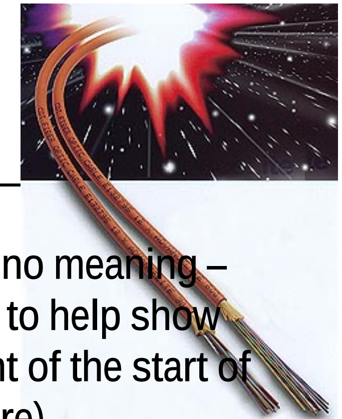

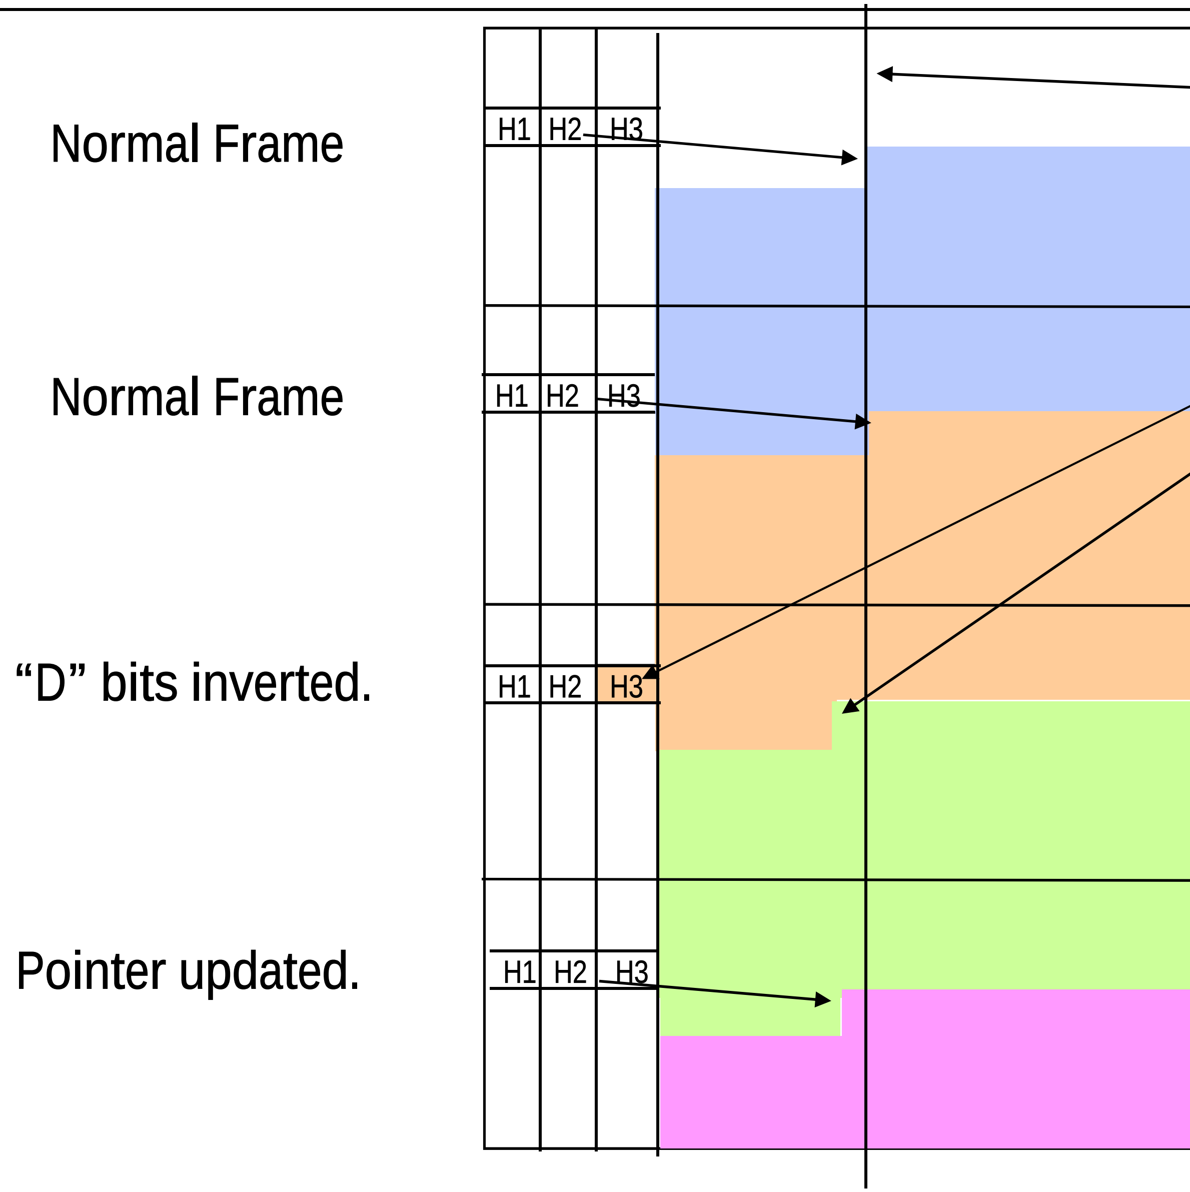

This line has no meaning – it’s just there to help show the movement of the start of the frame (here).

Note that H3 is being used to carry a payload octet

The pointer is invalid here because of the “D” bit inversion. The value without inversion is the same as the previous frame.

P. Michael Henderson, mike@michael-henderson.us

Details of H1, H2, H3 octets during justification

Negative Justification using Inverted "D" Bits

Frame status	New Data Flag				Unused		I	D	I	D	I	D	I	D	I	D
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	0
Invert "D" Bits	0	1	1	0	X	X	0	1	0	0	1	0	1	0	1	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1

A Positive Pointer Adjustment

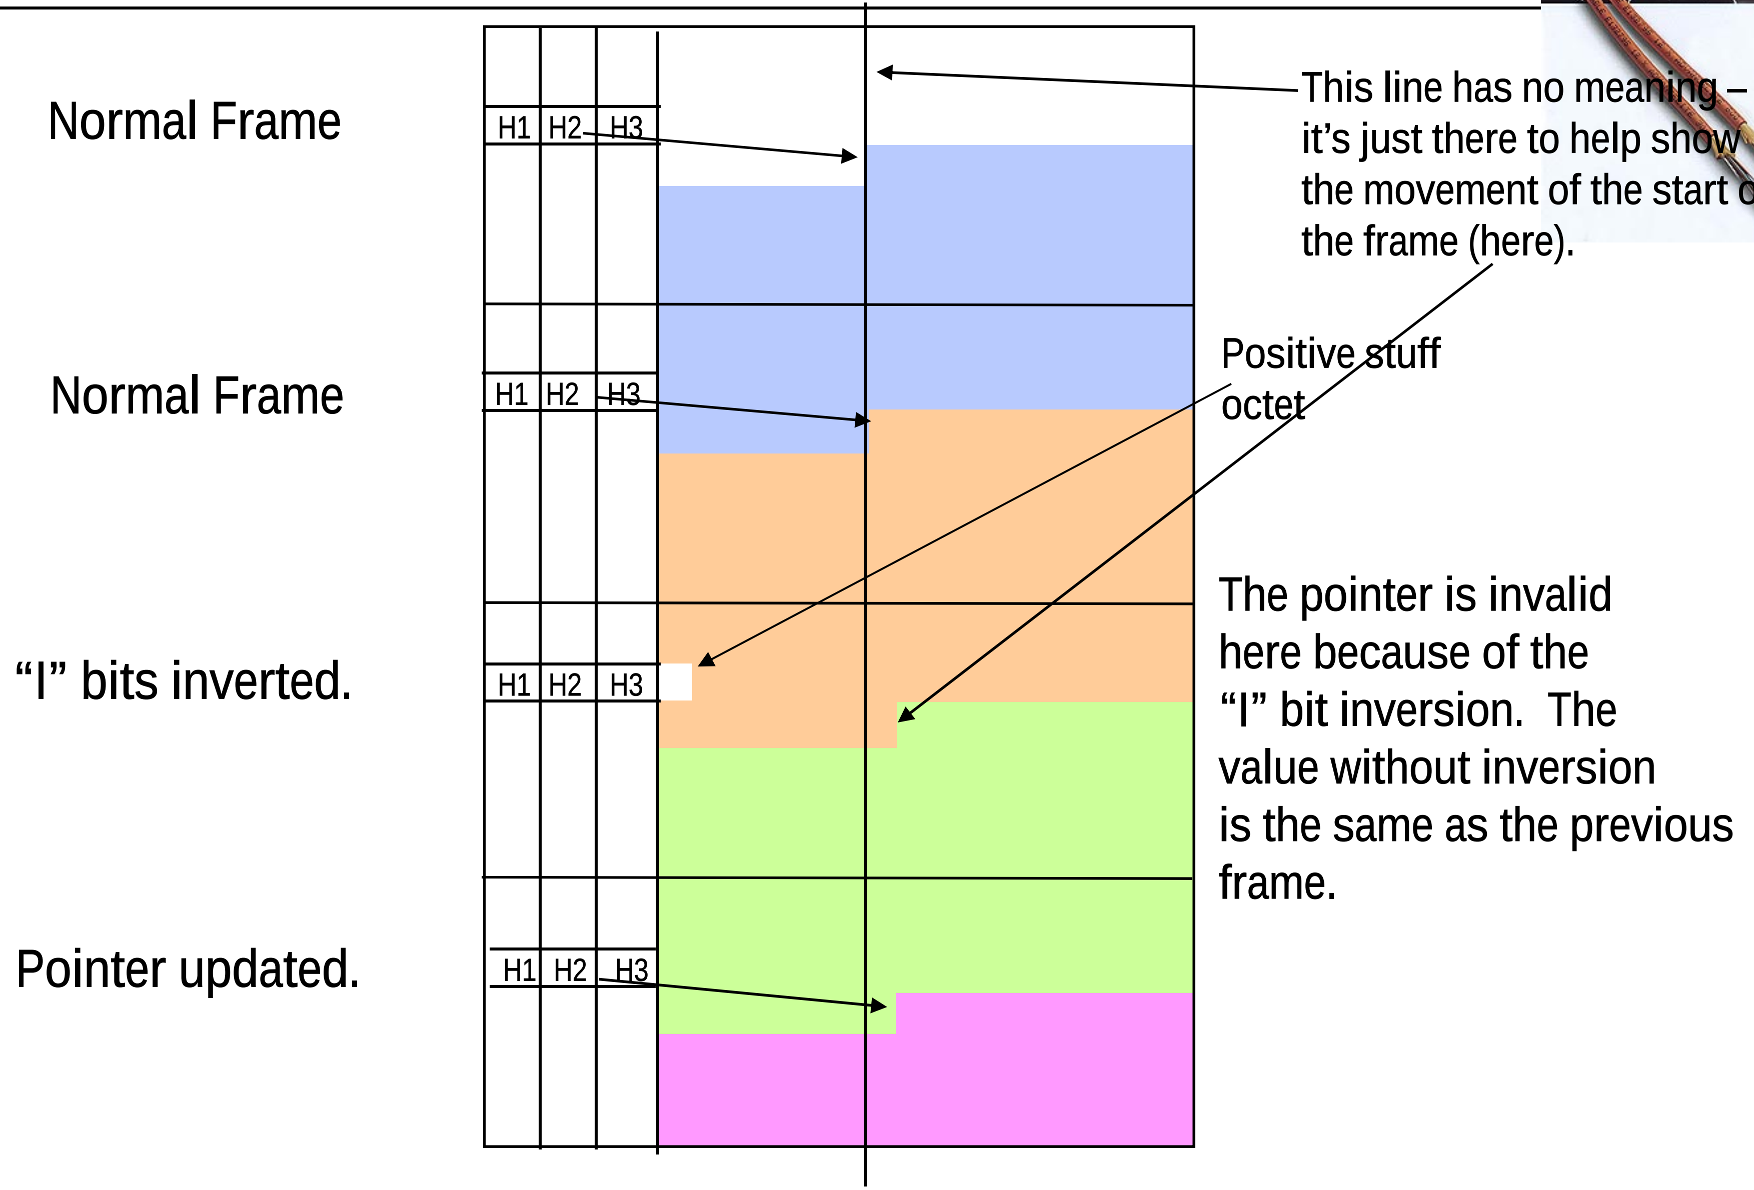

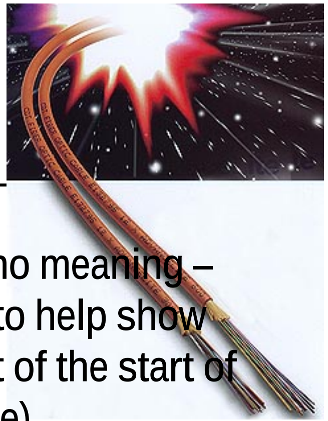

P. Michael Henderson, mike@michael-henderson.us

Details of H1, H2, H3 octets during justification

Positive Justification using Inverted "I" Bits

Frame status	New Data Flag				Unused		I	D	I	D	I	D	I	D	I	D
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	0
Invert "I" Bits	0	1	1	0	X	X	1	0	1	1	0	1	0	1	0	0
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1

Details of H1, H2, H3 octets when using NDF

One Octet Negative Adjustment using the NDF

Frame status	New Data Flag				Unused		I	D	I	D	I	D	I	D	I	D
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	0
NDF indicator	1	0	0	1	X	X	0	0	0	1	1	1	1	1	0	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	0	1

Details of H1, H2, H3 octets when using NDF

One Octet Positive Adjustment using the NDF

Frame status	New Data Flag				Unused		I	D	I	D	I	D	I	D	I	D
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	0
NDF indicator	1	0	0	1	X	X	0	0	0	1	1	1	1	1	1	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1
New ptr value	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1
Normal frame	0	1	1	0	X	X	0	0	0	1	1	1	1	1	1	1

Payload Overhead

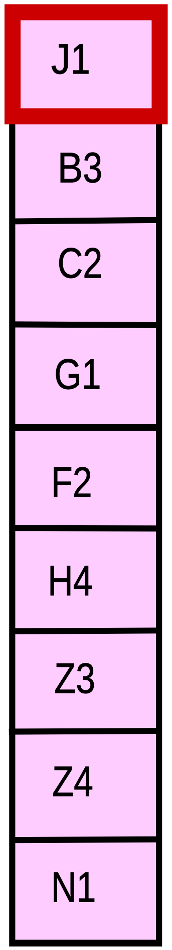

- Path trace (J1) used to check continued connection between the path devices (end-to-end).
- 64 octet for SONET and 16 octet for SDH message.

P. Michael Henderson, mike@michael-henderson.us

Error Checking

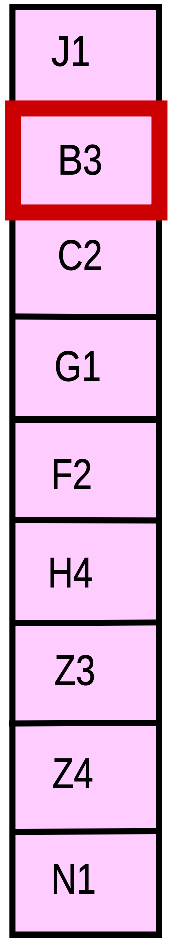

- BIP-8 parity check over the payload, only.
- Even parity.
- Parity is for *previous* payload, not this one.

P. Michael Henderson, mike@michael-henderson.us

Path Signal Label

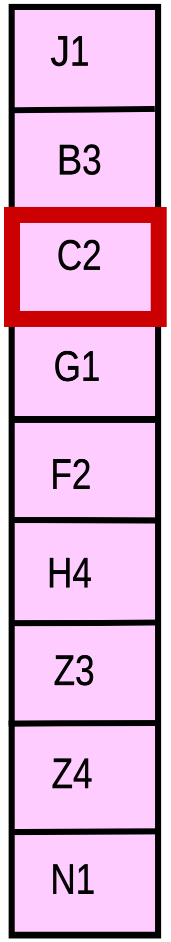

- Indicates the type of payload carried in the SPE.
- For our discussion, we'll be interested in a value of 0x02 for floating virtual tributary (VT) mode, 0x04 for asynchronous mapping of DS-3, 0x13 for mapping of ATM, 0x16 for packet over SONET (POS), and 0x1b for generic framing procedure (GFP).

Path Status and Path User Channel

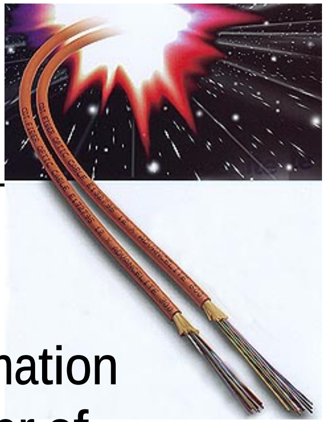

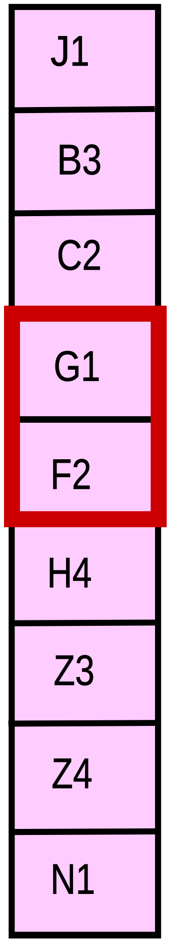

- The G1 octet (Path Status) sends information back to the sender indicating the number of parity errors, or if a complete failure is detected.
- The F2 octet is a user channel for Path applications (end-to-end). It is not subject to standardization (anything can be put into it).

Multi-Frame Indicator

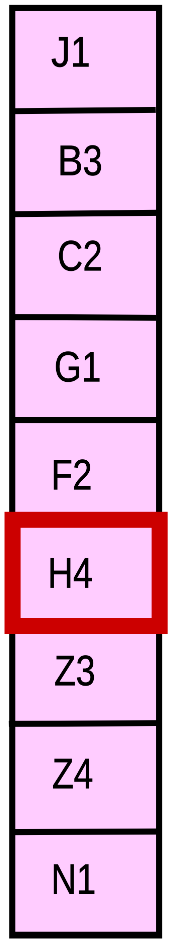

- The last two bits of this octet counts from 00 to 11 continuously to provide a multi-frame indicator for VT payloads (to be explained later).
- Some of the other bits are being defined for use with virtual concatenation but will not be described here.

P. Michael Henderson, mike@michael-henderson.us

TCM and Reserved octets

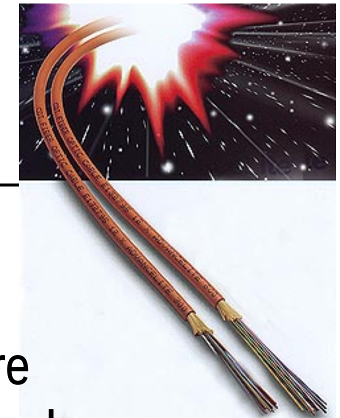

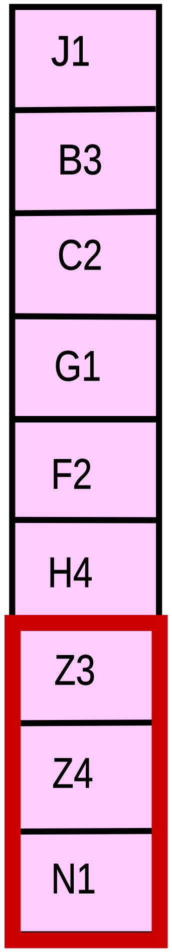

- The Z3, Z4 octets are reserved for future standardization and have no meaning today.
- The N1 octet is for tandem connection monitoring and is fairly complex. No attempt will be made here to explain it.

