Demonstration of COTS Hardware for Capture, Playback and Processing of SMPTE ST 2022-6 Media Streams

Thomas Edwards

FOX Network Engineering & Operations

Introduction

The broadcast industry is considering a change from carriage of uncompressed HD video in professional broadcast plants from SDI (serial digital interface, mainly over 75Ω coaxial cable) to an Ethernet networked solution.

SMPTE ST 2022-6 is a standard for IP carriage of SDI bit streams, and has been used in this demo.

The high bit rate and QoS requirements for professional media can be challenging for typical COTS PC systems. Use of TCP/IP stack bypass and DMA-enabled NIC drivers can allow such systems to capture and playback these high speed signals.

This work demonstrates a solution to using COTS PC hardware for Capture, Playback and Processing of uncompressed HD video streams.

Today's Broadcast Plant

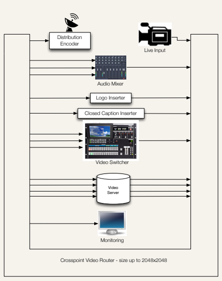

HD-SDI: 1.485 Gbps Uncompressed HD serial digital video & audio on 75Ω coax cable

Huge crosspoint switches ("routers") connect output of every device with input of every device, up to 2k x 2k

Devices to cut between or mix video streams ("switchers")

A Networked Broadcast Plant?

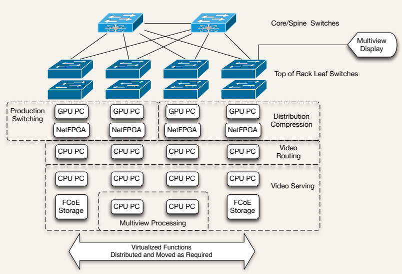

Commodity, programmable components support functional virtualization

Video processing distributed as required by business needs on demand

Advantages of Professional Media Networking

- Enhanced flexibility and agility of broadcast plant
- Cable reduction through statistical multiplexing of multiple signals
- Reduced cost from use of more Common, Off-the-Shelf (COTS) PC & IT equipment
- Easier monitoring of video signals
- Making signal metadata a "first class citizen"
- A single Ethernet fabric to operate for all data, including video, audio, comms, IFB, automation, file transfer, etc.
- Distributing video processing over private or possibly public cloud infrastructures
- A single multiresolution fabric to carry SD, HD, compressed and uncompressed, with an easy upgrade path to 4K, 8K resolutions

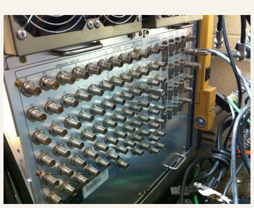

vs.

Cable Aggregation

Ethernet Type	JPEG 2000 HD Streams @ 150 Mbps*	Uncompressed HD streams @ 1.5 Gbps
1 GbE	6	N/A
10 GbE	66	6
40 GbE	266	26
100 GbE	666	66

*JPEG 2000 @ 150 Mbps – "Visually lossless" to pro eyes at 5 encode/decode generations, but has more latency than uncompressed video

SMPTE ST 2022-6

The SMPTE ST 2022-6 standard provides for encapsulation of the payloads of a variety of SMPTE serial digital video standards.

2022-6 carries the SDI video data in RTP datagrams. Each datagram has a short HBRM (High Bit Rate Media) header along with 1376 octets of video payload. The last datagram of a frame has zero padding beyond the end of video data in the frame, and also has its RTP market bit set.

Below is a Wireshark dissection of a 2022-6 datagram:

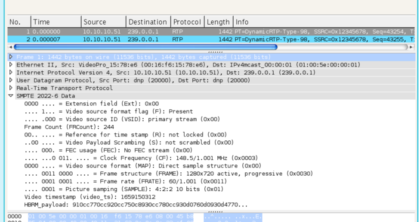

(Lua 2022-6 Wireshark dissector available at https://github.com/FOXNEOAdvancedTechnology/smpte2022-6-dissector)

The diagram below shows how a 720p59.94 SDI video signal is split into 2022-6 datagrams (boundaries are green), with approximately 3 datagrams per video line:

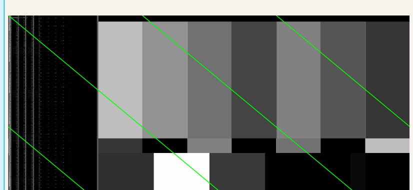

SDI to 2022-6 conversion done by a Nevion VS902 card in this demo.

