

Hardware UDP



High Performance Simultaneous Data Acquisition

Reference document for details relating to D-TACQ support for Hardware UDP on
ACQ2106 Carrier via SFP Port D on MGT482-SFP
ACQ2206 Carrier via SFP Port D on MGT483-SFP (10G support)

Contents

1	Introduction	5
1.1	1G	5
1.2	10G	5
2	Setup	6
2.1	Python - ACQ400 HAPI Installation	6
2.2	C - udpperf Installation	6
2.3	Optimising Host for best performance	6
2.3.1	Enable Jumbo frames	6
2.3.2	Network Interface Optimisations	6
3	Streaming Sampled data to host via UDP	7
3.1	Configure the HUDP Module	7
3.1.1	ARP Request Packet	8
3.2	1G Streaming	9
3.3	10G Streaming with udprx.cpp	10
3.3.1	Configure the HUDP Module	10
3.3.2	Validate performance with sequence check	10
3.3.3	Receive data and store to disk	10
3.3.4	Receive data and store to disk - Max Samples argument	10
3.3.5	(Optional) Monitor progress while writing to disk	11
3.4	Plotting Sampled Data on Host	12
3.5	Multiple Samples per Packet	13
3.6	Decimation	14
3.7	Streaming Multiple Sites of Data	15
4	Streaming UDP Data to acq2106 XO Modules	16
5	Monitoring Packets with Wireshark	17
6	Low Latency Demo Setup	18
6.1	HW	18
6.2	HUDP Core Configuration	19
6.3	CSS OPI	19
6.4	Discontinuity Counter	19
6.5	Clock and Excite Signal Phase Relationship	20
6.6	Latency Measurements	21
6.6.1	Point-to-Point	21
6.6.2	Through Switch	21
6.7	HUDP - Internal workings	22
7	Frame Specifications	23
7.1	Default MTU	23
7.1.1	1G	23
7.2	Jumbo Frames	24
7.2.1	1G	24
7.2.2	10G	25
8	Troubleshooting	26
8.1	sfp webpage	26
8.2	hudp webpage	26
8.3	tcpdump	27
A	Example run on D-TACQ Host (naboo)	28
A.1	Set ACTIVE_NIC	28
A.2	Enable Jumbo frames and set host optimisations	28

A.3	Set up a full ACQ424 box and maximise packet length	28
A.4	Force a negotiation (optional, if LINK_STATUS is GOOD then skip)	28
A.5	Spin up udprx on the host to receive and validate packets	28
A.6	Start a stream and observe packet reception statistics	28
B	Reliability Testing	29
B.1	Interrogate HUDP stresser logs	29
B.2	Optimising host for very high rates	30
B.2.1	Isolate a CPU	30
C	Validated Hardware	31
C.1	NICs	31
C.2	SFP+ Transceiver Modules	31
C.3	Optical Fiber	31
C.4	SFP+ Active Optical Fiber Cables - Transceiver and fiber in one fixed assembly	31
C.5	Copper	31
C.6	NIC / Transceiver Vendor Lock	31
C.7	How to query NIC and transceiver info	31
D	Loopback to D-TACQ Fiber Ethernet Port	32
E	Resource Usage	33
F	HUDP 1G Logic Schematic View	33

Revision History

Revision	Date	Author(s)	Description
1.0	26/05/2022	SR	Created
2.0	06/07/2022	SR	New diagrams. More prose. Added HUDP to Host section
3.0	15/07/2022	SR	Section better specifying packet structure and max limits
4.0	18/01/2023	SR	Some formatting fix-ups
5.0	27/04/2023	SR	More usage examples for end users
6.0	12/12/2023	SR	Restructure
6.1	12/12/2023	SR	Update to include 10G, minor corrections
6.2	01/02/2024	SR	Add Reliability Testing; optimising host for high rates; expand Troubleshooting
7	26/04/2024	SR	Add Validated Hardware section
8	13/01/2025	SR	Documenting dummy data sent during forced ARP request

1 Introduction

The Hardware UDP (HUDP) capability in D-TACQ carriers uses an FPGA IP core to enable high-speed data movement over standard Fiber Ethernet infrastructure.

High data rates OR relatively low packet latency can be achieved.

1.1 1G

Available on ACQ2106 or ACQ2206 carriers.



Figure 1: View of ACQ2106 MGT482-SFP Ports. Port D can be used for HUDP

- Max Throughput
 - with default MTU => 119.6 MB/s
 - with Jumbo frames, MTU=9000 => **124 MB/s**
 - *All devices in the network path (e.g. network switches) must support Jumbo frames.*
- Supports Low-Latency Control applications
- 1G UDP FPGA personalities are denoted by the "UDP" filename suffix

1.2 10G




Figure 2: View of ACQ2206 MGT483-SFP Ports. Port D can be used for HUDP

10G is only supported on ACQ2206 carriers.

Data throughput may be limited by a host's ability to receive high packet rates. 10G UDP must be paired with 10G capable NIC and SFP+ modules.

- Max Throughput
 - with Jumbo frames, MTU=9000 => **1.24 GB/s**
 - *All devices in the network path (e.g. network switches) must support Jumbo frames.*
- 10G UDPX FPGA personalities are denoted by the "UDP" filename suffix
- Ports
 - Initial release, Port D
 - Future multi-port A, B option under development

2 Setup

2.1 Python - ACQ400 HAPI Installation

Python code and full install instructions can be retrieved from : https://github.com/D-TACQ/acq400_hapi

To install :

```
mkdir PROJECTS; cd PROJECTS
git clone https://github.com/D-TACQ/acq400_hapi.git
cd acq400_hapi
```

on Linux , run `source ./setpath`

on Windows, run `SETPYTHONPATH.BAT`

2.2 C - udpperf Installation

C code and full install instructions can be retrieved from : <https://github.com/D-TACQ/udpperf.git>

To install :

```
mkdir PROJECTS; cd PROJECTS
git clone https://github.com/D-TACQ/udpperf.git
cd udpperf
cmake -DCMAKE_BUILD_TYPE=Release
make
```

2.3 Optimising Host for best performance

In order to accommodate high packet rates and Jumbo frames it is necessary to prepare the host network interface.

2.3.1 Enable Jumbo frames

```
sudo ip link set NETWORK-INTERFACE mtu 9000
```

Here NETWORK-INTERFACE should be replaced with the user's network interface name, e.g.

```
sudo ip link set enp10s0f0 mtu 9000
```

This assumes that a user has already configured an IP address and subnet mask for the NIC in question.

2.3.2 Network Interface Optimisations

```
sudo sysctl -w net.core.rmem_max=26214400 # Receive queue
sudo sysctl -w net.core.wmem_max=12582912 # Transmit queue
sudo sysctl -w net.core.netdev_max_backlog=5000
sudo ifconfig NETWORK-INTERFACE mtu 9000 txqueuelen 10000 up
```

3 Streaming Sampled data to host via UDP

3.1 Configure the HUDP Module

We begin by running the `hudp_setup` script from the host. Usage and arguments below. There is also extensive explanation provided in the comments at the beginning of the script.

```
[dt100@naboo acq400_hapi]$ python3 ./user_apps/acq2106/hudp_setup.py -h
usage: hudp_setup.py [-h] [--netmask NETMASK] [--tx_ip TX_IP] [--rx_ip RX_IP]
                  [--gw GW] [--port PORT] [--run0 RUN0] [--play0 PLAY0]
                  [--broadcast BROADCAST] [--disco DISCO] [--spp SPP]
                  [--hudp_decim HUDP_DECIM]
                  txuut rxuut

hudp_setup

positional arguments:
  txuut                transmit uut
  rxuut                transmit uut

optional arguments:
  -h, --help            show this help message and exit
  --netmask NETMASK    netmask (default: 255.255.255.0)
  --tx_ip TX_IP        rx ip address (default: 10.12.198.128)
  --rx_ip RX_IP        tx ip address (default: 10.12.198.129)
  --gw GW              gateway (default: 10.12.198.1)
  --port PORT          port (default: 53676)
  --run0 RUN0          set tx sites+spad (default: 1 1,16,0)
  --play0 PLAY0        set rx sites+spad (default: 1 16)
  --broadcast BROADCAST broadcast the data (default: 0)
  --disco DISCO        enable discontinuity check at index x (default: None)
  --spp SPP            samples per packet (default: 1)
  --hudp_decim HUDP_DECIM hudp decimation, 1..16 (default: 1)
```

This configures all of the necessary HUDP variables on the box. Many of these can be monitored using the HUDP tab in the `capture.opi`.

