Voice meets data – where did my 10% capacity go?

Yi Chu

IP Engineering

Sprint

January 19, 2010

Where did 10% of my POS IMT capacity go?

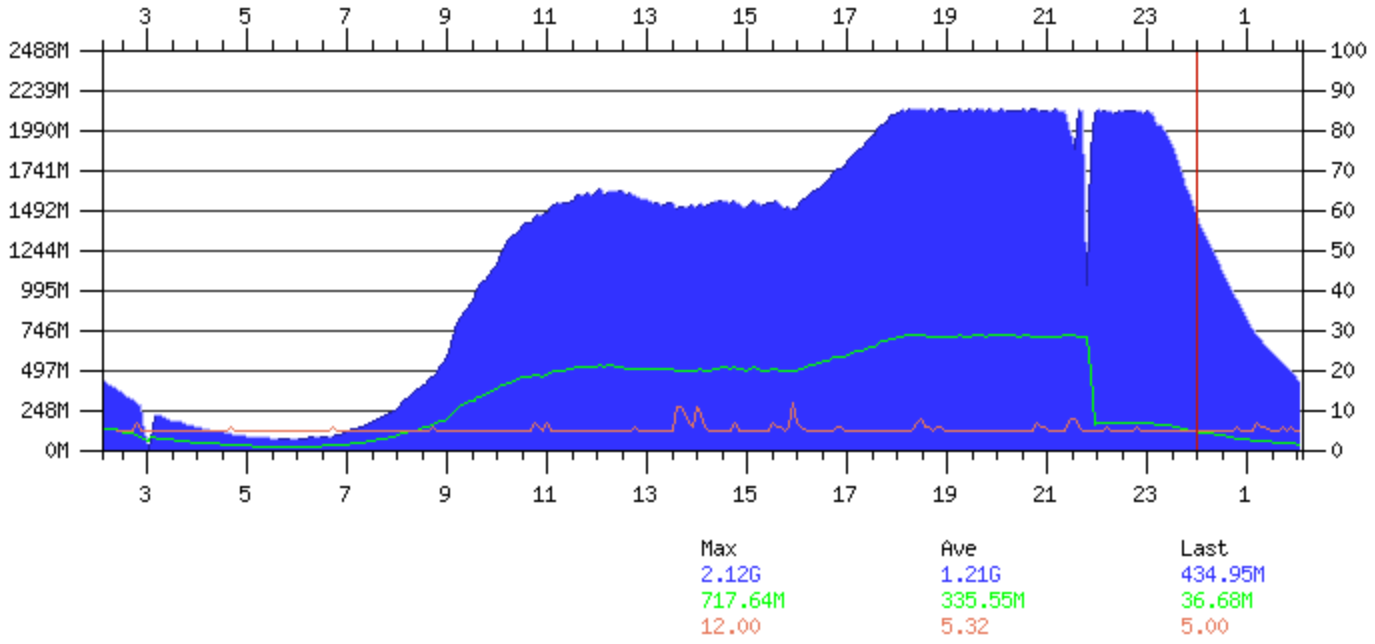

Background

- Mostly (90%+) VoIP traffic – burstiness not expected
- G.711 encoding
- No silence suppression
- 20msec packetization
- Egress traffic load-shared with another link
- Encap is HDLC on OC48 POS link
- So where did the 10% go?
- Peter Lothberg pointed me to HDLC framing

HDLC Frame

HDLC Frame

Flag, 7E hex

Escape: 7E => 7D-5E
7D => 7D-5D

‘User data which contains 7E is resolved using an escape sequence which converts 7E to 7D-5E [with 7D being the escape character]. If 7D is used in the data stream it again is converted into 7D-5D. ’