P. Michael Henderson, mike@michael-henderson.us

Virtual Tributaries

P. Michael Henderson, mike@michael-henderson.us

Plesiochronous Data Rates of Interest

Type of Digital Circuit	Bit Rate (Mbps)
DS-1 (T1)	1.544
E1	2.048
DS-1C	3.152
DS-2	6.312
DS-3 (T3)	44.736

STS-1 Payload Mapping

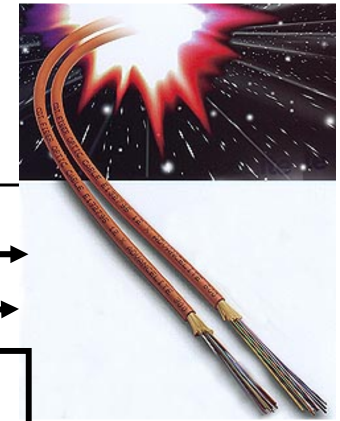

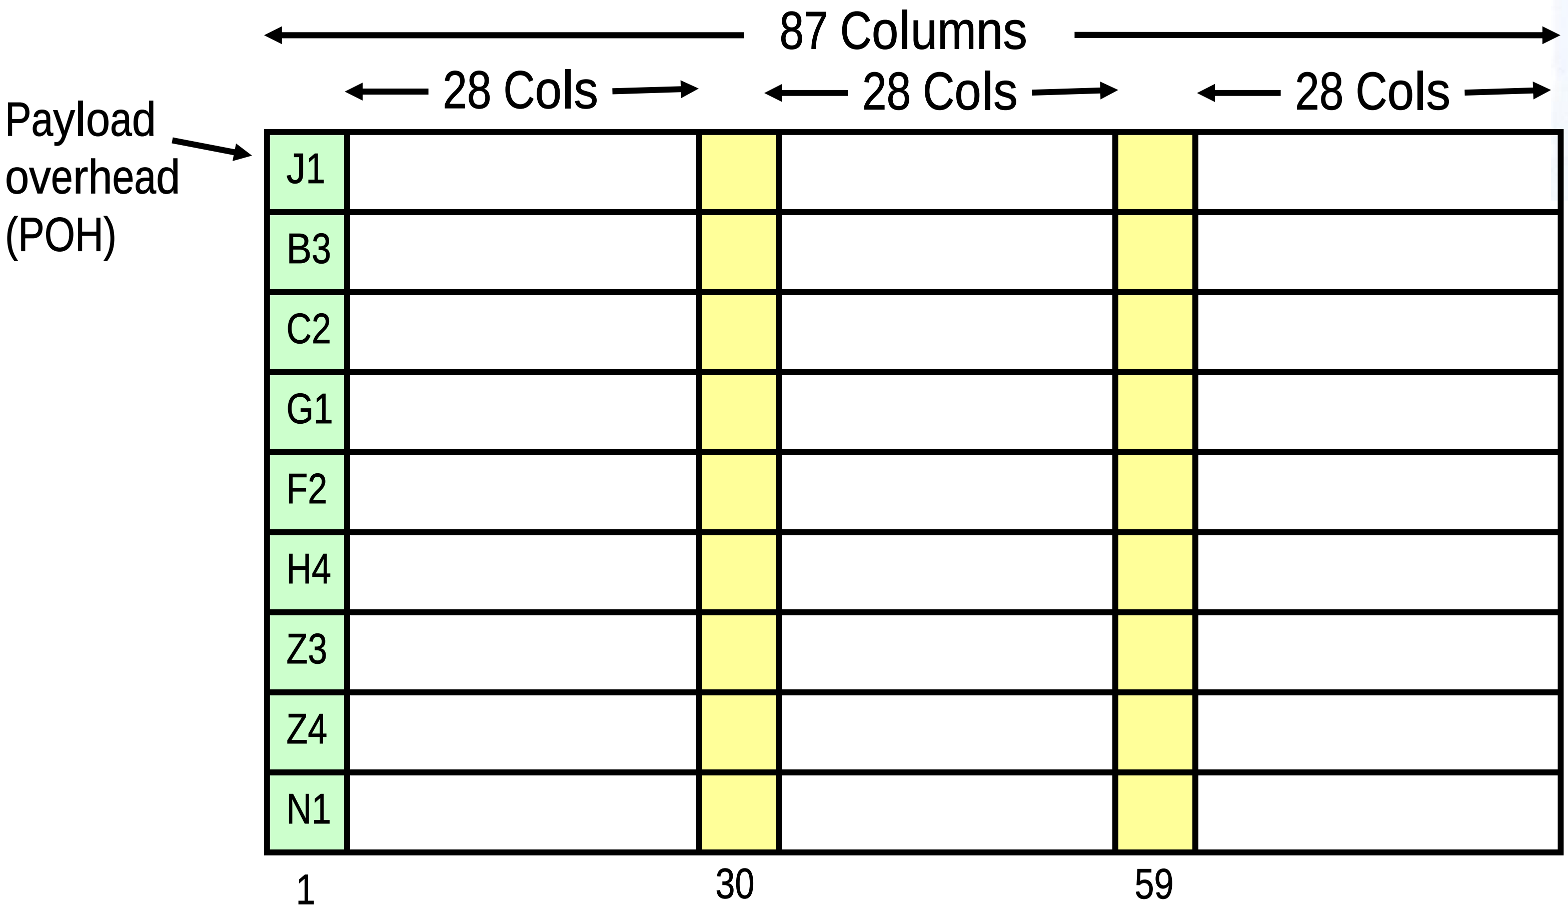

All mappings of payload into an STS-1 payload have these two columns taken as stuff columns.

P. Michael Henderson, mike@michael-henderson.us

Virtual Tributary Groups

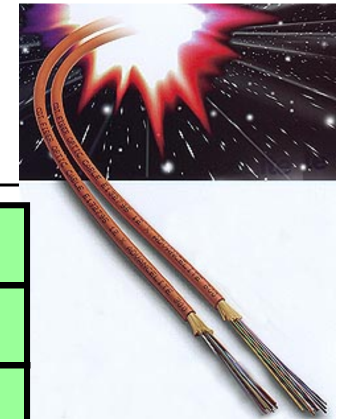

The “x” in Vx goes from 1 to 4

Order of transmission

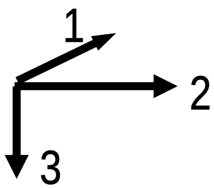

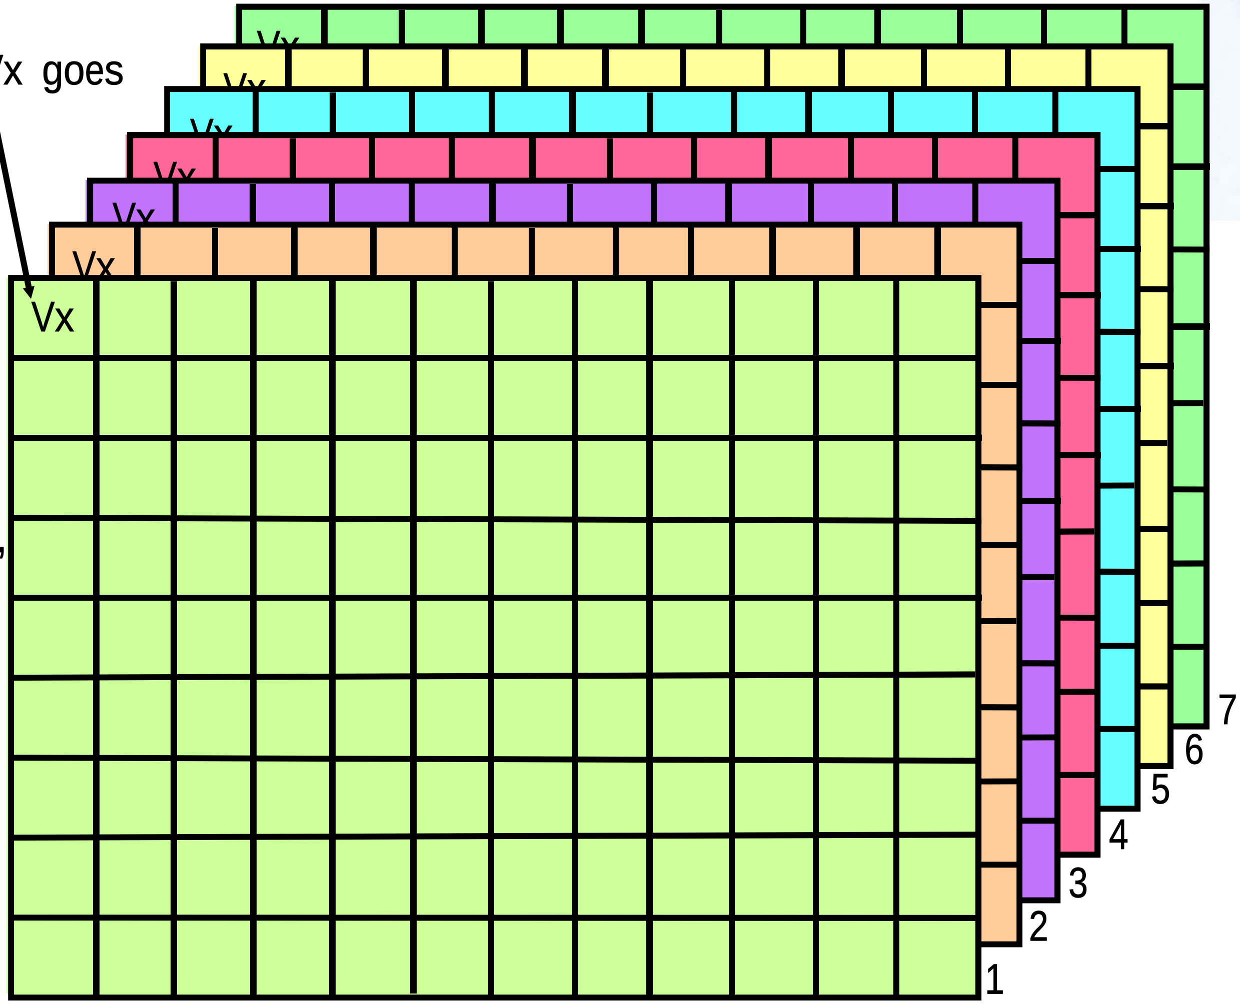

When interleaving, the first column of VTG#1 goes next to the POH. Columns 30 & 59 of the SPE are skipped.

The 84 usable payload columns are divided into seven groups of twelve columns. Each set of twelve is called a “Virtual Tributary Group” (VTG)

Virtual Tributaries

- A Virtual Tributary Group is further subdivided into Virtual Tributaries (VT).
  - Three columns makes a VT-1.5 (1.728 Mbps gross, good for a DS-1 at 1.544 Mbps).
  - Four columns makes a VT-2 (2.304 Mbps gross, good for an E1 at 2.048 Mbps).
  - Six columns makes a VT-3 (3.456 Mbps gross, good for a DS-1C at 3.152 Mbps).
  - Twelve columns makes a VT-6 (6.912 Mbps gross, good for a DS-2 at 6.312 Mbps).
- A VTG contains four VT-1.5s, or three VT-2s, or two VT-3s, or one VT-6.
- A VTG can only contain one kind of VT. Different VTGs in an SPE can contain different VT types, but within a VTG, there can only be one kind of VT.

Virtual Tributaries – The VT-1.5

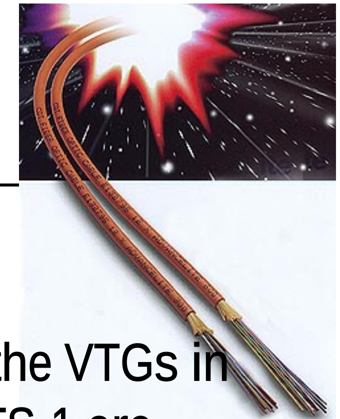

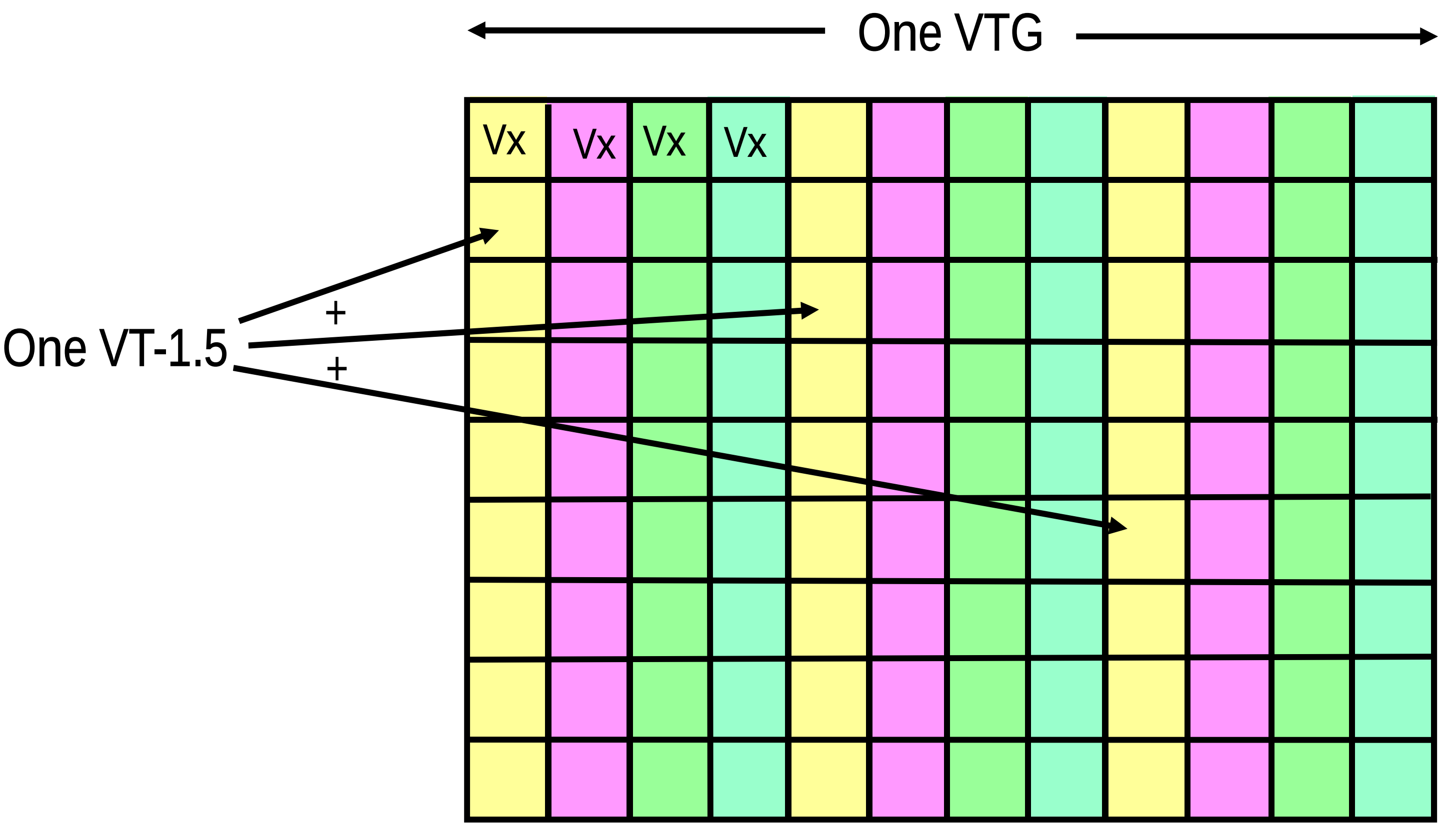

If all the VTGs in an STS-1 are carrying VT-1.5s, the capacity of the STS-1 SPE is 28 DS-1s, or the same as a DS-3. (Seven VTGs times four VT-1.5s per VTG)

The four VT-1.5s are interleaved into the VTG. Note the colors – there are three of each color, with each color indicating one VT-1.5 (for four inside the VTG. Each VT is 27 octets, with one octet taken for the Vx octet.

Superframes

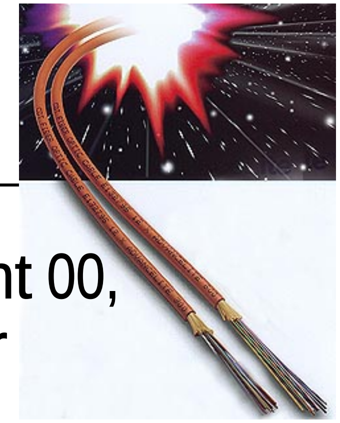

- Remember that the last two bits of H4 in the POH count 00, 01, 10, 11, 00, etc. This produces a superframe of four frames.
- The frame after the SPE with an H4 value of 00 will have the first octet in the VT identified as V1. The one after value 01 will be identified as V2, etc.
- V1 and V2 form a pointer, exactly like the H1, H2 octets.
  - Bits 5 and 6 are used to indicate the type of VT (unused in H1, H2) (VT-6=00, VT-3=01, VT-2=10, VT-1.5=11)
- V3 is the negative stuff opportunity, exactly like H3.
- V4 is reserved for future standardization.

