

# Bristol NVMe Firmware

---

ADAM GILLARD

ag17009@bristol.ac.uk

# HiTech Global K800

---

NVMe firmware port to HiTech Global K800 initial testing complete

- Error in reading *Peak Latency* values in test software are suspected to be due to a clock domain crossing error
- This is further backed up by 14 critical warnings in Vivado CDC Report
- “Bug Report” has been written up and sent to BEAM

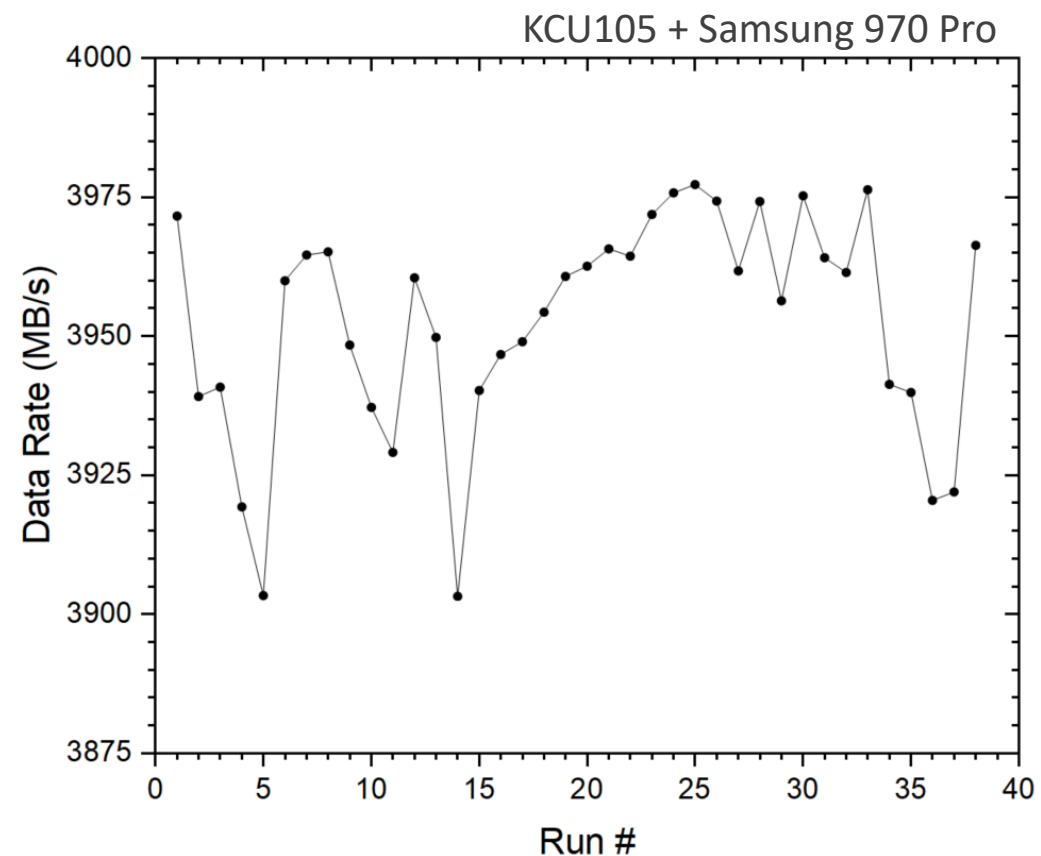
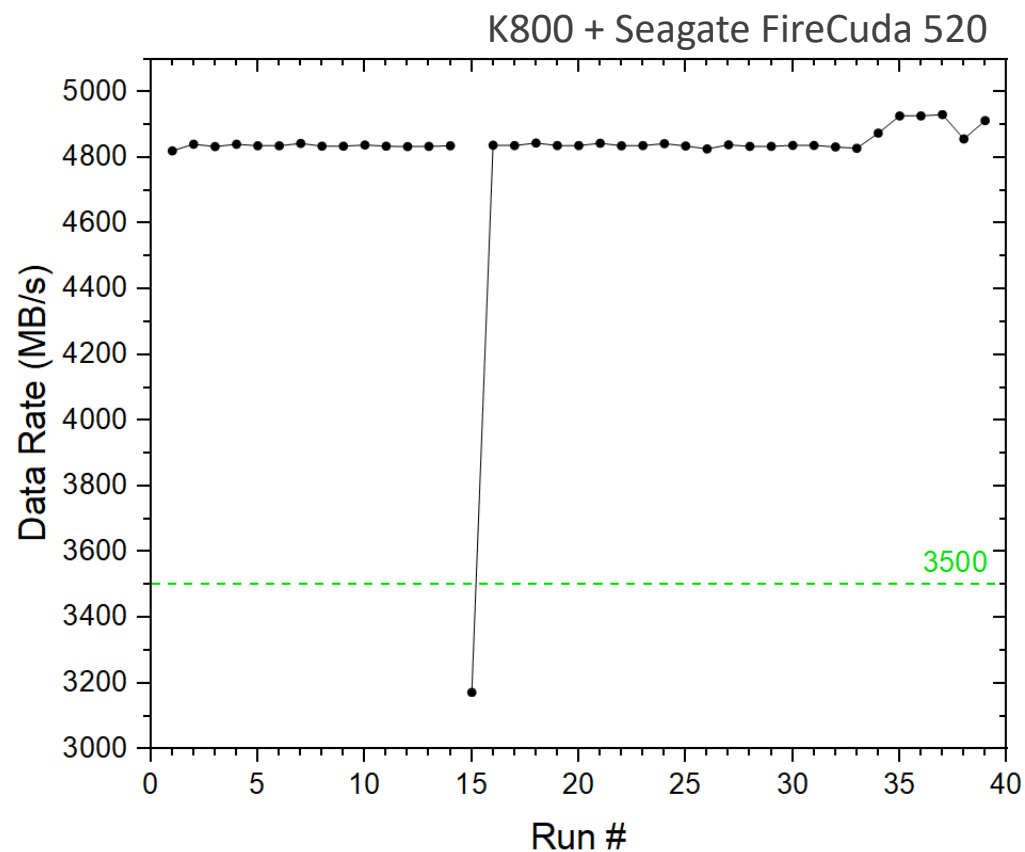
Build files can be found at <https://github.com/DUNE/pl-nvme/tree/HTG-K800>