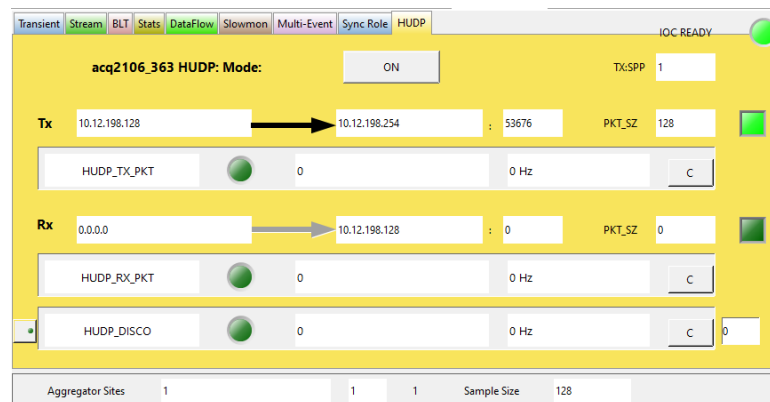


Figure 3: HUDP OPI after configuration by `udp_setup.py`

Here we give the `acq2106` an address of `10.12.198.128` and we plan to receive packets on host (`naboo:10.12.198.254`). Our port on both ends is the default `53676`.

```
[dt100@naboo acq400_hapi]$ UUT=acq2106_363
[dt100@naboo acq400_hapi]$ python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 \
--run0='1 1,16,0' $UUT none
txuut acq2106_363
enable disco at 16
```

The `run0='1 1,16,0'` argument tells us that we have requested data from Site 1, and that we have enabled a Scratchpad (SPAD) of length 16 LWords.

The Scratchpad is a series of writable elements that we append to the end of each sample of analog data. The first of these is always a sample counter. This is very useful to verify that no packets have been lost on the wire.

3.1.1 ARP Request Packet

As part of the `hudp_setup.py` a single small data packet is transmitted from the UUT. This forces the underlying HUDP logic to perform an ARP request. This ensures that when we begin sending high throughput streaming data there are no interruptions to the raw data transmission which could result in loss of data. This packet will share the address and port attributes of any other data packet but can be identified by its size and payload.

It is recommended that `hudp_setup` is executed before monitoring for data arriving on the host.

This eliminates the need for a host-side program to filter or ignore this dummy packet.

The packet has the minimum Ethernet frame length of 60 bytes. The UDP payload is one byte of 0x55.

No.	Time	Source	Destination	Protocol	Length	Info
3	0.040436397	DTACQSolutio_d3:00:85	Broadcast	ARP	60	Who has 10.12.198.254? Tell 10.12.198.16
5	0.040503032	10.12.198.16	10.12.198.254	UDP	60	53676 → 53676 Len=1
21	5.247483999	DTACQSolutio_d3:00:85	Intel_b9:22:44	ARP	60	10.12.198.16 is at 00:21:54:d3:00:85


```

Frame 5: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface enp10s0f0, id 0
Section number: 1
  Interface id: 0 (enp10s0f0)
  Encapsulation type: Ethernet (1)
  Arrival Time: Jan  7, 2025 16:12:19.055229439 GMT
  UTC Arrival Time: Jan  7, 2025 16:12:19.055229439 UTC
  Epoch Arrival Time: 1736266339.055229439
  [Time shift for this packet: 0.000000000 seconds]
  [Time delta from previous captured frame: 0.000038813 seconds]
  [Time delta from previous displayed frame: 0.000066635 seconds]
  [Time since reference or first frame: 0.040503032 seconds]
  Frame Number: 5
  Frame Length: 60 bytes (480 bits)
  Capture Length: 60 bytes (480 bits)
  [Frame is marked: False]
  [Frame is ignored: False]
  [Protocols in frame: eth:ethertype:ip:udp:data]
  [Coloring Rule Name: UDP]
  [Coloring Rule String: udp]
  Ethernet II, Src: DTACQSolutio_d3:00:85 (00:21:54:d3:00:85), Dst: Intel_b9:22:44 (3c:fd:fe:b9:22:44)
    Destination: Intel_b9:22:44 (3c:fd:fe:b9:22:44)
      ... ..0. .... = LG bit: Globally unique address (factory default)
      ... ..0. .... = IG bit: Individual address (unicast)
    Source: DTACQSolutio_d3:00:85 (00:21:54:d3:00:85)
      ... ..0. .... = LG bit: Globally unique address (factory default)
      ... ..0. .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
    [Stream index: 3]
    Padding: 00000000000000000000000000000000
  Internet Protocol Version 4, Src: 10.12.198.16, Dst: 10.12.198.254
    0100 ... = Version: 4
    ... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 29
    Identification: 0x0000 (0)
    010. .... = Flags: 0x2, Don't fragment
    ... 0 0000 0000 0000 = Fragment Offset: 0
    Time to Live: 64
    Protocol: UDP (17)
    Header Checksum: 0x99a9 [validation disabled]
    [Header checksum status: Unverified]
    Source Address: 10.12.198.16
    Destination Address: 10.12.198.254
    [Stream index: 1]
  User Datagram Protocol, Src Port: 53676, Dst Port: 53676
    Source Port: 53676
    Destination Port: 53676
    Length: 9
    Checksum: 0x665b [unverified]
    [Checksum Status: Unverified]
    [Stream index: 0]
    [Stream Packet Number: 1]
  [Timestamps]
  UDP payload (1 byte)
  Data (1 byte)
    Data: 55
    [Length: 1]
    
```

Figure 4: Example ARP Request Data Packet

3.2 1G Streaming

This example shows the streaming procedure for 1G data rates. For higher data/packet rates a more tailored receiver program may be required. See Section 3.3.

Packets will be received on the host by the netcat program and redirected to a file for later analysis.

We start the `nc` process on the host, then begin a capture on the `acq2106`. In this way we always catch the beginning of the shot.

Open two terminals on your host. From the first we will run the `nc` command and from the other we will control the start/stop of the `acq2106` stream process.

```
$ nc -ulv 10.12.198.254 53676 | pv > shot_data
Ncat: Version 7.50 ( https://nmap.org/ncat )
Ncat: Listening on 10.12.198.254:53676
Ncat: Connection from 10.12.198.128.
113MiB 0:00:15 [23.2MiB/s] [ <=>
]
```

```
$ python3 ./user_apps/acq400/acq400_continuous.py \
--run=1 $UUT
$
$ # Wait for the nc > file to fill with data
$
$ python3 ./user_apps/acq400/acq400_continuous.py \
--stop=1 $UUT
```

We can use `hexdump` to take a quick look at the data. Here we have received data comprised of 32 LWord samples, where the first 16 LWords are analog data and the remaining 16 are SPAD values. The SPAD values begin on the 17th entity (counting from 1) and we can see this clearly by looking for the sample counter. I have forced some fake headings into the `hexdump` output below to illustrate this.

The use of `cut` here means we don't have to look at all 32 values, thus saving ourselves some space in our terminal.

```
[dt100@naboo UDP_TEST]$ hexdump -e '32/4 "%08x " "\n"' shot_data | cut -d ' ' -f 1-4,17-20 | head
ANALOG ANALOG ANALOG ANALOG COUNTER XX XX SIGNATURE
ea631a50 f571ff61 ffeffff9 fff10000 00000001 00000000 00561657 00000000
ea651a21 f571ff5f fffffff8 fff10000 00000002 00000000 0056171f 00000000
ea6219f0 f572ff5f fffffff9 fff10000 00000003 00000000 005617e7 00000000
ea6319c2 f572ff5f fffffff9 fff10000 00000004 00000000 005618af 00000000
ea671993 f572ff5f fffffff8 fff10000 00000005 00000000 00561977 00000000
ea691964 f572ff61 fffffff9 fff10001 00000006 00000000 00561a3f 00000000
ea691935 f572ff60 fffffff9 fff10000 00000007 00000000 00561b07 00000000
ea641904 f572ff63 fffffff8 fff00000 00000008 00000000 00561bcf 00000000
ea6718d5 f572ff63 fffffff9 fff10000 00000009 00000000 00561c97 00000000
ea6518a6 f572ff64 fffffff8 fff10000 0000000a 00000000 00561d5f 00000000
```

We will plot this data graphically in the Section 3.4.

3.3 10G Streaming with udprx.cpp

For updperf install instructions see Section 2.2.