V1	V2	V3	V4

The V1, V2 Pointers

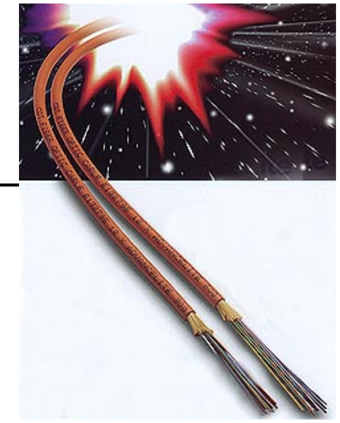

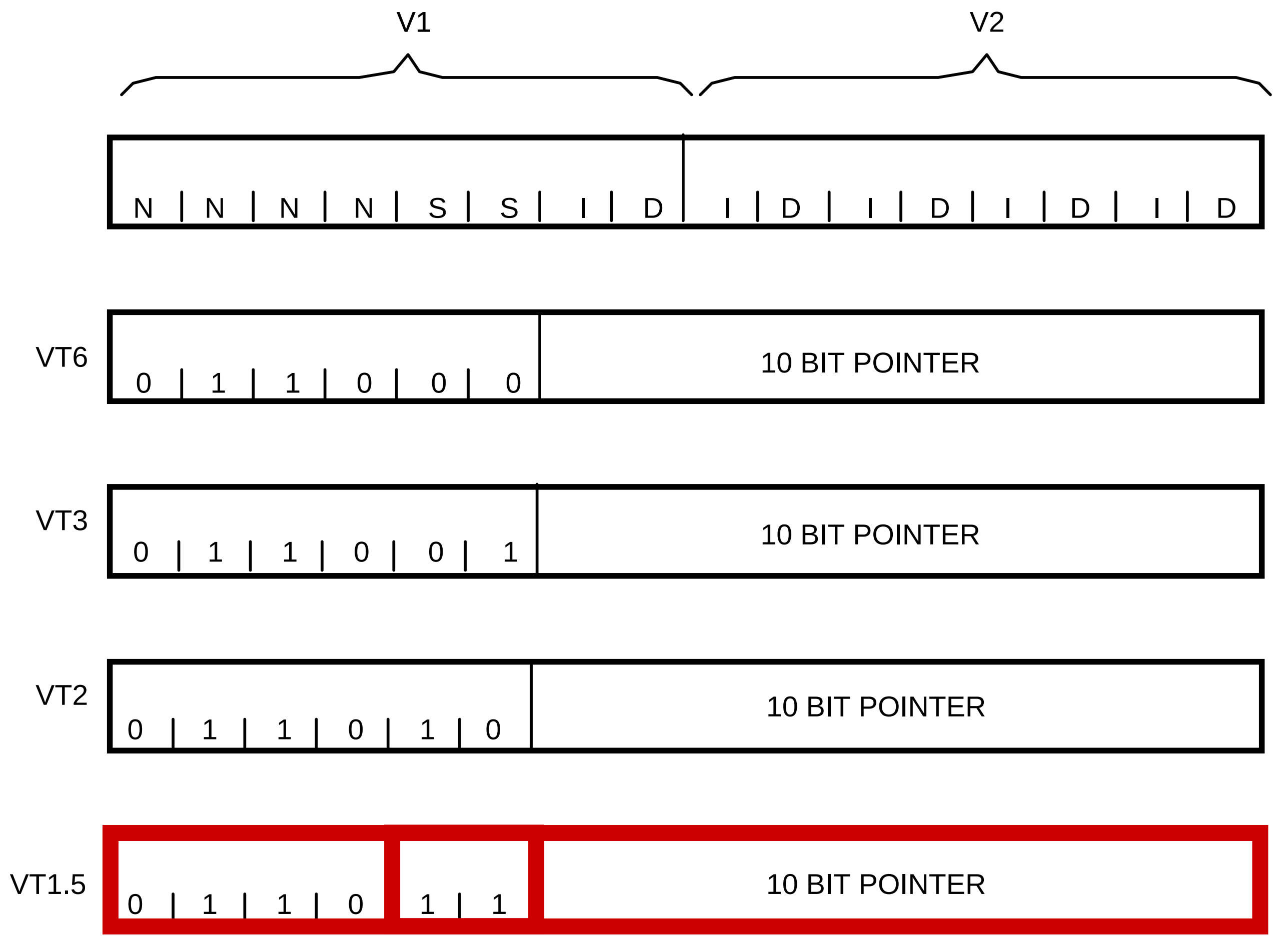

NEW DATA FLAG – INVERT 4 N BITS
NEGATIVE STUFF – INVERT 5 D BITS
POSITIVE STUFF – INVERT 5 I BITS

I – Increment Bit
D – Decrement Bit
N – New Data Flag BIt
S - VT Size Bit

P. Michael Henderson, mike@michael-henderson.us

Pointing to the Start of the VT Payload

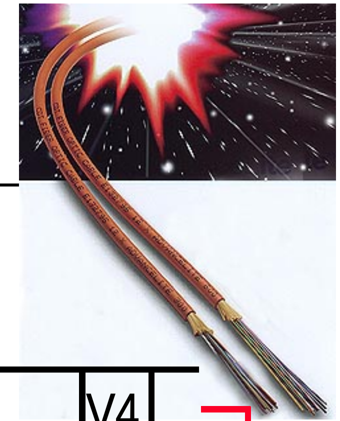

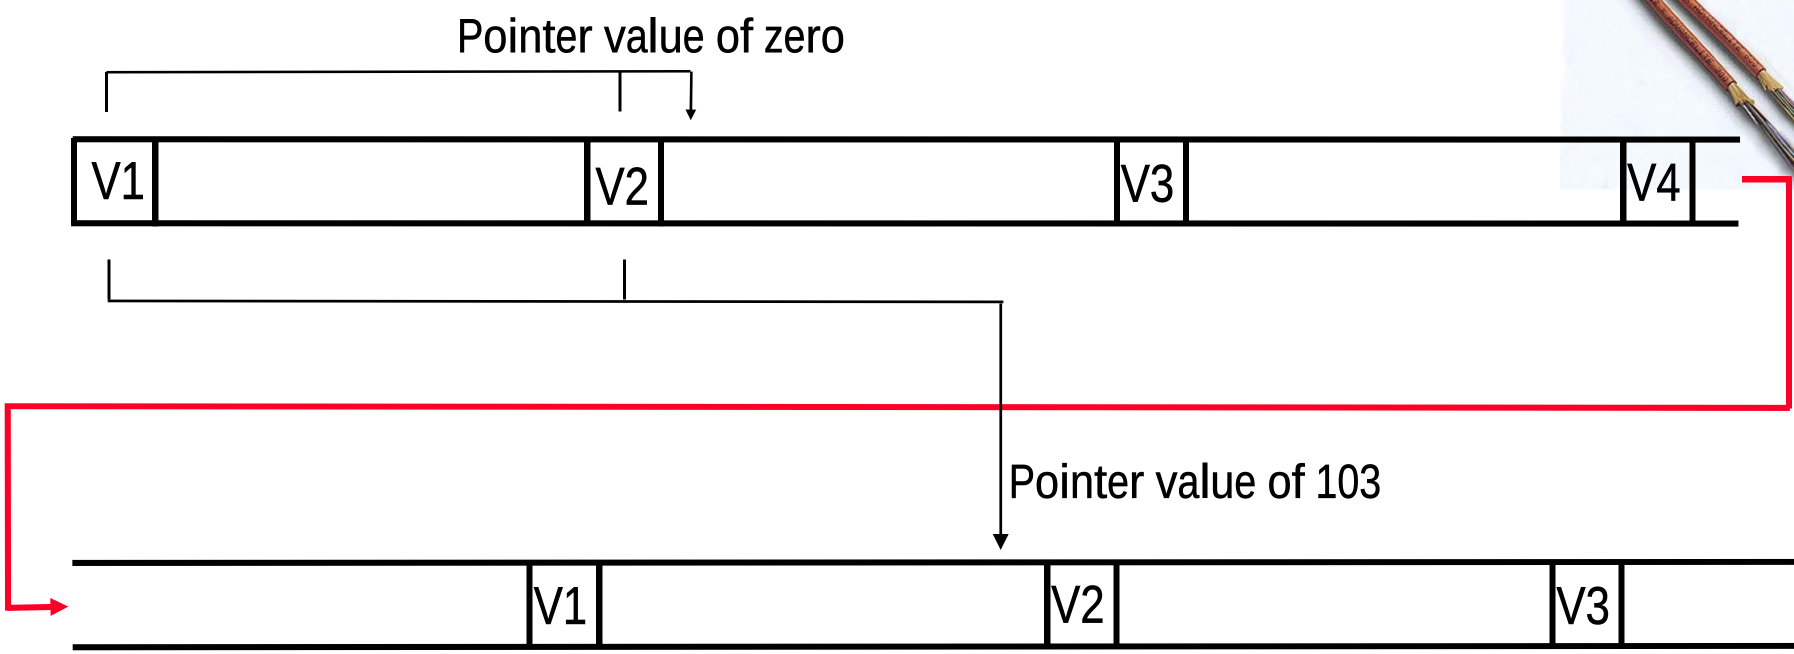

A VT-1.5 is made up of three columns of 9 octets or 27 octets. Four frames make a supreframe, or 108 octets. Once the V1, V2, V3, V4 octets are removed from the count, we have 104 octets remaining. Since the pointer counts from zero, the highest value of the pointer is 103.

P. Michael Henderson, mike@michael-henderson.us

The VT-1.5 Payload

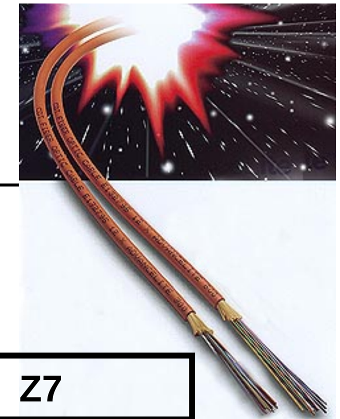

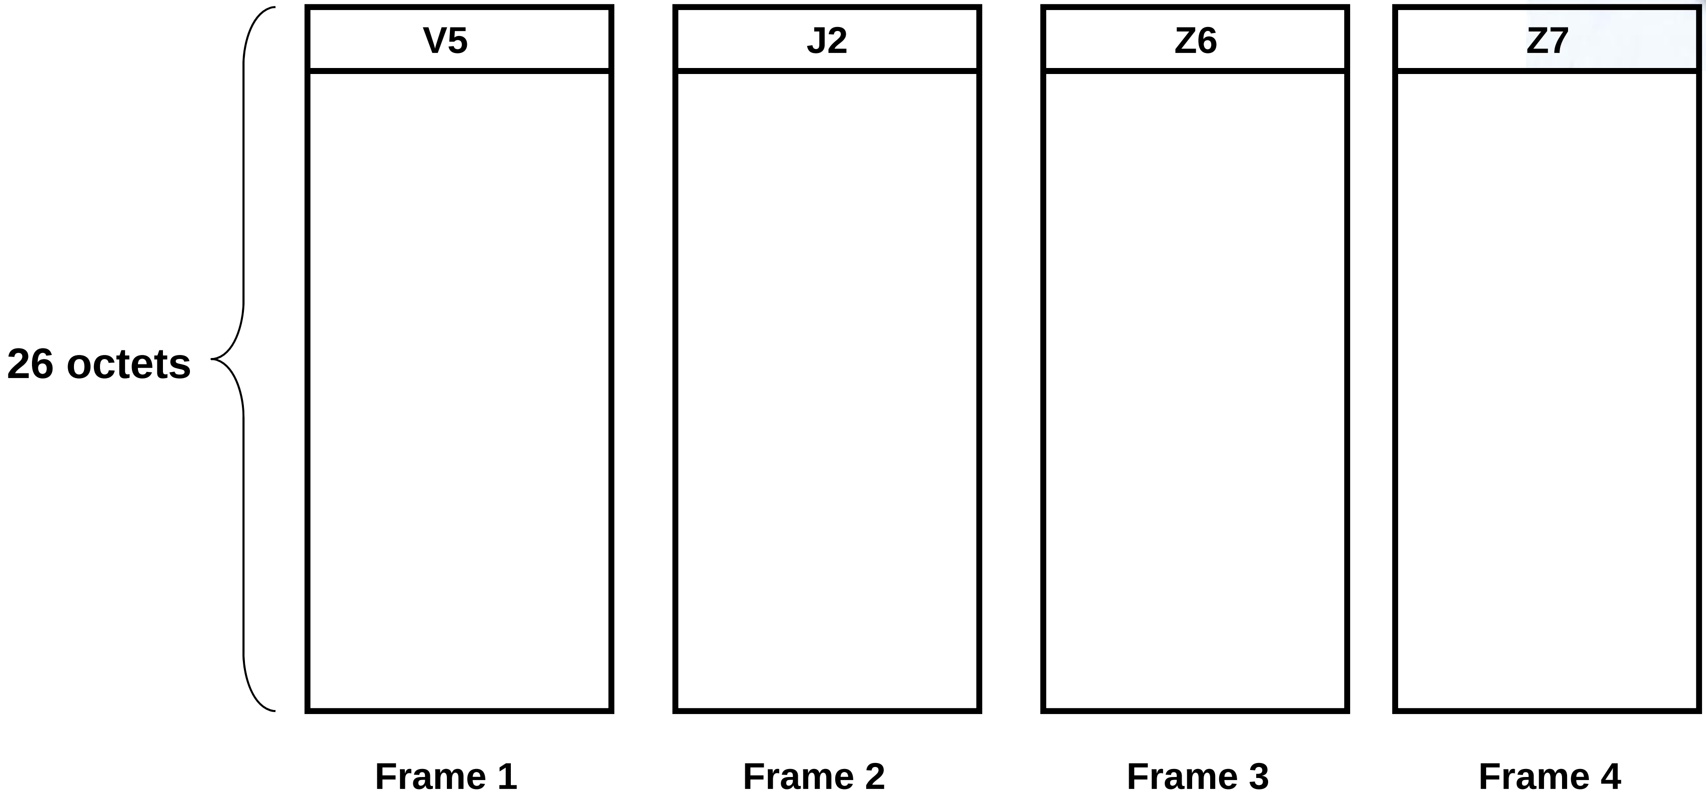

P. Michael Henderson, mike@michael-henderson.us

Byte Synchronous and Asynchronous VTs

- We're going to look at two ways that DS-1 traffic can be carried in a VT payload – byte synchronous and asynchronous.
- Byte synchronous preserves the location of the payload octets in a T1 frame (each speech sample).
  - Used primarily to transport channelized T1s which are carrying voice calls.
- Asynchronous simply transports the 1.544 Mbps stream without concern for which byte is which.
  - Used to carry T1s which are carrying data.
- Bits 5, 6, & 7 of V5 indicate which is being carried (010 = asynchronous, 100 = byte synchronous).

Byte Synchronous Mapping – VT-1.5

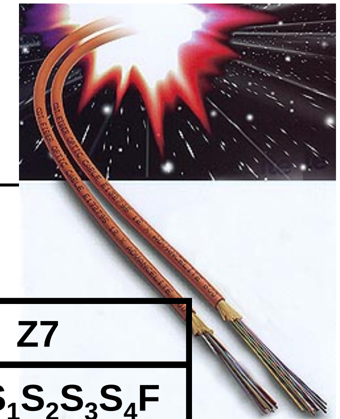

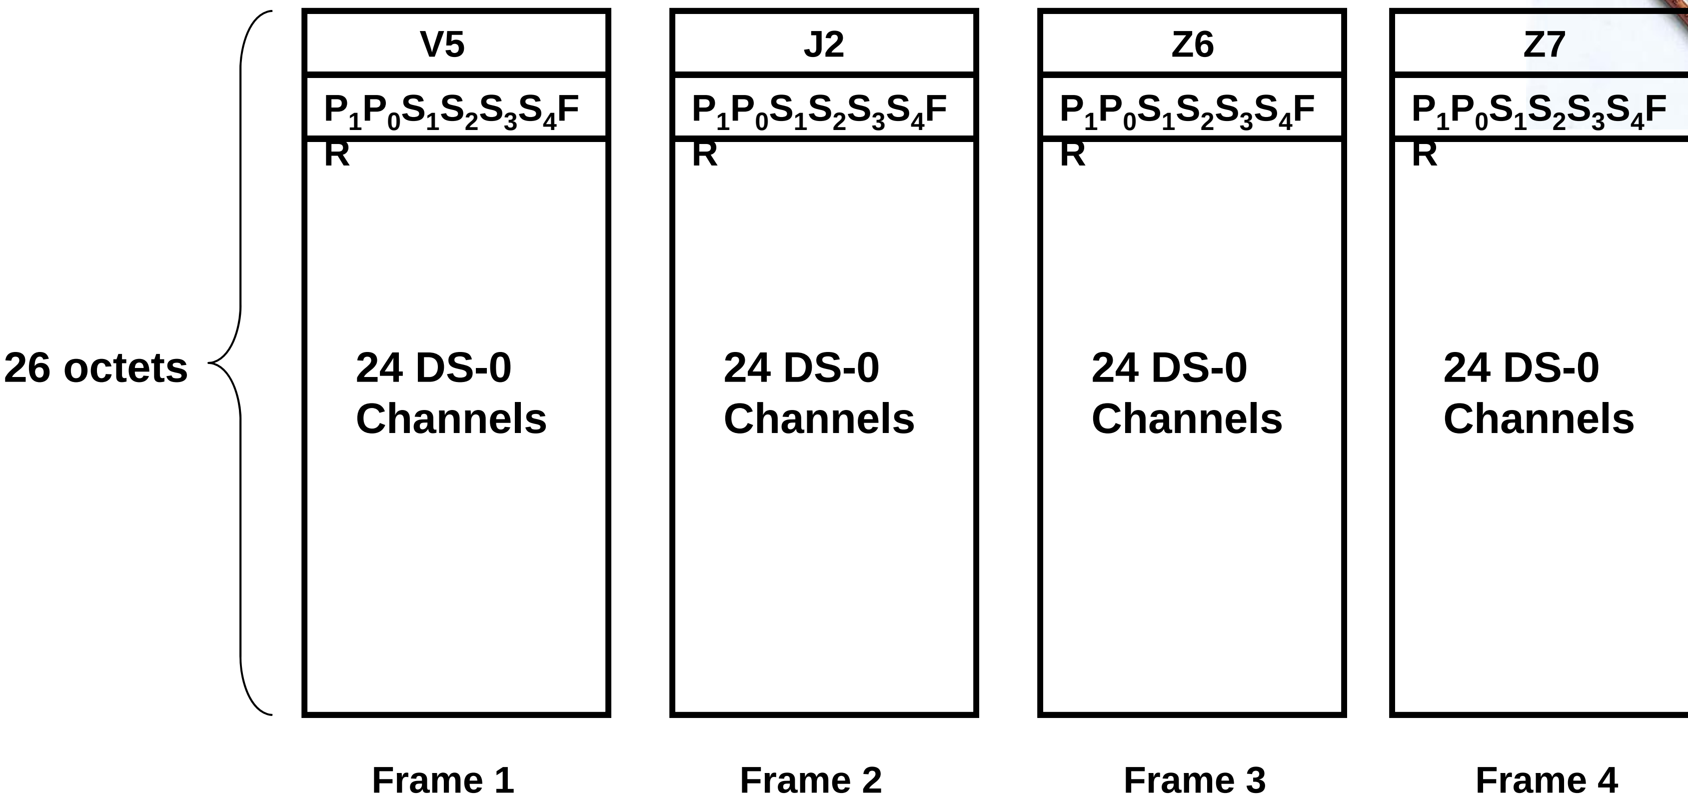

R = Fixed stuff bit

F = DS-1 framing bit

Sx = Signaling bits, A, B, C, D bits in a DS-1

Px = Phase bits, indicates which frame set for SF or ESF framing

Asynchronous Mapping – VT-1.5

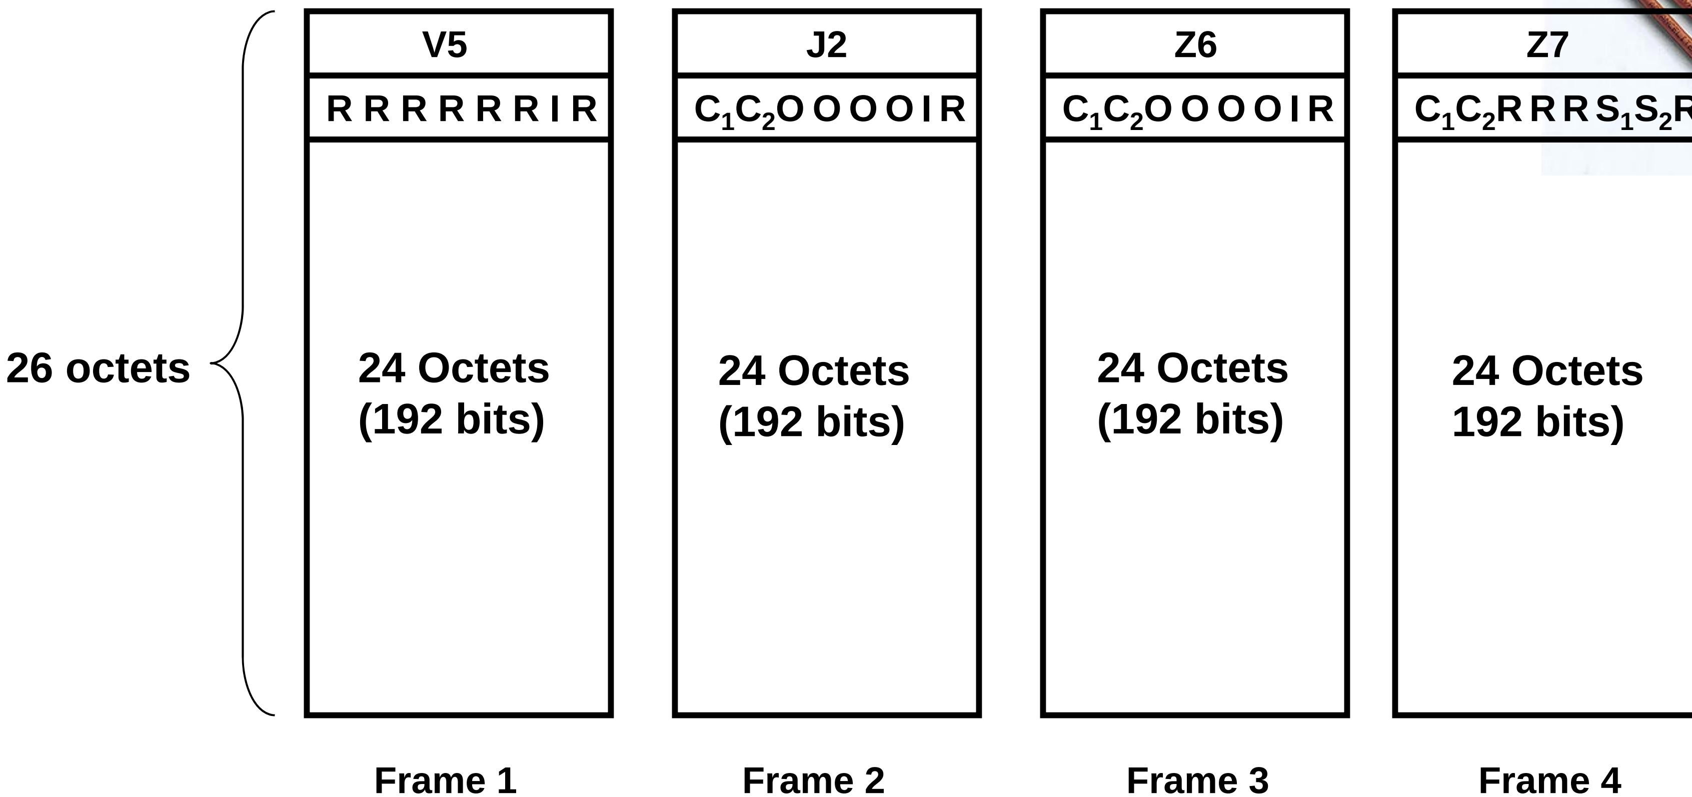

R = Stuff bits – no meaning

O = Future standardization – unused

I = Information bit. Makes the 193rd bit of a DS-1 frame

Cx = Controls use of the S bits.