G.711 Encoding - ethereal capture

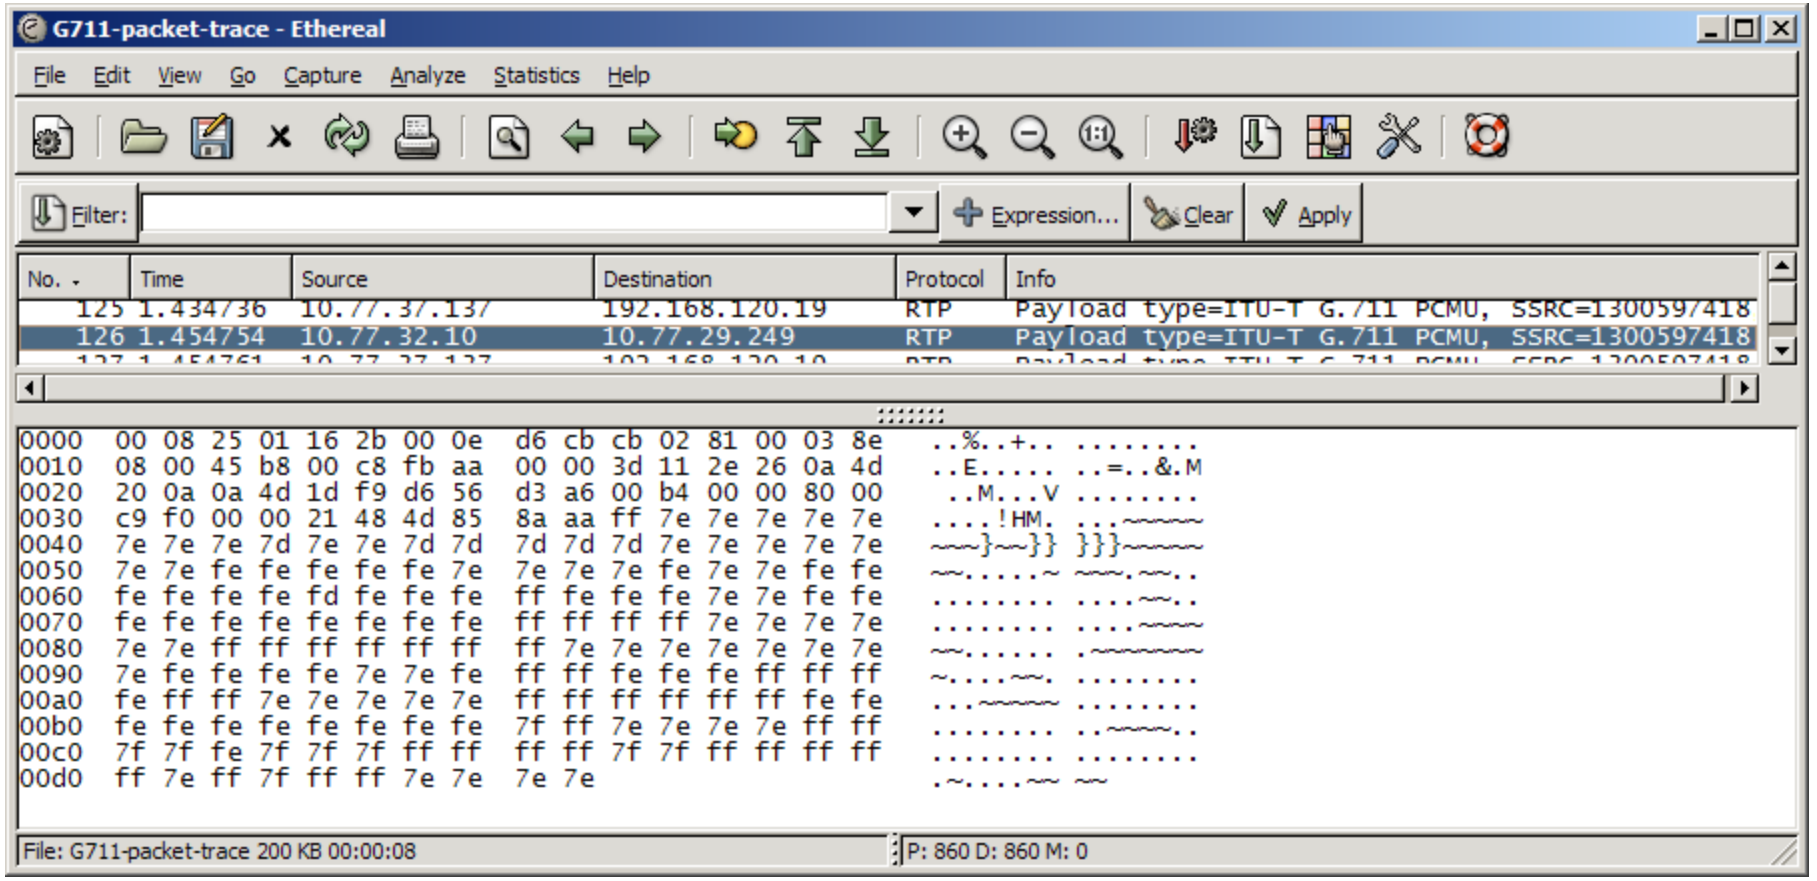

Notice many the '~' and '}'? But why?

G.711 μ-law codec

Ask people who know:
Here is a correspondence with Prof. Dave Petr (http://www.ittc.ku.edu/~dwp/):

all of the bytes are just pcm codewords (digital representations of the voice sample values), and it just turns out that 7E is one of the codewords representing a very small signal amplitude (hence the frequency of this codeword in silence portions).

i would guess that you also see a bunch of 7F, FE and FF codewords as well -- these 4 represent the 2 smallest amplitudes, both + and -. looking at the trace just now, i see that this is indeed the case.

G.711 spec

Well, if you care to investigate further, here is the URL:
http://www.itu.int/rec/T-REC-G.711/en