In order to support high data/packet rates we have to use some lower level code. udprx provides this. The program supports data reception and optional sequence number verification.

```
dt100@staffa:~/PROJECTS/udpperf$ ./udprx -h
UDP receiver with 32 bit sequence number check.
Usage: ./udprx [OPTIONS]

Options:
  -h,--help                Print this help message and exit
  -p,--port INT            UDP receive port
  -b,--socket_buffer_size INT socket buffer size (bytes)
  --spp INT                Samples per packet
  --ssb INT                Sample size (bytes)
  -c,--count_column INT    Count column (indexed from 0)
  -t,--step INT            Count step (default:1), but may be decimated
  -R,--rt_prio INT         set POSIX RT priority (0: no set)
  -o,--output INT          1: output data to stdout
  -q,--quiet INT           1: stop reporting
  -S,--max_samples UINT    stop after this many samples, 0: no limit
  -M,--max_errs INT        stop after this many errors
  -a,--local_address TEXT  optional local address
                           eg multiple NICs, one port
```

From here we assume the user has already understood the function and options provided in hudp_setup.py detailed in Section 3.

3.3.1 Configure the HUDP Module

```
dt100@staffa:~/PROJECTS/acq400_hapi$ UUT=acq2206_014
dt100@staffa:~/PROJECTS/acq400_hapi$ python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 \
--run0='1,2,3,4,5,6 1,16,0' --spp=20 $UUT none
txuut acq2206_014
TX configured. ssb:448 spp:20 tx_pkt_size 8960
```

3.3.2 Validate performance with sequence check

Prime udprx to receive packets then start a stream on the UUT.

```
dt100@staffa:~/PROJECTS/udpperf$ sudo ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -c 96
0x00000001  1 ini
0x00000002  2 ini
0x00000003  3 ini
0x00000004  4 ini
0x00000005  5 ini
Rx rate: 0.19 Mbps, rx 0 MB (total: 0 MB), Elapsed 00:00:00, ErrCount = 0, PER 0.000e+00
Rx rate: 3587.64 Mbps, rx 448 MB (total: 4486 MB), Elapsed 00:00:10, ErrCount = 0, PER 0.000e+00
Rx rate: 3587.66 Mbps, rx 448 MB (total: 8971 MB), Elapsed 00:00:20, ErrCount = 0, PER 0.000e+00
Rx rate: 3587.69 Mbps, rx 448 MB (total: 13457 MB), Elapsed 00:00:30, ErrCount = 0, PER 0.000e+00
```

3.3.3 Receive data and store to disk

Be careful when dealing with high data rates! You don't want to fill your disk!

```
sudo ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -q 1 -o 1 > shot_data
```

3.3.4 Receive data and store to disk - Max Samples argument

Capture exactly 1 million samples

```
sudo ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -q 1 -o 1 -S 1000000 > shot_data_1M
```

3.3.5 (Optional) Monitor progress while writing to disk

Pipe data through `pv` to monitor progress while writing to disk

```
dt100@staffa:~/PROJECTS/udpperf$ sudo ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -q 1 -o 1 | pv > shot_data  
6.66GiB 0:00:27 [ 427MiB/s] [ <=> ]
```

3.4 Plotting Sampled Data on Host

```
[dt100@naboo acq400_hapi]$ python3 ./user_apps/analysis/host_demux.py --save=DATA --pchan=1,33 \
--src="/home/dt100/UDP_TEST/shot_data" $UUT --pses=1:20000:1
INFO: plottext available as plot backup when no graphics
data_type 16 np <class 'numpy.int16'>
args.pc_list [0, 32]
data saved to directory: /home/dt100/UDP_TEST/DATA
Plotting with Matplotlib. Subrate = 1
plot_mpl num 1
```

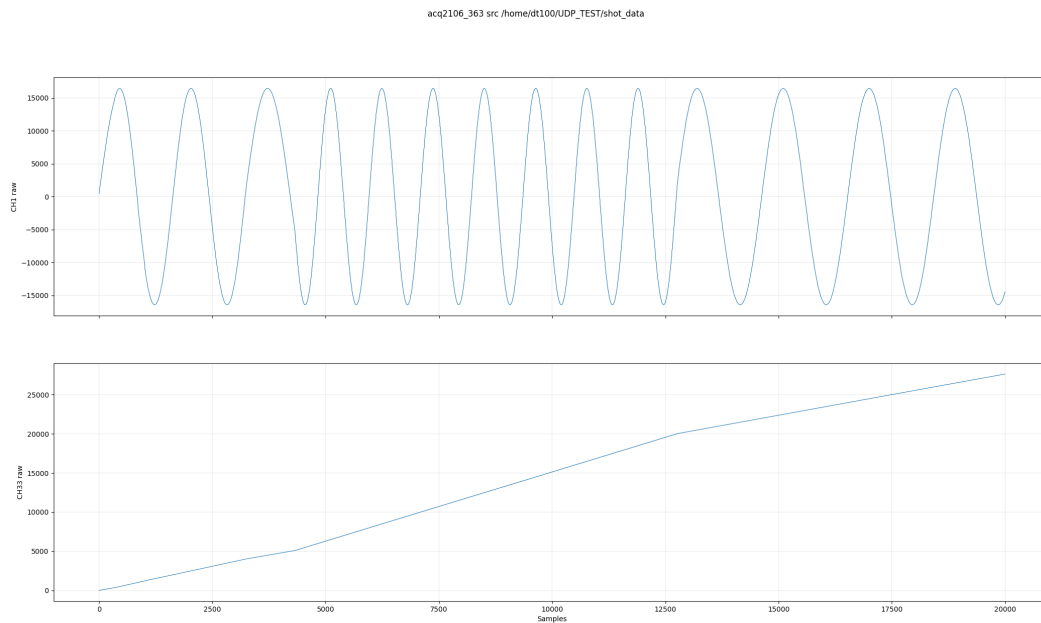


Figure 5: An example of packet loss when we overwhelm the host with too many packets per second

Note the unstable sine wave and non-contiguous sample counter indicating that packets have been lost. In this case we know through previous testing that this particular host can only cope with 80,000 samples per second, and here we are sending 200,000 per second.

For low latency applications we want to minimise our packet size so that we do not incur a delay by waiting for a packet to fill with user data. For streaming applications, this is obviously terribly inefficient.

We have two ways of dealing with this problem.

- Firstly, we can pack multiple samples into a packet; hence reducing the packet rate.
- Secondly, we can apply decimation in the HUDP module by discarding every X samples.

3.5 Multiple Samples per Packet

This time we run our setup script but include an argument for samples per packet. `--spp 10`

```
[dt100@naboo acq400_hapi]$ python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 \
--run0='1 1,16,0' acq2106_363 none --spp 10
txuut acq2106_363
TX configured. ssb:128 spp:10 tx_pkt_size 1280
```