Sx = Stuff bits. May or may not carry information bits, depending upon values in Cx

Finished with Virtual Tributaries!!!

P. Michael Henderson, mike@michael-henderson.us

Support for DS-3 signals

- A DS-3 signal at 44.736 Mbps takes the whole STS-1 SPE.
  - No virtual tributaries are used.
- A DS-3 can only be carried with asynchronous mapping.
  - There is no byte synchronous mapping for a DS-3.

P. Michael Henderson, mike@michael-henderson.us

Support for DS-3 signals

		28 Octets					28 Octets					28 Octets		
STS Path Overhead	R	R	C1	25 I	Fixed Stuff	R	C2	I	25 I	Fixed Stuff	R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I
	R	R	C1	25 I		R	C2	I	25 I		R	C3	I	25 I

Octets:

		8 bits
I	=	i i i i i i i i
R	=	r r r r r r r r
C1	=	r r c i i i i i
C2	=	c c r r r r r r
C3	=	c c r r o o r s

Bits:

i = information bit
r = fixed stuff bit
c = justification control bit
s = justification opportunity bit
o = overhead bit

P. Michael Henderson, mike@michael-henderson.us

Support for DS-3 signals

- There are 77 full information payload octets, or 616 information bits.
- The C1 octet has 5 information bits, making the total number of information bits per row 621.
- Nine of these per frame, times 8,000 frames per second gives 44.712 Mbps.
  - So how do we carry a 44.736 Mbps signal in 44.712 Mbps?
- The answer is in the “s” bit in C3. Sometimes it carries an information bit.
  - If it carried an information bit every frame, the data rate would be 44.784 Mbps.

P. Michael Henderson, mike@michael-henderson.us

Support for ATM, POS, and GFP

- Asynchronous Transfer Mode (ATM), Packet over SONET (POS) and Generic Framing Procedure (GFP) are simply mapped into the SPE as a serial data stream, octet aligned with the SONET/SDH octets.
- When traffic is mapped into an STS-1 payload, columns 30 and 59 are not used for payload (fixed stuff).
- ATM, POS and GFP can be mapped into higher speed concatenated payloads.
  - Mapping is simply done by putting the ATM, POS or GFP octets into the concatenated SPE. SONET/SDH does not examine the payload octets in any way.

Generic Framing Procedure

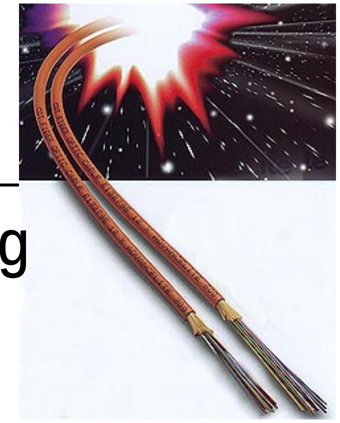

- A way of framing variable length data without a framing character.
  - To overcome the problem of “shielding” in POS.
- GFP is being defined in the ANSI T1X1.5 committee.
  - Can be thought of a variable length ATM type of framing.
  - Header has header check octets and payload length. GFP frame delineation is similar to finding ATM header.

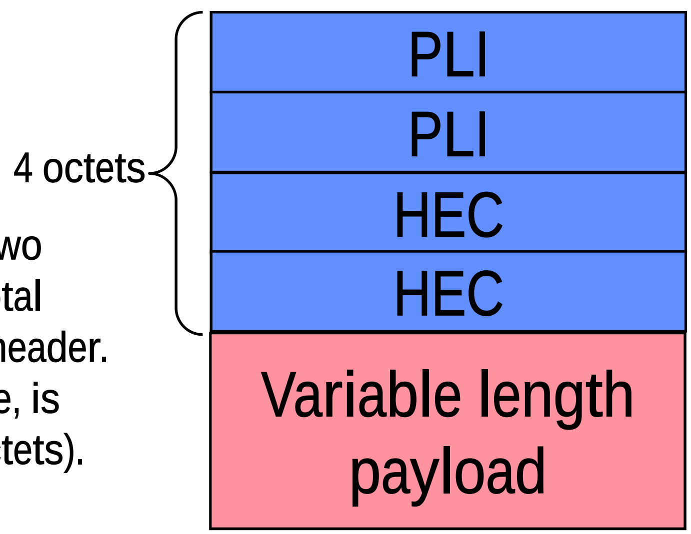

Length specified in the two octets of the PLI is for total frame length, including header. Max frame size, therefore, is (64K-1) octets (65,535 octets).

PLI = Payload length indicator
HEC = Header error control (CRC-16)

P. Michael Henderson, mike@michael-henderson.us

Automatic Protection Switching

P. Michael Henderson, mike@michael-henderson.us

A Unidirectional Ring

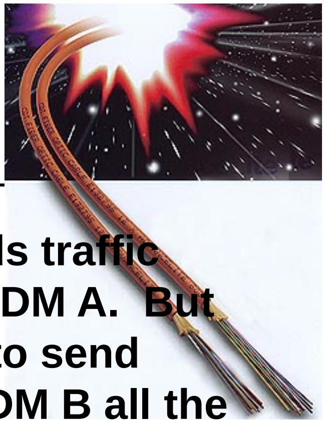

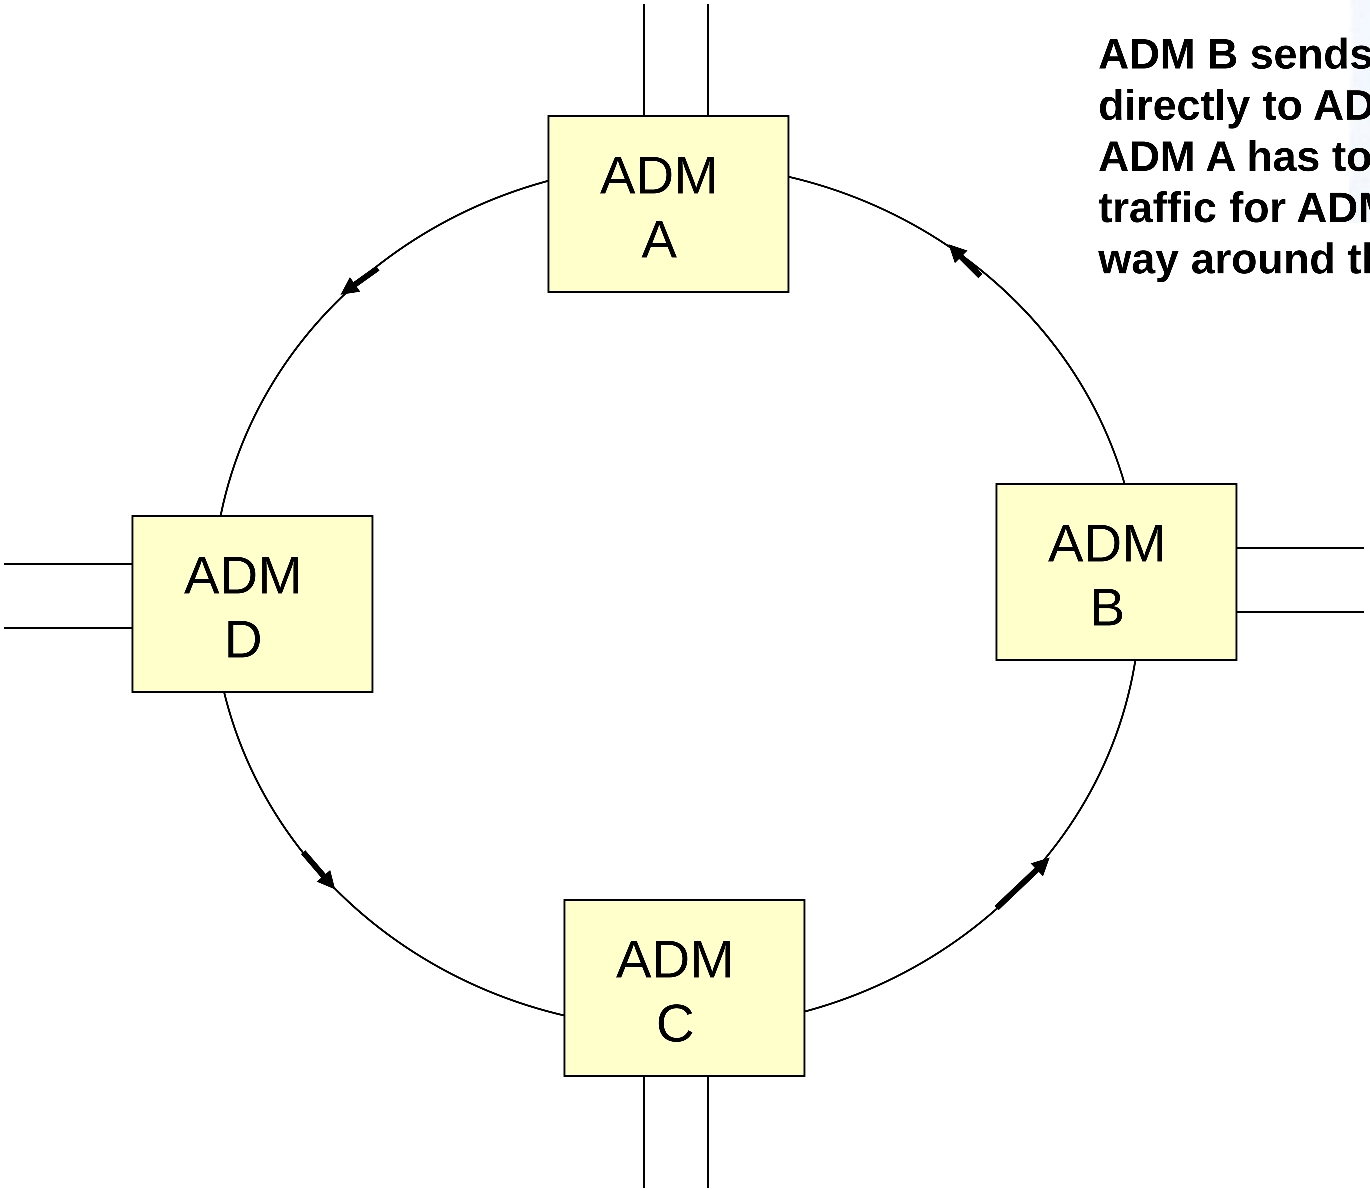

ADM B sends traffic directly to ADM A. But ADM A has to send traffic for ADM B all the way around the ring.

P. Michael Henderson, mike@michael-henderson.us

Backup for a Unidirectional Ring

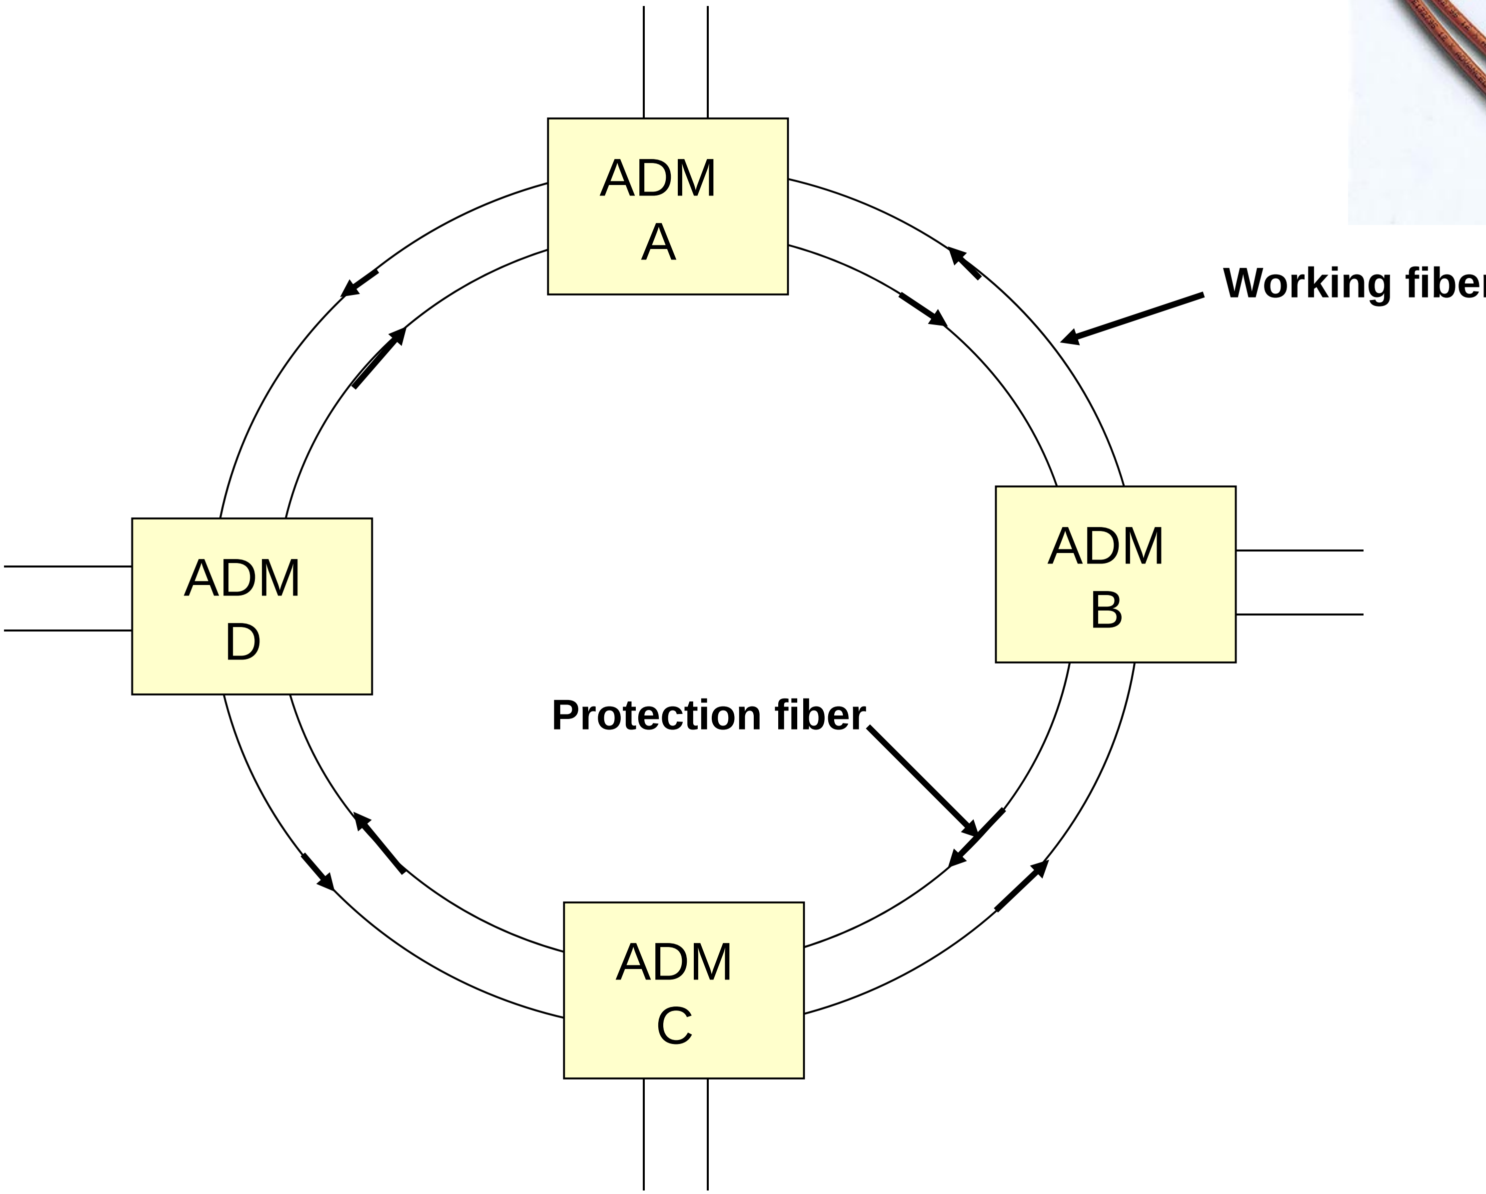

A Bi-directional Ring

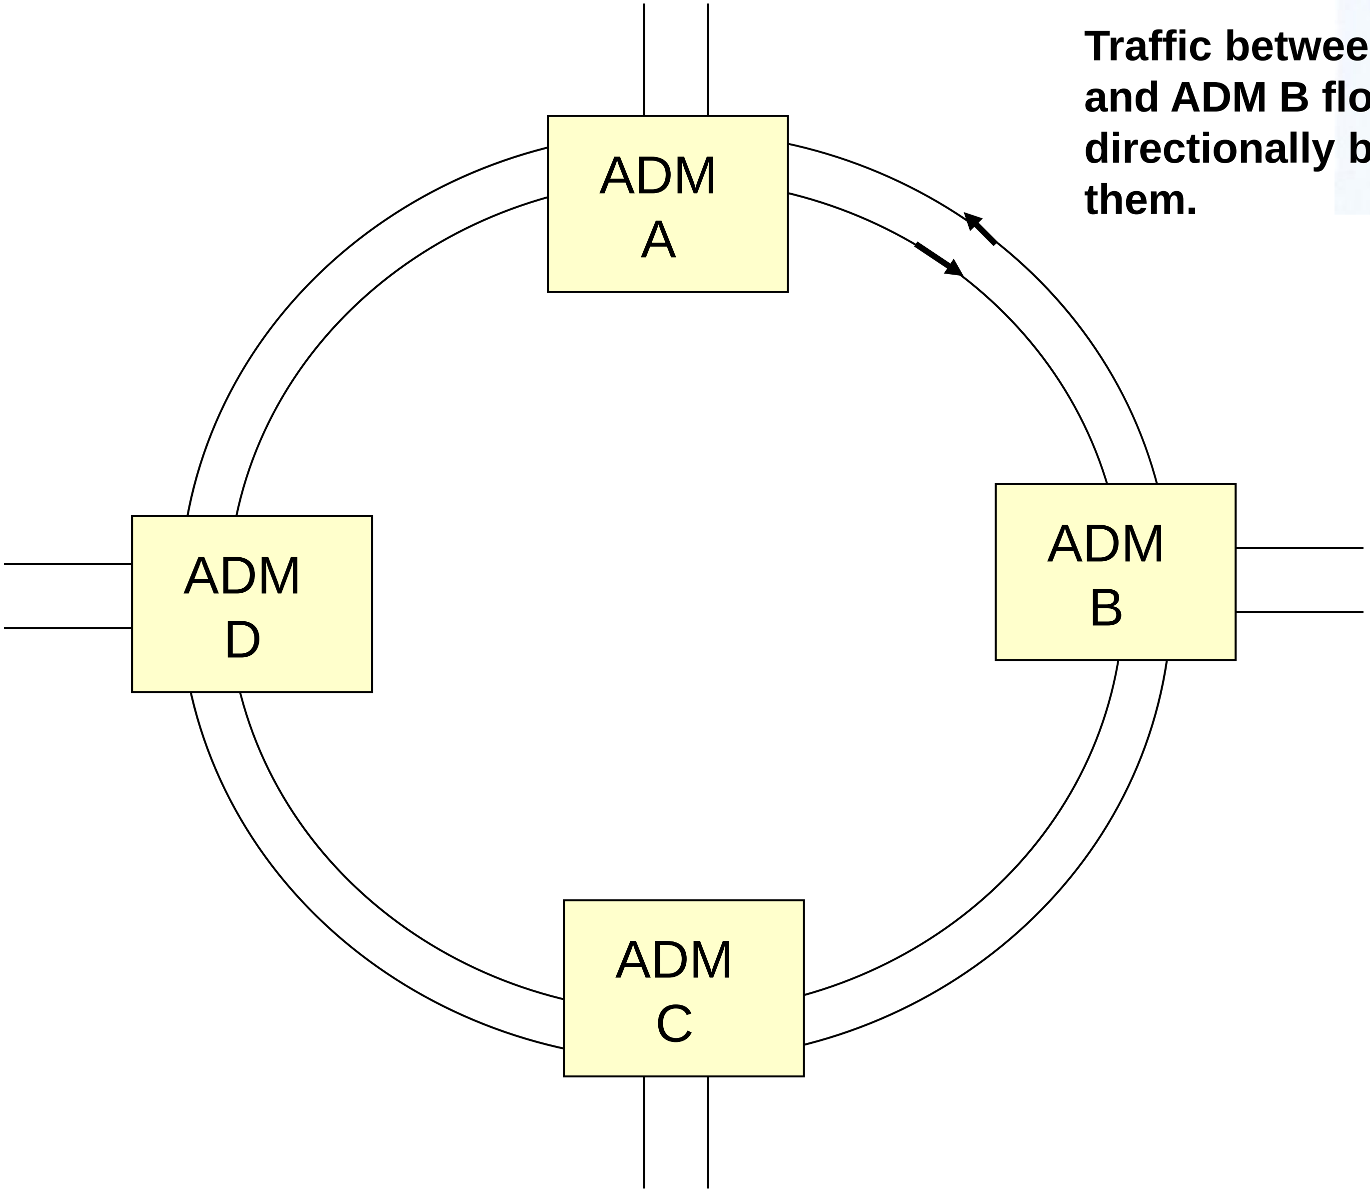

Traffic between ADM A and ADM B flows bi-directionally between them.