- Includes build instructions
- *Full step-by-step instructions from KCU105 to HTG-K800 can be provided upon request*

Received test data from Rob Hallsal (RAL) using KCU105 + Samsung 970 Evo Drives

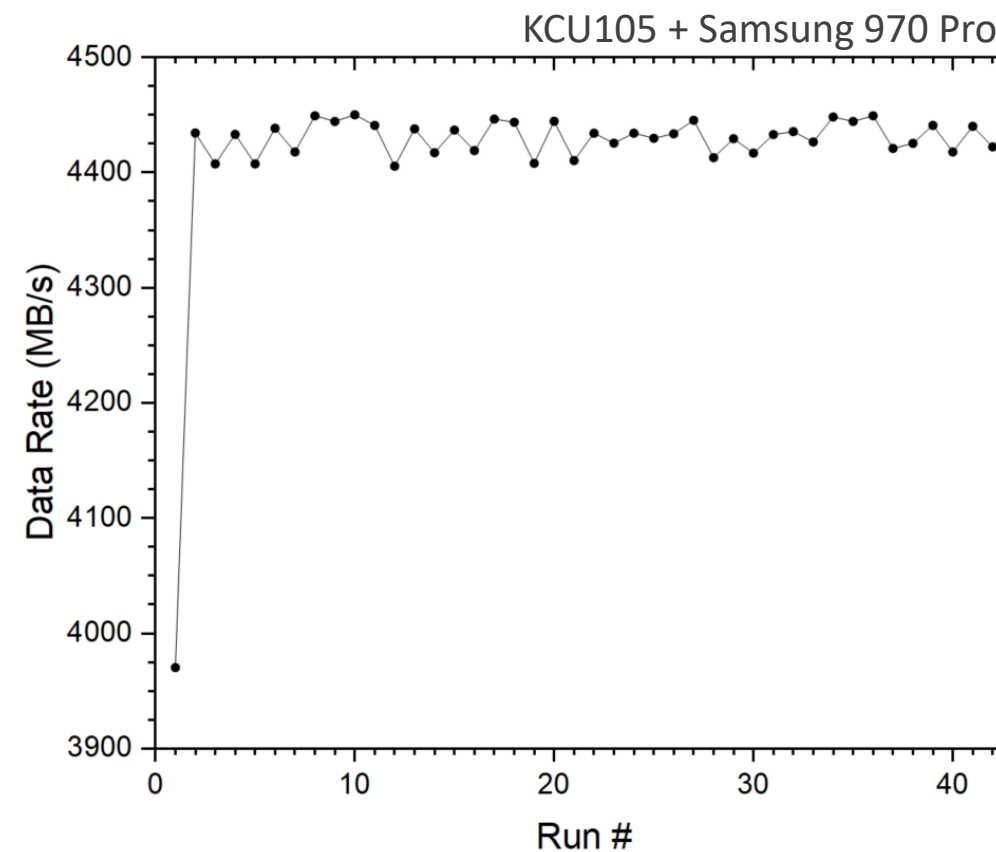
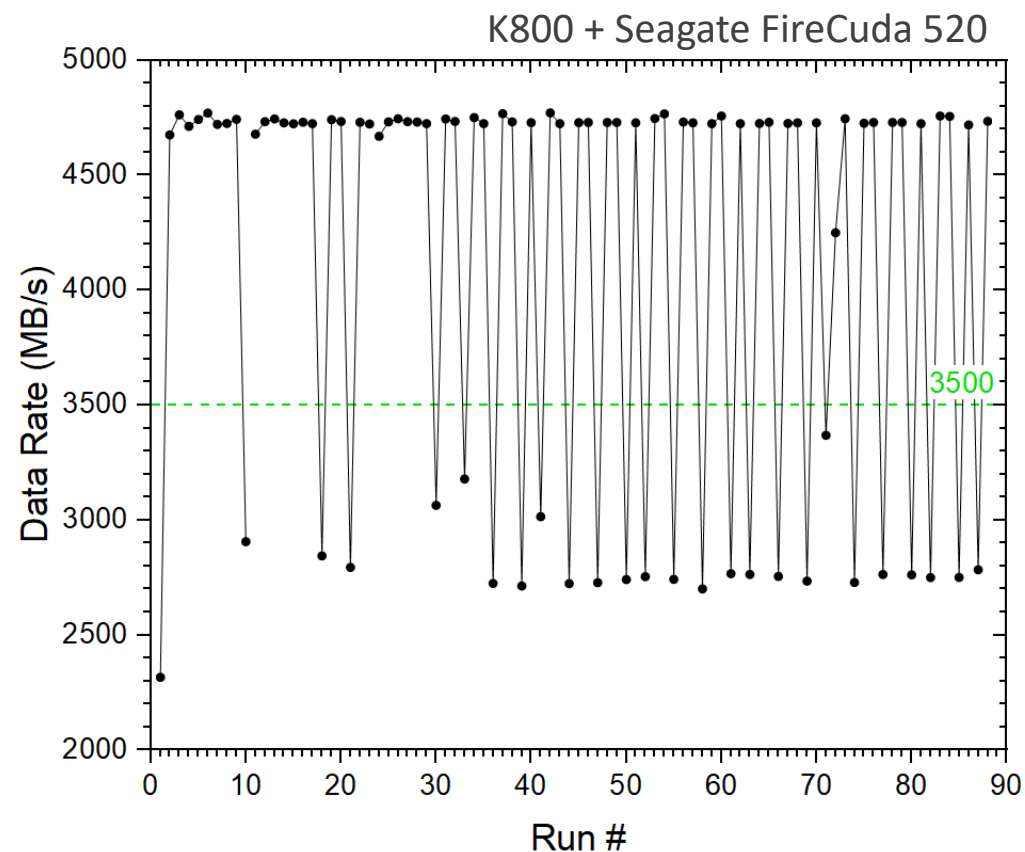
# Testing Data – 20GB Write