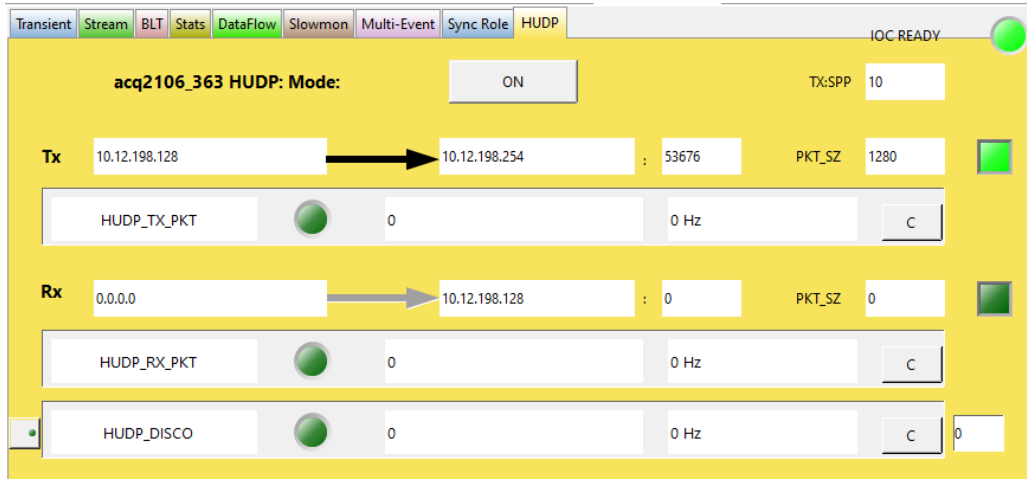


Figure 6: HUDP OPI with Multiple Samples per Packet showing new PKT_SZ

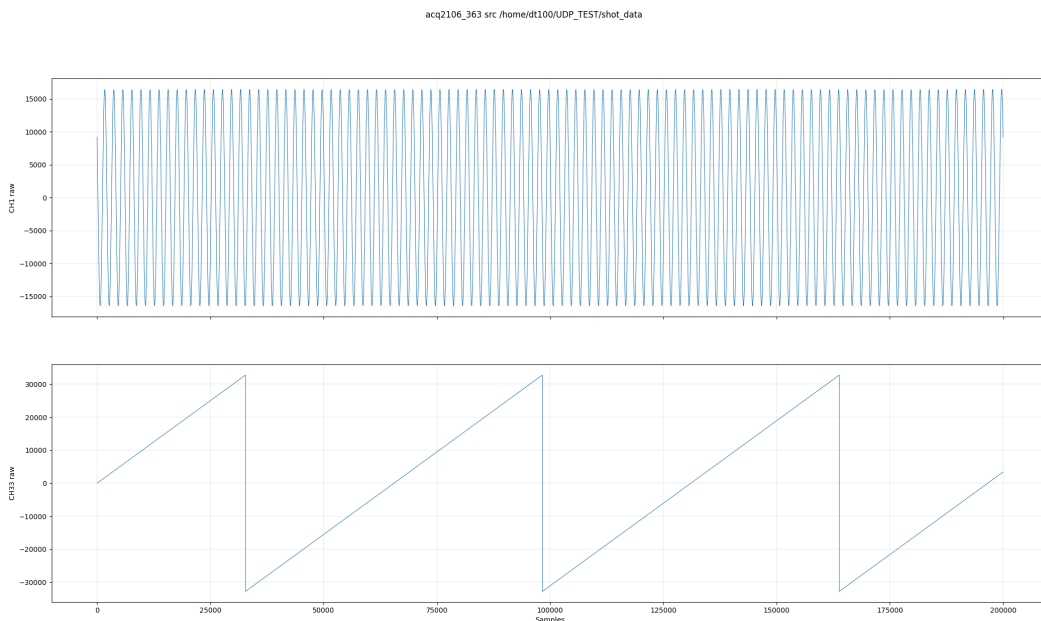


Figure 7: One whole seconds worth of data with zero packet loss

CH33 contains the bottom 16 bits of an embedded sample counter so it wraps at 32767, but what we are focusing on here is the lack of any other unexpected jumps across the entire shot. There are more robust ways of validating this transmission but this test is presented here as an easily understood, basic example.

3.6 Decimation

We re-run the setup script again to return to 1 sample per packet. But now we will apply sample decimation in the HUDP module to lower the packet rate **AND** the data rate, instead. `--hudp_decim 10`

```
[dt100@naboo acq400_hapi]$ python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 \
--run0='1 1,16,0' acq2106_363 none --hudp_decim 10
txuut acq2106_363
TX configured. ssb:128 spp:1 tx_pkt_size 128
```

We requested a decimation of 10. This means we discard 9 samples between outputs, a "decimation" of factor 10.

We can see here that the MiB/s rate has dropped by a factor of 10 relative to full rate example in Section 3.2.

```
[dt100@naboo UDP_TEST]$ nc -ulv 10.12.198.254 53676 | pv > shot_data
Ncat: Version 7.50 ( https://nmap.org/ncat )
Ncat: Listening on 10.12.198.254:53676
Ncat: Connection from 10.12.198.128.
50.5MiB 0:00:31 [2.44MiB/s] [ <=>
```

The effective sample rate has now decreased so our sampled waveform has a higher apparent frequency than the plots earlier in this document.




Figure 8: Decimated data showing effective decreased sample rate

Note also that the sample counter now jumps in increments of 10 samples.

```
[dt100@naboo UDP_TEST]$ hexdump -e '32/4 "%08x " "\n"' shot_data | cut -d ' ' -f 1-4,17-20 | head
ANALOG ANALOG ANALOG ANALOG COUNTER XX XX SIGNATURE
e5c0c068 f490ffd7 ffbfff3 ffe0006 00000001 00000000 204ac4af 00000000
e5b3c0b4 f490ffd7 ffbfff3 fff00007 0000000b 00000000 204acc7f 00000000
e5b1c113 f491ffd6 fffcfff2 ffe0006 00000015 00000000 204ad44f 00000000
e5b6c184 f491ffd6 fffaaff2 fff00006 0000001f 00000000 204adc1f 00000000
e5aac202 f490ffd5 fffaaff3 ffe0006 00000029 00000000 204ae3ef 00000000
e5a5c290 f491ffd5 ffbfff4 fff00006 00000033 00000000 204aebbf 00000000
e5a9c32f f491ffd4 ffbfff3 fff00006 0000003d 00000000 204af38f 00000000
e5a2c3dd f491ffd4 ffbfff3 fff00005 00000047 00000000 204afb5f 00000000
e592c498 f491ffd2 ffbfff3 fff00005 00000051 00000000 204b032f 00000000
e58ec562 f491ffd1 fffcfff4 fff00005 0000005b 00000000 204b0aff 00000000
```

3.7 Streaming Multiple Sites of Data

```

1 D-TACQ Solutions ACQ423ELF ACQ423ELF-32-200-16-FFC N=32 M=09 E42310229
2 D-TACQ Solutions ACQ423ELF ACQ423ELF-32-200-16-FFC N=32 M=09 E42310230
6 D-TACQ Solutions DI0482ELF DI0482ELF N=32 M=6B E48220160

```

Let us now include all of the “input” capable sites in the Aggregator set. We will also reduce the length of the SPAD to give a nice round number. Note, `run0='1,2,6 1,15,0'`

We'll also decimate by 10 to keep the packet rate to the host down.

```

[dt100@naboo acq400_hapi]$ python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 \
--run0='1,2,6 1,15,0' acq2106_363 none --hudp_decim 10
txuut acq2106_363
TX configured. ssb:192 spp:1 tx_pkt_size 192

```

The output from the `hudp_setup` script helpfully tells us our samples size in bytes (ssb). So we can see that our sample size in LWords is now 48 (192÷4). The new hexdump is presented below. Again, context headings have been added to aid understanding.

```

[dt100@naboo UDP_TEST]$ hexdump -e '48/4 "%08x " "\n" shot_data | cut -d ' ' -f 1,2,17,18,33-37 | head
1      2      17      18      33      34      35      36      37
ACQ423_1 ACQ423_1 ACQ423_2 ACQ423_2 DI0482 COUNTER XX      XX      SIGNATURE
ffff0000 ffffffff 00060005 00000000 00000000 00000001 00000000 906accb5 00000000
ffffffff 00000000 00060006 00000000 00000000 0000000b 00000000 906ad485 00000000
0000ffff ffffffff 00050005 00000000 00000000 00000015 00000000 906adc55 00000000
fff00000 ffff0000 00060006 00000000 00000000 0000001f 00000000 906ae425 00000000
0000ffff ffffffff 00060005 ffff0000 00000000 00000029 00000000 906aebf5 00000000
ffffffff 00000000 00050005 0000ffff 00000000 00000033 00000000 906af3c5 00000000
ffffffff 00000000 00060005 00000000 80000000 0000003d 00000000 906afb95 00000000
ffffffff 00000000 00060005 00000000 80000000 00000047 00000000 906b0365 00000000
ffffffff 0000ffff 00060005 00000000 80000000 00000051 00000000 906b0b35 00000000
ffffffff 00000000 00060006 00000000 80000000 0000005b 00000000 906b1305 00000000

```

4 Streaming UDP Data to acq2106 XO Modules

In this Section we demonstrate sending data from the host to the acq2106 over UDP and playing it out on the AO424 module. For simplicity we will use previously sampled data to replay on the AO424. Process below :

- Sample analog data on an ADC module
- Stream this data over UDP to the host
- Store the data in a file
- Stream this data file back to the box over UDP to be played out by the DAC module

We will limit the sample/update rate to 20 kHz so as not to stress the host through high frequency packet generation. A C program has been created to take a source file and send fixed length packets to a specific IP address and port. This is available here : <https://github.com/petermilne/HUDP>

```
# Set the sample rate to 20 kHz
python3 ./user_apps/acq400/sync_role.py --fclk=20k --toprole=master,strg acq2106_363

# Configure the acq2106 HUDP TX to send us 32 channels of analog data with 16 LWords of SPAD
python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 --run0='1 1,16,0' acq2106_363 none

# Capture some data to file
nc -ulv 10.12.198.254 53676 | pv > shot_data

# Configure the acq2106 HUDP RX to accept packets from IP addr and port
# Also, set up the distributor to send 32 channels of data to Site 3 and discard the SPAD LWords
python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.254 --rx_ip 10.12.198.128 --play0='3 16' none acq2106_363

# Stream the data file back to the acq2106
# Here 50 usecs equates to 20 kHz, matching our ADC sample rate
[dt100@naboo HUDP]$ ls -lh shot_data
-rw-r--r-- 1 dt100 d-tacq 104M Nov 23 11:00 shot_data
[dt100@naboo HUDP]$
[dt100@naboo HUDP]$ ./udp_client
usage: udpclient <IP address> <port> <sz> <file> <usecs>
[dt100@naboo HUDP]$ ./udp_client 10.12.198.128 53676 128 shot_data 50
```