Backup for a Bi-directional Ring

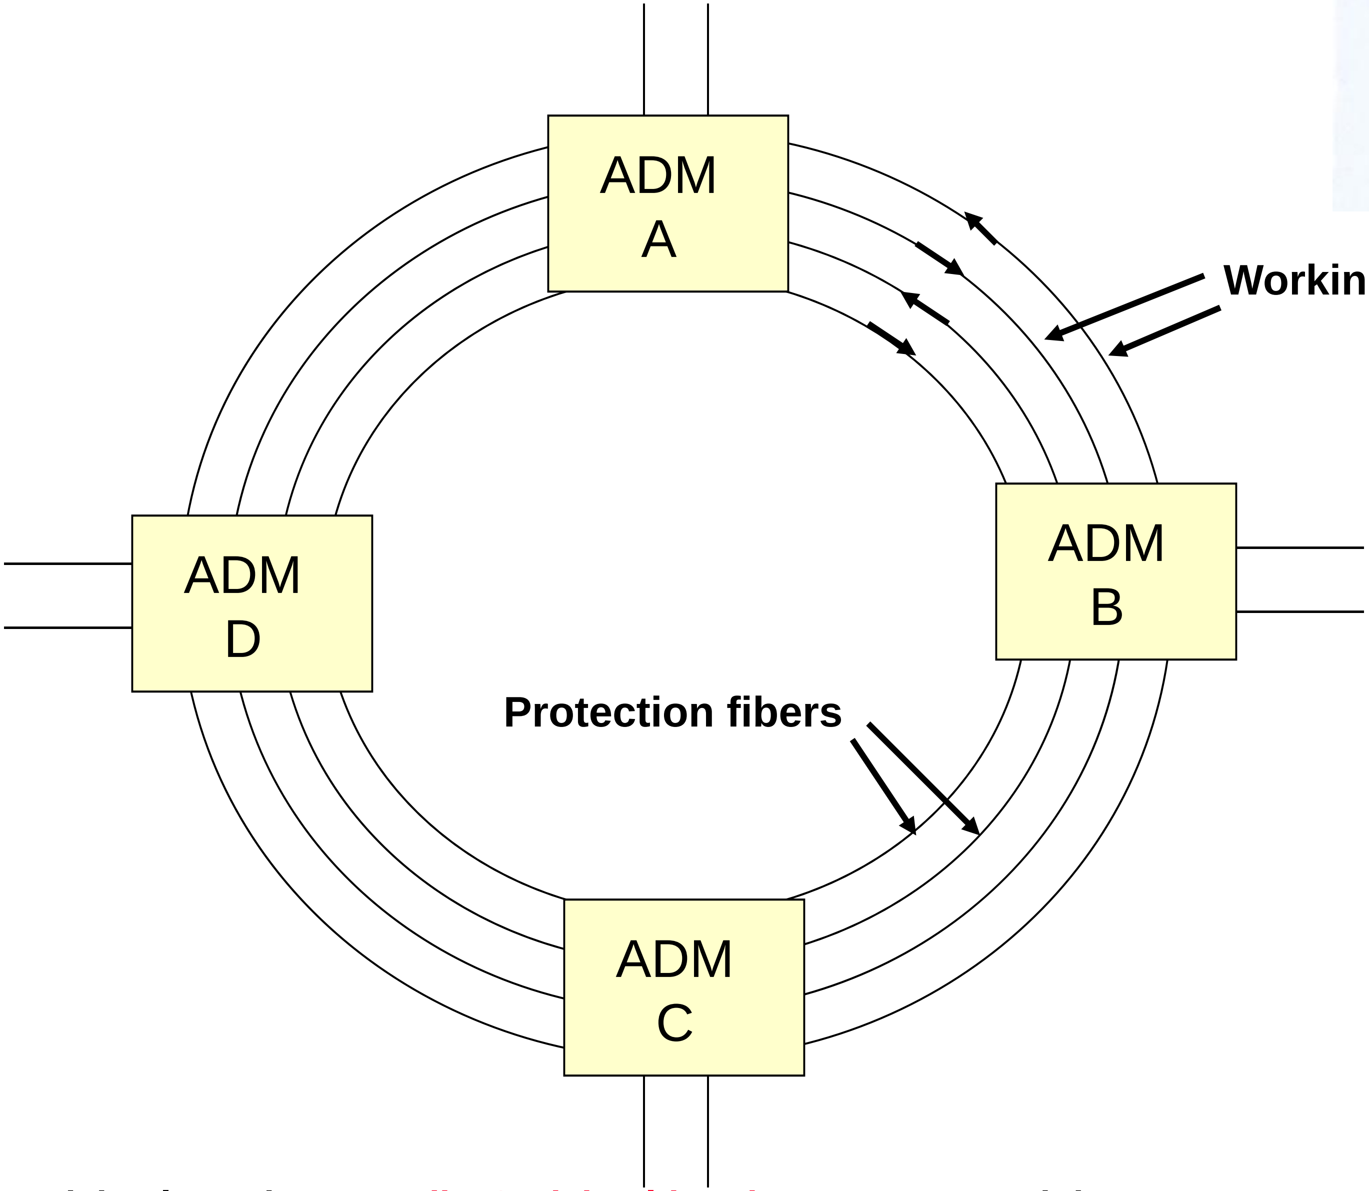

Path Switching

- Today, path switching is only used on unidirectional rings – hence the name Unidirectional Path Switched Rings (UPSR).
- What does Path Switching mean?
  – At the exit node, both fibers are monitored and the path traffic extracted.
  – Based on several factors, especially error rate, the “best” traffic is selected to be handed off to the customer.
- To a large degree, path switching can be done completely by the receiver, without coordination with the sender.
- But, path switching requires two sets of electronics to extract the path information.

Path Switching – How does it work?

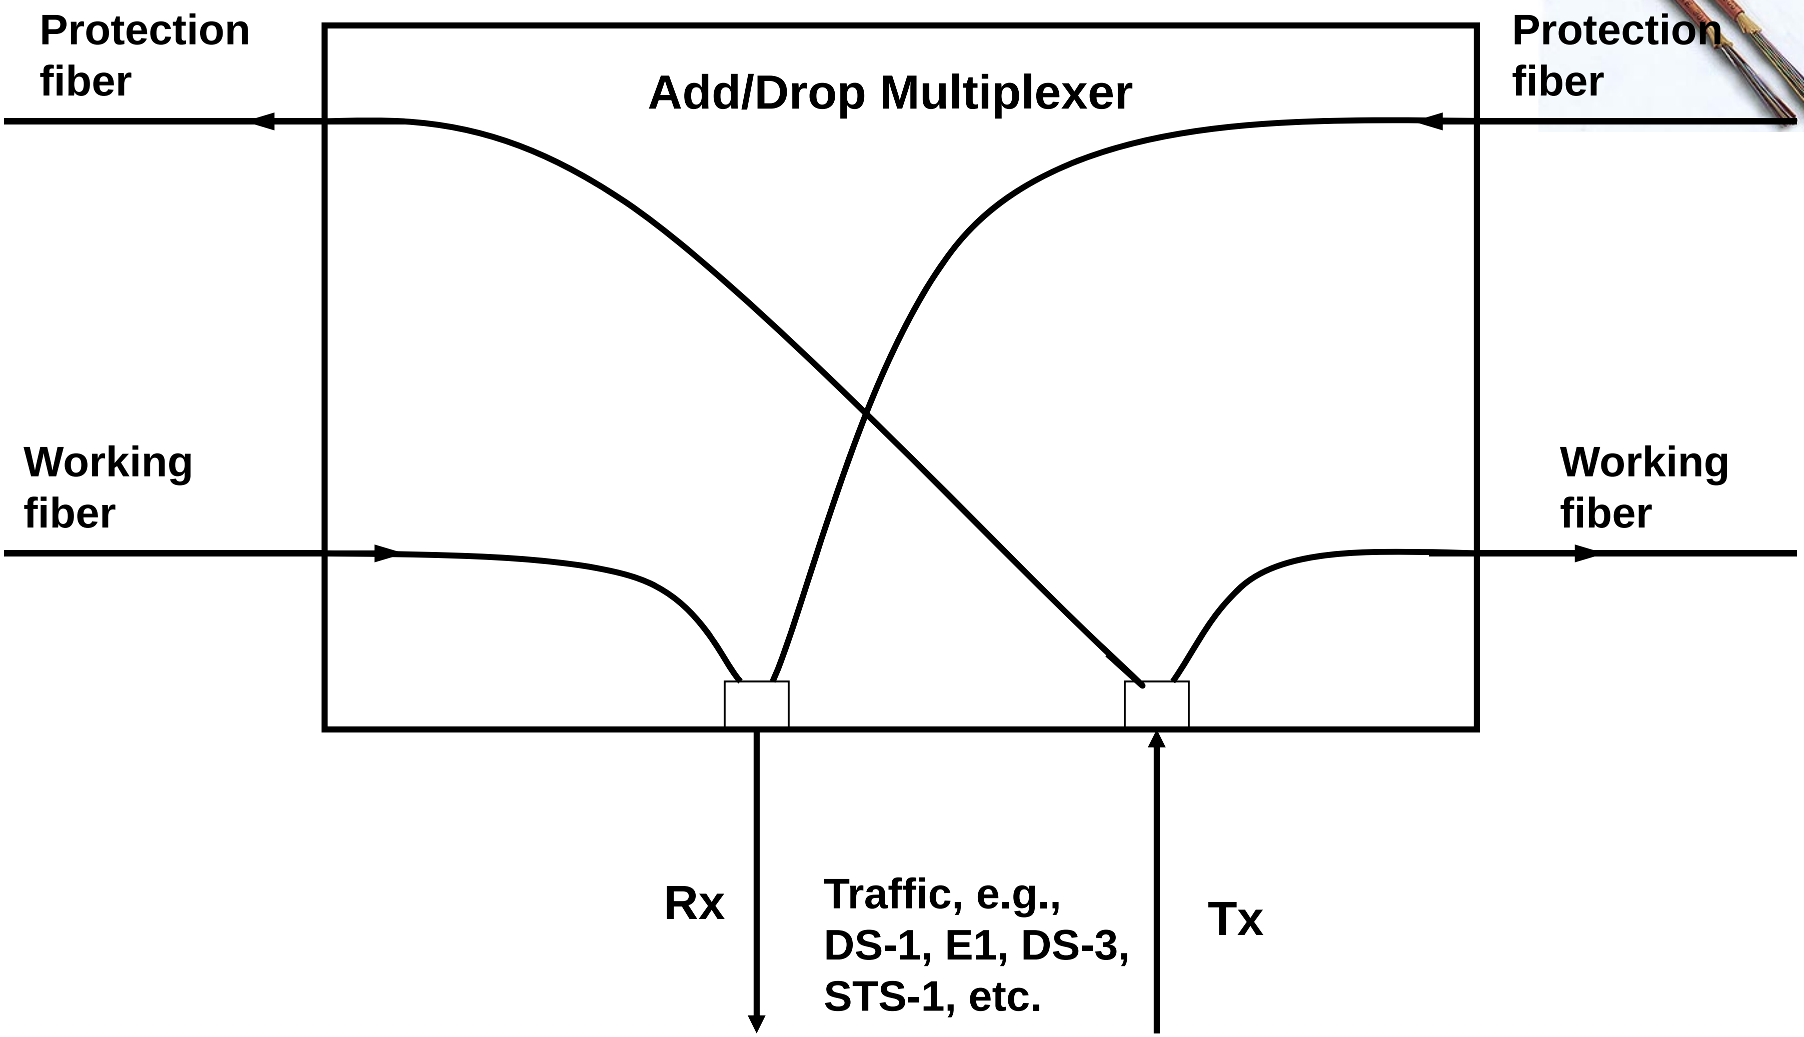

P. Michael Henderson, mike@michael-henderson.us

Line Switching

- Line Switching is done on bi-directional rings, either two fiber or four fiber. Thus the name Bi-directional Line Switched Ring (BLSR).
- The total capacity of the fibers must be twice the carried traffic. For four fiber systems, this means two fibers are dedicated for protection.
- What does Line Switching mean?
  - Two adjacent nodes monitor the traffic between them.
  - If one detects “failure” on a fiber, it signals the other.
  - The two nodes coordinate switching to the protection fiber(s).
- The K1, K2 octets in the transport overhead are used for this signaling.

Line Switching – How is it done?

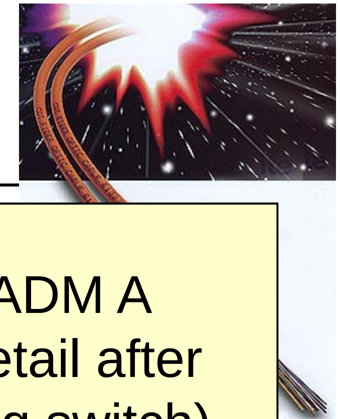

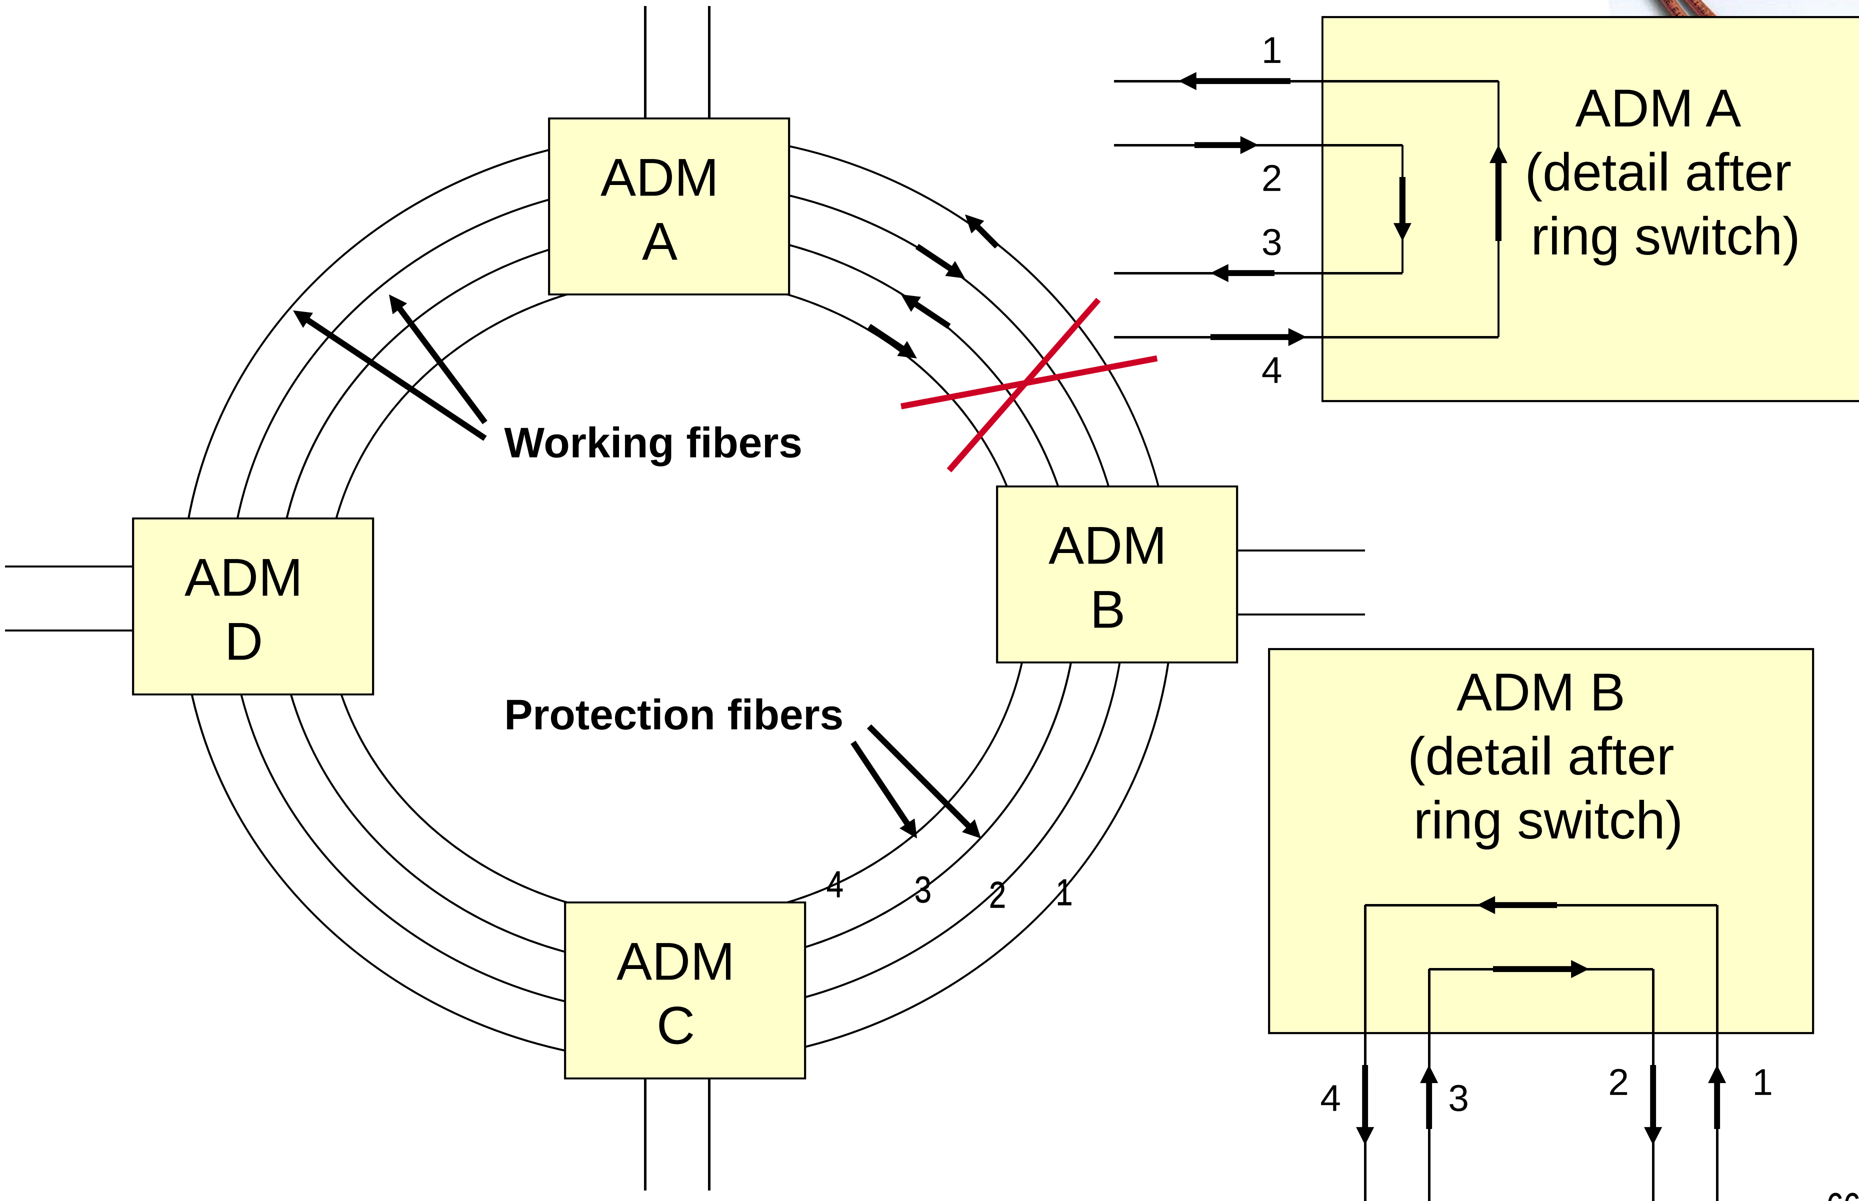

Summary

- SONET/SDH is a complex subject, with many layers.
- This presentation has attempted to take you from the SONET/SDH frame to the virtual tributaries.
- There's a lot more to SONET/SDH than was covered in this presentation – many things were ignored to avoid hopeless complexity.
- I hope this introduction will stimulate some of you to study the subject in more detail, perhaps by reading the standards.