Written to 4K Blocks



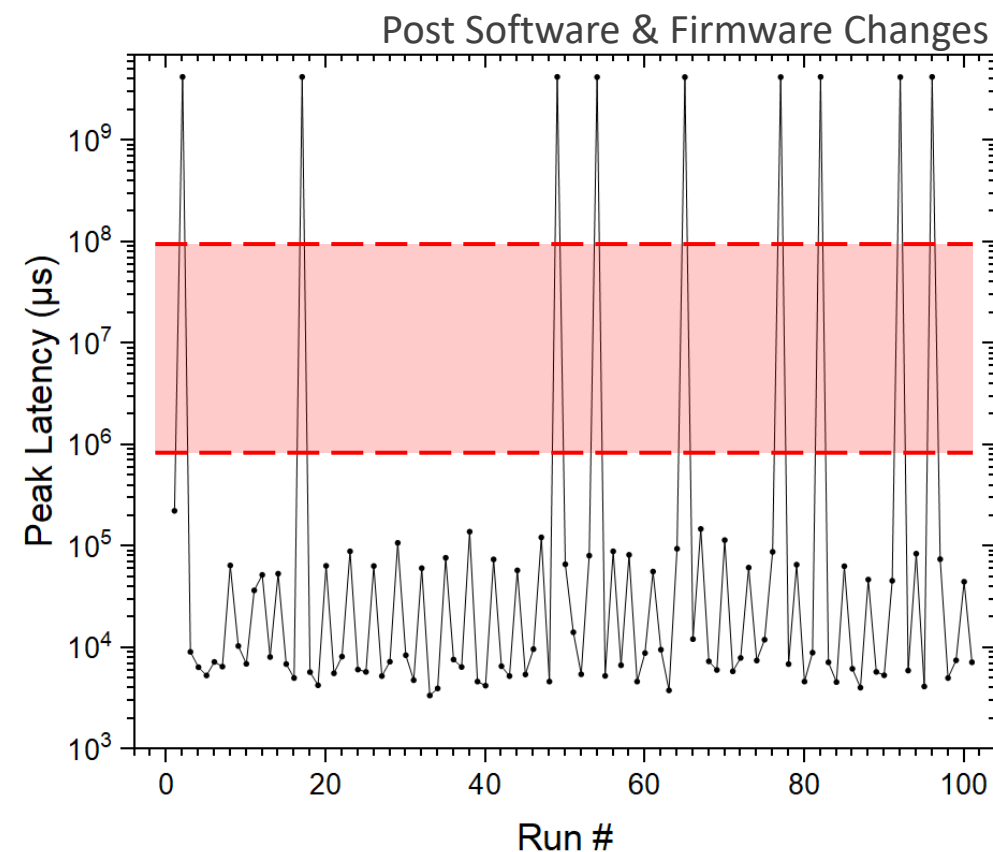

# Testing Data – 200GB Write

Written to 4K Blocks



# Testing Data – Peak Latency (200GB)

Written to 4K Blocks



# NVMe Readout

---

What do we need to do?

1. Take data (already in 4kByte blocks) – using incrementing counter at this stage
2. Write data to drives in parallel
3. Read data out to host PC
  1. Confirm all blocks written to drives
  2. Confirm all blocks have 4k Bytes
  3. Confirm data is written in correct sequence
4. Stitch blocks back together in the correct order to give a single data-stream
  1. Requires special treatment for missing blocks
5. Check the data-stream is correct (no missing blocks)

# NVMe Readout

## Initial sanity check

- Use BEAM test software to read 2 blocks to binary output file (.bin)
- Very primitive Python script to convert output to list of Bytes in integer format
- Output consisted of 8192 elements (Bytes) as expected
- Counter uses 4 Bytes: each of values 0 to 255

```
f=open("testoutput.bin", "rb")
num=list(f.read())
print(num)
f.close()
```