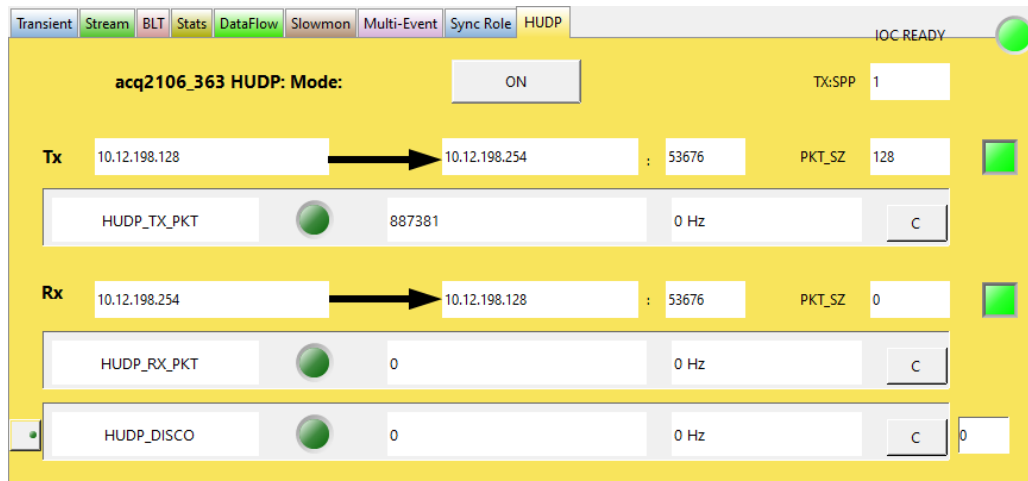


Figure 9: HUDP OPI showing TX packet count and RX variables set

After we run the last command in the listing above, the RX packet counter should be counting and the DAC should now be updating.

We can see the result of this in Figure 10. The yellow trace shows us what the signal generator is producing, and hence what the ADC had captured. The blue trace is live output from the DAC showing the signal faithfully recreated.



Figure 10: Oscilloscope trace showing data from signal generator in yellow and similar data from DAC in blue

5 Monitoring Packets with Wireshark

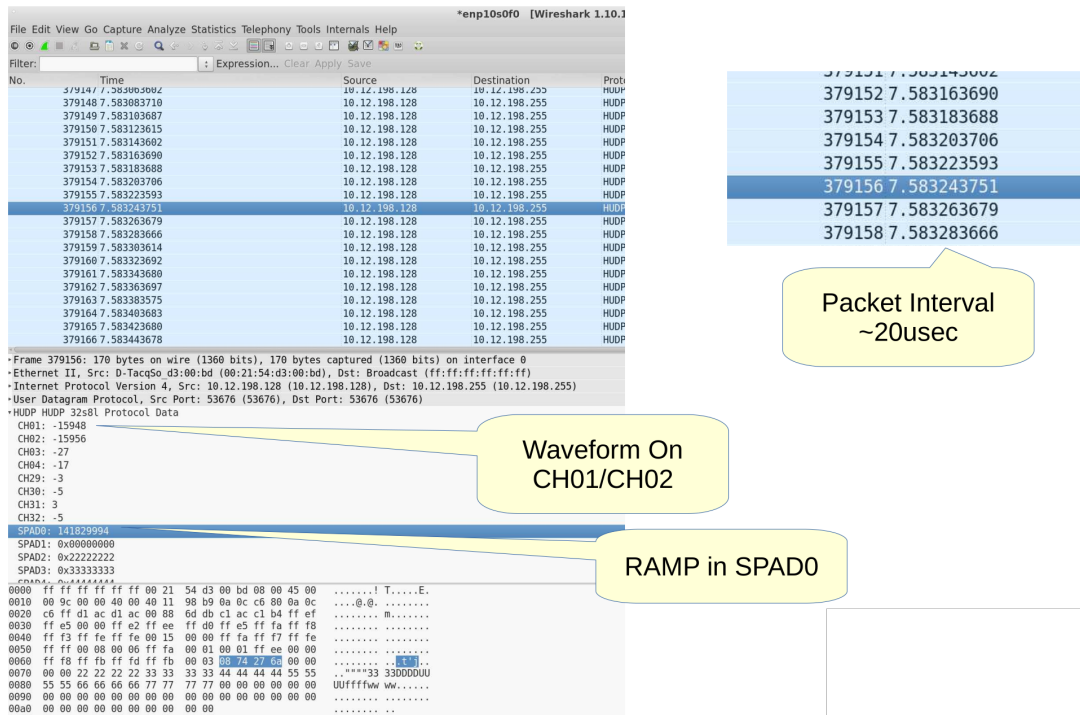


Figure 11: Monitoring packets with Wireshark plugin

https://github.com/D-TACQ/CUSTOM_WRPGB/blob/master/CARE/dot.wireshark.plugins.hudp32s_8l.lua

6 Low Latency Demo Setup

6.1 HW

- AI Box : acq2106_189
- AO Box : acq2106_274




Figure 12: HUDP Demo Hardware Configuration

We set the AI box clock to bypass so that there is no phase delay between the clock from the signal generator and the clock which the ADC sees. The AO is in Low Latency mode (update as soon as you have data) so the clock on the AO box is less important.

Bring the clock back out of HDMI from the 2106 before displaying on scope. This tells us for certain that the Si5326 is successfully in bypass mode and has the added benefit of showing the user where the sample clock lies in relation to their excite signal.

The signal generator which is providing the clock signal, is also driving the External Trigger input of a secondary function generator. This secondary function generator outputs a square wave burst on the rising edge of the clock (not on every edge, burst period is longer than clock period).

We can then control the phase of the burst signal relative to the rising edge of the clock to account for the aperture delay and slew rate of ACQ424 input.

See figures in Section 6.6 for a detailed look at the oscilloscope traces and latency measurements.

6.2 HUDP Core Configuration

D-TACQ have provided a handy Python interface to help when configuring the HUDP Core. This is further discussed in Section 3.1.

For our example above (Figure 12), the configuration is as follows.

We execute a `sync_role` command to place the AI box's clock synthesiser in bypass.

```
./user_apps/acq400/sync_role.py --fin=50k --fclk=50k --si5326_bypass 1 --toprole=fpmaster,strg acq2106_189
```

Then we set up a TX-RX pair between box `acq2106_189` and `acq2106_274`. Here we use the broadcast address `.255` so that the PC which is also in the network can easily spy on the packets as they fly past.

```
./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.255 --run0='1 1,16,0' --play0='1 16' \
--broadcast=1 --disco=16 acq2106_189 acq2106_274
```

6.3 CSS OPI

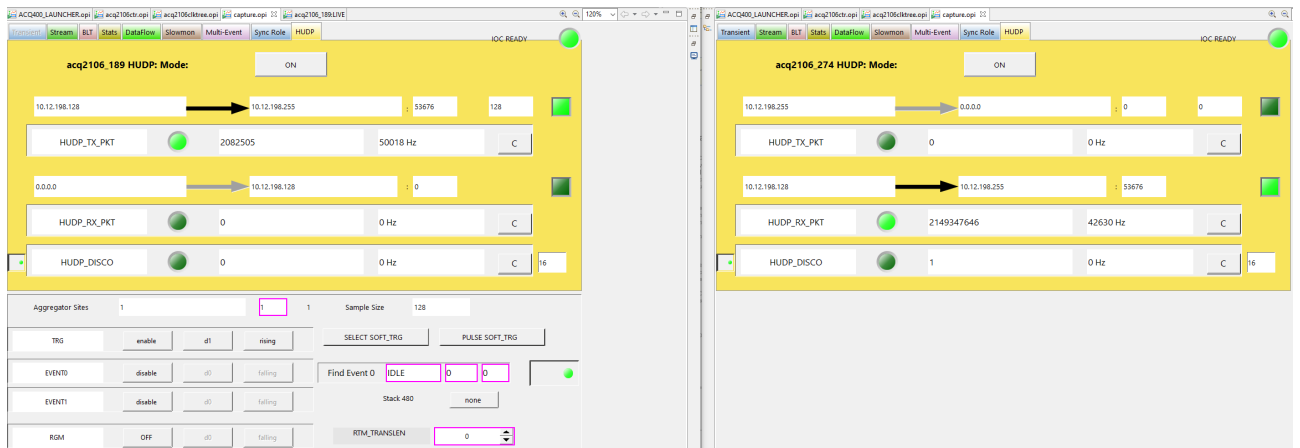


Figure 13: HUDP OPIs : A TX-only box on the left and an RX-only box on the right

6.4 Discontinuity Counter

The discontinuity (disco) counter uses the embedded sample counter (SPAD0) contained within every sample/packet to verify that we have had no skips or lost packets throughout the course of a run.