P. Michael Henderson, mike@michael-henderson.us

Questions?

Backup Material

P. Michael Henderson, mike@michael-henderson.us

Carrying an E1 in a Virtual Tributary

P. Michael Henderson, mike@michael-henderson.us

Virtual Tributaries – The VT-2

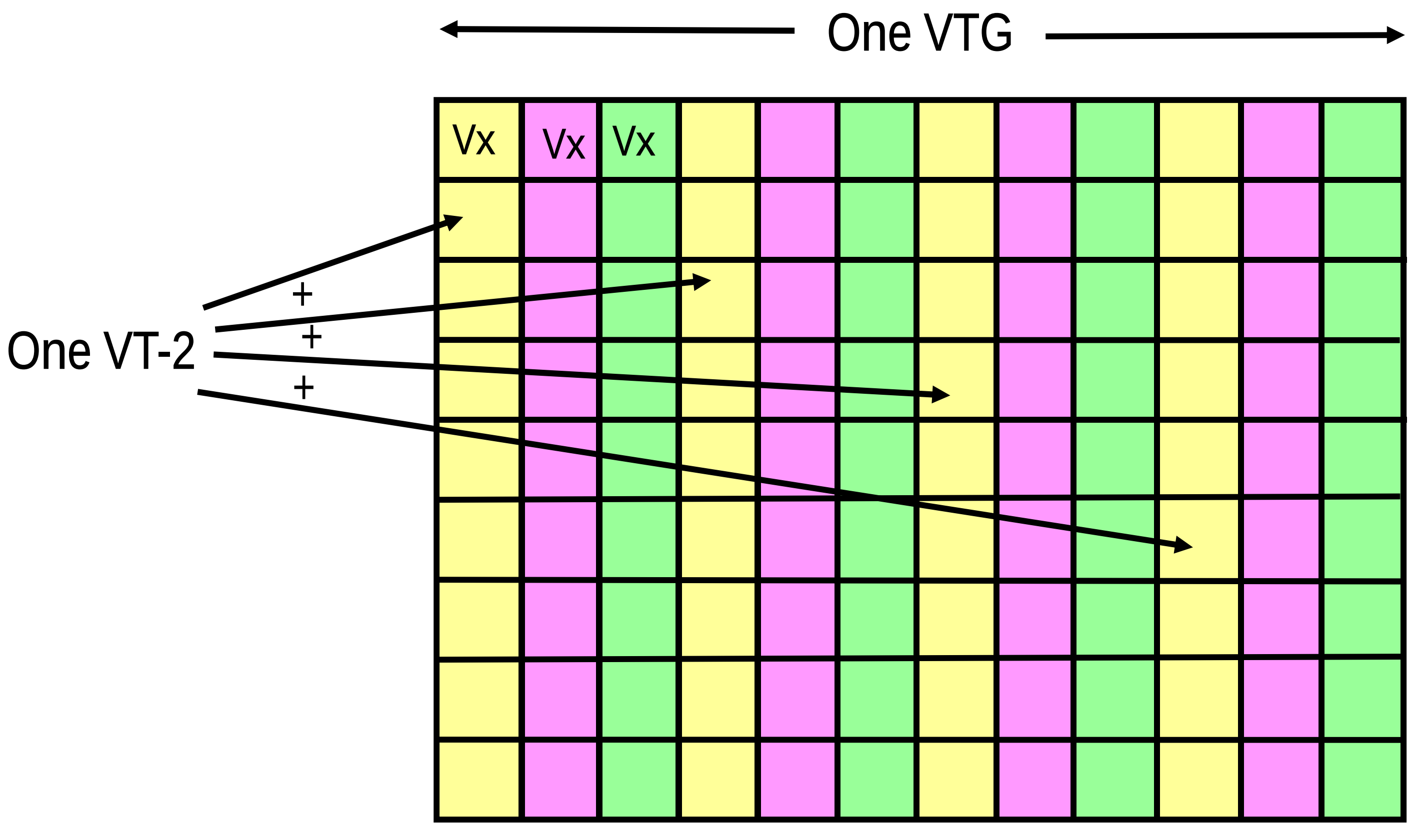

The three VT-2s are interleaved into the VTG. Note the colors – there are four of each color, with each color indicating one VT-2 (for three inside the VTG. Each VT is 36 octets, with one octet taken for the Vx octet.

P. Michael Henderson, mike@michael-henderson.us

Superframes

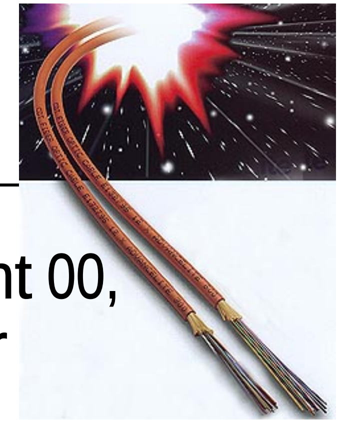

- Remember that the last two bits of H4 in the POH count 00, 01, 10, 11, 00, etc. This produces a superframe of four frames.
- The frame after the SPE with an H4 value of 00 will have the first octet in the VT identified as V1. The one after value 01 will be identified as V2, etc.
- V1 and V2 form a pointer, exactly like the H1, H2 octets.
  - Bits 5 and 6 are used to indicate the type of VT (unused in H1, H2) (VT-6=00, VT-3=01, VT-2=10, VT-1.5=11)
- V3 is the negative stuff opportunity, exactly like H3.
- V4 is reserved for future standardization.

V1	V2	V3	V4

The V1, V2 Pointers

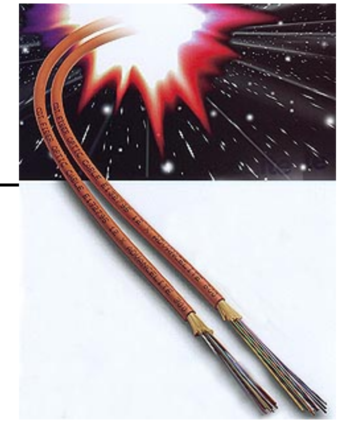

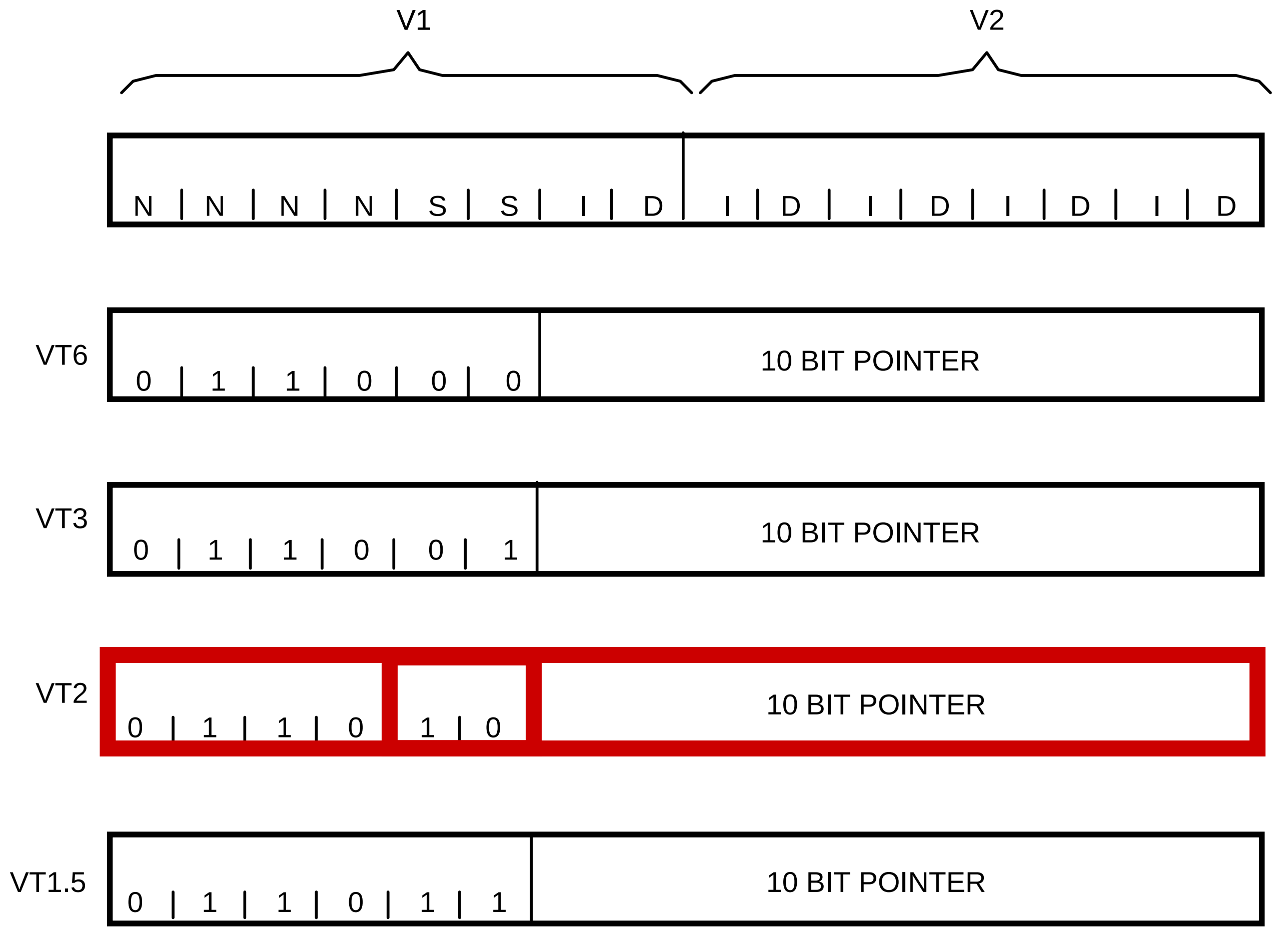

NEW DATA FLAG – INVERT 4 N BITS
NEGATIVE STUFF – INVERT 5 D BITS
POSITIVE STUFF – INVERT 5 I BITS

I – Increment Bit
D – Decrement Bit
N – New Data Flag BIt
S - VT Size Bit

Pointing to the Start of the VT Payload

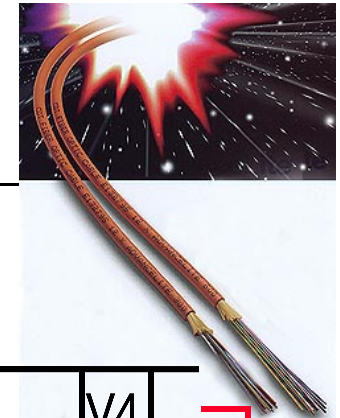

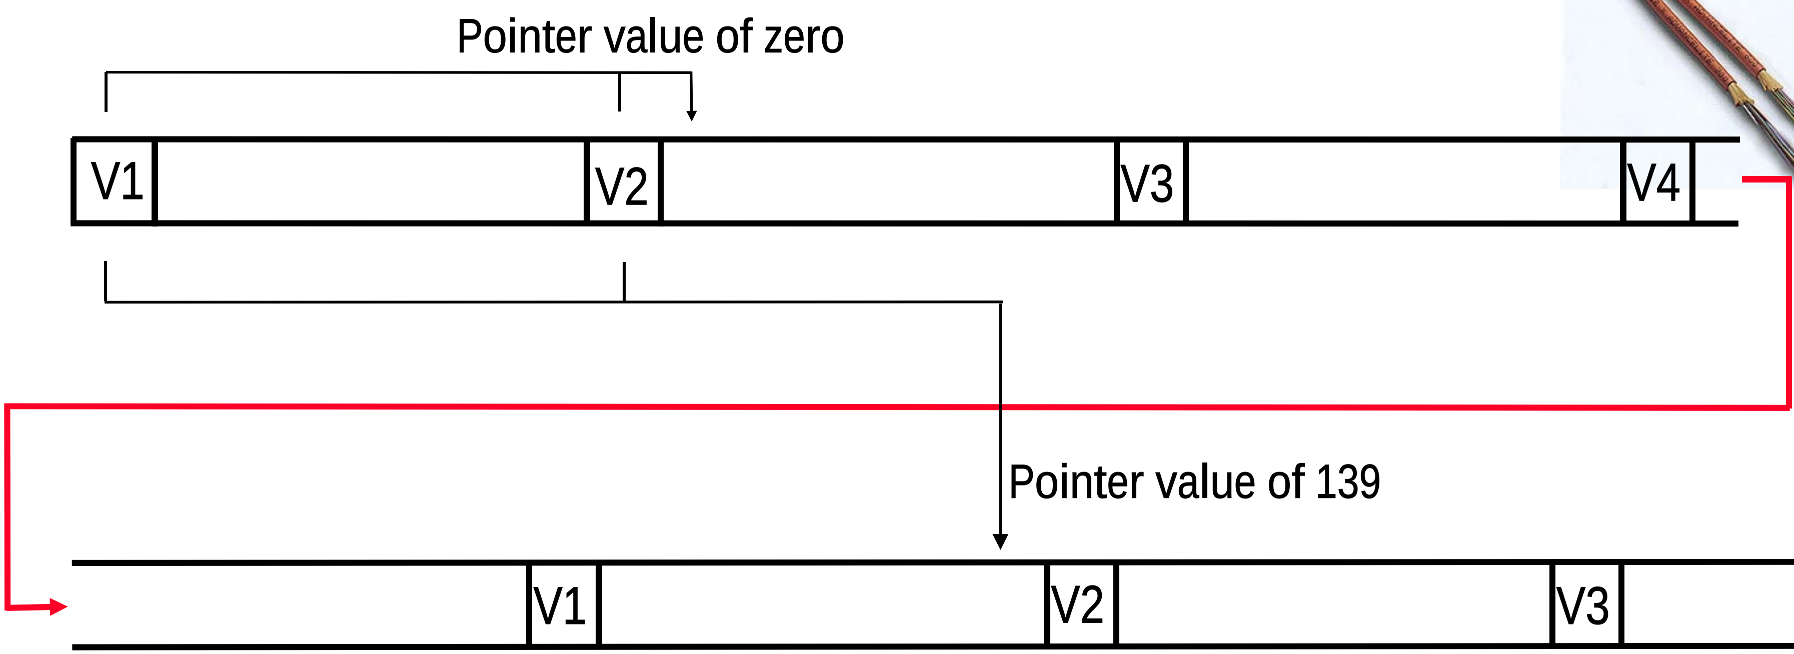

A VT-2 is made up of four columns of 9 octets or 36 octets. Four frames make a supreframe, or 144 octets. Once the V1, V2, V3, V4 octets are removed from the count, we have 140 octets remaining. Since the pointer counts from zero, the highest value of the pointer is 139.

