Native IP decoding of MPEG-TS video to Uncompressed IP (and vice versa) on COTS hardware

Kieran Kunhya <kierank@obe.tv>

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

Introduction

- Many facilities have large numbers of outside sources, hundreds of single 1U encoders and decoders, single channel
  - Often not much in the way of switchers
  - Basic processing, graphics, commentary
- Using entirely COTS hardware for encoding and decoding to/from Uncompressed IP
  - Complex software engineering problems
  - Multidisciplinary: Broadcast, Networking, Multimedia
- First half of presentation about operational challenges, second part is deep dive into Software Engineering problems

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

Who we are?

- Company specialising in software-based encoders and decoders for Sport, News and Channel contribution (B2B)
- Based in Central London
- Build everything in house:
  - Hardware, firmware, software

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

Where it all began?

- Major UK News Channel moving building
  - Approached in 2015 about possibility of future update to 2022-6
  - Hundreds of encoders and decoders for contribution

- In the end run as SDI for 1.5 years, then “lift-and-shift”
  - Designed process for operators to migrate equipment overnight

Why do this?

- Reduce use of encapsulators
- Beyond historical-re-enactment
  - Using native IP devices to process video at scale
  - Don’t invent the internet to use it to send more faxes

- Still a lot of 2022-6, not going away. Useful in linear environments with freerunning sources, no need for PTP

Challenges

- Selection of Network Interface Card (NIC) in 2015!
  - Limited choice (10GbE) and capabilities
  - Chose Intel 82599
- No analysis equipment until ~2016-17
- Limited availability of 2022-6 equipment
- Mixed SD/HD and PAL/NTSC including 720p59.94

Broadcast as an IT service

- Large amount of compute
  - Hundreds of cores (small supercomputer)
  - Usable for other things (e.g transcode)
- Managed using Ansible
  - “open-source software provisioning, configuration management, and application-deployment tool”
- COTS hardware to reduce costs, time to deployment

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

Hardware provisioning

- Well known methods for large scale hardware provisioning of IT hardware
  - Ubuntu Preseed to flash precise software onto hardware, inc. IP addresses etc
  - Ship anywhere in the world, plug-in, go
- Used by CDNs which colocate equipment on other people's network

Spin up and spin down

- Goal to allow spin up and spin down production
  - e.g Sports, Elections
- Weekend operators did this by themselves over the weekend!
  - Doubled decode capacity on same hardware
- Cloud-style usage models
  - Pay per usage

Activation log

IP address	Extra data	Type	Date activated
[redacted]		Activation	September 29, 2019 — 2:42 pm (+01:00)
		Activation	September 29, 2019 — 2:23 pm (+01:00)
		Activation	September 29, 2019 — 1:41 pm (+01:00)
		Activation	September 29, 2019 — 11:41 am (+01:00)
		Activation	September 29, 2019 — 11:18 am (+01:00)
		Activation	September 29, 2019 — 10:58 am (+01:00)
		Activation	September 29, 2019 — 10:33 am (+01:00)
		Activation	September 29, 2019 — 9:59 am (+01:00)
		Activation	September 28, 2019 — 2:03 pm (+01:00)
		Activation	September 28, 2019 — 1:38 pm (+01:00)
		Activation	September 28, 2019 — 1:12 pm (+01:00)
		Activation	September 28, 2019 — 9:16 am (+01:00)

Step forward to 2019/20

- Well known consumer brand doing sports
  - Based at large production house
  - 2.5 months from commission to deploy (!!)
- World Feed of Sport decoded to 2110 for downstream processing
  - 15-20Mbps H.264, 4-8 pairs of audio
- 1080i25 or 29.97, 16ch single flow @ 125us
- Up to 15 simultaneous matches

Incredible density vs SDI

- 2 full racks of decoders or encoders replaced with 2U of blade servers
- Add CPUs/RAM to nodes as needed
- Envisage cloud like racks of this in near-future
- 64+ Channels in 2U

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

The software engineering

High Level Overview of Challenges

- High data rates of UDP traffic
  - Many Gbps
- Unusual pixel formats, software works in 8/16-bit, packed 10-bit in 2022-6/2110
- Very tight timing requirements:
  - Order of us
- 2022-7 skew
- Internal multicast routing