Check that the latest received sample count is equal to the previously decoded sample count plus 1. Note, this will fire once at the start of every shot (when our previous sample count is assumed to be zero and the first received sample count will also be zero.) It will also fire spuriously on the rollover of the 32-bit counter (once a day at 50 kHz). We could enhance the logic to account for these exceptions.

6.5 Clock and Excite Signal Phase Relationship

Here we demonstrate the effect of sampling the analogue rising edge at different points in the low-to-high transition. Even though the signal generator produces a very sharp rising edge, a combination of the bandwidth and slew rate limitations of the front end of the ADC mezzanine, "flatten out" the rising edge somewhat.

- With the sample clock almost coincident with the analogue rising edge (far left) the ADC only manages to capture the very beginning of the rising edge, resulting in a very low amplitude update on the DAC.
- Pulling the excite signal forward in time (middle section) means that the sample clock is in effect, delayed, and thus samples further up the slope of the rising edge. This results in a higher amplitude update from the DAC.
- With the rising edge now having almost completed its low-to-high transition before the sample clock arrives, the ADC samples much further up the slope and the output from the DACs is almost full scale.




Figure 14: The result of the sample clock sampling the rising edge at various points during the signal slew

■ ADC sample clock ■ Sig-gen ■ LLC Filter DAC Output ■ Std. Filter DAC Output

6.6 Latency Measurements

6.6.1 Point-to-Point

Best possible result. AI box straight to AO box. Very stable. 12.6 μ s. We should repeat this test using a 2 MHz ACQ425 (smallest possible aperture delay) and an AO424-LLC in 16CH mode to minimise both input and output delay as well as the smallest possible packet. This could then be compared to the Aurora latency record of 5.1 μ s from the D-TACQ LLC White Paper.



Figure 15: HUDP Point-to-Point Latency

6.6.2 Through Switch

Note the extra jitter incurred by the trip through the switch in addition to the further \approx 5 μ s of latency. 17.4 μ s.




Figure 16: HUDP Latency through Switch

6.7 HUDP - Internal workings

Here we have a more detailed diagram which helps visualise data flows, both within an individual box employing the HUDP core, and in a larger system utilising HUDP as the backbone for data transfer.

Direct Data Feed With HUDP

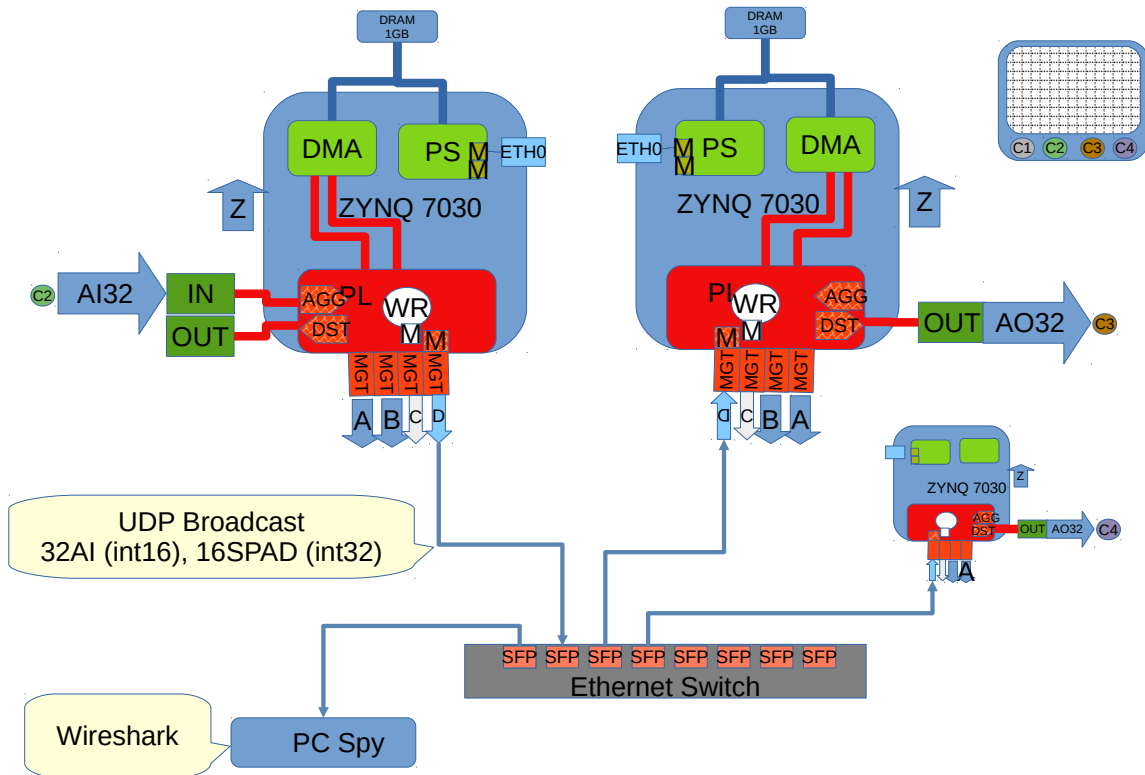


Figure 17: More detailed example of HUDP dataflows

This helps to visualise the entire datapath. Beginning at C2 AI32 we see the data flow into the site, through the Aggregator, down to the HUDP core and associated Multi-Gigabit Transceiver (MGT) port. We push packets out on SFP, through the Ethernet switch and back to the RX box. These packets hit the HUDP core on the AO box, data is stripped and passed down to the Distributor which is responsible for sharing out the data to selected sites. From here we carry on to the AO32 site and eventually back out into the analogue domain to be sampled by the oscilloscope (C3).

7 Frame Specifications

7.1 Default MTU

Frame Section	Size (Bytes)	Comment	Size (Bytes)
Inter-Frame Gap (IFG) (96 ns)	12	(125 MB/s × 96 ns = 12B)	12
MAC Preamble	8	including Start of Frame Delimiter (SFD)	8
MAC Destination Address	6		6
MAC Source Address	6		6
EtherType (or Length)	2		2
IPv4 Header	20	Payload MTU 1500B	20
UDP Header	8		8
User Data (CALC_PKT_SIZE)	1472		1472
Frame Check Sequence (FCS)	4		4
Total Frame Size	1518	Total w/ overheads (includes IFG and Preamble)	1538

Table 1: Breakdown of an Ethernet Frame from HUDP Core

7.1.1 1G

Gigabit Ethernet raw line rate = 1.25 Gb/s,
accounting for 8B/10B encoding leaves us with 1 Gb/s,
this equates to 125 MB/s.

So we can transfer 125 million bytes per second, and our maximum frame size is 1538 bytes. This gives us a maximum packet rate (with a 1538B frame size) of :

$$\frac{125 \times 10^6}{1538} \approx 81,274 \text{ frames per second}$$

This means that our maximum rate for transferring (useful) user data is as follows :

$$81,274 \times 1472\text{B} = 119.635 \text{ MB/s} \\ \approx 114 \text{ MiB/s}$$

From the above we can see that our total overhead is equal to :

$$1538 - 1472 = 66\text{B}$$

To calculate our max packet rate for any sample size, and hence our maximum sample rate in a low latency system, we would use the following equation :

$$\frac{125 \times 10^6}{(\text{Sample Size} + 66)} = \text{Max. packets per second} \quad (1)$$

7.2 Jumbo Frames

Frame Section	Size (Bytes)	Comment	Size (Bytes)
Inter-Frame Gap (IFG) (96 ns)	12	(125 MB/s × 96 ns = 12B)	12
MAC Preamble	8	including Start of Frame Delimiter (SFD)	8
MAC Destination Address	6		6
MAC Source Address	6		6
EtherType (or Length)	2		2
IPv4 Header	20	Payload MTU 9000B	20
UDP Header	8		8
User Data (CALC_PKT_SIZE)	8972		8972
Frame Check Sequence (FCS)	4		4
Total Frame Size	9018	Total w/ overheads (includes IFG and Preamble)	9038

Table 2: Breakdown of an Ethernet Frame from HUDP Core

Ethernet Jumbo Frames commonly carry payloads of up to 9038 bytes.

7.2.1 1G

$$\frac{125 \times 10^6}{9038} \approx 13,830 \text{ frames per second}$$

We still have the same overall overhead, so our max user data length is :

$$9038 - 66 = 8972\text{B}$$

This means that our maximum rate for transferring (useful) user data is as follows :

$$13,830 \times 8972\text{B} = 124.08 \text{ MB/s}$$

$$\approx 118 \text{ MiB/s}$$

7.2.2 10G

10 Gigabit Ethernet raw line rate = 10.3125 Gb/s,
accounting for 64B/66B encoding leaves us with 10 Gb/s,
this equates to 1.25 GB/s.