The VT-2 Payload

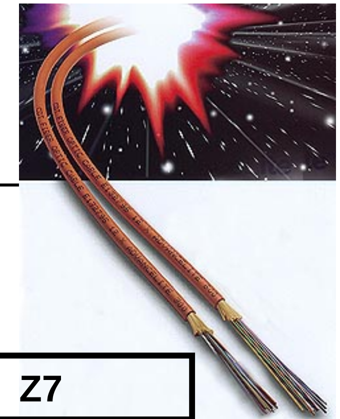

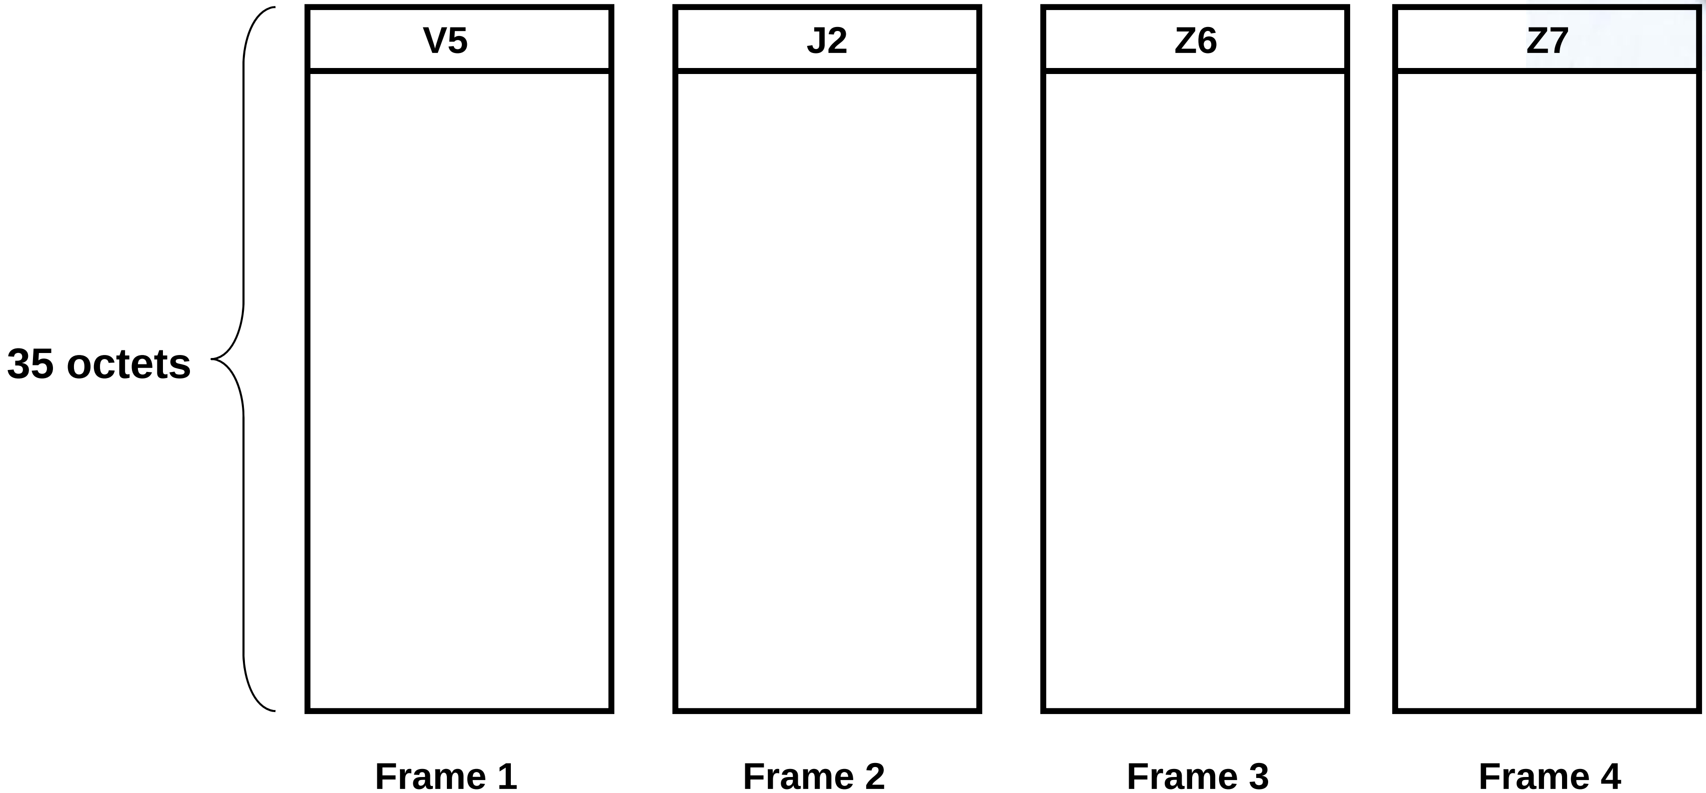

P. Michael Henderson, mike@michael-henderson.us

Byte Synchronous and Asynchronous VTs

- We're going to look at two ways that E-1 traffic can be carried in a VT payload – byte synchronous and asynchronous.
- Byte synchronous preserves the location of the payload octets in an E1 frame (each speech sample).
  - Used primarily to transport channelized E1s which are carrying voice calls.
- Asynchronous simply transports the 2.048 Mbps stream without concern for which byte is which.
  - Used to carry E1s which are carrying data.
- Bits 5, 6, & 7 of V5 indicate which is being carried (010 = asynchronous, 100 = byte synchronous).

Byte Synchronous Mapping – 30 Channels

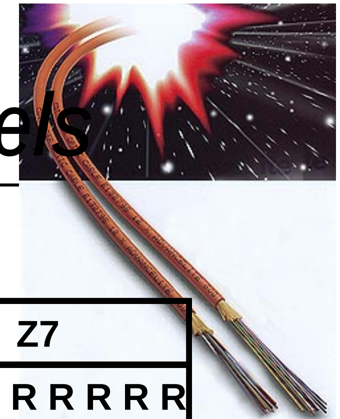

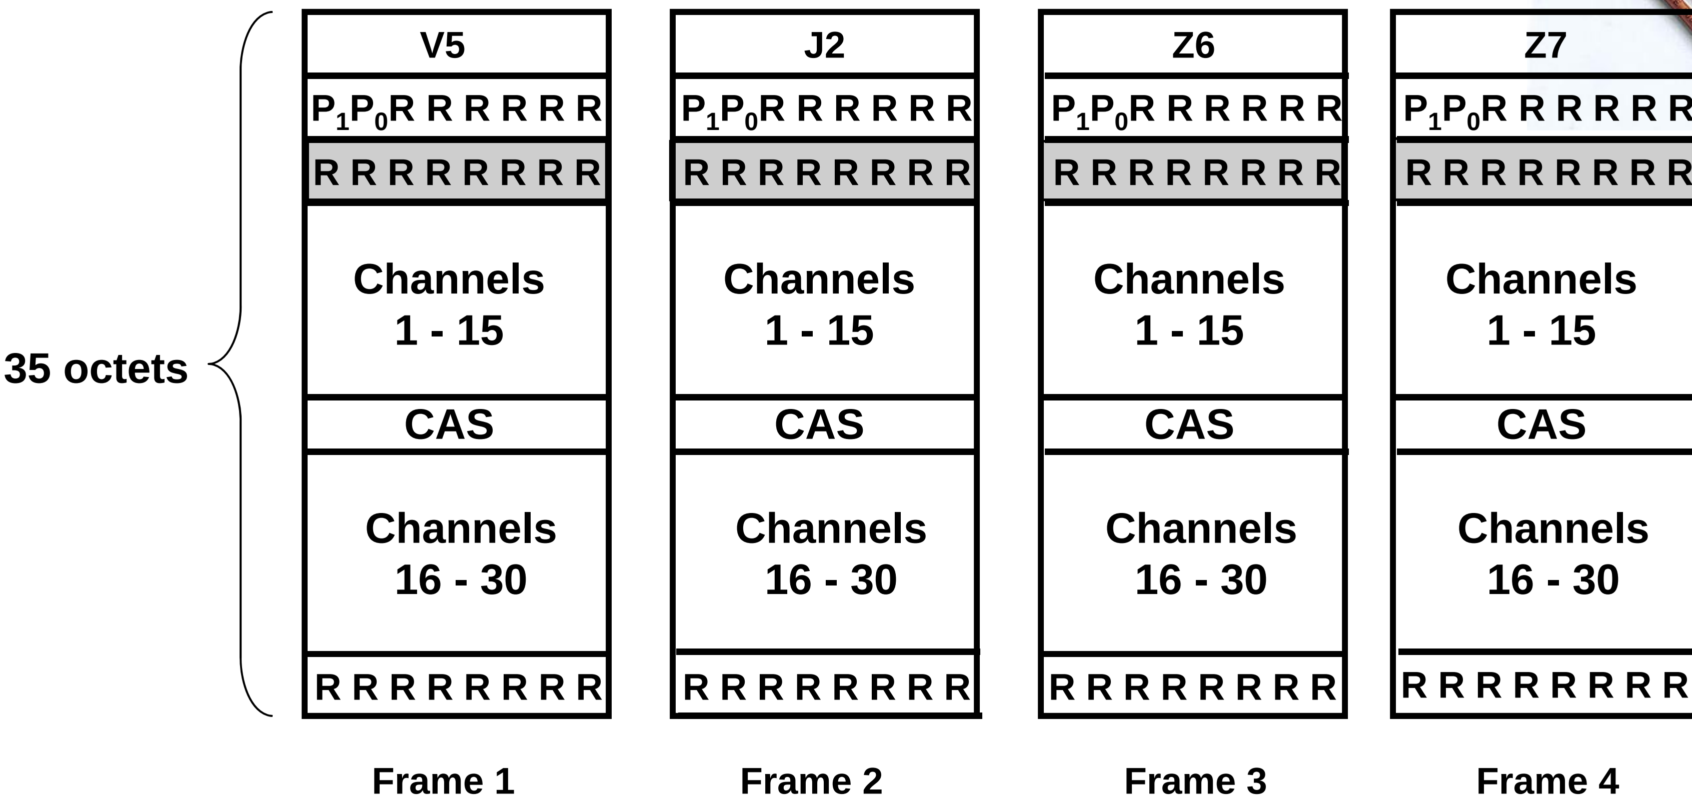

R = Fixed stuff bit
R = may be used for channel 0, if required
Px = Phase bits, indicates which channel the CAS bits apply to

Byte Synchronous Mapping – 31 Channels

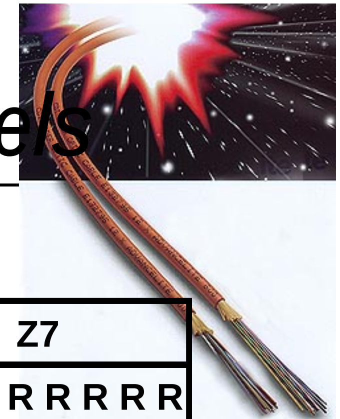

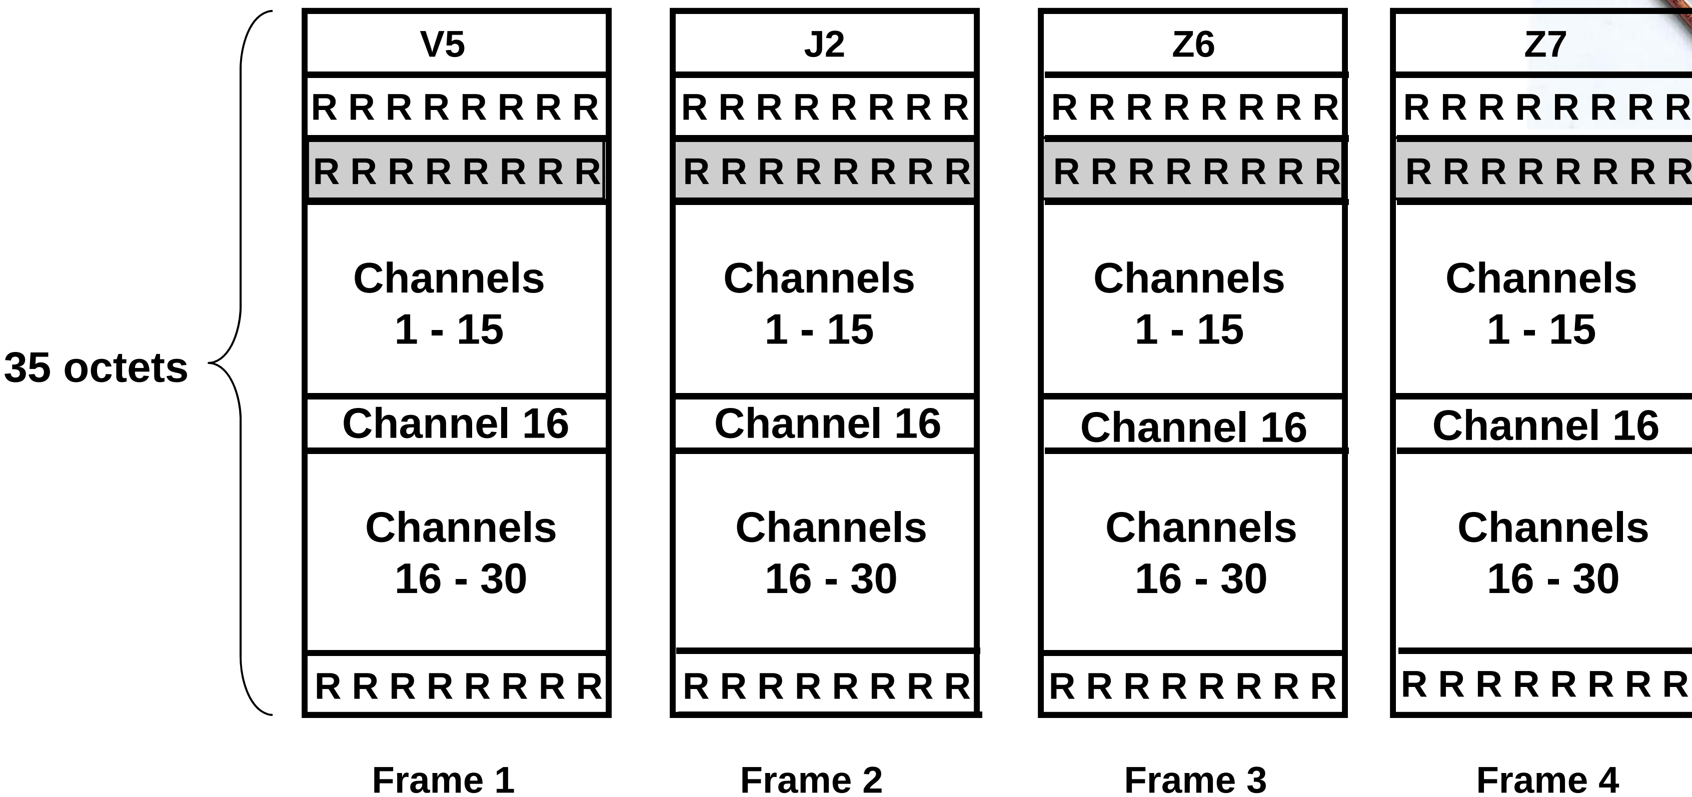

R = Fixed stuff bit

R = Indicates may be used for channel 0

P. Michael Henderson, mike@michael-henderson.us

Asynchronous Mapping – VT-2

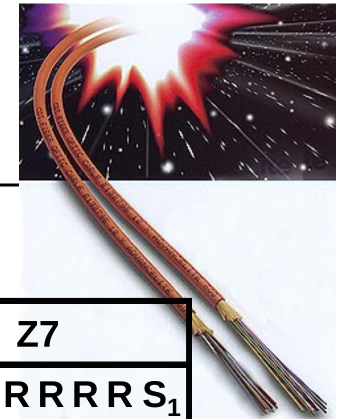

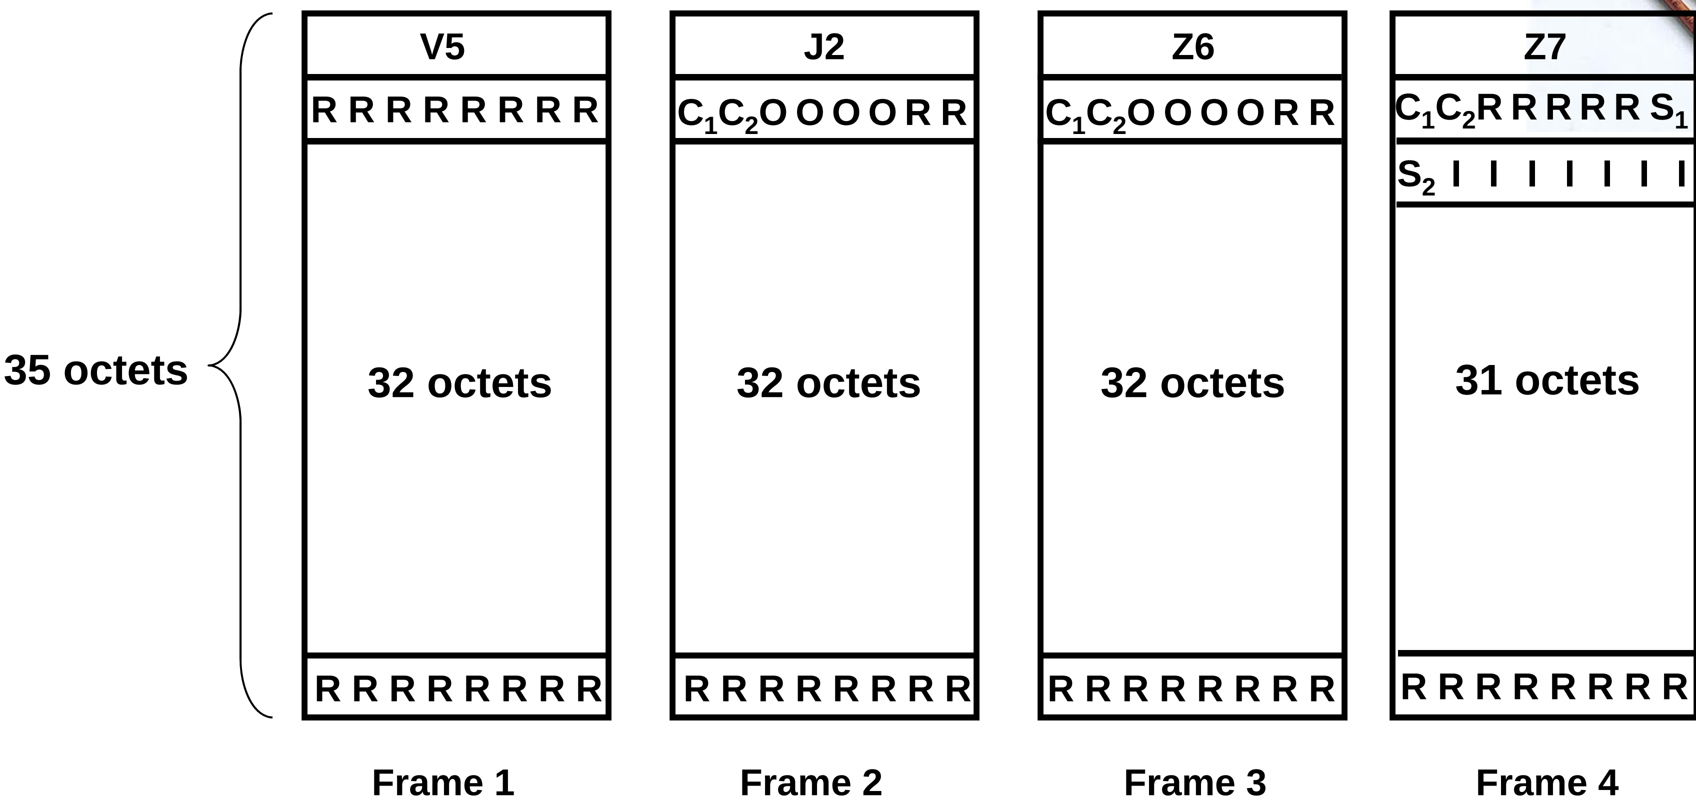

I = Information bit (Data)
R = Fixed stuff bit
O = Overhead – same as R

Cx = Stuff control
Sx = Stuff opportunity

Synchronous Digital Hierarchy (SDH)

P. Michael Henderson, mike@michael-henderson.us

SDH Terminology

- SDH terminology is different from SONET and several new concepts are introduced.
- One primary difficulty is that SDH starts with the equivalent of an STS-3 signal but has to map the same PDH signals.
- SDH makes much more use of the “group” terminology.
  - Compare to Virtual Tributary Groups (VTGs) in SONET.
- Let’s look at the STM-1 frame and the ways that PDH signals get mapped.

Administrative Unit Groups and AUs

- An Administrative Unit Group (AUG) consists of the payload portion of the frame, plus row four of the overhead.
- SONET has no equivalent of the AUG.
- The AUG can contain either:
  - Three Administrative Unit level 3 (AU-3), or
  - One AU-4.
- We'll look at the AU-3 first and then at the AU-4.

The STM-1 frame (not to scale)

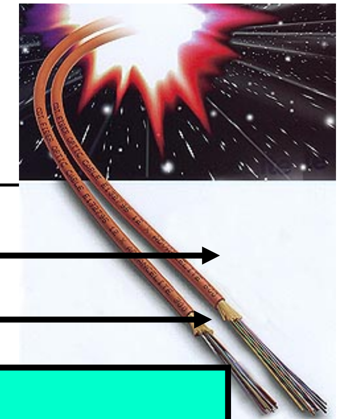

270 Columns

261 Columns

A1	A1	A1	A2	A2	A2	J0	Z0	Z0	
B1	X	X	E1	X	X	F1	X	X	
D1	X	X	D2	X	X	D3	X	X	
H1	H1	H1	H2	H2	H2	H3	H3	H3	
B2	B2	B2	K1	X	X	K2	X	X	
D4	X	X	D5	X	X	D6	X	X	
D7	X	X	D8	X	X	D9	X	X	
D10	X	X	D11	X	X	D12	X	X	
S1	Z1	Z1	M0/1	Z2	M2	E2	X	X	

Section overhead

Administrative Unit Group (AUG)

Administrative Unit - 3

- The AUG is larger than needed to carry PDH traffic. Plus, some technique must be provided to maintain some compatibility with SONET STS-1.
- Three AU-3s fit inside the STM-1 AUG.
- Begins to look a lot like an STS-3 carrying three STS-1 signals.