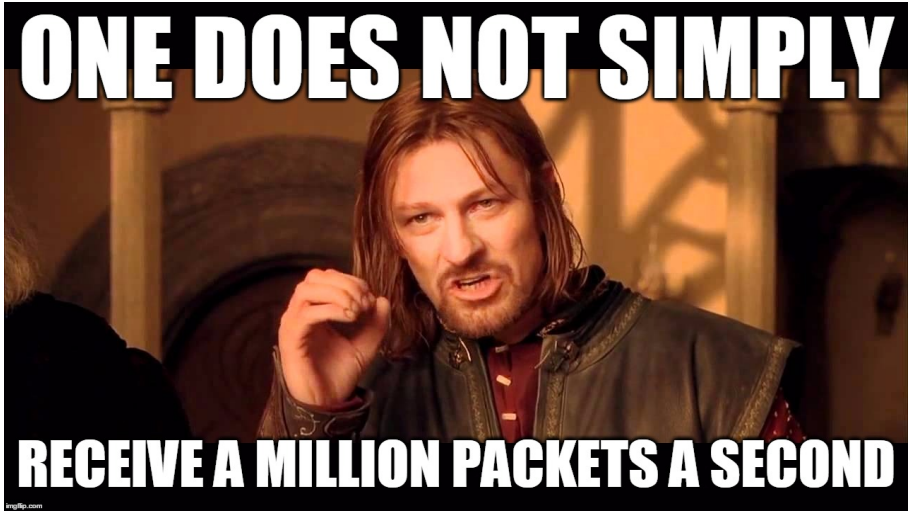

OPEN
BROADCAST SYSTEMS

VSF
VIDEO SERVICES FORUM

Kernel Bypass

- Standard OSs not designed for volume of UDP traffic
  - Might change in future as HTTP3 deployed which uses UDP
- Direct path to NIC memory
  - Craft packets yourself IP/UDP headers
- We chose netmap, support Intel and Mellanox cards
  - Can't use vendor specific libraries

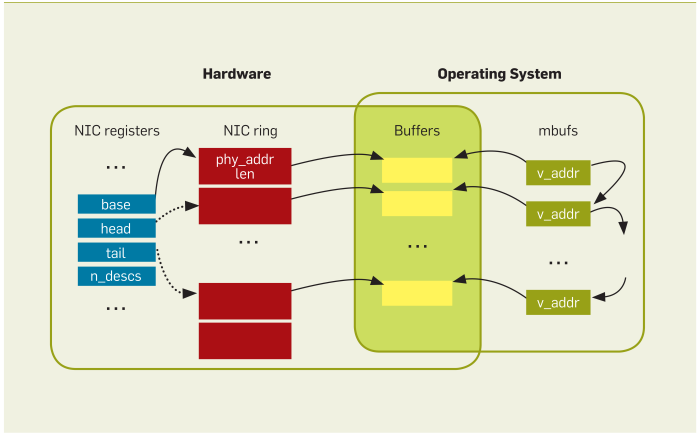

Packet pacing

- Need to space packets evenly
- Canonical method of NIC assisted pacing
  - Send using a TX queue, ratelimit queue
  - Grace et. Al, BBC R&D 2016
  - Ideally let the OS do PTP, ping, IGMP
- Fix all the hardware bugs in pacer
  - Mellanox and Intel had large drift (sometimes 1000ppm)
  - Fractional framerate compensation

2022-7 skew

- Many problems
  - Cannot send two packets exactly at the same time
  - After loss of a path, queue empties
  - Tricky to realign queue, continually moving target

- Would love to have hardware duplicate packets
  - Send once, header rewrite using NIC eSwitch

Accelerated pixel processing

- Hand written SIMD code, multiple pixels processed at once
  - Many permutations pixel formats, dozens of functions
  - Around 25 times faster than naïve implementations!
- Pixel transformations straight into NIC memory
  - SIMD + Kernel Bypass = 100-200x speed improvement!
- More detailed explanation at Demuxed 2017 presentation: www.youtu.be/watch?v=A4L5xEXXlas

Implementing SDI (2022-6) in software

- SDI CRC is CPU intensive
  - 10-bit data, 25-bit polynomial
  - Useless in IP
- Pixel spanning multiple packets…
- SDI Level B, painful to handle
- Cursed mixed cadence in NTSC
  - 1601/1602 on each audio group
- **However, single code base across baseband and 2022-6**
  - Regression tests based on real-world captures

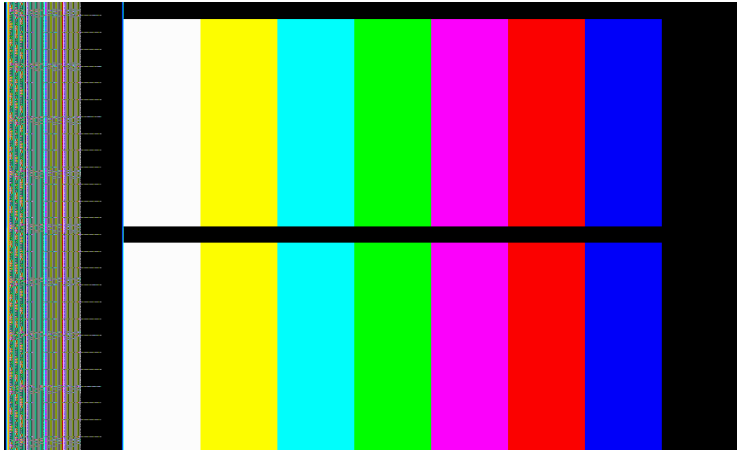

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

2110 software challenges (video)

- 3 coupled variables
  - Frequency
  - Phase
  - Packet Bursts
    - NIC dependent, difference between Narrow and Wide
- Change one and the others move from target!