So we can transfer 125 billion bytes per second, and our maximum (jumbo) frame size is 9038 bytes. This gives us a maximum packet rate (with an 9038B frame size) of :

$$\frac{1.25 \times 10^9}{9038} \approx 138,304 \text{ frames per second}$$

We still carry the standard 66 byte overhead per frame, thus our useful data per frame is 8972B. This means that our maximum (theoretical) rate for transferring (useful) user data is as follows :

$$\begin{aligned} 138,304 \times 8972B &\approx 1.24 \text{ GB/s} \\ &\approx 1.15 \text{ GiB/s} \end{aligned}$$

8 Troubleshooting

Things to check :

- Firewall configuration
- SFP Plugged in? See Figure 18
- Link Status
 - `sudo ethtool NETWORK-INTERFACE`
 - LINK_STATUS on the hudp system webpage (`http://$UUT/d-tacq/#hudp`)
 - If no link UP, try a `set.site 10 hudp_gt_reset P # P for pulse`
- Valid ARP Response? See Figure 19
- Do you have a switch in the datapath?
 - Has it been configured to support jumbo frames?

8.1 sfp webpage

The HUDP SFP is plugged into SFP4/D and we can see both Tx and Rx power.

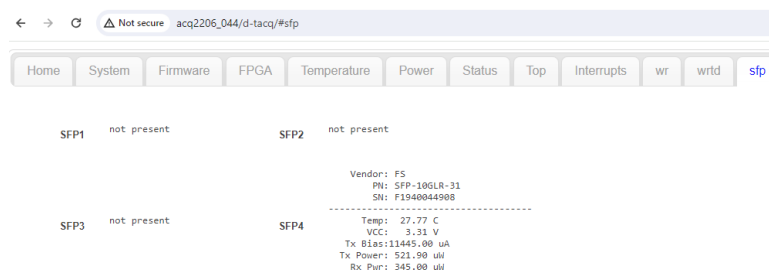


Figure 18: Example of SFP status on webpage

8.2 hudp webpage

shows the current HUDP setup. Note that with `arp_max_resp`, the destination mac address is only probed via ARP at the start of capture. If a capture fails to transfer data, and `arp_max_resp` is zero, the ARP process has failed.

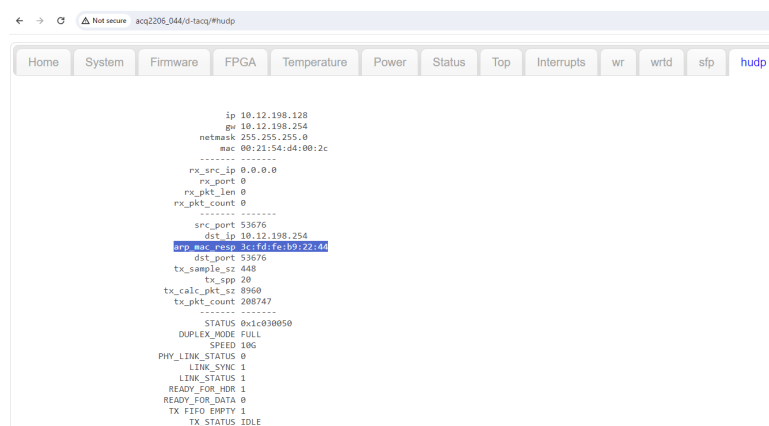


Figure 19: Example of statuses exposed on the hudp webpage

8.3 tcpdump

Use tcpdump to capture one packet arriving to port 53676 on interface enp10s0f0.

This is before we hit the firewall so is useful to verify if packets are flowing at a low level.

```
dt100@naboo:~/PROJECTS/udpperf$ sudo tcpdump -i enp10s0f0 -n udp port 53676 -X -c 1
tcpdump: verbose output suppressed, use -v[v]... for full protocol decode
listening on enp10s0f0, link-type EN10MB (Ethernet), snapshot length 262144 bytes
11:13:05.121632 IP 10.12.198.128.53676 > 10.12.198.254.53676: UDP, length 8960
    0x0000:  4500 231c 0000 4000 4011 763a 0a0c c680  E.#...@.v:....
    0x0010:  0a0c c6fe d1ac d1ac 2308 c8af f407 e507  .....#.....
    0x0020:  fa07 f407 f907 eb07 f307 eb07 f207 eb07  .....
    0x0030:  f907 ea07 f607 f307 fb07 f207 f607 f507  .....
    0x0040:  ee07 f107 ef07 e707 ec07 ed07 f207 e907  .....
    0x0050:  ec07 f107 f307 f507 f707 e707 f2ff f9ff  .....
```

A Example run on D-TACQ Host (naboo)

Example UUT has 6 sites

- Each has 32 channels of 16-bit ADC = $6 \times 32 \times 2B = 384 B$
- Plus $16 \times 4B = 64 B$ of Scratchpad (SPAD) instrumentation
- Total sample size of $384 + 64 = 448 B$

SPAD0 holds an included sample counter. In this example, breaking the sample down into 32b (4B, 1 LWord) quantities, the sample count is the 96th element (counting from zero). This provides us with the `-c 96` argument to the validation script `udprx`.

Various parameters may need to change if your HW configuration does not match the example precisely.

A.1 Set ACTIVE_NIC

```
ACTIVE_NIC=enp10s0f0
```

A.2 Enable Jumbo frames and set host optimisations

```
sudo sysctl -w net.core.rmem_max=26214400
sudo sysctl -w net.core.wmem_max=12582912
sudo sysctl -w net.core.netdev_max_backlog=5000
sudo ifconfig $ACTIVE_NIC 10.12.198.254 netmask 255.255.255.0 mtu 9000 txqueuelen 10000 up
```

A.3 Set up a full ACQ424 box and maximise packet length

```
cd ~/PROJECTS/acq400_hapi
source setpath
```

```
python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.254 --run0='1,2,3,4,5,6 1,16,0' \
--spp=20 acq2206_013 none
```

A.4 Force a negotiation (optional, if LINK_STATUS is GOOD then skip)

Can't seem to enable auto-negotiation on the host interface... this is required to settle the block lock state machine in the PCS/PMA core. See **LINK_STATUS** knob on hudp webpage before and after forcing a negotiation.

10G Ethernet PCS/PMA (10GBASE-R) does not support auto-negotiation. This is probably the issue.

```
# ethtool -r|--negotiate devname
sudo ethtool -r $ACTIVE_NIC
```

A.5 Spin up udprx on the host to receive and validate packets

```
cd ~/PROJECTS/udpperf
```

```
sudo ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -c 96
```

A.6 Start a stream and observe packet reception statistics

```
Rx rate: 0.40 Mbps, rx 0 MB (total: 0 MB), Elapsed 00:00:00, ErrCount = 0, PER 0.000e+00
Rx rate: 3583.96 Mbps, rx 448 MB (total: 4480 MB), Elapsed 00:00:10, ErrCount = 0, PER 0.000e+00
Rx rate: 3584.02 Mbps, rx 448 MB (total: 8961 MB), Elapsed 00:00:20, ErrCount = 0, PER 0.000e+00
Rx rate: 3583.98 Mbps, rx 448 MB (total: 13441 MB), Elapsed 00:00:30, ErrCount = 0, PER 0.000e+00
Rx rate: 3584.04 Mbps, rx 448 MB (total: 17922 MB), Elapsed 00:00:40, ErrCount = 0, PER 0.000e+00
Rx rate: 3583.98 Mbps, rx 448 MB (total: 22402 MB), Elapsed 00:00:50, ErrCount = 0, PER 0.000e+00
Rx rate: 3583.98 Mbps, rx 448 MB (total: 26882 MB), Elapsed 00:01:00, ErrCount = 0, PER 0.000e+00
```

B Reliability Testing

Example below for an ACQ425 system running at 2 MHz

```
python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.199.128 --rx_ip 10.12.199.254 \
--run0='1,2,3,4,5 1,8,0' --spp=45 acq2206_013 none
```

```
dt100@naboo:~/PROJECTS/udpperf$ cat hudp_stresser
UUT=$1
echo $UUT
runs=${2:-2}
minutes=${3:-10}
samples=$(( $minutes*60*2000000 ))
echo Capture $samples samples

for i in $(seq -w 1 $runs)
do
    echo Run $i
    echo Run $i > logs/run${i}.log
    date >> logs/run${i}.log
    sudo taskset -c 20 ./udprx -R 40 -p 53676 --spp 45 --ssb 192 -c 40 -S ${samples} >> logs/run${i}.log 2>&1 & PID_RX=$!
    python3 ../acq400_hapi/user_apps/acq400/acq400_continuous.py --run=1 $UUT
    wait $PID_RX
    python3 ../acq400_hapi/user_apps/acq400/acq400_continuous.py --stop=1 $UUT
done
```

```
dt100@naboo:~/PROJECTS/udpperf$ ./hudp_stresser acq2206_013 50 2
acq2206_013
Capture 240000000 samples
Run 01
Run 02
...
Run 49
Run 50
```