Python ▾

```
[0, 0, 0, 0, 1, 0, 0, 0, 2, 0, 0, 0, 3, 0, 0, 0, 4, 0, 0, 0, 5, 0, 0, 0, 6, 0, 0, 0,
7, 0, 0, 0, 8, 0, 0, 0, 9, 0, 0, 0, 10, 0, 0, 0, 11, 0, 0, 0, 12, 0, 0, 0, 13, 0, 0,
0, 14, 0, 0, 0, 15, 0, 0, 0, 16, 0, 0, 0, 17, 0, 0, 0, 18, 0, 0, 0, 19, 0, 0, 0, 20,
0, 0, 0, 21, 0, 0, 0, 22, 0, 0, 0, 23, 0, 0, 0, 24, 0, 0, 0, 25, 0, 0, 0, 26, 0, 0, 0,
27, 0, 0, 0, 28, 0, 0, 0, 29, 0, 0, 0, 30, 0, 0, 0, 31, 0, 0, 0, 32, 0, 0, 0, 33, 0,
0, 0, 34, 0, 0, 0, 35, 0, 0, 0, 36, 0, 0, 0, 37, 0, 0, 0, 38, 0, 0, 0, 39, 0, 0, 0, 4
0, 0, 0, 41, 0, 0, 0, 42, 0, 0, 0, 43, 0, 0, 0, 44, 0, 0, 0, 45, 0, 0, 0, 46, 0, 0,
0, 47, 0, 0, 0, 48, 0, 0, 0, 49, 0, 0, 0, 50, 0, 0, 0, 51, 0, 0, 0, 52, 0, 0, 0, 53,
0, 0, 0, 54, 0, 0, 0, 55, 0, 0, 0, 56, 0, 0, 0, 57, 0, 0, 0, 58, 0, 0, 0, 59, 0, 0,
0, 60, 0, 0, 0, 61, 0, 0, 0, 62, 0, 0, 0, 63, 0, 0, 0, 64, 0, 0, 0, 65, 0, 0, 0, 66, 0,
0, 0, 67, 0, 0, 0, 68, 0, 0, 0, 69, 0, 0, 0, 70, 0, 0, 0, 71, 0, 0, 0, 72, 0, 0, 0, 7
3, 0, 0, 0, 74, 0, 0, 0, 75, 0, 0, 0, 76, 0, 0, 0, 77, 0, 0, 0, 78, 0, 0, 0, 79, 0, 0,
0, 80, 0, 0, 0, 81, 0, 0, 0, 82, 0, 0, 0, 83, 0, 0, 0, 84, 0, 0, 0, 85, 0, 0, 0, 86,
0, 0, 0, 87, 0, 0, 0, 88, 0, 0, 0, 89, 0, 0, 0, 90, 0, 0, 0, 91, 0, 0, 0, 92, 0, 0, 0,
93, 0, 0, 0, 94, 0, 0, 0, 95, 0, 0, 0, 96, 0, 0, 0, 97, 0, 0, 0, 98, 0, 0, 0, 99, 0,
0, 0, 100, 0, 0, 0, 101, 0, 0, 0, 102, 0, 0, 0, 103, 0, 0, 0, 104, 0, 0, 0, 105, 0, 0,
0, 106, 0, 0, 0, 107, 0, 0, 0, 108, 0, 0, 0, 109, 0, 0, 0, 110, 0, 0, 0, 111, 0, 0, 0,
112, 0, 0, 0, 113, 0, 0, 0, 114, 0, 0, 0, 115, 0, 0, 0, 116, 0, 0, 0, 117, 0, 0, 0, 11
8, 0, 0, 0, 119, 0, 0, 0, 120, 0, 0, 0, 121, 0, 0, 0, 122, 0, 0, 0, 123, 0, 0, 0, 124,
0, 0, 0, 125, 0, 0, 0, 126, 0, 0, 0, 127, 0, 0, 0, 128, 0, 0, 0, 129, 0, 0, 0, 130, 0,
0, 0, 131, 0, 0, 0, 132, 0, 0, 0, 133, 0, 0, 0, 134, 0, 0, 0, 135, 0, 0, 0, 136, 0, 0,
0, 137, 0, 0, 0, 138, 0, 0, 0, 139, 0, 0, 0, 140, 0, 0, 0, 141, 0, 0, 0, 142, 0, 0, 0,
143, 0, 0, 0, 144, 0, 0, 0, 145, 0, 0, 0, 146, 0, 0, 0, 147, 0, 0, 0, 148, 0, 0, 0, 14
9, 0, 0, 0, 150, 0, 0, 0, 151, 0, 0, 0, 152, 0, 0, 0, 153, 0, 0, 0, 154, 0, 0, 0, 155,
0, 0, 0, 156, 0, 0, 0, 157, 0, 0, 0, 158, 0, 0, 0, 159, 0, 0, 0, 160, 0, 0, 0, 161, 0,
0, 0, 162, 0, 0, 0, 163, 0, 0, 0, 164, 0, 0, 0, 165, 0, 0, 0, 166, 0, 0, 0, 167, 0, 0,
0, 168, 0, 0, 0, 169, 0, 0, 0, 170, 0, 0, 0, 171, 0, 0, 0, 172, 0, 0, 0, 173, 0, 0, 0,
174, 0, 0, 0, 175, 0, 0, 0, 176, 0, 0, 0, 177, 0, 0, 0, 178, 0, 0, 0, 179, 0, 0, 0, 18
```