PF_RING: Bypass normal TCP/IP Stack

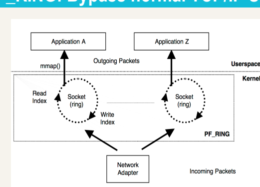

PF_RING (from ntop.org) is a new type of network socket to improve packet capture & transmit speed

No "bells and whistles" of regular network stack (e.g. crafting UDP/IP packets, routing, error detection)

10 Gbps capture & hardware packet filtering using commodity NICs (e.g. Intel 82599-based)

Still requires: Kernel copy packet from NIC to ring, Userland: read packet from ring

DNA (Direct NIC Access) Driver

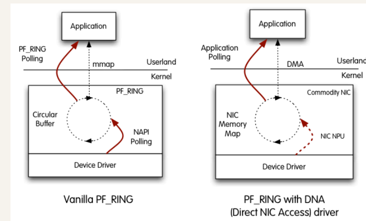

PF_RING DNA maps NIC memory and registers to userland with DMA transfer done by the NIC NPU (Network Process Unit).

CPU cycles are not used to move packets off adapter. DNA provides 10 Gbps "wire rate" TX/RX on Sandy Bridge and later chipset systems. It works with typical Intel 10GbE NICs.

"nBox""

COTS SuperMicro chassis with X9SPU-F MB

- 1 x Intel Xeon E3-1230 V2 @ 3.30 Ghz, 8GB RAM
- Intel 82599EB-based dual 10 Gbps SFP+ NIC
- 4 x WD Velociraptor 10KRPM 1TB SATA HDs
- Ubuntu 12.04.2 LTS (GNU/Linux 3.2.0-44-generic x86_64)
- PF_RING module & DNA driver

PF_RING-based Applications

- Disk2n: Captures packets from 10 GbE NIC to PCAP files
- N2disk: Replays packets from PCAP files to 10 GbE NIC **with cadence specified by PCAP file packet timestamps**

Capture & Playback Results:

- 13584 frame (~226s) 720p59.94 test clip on 2 HD-SDI streams of 2022-6 was recorded by disk2n
- **Perfect playback by n2disk** of SDI as converted by VS902
- Every pixel tested "OK" by Video Clarity

SMPTE RP 168 Switching & 2022-6

SMPTE RP 168 is the "Definition of Vertical Interval Switching Point for Synchronous Video Switching". The recommended switch point is located so "the effects of any signal discontinuity in the chain are minimized", is early in the vertical blanking interval to keep ANC data with its intended frame, and SMPTE standards ensure no vital ANC data is located on the line after switch point.

As suggested in [M. Laabs, EBU Technical Review, 2012 Q4] one can count an integral number of 2022-6 RTP datagrams into the frame to find the RP-168 valid switch point. Software was written for this demo to do clean switching in this fashion based on PCAP files captured by n2disk, for later playback of switched streams with disk2n.

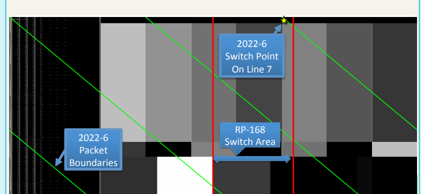

Logo Insertion

Logos and other graphical insertion can also be done using PCAP files of captured 2022-6 streams. Care needs to be taken to recalculate the HD-SDI Cyclic Redundancy Check (CRC) on lines where the active video changes. This is not an issue for switching because line 7 is in the vertical interval, and all active video pixels are always black.

Contact Information

Thomas Edwards
thomas.edwards@fox.com
+1-310-369-6696
VP Engineering & Development, FOX NE&O
10201 W Pico Blvd, Los Angeles, CA 90035 USA