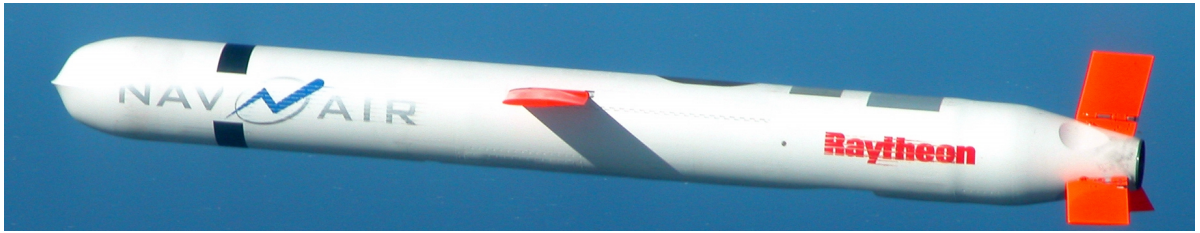

2110 software challenges (video)

- Software senders can’t just create a gap for the vertical blanking
  - Stuff with other traffic (audio etc)
- Phase is hard, don’t know where you start
  - Start sending and drift your way back to PTP epoch
  - Start too late and some devices will never recover!
  - Many NICs have hardware transmit timestamps
  - Filter some timestamps that go back in the past..
- 50/59.94fps harder, less time to react

2110 software challenges (audio)

- 125us audio is tight. Bypass is overkill for 30Mbps of traffic
- Seen one vendor send two packets at the same time (250us gap)

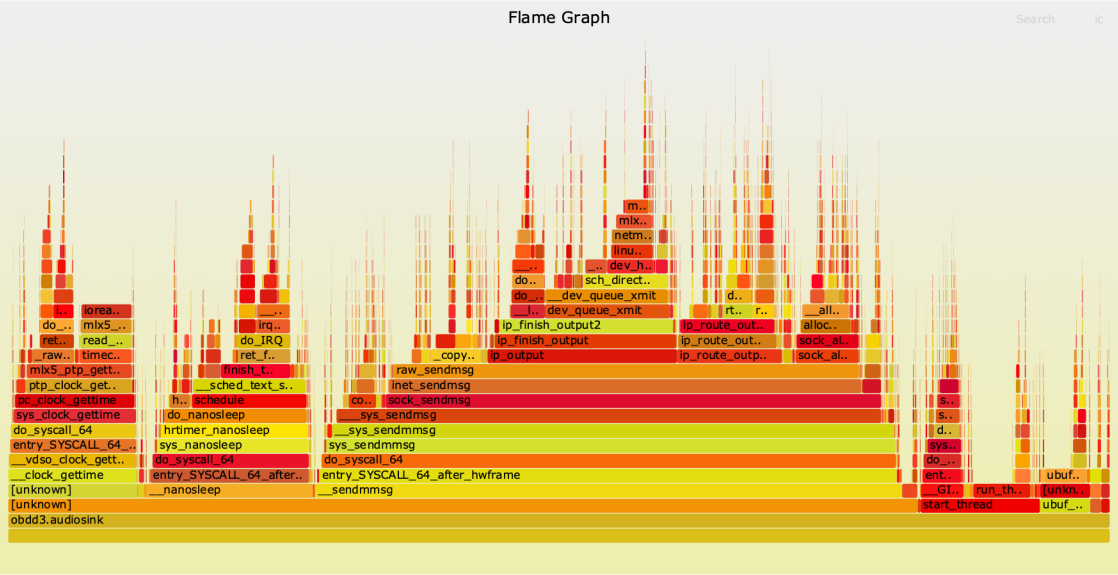

OPEN BROADCAST SYSTEMS

VSF VIDEO SERVICES FORUM

General software comments

- Challenging to maintain a modular codebase with Uncompressed IP
  - Technical debt accumulation
- Hard to use standard paradigms such as reference counting (especially with tiny 125us buffers)
- Very tight integration with application

Interoperability…

- One decapsulator expects source port to match destination port…
- Many don’t verify IP checksums
- Not all receivers can handle 16ch flows
- Some analysers and receivers just crash
  - Sometimes within minutes, sometimes days
  - Expecting SDI-like flows…
- Clear that some receivers using packet arrival time for audio and not timestamps
- Some transmitters miles off the epoch
  - Operators route freerunning sources, causes problems for native IP devices

Where next?

- Packet Pacing is a crutch
  - Hardware packet pacers are buggy
  - Easier to send line rate and multiplex ourselves
    - Not container/VM friendly
- Docker containers
  - How do have different 2110 implementations coexist?
- True remote production
  - Use off-prem compute hardware, spin-up and down

Questions, also outside booth 13

OPEN
BROADCAST SYSTEMS

VSF
VIDEO SERVICES FORUM