and so on...

# NVMe Readout

Written simple Python script which piggybacks off of Beam *test\_nvme* software

1. Uses *test\_nvme* to write 20GB data
2. Uses *test\_nvme* to read a specified number of 4kByte blocks to .bin output file
3. Interpret binary data to list of bytes in integer format
4. Checks the list contains expected number of Bytes
5. Compiles Bytes into a single data-stream string

```
1 #import numpy as np
2 import subprocess
3
4
5 def readFile(filename):
6     """ Read a binary file and output elements as a list
7     """
8     f = open(filename, "rb")
9     num = list(f.read()) # List of all elements
10    f.close()
11
12    return num
13
14 def dataStream(elementList, numBlocksRead):
15     """ Takes in list of binary output and compiles to single datastream
16     """
17     elementStream = ''.join(map(str, elementList))
18     if numBlocksRead <= 2:
19         print(elementStream)
20
21     return elementStream
22
23 def bashProcess(command):
24     """ Run a bash command to terminal
25     """
26     subprocess.run(command, shell=True)
27     print("Command completed")
28
29     return
30
31 def lenCheck(elementList):
32     #if numBlocksRead <= 2:
33     #    print(elementList)
34
35     if len(elementList) == 4096 * numBlocksRead:
36         print("No blocks missing. {} kBytes read.".format(len(elementList)/1024))
37     else:
38         print("Blocks missing")
39
40     return
41
42
43 filename = "output.bin"
44 numBlocksRead = 2
45 command = ""
46         ./test_nvme -d 2 -s 0 -n 5242880 capture
47         ./test_nvme -d 2 -s 0 -n {0} -o {1} read"".format(numBlocksRead, filename)
48
49 bashProcess(command)
50 elementList = readFile(filename)
51 lenCheck(elementList)
52
53 elementStream = dataStream(elementList, numBlocksRead)
```



# NVMe Readout

Example output for 10 block readout

```
[ag17009@kairos test]$ python new_readout.py
nvmeCapture: Write FPGA data stream to Nvme devices. nvme: 2 startBlock: 0 numBlocks: 5242880
10:39:08.090: ErrorStatus: 0x0, StartBlock: 0, DataRate: 4832.916 MBytes/s, PeakLatency: 3028 us
NvmeRead: nvme: 2 startBlock: 0 numBlocks: 10
Read complete at: 10 blocks
NvmeRead: rate: 4.684355 MBytes/s
Command completed
No blocks missing. 40.0 kBytes read.
```

Problem here is that the *test\_nvme* software does its own error checking for missing blocks when it reads, so you'll never produce an error case when piggybacking off of it in this way.

- Need to write the readout solution into dedicated software that can perform the same tasks as Beam test software
- Need custom data sets with intentional errors for testing scenarios where blocks are missing/corrupt

# NVMe Readout

## Example of 8kByte data-stream output

[illegible]

These are still very primitive readout solutions performed mostly for sanity check and proof of concept purposes.