B.1 Interrogate HUDP stresser logs

```
dt100@naboo:~/PROJECTS/udpperf/logs$ for i in $(seq -w 1 50); do echo Run $i $(tail -1 run${i}.log);done
Run 01 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 02 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 03 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 04 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 05 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 06 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 07 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 08 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 09 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 10 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 11 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 12 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 13 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 14 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 15 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 16 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 17 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 18 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 19 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 20 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 21 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 22 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 23 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 24 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 25 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 26 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 27 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 28 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 29 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 30 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 31 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 32 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 33 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 34 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 35 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 36 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 37 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 38 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 39 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 40 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 41 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 42 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 43 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 44 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 45 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 46 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 47 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 48 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 49 RxRate: 384 MB/s (total: 42247 MB) Elapsed 00:01:50 PktRec = 4889800 ErrCount = 0 PktsLost = 0 PER 0.000e+00
Run 50 RxRate: 384 MB/s (total: 42246 MB) Elapsed 00:01:50 PktRec = 4889600 ErrCount = 0 PktsLost = 0 PER 0.000e+00
```

B.2 Optimising host for very high rates

B.2.1 Isolate a CPU

In this example we isolate CPUs 20 through 23 meaning that the operating system will not, by default, schedule any tasks on them. This leaves the user to specifically target these CPUs for uninterrupted performance using the `taskset` command.

```
vi /etc/default/grub
GRUB_CMDLINE_LINUX_DEFAULT="isolcpus=20,21,22,23"

sudo update-grub

# Check it worked
cat /proc/cmdline
BOOT_IMAGE=/vmlinuz-5.15.0-91-generic root=/dev/mapper/ubuntu--vg-ubuntu--lv ro isolcpus=20,21,22,23
```

Example of using `taskset` to tie the `udprx` process to CPU #20.

```
sudo taskset -c 20 ./udprx -R 40 -p 53676 --spp 20 --ssb 448 -c 96
```

C Validated Hardware

C.1 NICs

Vendor	Controller	Card Description
Intel	Intel X710	DA2 10GbE SFP+
-	Intel XL710	DA4
-	Intel E810	DA4 25GbE OCP3.0 - Only supported to 10GbE rates

Table 3: Table Comparing SNR figures for various input signals

C.2 SFP+ Transceiver Modules

- [FS SFP+ 10GBASE-LR Optical Transceiver Module](#)
- WTD RTX228-461
- Fiberstore GE-LC-1310

C.3 Optical Fiber

- [FS LC UPC Fiber Optic Patch Cable](#)

C.4 SFP+ Active Optical Fiber Cables - Transceiver and fiber in one fixed assembly

- [Gigalight GSS-MDO100-00XC SFP+ Active Optical Cable](#)

C.5 Copper

- N.B. **Passive** Direct Attach Copper cables are not supported.
- Active Direct Attach Copper : We haven't validated any Active DAC cables, please contact D-TACQ if you have a requirement.

C.6 NIC / Transceiver Vendor Lock

Please be aware some NICs can be very particular about which SFP+ modules they will communicate with. Here is an example of a driver message from the X710 when connected to the Gigalight GSS-MDO100-00XC.

```
dt100@naboo:~//$ sudo dmesg
[1386284.467239] i40e 0000:0a:00:01: Rx/Tx is disabled on this device because an unsupported SFP module type was detected.
[1386284.467251] i40e 0000:0a:00:01: Refer to the Intel(R) Ethernet Adapters and Devices User Guide for a list of supported modules.
```

C.7 How to query NIC and transceiver info

```
dt100@naboo:~/PROJECTS/udpperf$ sudo lshw -C network | grep product -A 1
product: RTL8125 2.5GbE Controller
vendor: Realtek Semiconductor Co., Ltd.
--
product: I211 Gigabit Network Connection
vendor: Intel Corporation
--
product: Ethernet Controller X710 for 10GbE SFP+
vendor: Intel Corporation
--
product: Ethernet Controller X710 for 10GbE SFP+
vendor: Intel Corporation
dt100@naboo:~/PROJECTS/udpperf$ sudo ethtool -m enp10s0f1 | grep Vendor
Vendor name      : FS
Vendor OUI      : 00:1b:21
Vendor PN       : SFP-10GLR-31
Vendor rev      : A
Vendor SN       : F1940044909
```

D Loopback to D-TACQ Fiber Ethernet Port

In an HUDP FPGA image, port D on the MGT482 module is employed as the HUDP endpoint, but in standard D-TACQ FPGA images, port D can be used as a second 1G Ethernet port for the Linux operating system.



Figure 20: View of MGT482-SFP Ports

Ports are labelled left to right “A B CWR D/ETH”

Users can loop an “HUDP” port back to a “Fiber Ethernet” port as a minimum viable test.

- Connect Port D on the HUDP system to Port D on the Fiber Ethernet system
- Configure the Fiber Ethernet
 - `ifconfig eth1 XXX.XXX.XXX.XXX up`
 - e.g. `ifconfig eth1 10.12.198.5 up`
- Program the HUDP module to target your newly configured IP address

```
python3 ./user_apps/acq2106/hudp_setup.py --tx_ip 10.12.198.128 --rx_ip 10.12.198.5 --run0='1 1,16,0' acq2106_367 none
```

- Prepare to receive data on the Fiber Ethernet system

```
nc -ulv -s 10.12.198.5 -p 53676 | pv > /tmp/shot_data
```

- Start streaming on the HUDP box, this will automatically forward data out of the HUDP port and it should be received by the Fiber Ethernet system

```
acq2106_130> nc -ulv -s 10.12.198.5 -p 53676 | pv > /tmp/shot_data
listening on 10.12.198.5:53676 ...
connect to 10.12.198.5:53676 from 10.12.198.128:53676 (10.12.198.128:53676)
7.01MiB 0:00:14 [1.67MiB/s] [
```

We can then hexdump the received data, as described in Section 3.2.

```
acq2106_130> hexdump -e '48/4 "%08x " "\n" ' /tmp/shot_data | cut -d ' ' -f 1-4,33-35 | head
00038220 00073721 fffc3022 fff04323 00000001 00000000 b0310454
00037b20 00074d21 fffc5322 fff01023 00000002 00000000 b0310654
00036d20 00075321 fffc3d22 fff02823 00000003 00000000 b0310854
00035420 00075721 fffc6522 fff00723 00000004 00000000 b0310a54
00036b20 00073a21 fffc7722 fff03a23 00000005 00000000 b0310c54
00033920 00074b21 fffc6322 fff02423 00000006 00000000 b0310e54
00034220 00074421 fffc3c22 fff02a23 00000007 00000000 b0311054
00035620 00074c21 fffc3722 fff04523 00000008 00000000 b0311254
00033b20 00076e21 fffc1f22 fff02423 00000009 00000000 b0311454
00032220 00075121 fffc3a22 fff04523 0000000a 00000000 b0311654
```


E Resource Usage

UDP Standard = (FIFO Depth 4096)

UDP Jumbo = (FIFO Depth 16384)

Site Type	Available	Used (No UDP)	Used (UDP Std.)	Std. Diff	Std. Util
Slice LUTs	78,600	36,794	40,019	3,225	4.10 %
Slice Regs	157,200	49,353	54,211	4,858	3.09 %
BRAM Tile	265	149.5	160	10.5	3.96 %

Table 4: Table Showing Resource Utilisation for UDP with Standard MTU

Site Type	Available	Used (No UDP)	Used (UDP Jumbo)	Jumbo Diff	Jumbo Util
Slice LUTs	78,600	36,794	40,050	3,256	4.14 %
Slice Regs	157,200	49,353	54,255	4,902	3.12 %
BRAM Tile	265	149.5	167	17.5	6.60 %

Table 5: Table Showing Resource Utilisation for UDP with Jumbo Frames

Site Type	Available	Std. Diff	Jumbo Diff	Diff	Diff %
Slice LUTs	78,600	3,225	3,256	31	0.04 %
Slice Regs	157,200	4,858	4,902	44	0.03 %
BRAM Tile	265	10.5	17.5	7	2.64 %

Table 6: Table Showing Extra Resource Required to implement Jumbo Frames

F HUDP 1G Logic Schematic View

Overleaf




Figure 21: 1G HUDP Logic Schematic View