Mapping of three AU-3s to an AUG

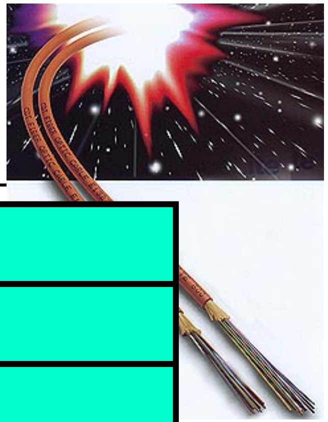

A1	A1	A1	A2	A2	A2	J0	Z0	Z0
B1	X	X	E1	X	X	F1	X	X
D1	X	X	D2	X	X	D3	X	X
B2	B2	B2	K1	X	X	K2	X	X
D4	X	X	D5	X	X	D6	X	X
D7	X	X	D8	X	X	D9	X	X
D10	X	X	D11	X	X	D12	X	X
S1	Z1	Z1	M0/1	Z2	M2	E2	X	X

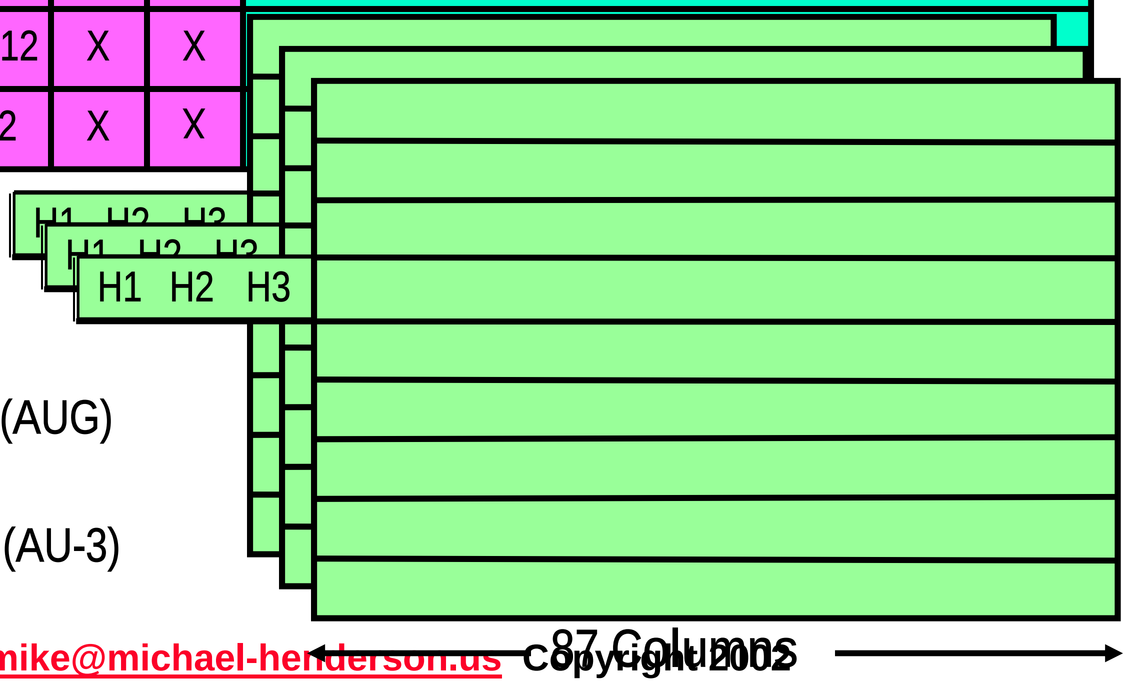

Section overhead

Administrative Unit Group (AUG)

Administrative Unit level 3 (AU-3)

P. Michael Henderson, mike@michael-henderson.us

Administrative Unit - 3

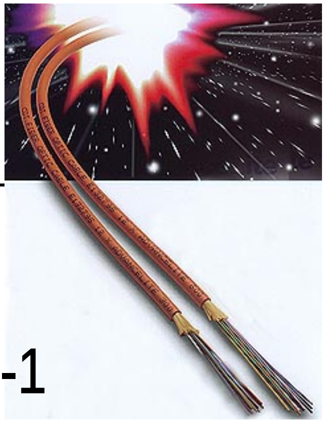

- But the AU-3 doesn't "float" like the payload of an STS-1 signal. How do we accommodate clock differences?
- We use a structure known as a Virtual Container – 3 (VC-3).
  – The AU-3 has fixed stuff columns at columns 30 and 59, just like SONET.
  – The VC-3 does not include these fixed stuff columns – they're taken into account when the VC-3 is mapped into the AU-3.
  – So the VC-3 has only 85 columns – 84 columns of payload and one column of payload overhead.

The "Floating" Virtual Container - 3

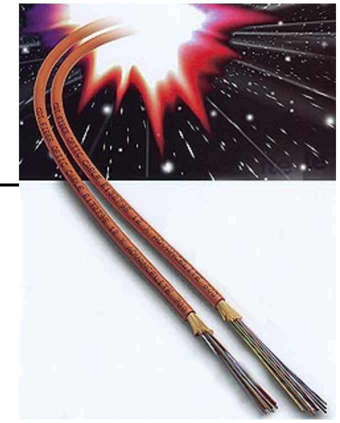

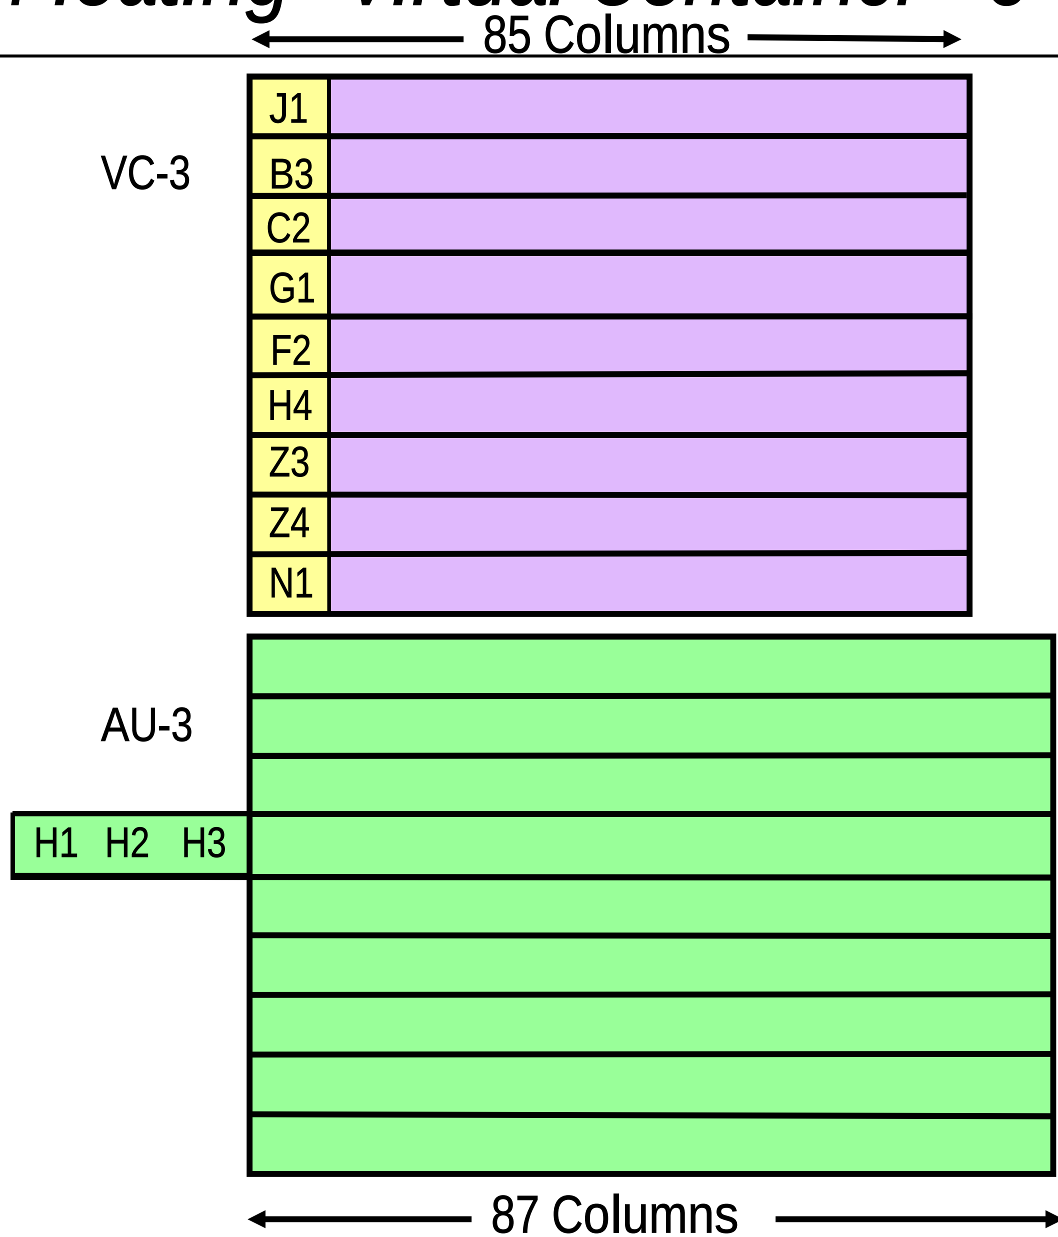

P. Michael Henderson, mike@michael-henderson.us

Mapping PDH Traffic

- When mapping PDH traffic, the VC-3 will contain seven Tributary Unit Group – 2s (TUG-2). These function *exactly* like the VTGs in SONET.
- Then, within the TUG-2, we have Tributary Units (TU-x) which carry the traffic exactly like SONET. The names are different.
  - Actually, each TU contains a Virtual Container which carries the traffic. See next slide for convention.

SONET Name	SDH Name
VT-1.5	TU-11/VC-11
VT-2	TU-12/VC-12
VT-6	TU-2/VC-2

SDH Naming Conventions

- Level 1 is the lowest level structure. Numbers get larger as the structure gets larger.
- Traffic is always carried in a container.
- A container becomes a Virtual Container (VC) with the addition of Path Overhead (POH)
- A Tributary Unit (TU) is a Virtual Container (VC) plus a pointer.
- A TU fits into a Tributary Unit Group (TUG).
- When a structure fills another structure, it gets the same number. Example, a TU-2 fills a TUG-2. A VC-3 fills an AU-3.

Administrative Unit - 4

- An alternate to the AU-3 is the AU-4.
- The AU-4 structure takes the entire payload of an STM-1
  - 261 columns plus the pointers in row four of the section overhead.
- Problem: How to map the PDH traffic in this larger frame?
- Hint: If we can get to a VC-3 structure, the TU structure will be exactly as described for the AU-3.

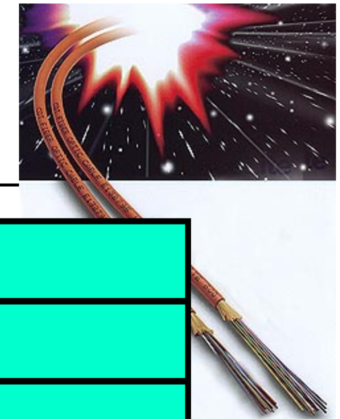

A1	A1	A1	A2	A2	A2	J0	Z0	Z0
B1	X	X	E1	X	X	F1	X	X
D1	X	X	D2	X	X	D3	X	X
B2	B2	B2	K1	X	X	K2	X	X
D4	X	X	D5	X	X	D6	X	X
D7	X	X	D8	X	X	D9	X	X
D10	X	X	D11	X	X	D12	X	X
S1	Z1	Z1	M0/1	Z2	M2	E2	X	X

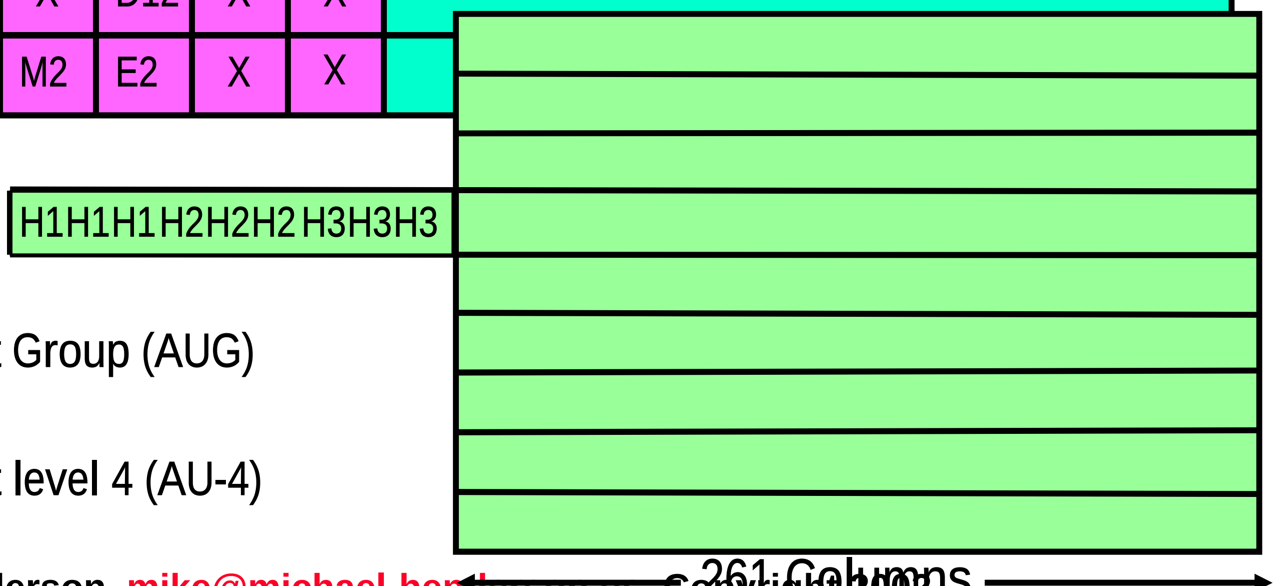

Section overhead

Administrative Unit Group (AUG)

Administrative Unit level 4 (AU-4)

P. Michael Henderson, mike@michael-henderson.us

Tributary Unit Group - 3

- Define a new structure known as the Tributary Unit Group – 3 (TUG-3).
- Consists of 86 columns, one larger than the 85 columns of a VC-3.
- The extra columns carries pointers to allow the VC-3 to "float."
- Three 86 column TUG-3s will take 258 columns, leaving three columns in the 261 column AU-4.
- One column is used for the payload overhead (POH) just like all other payloads.
- The two columns behind the POH are fixed stuff.
- So three TUG-3s fill an AU-4.

P. Michael Henderson, mike@michael-henderson.us

VC-3 Mapping into a TUG-3

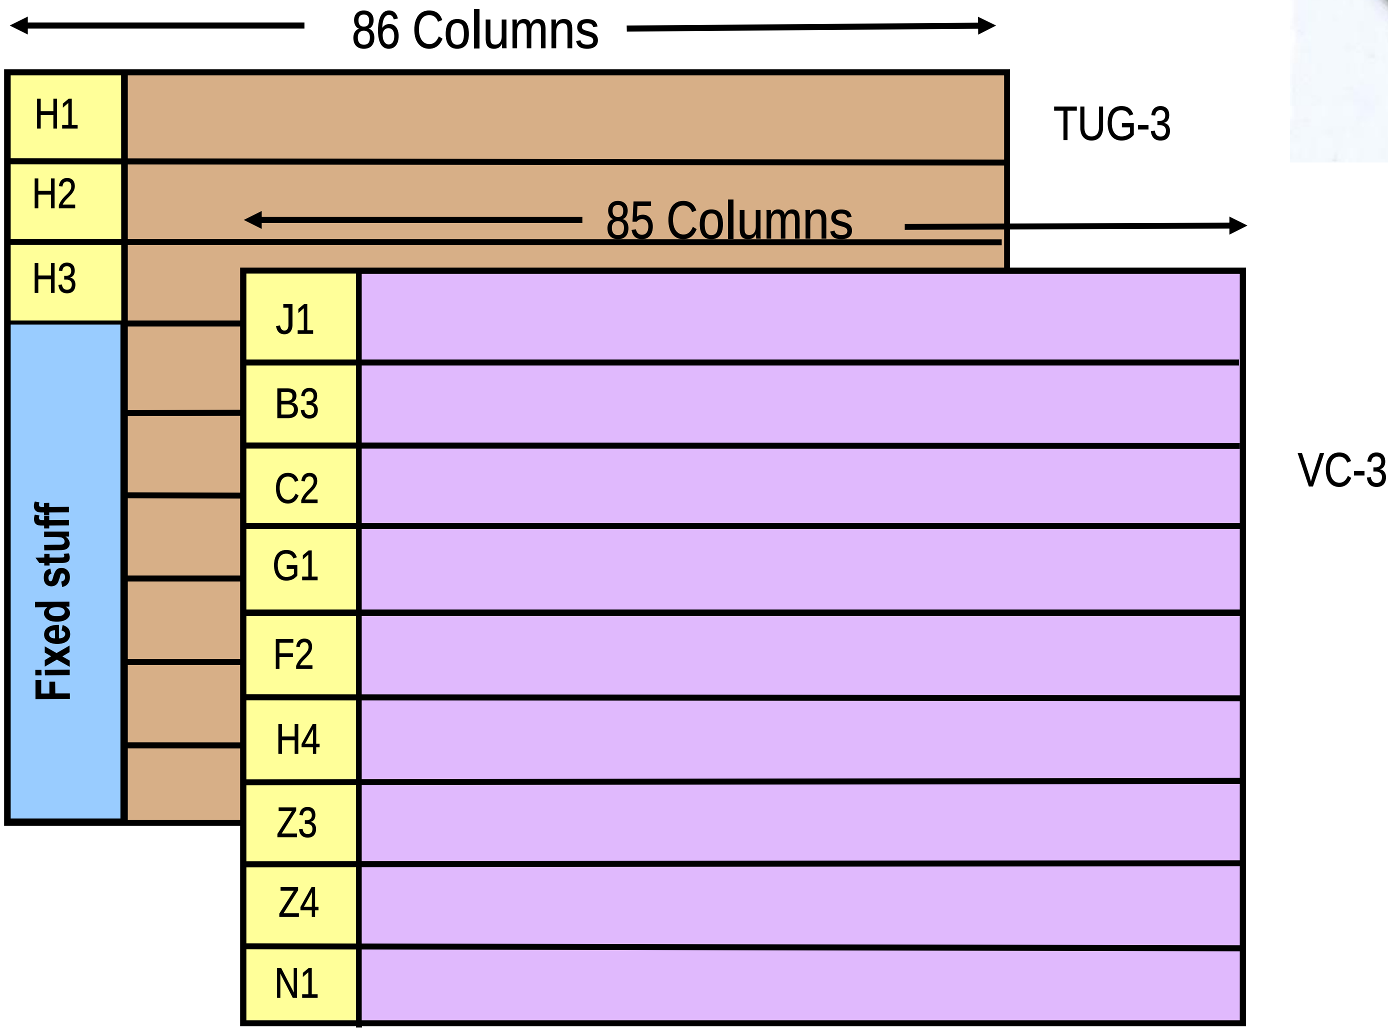

P. Michael Henderson, mike@michael-henderson.us

Virtual Container - 3

- Now that we have reached the VC-3 structure, all of the Tributary Unit mappings are the same as described earlier.

- Added note: There's another terminology in the TUG-3 case which I skipped for simplicity. The combination of the VC-3 plus the H1, H2, H3 pointers in the TUG-3 are called a Tributary Unit level 3 (TU-3). This is similar to the way the AU-3 consists of the payload plus the H1, H2, H3 pointers from the AUG.