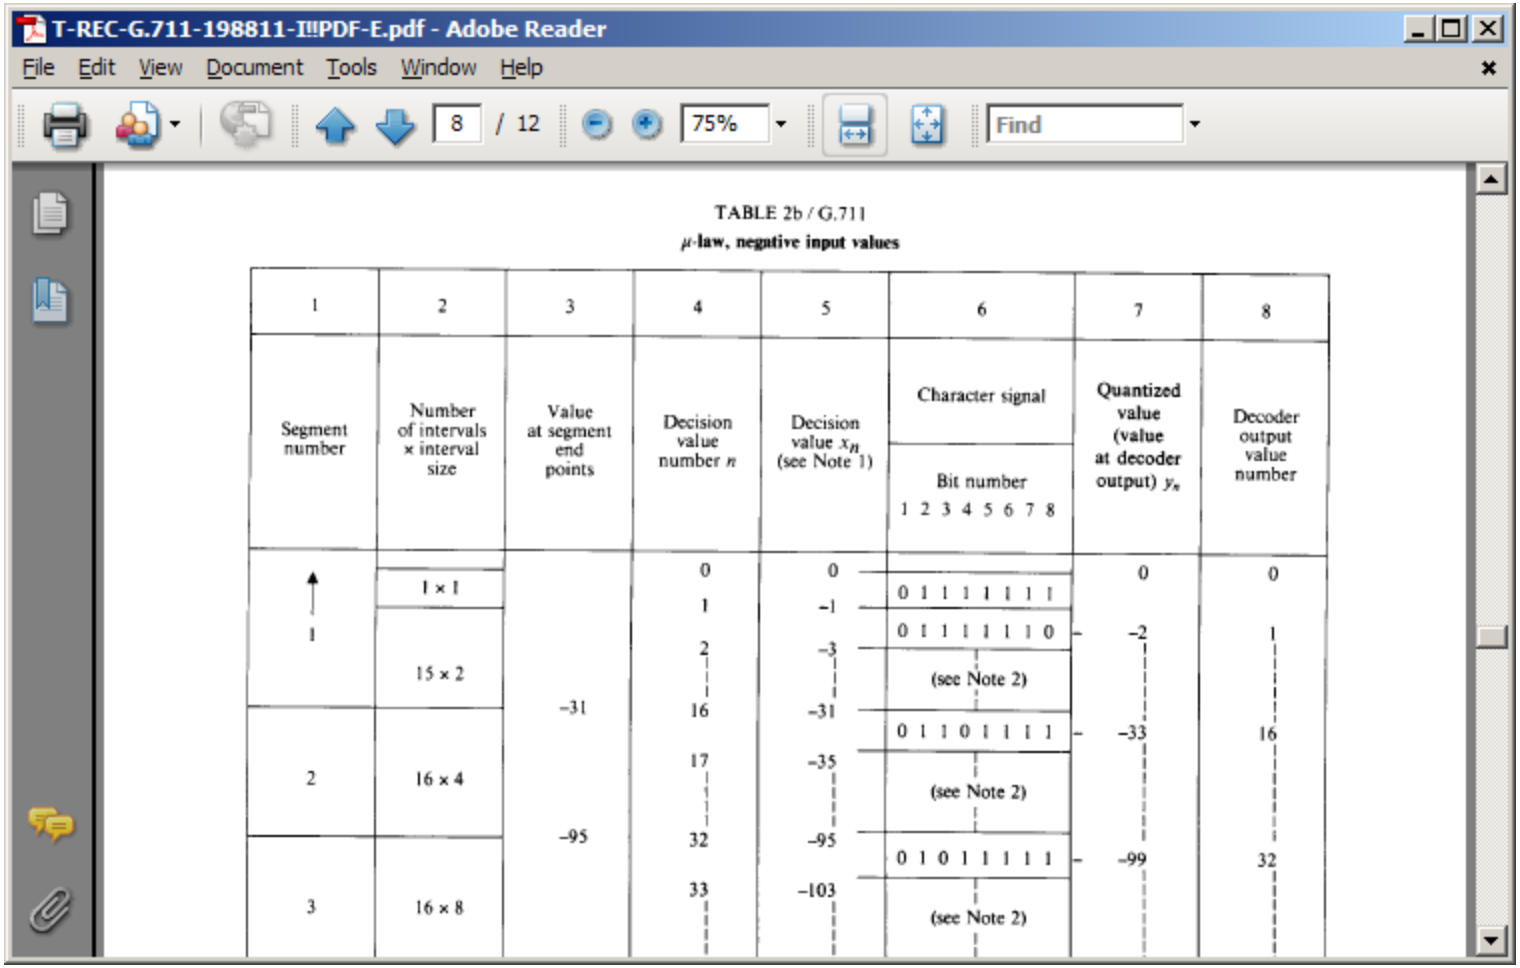

TABLE 2b / G.711

μ-law, negative input values

1	2	3	4	5	6	7	8
Segment number	Number of intervals × interval size	Value at segment end points	Decision value number n	Decision value x_n (see Note 1)	Character signal Bit number 1 2 3 4 5 6 7 8	Quantized value (value at decoder output) y_n	Decoder output value number
1	1 × 1		0	0	0 1 1 1 1 1 1 1	0	0
			1	−1	0 1 1 1 1 1 1 0	−2	1
	15 × 2		2	−3	(see Note 2)		
		−31	16	−31			
2	16 × 4		17	−35	0 1 1 0 1 1 1 1	−33	16
			32	−95	(see Note 2)		
		−95					
3	16 × 8		33	−103	0 1 0 1 1 1 1 1	−99	32
					(see Note 2)		

Conclusion

- 7E, 7D are just code words that represent low voice activity
- HDLC uses two bytes to transmit them
- So it happens that it takes more bits down the POS/HDLC wires when people are not talking
- The effect is not counted on router interface counters
- Resulting in seemingly lost capacity

Silence audio trace statistics

Byte Type	Count
7E	19,611
7D	1,995
Total packets	405
Payload Bytes per packet	200
Total bytes	81,000
Total 7E/7D bytes	21,606
extra %	27%

Here is the lab captured audio:

hindsight

- G.711 was developed in 1980’s
- HDLC in the 1990’s
- *[soap box] HDLC could have used another byte pattern for flag, if the data guys consulted the voice guys*