

# DUNE SPNG Updates

# May 1 2026

Tests in SGPU node

# Signal Processing Workflow

- Looking at the performance of the Signal Processing workflow that was executed on the SGPU node
- Workflow
  - Dense raw data → Multi View processing (Software filters, Formation of Initial ROIs and MP) → ROIs using Inference → Apply ROI → ROI applied data
- Jsonnet
  - test-subgraphs.jsonnet\*
- Events → 2
- Views → U, V and W [800,800, 480x2 channels] 2560 ticks readout
- Detector → pdhd
- APA → 1 APA

# Test candidate memory requirement

- Test candidate requires
  - ~15 GB GPU memory
  - <1 GB storage (tested for 2 events only)
- Condor jobs were submitted :
  - 1 workflow per condor job
  - 2 parallel workflow per condor job
  - 3 parallel workflow per condor job
  - 1 CPU and 1 GPU

```
universe = vanilla
executable = generated_condor/run_wirecell_worker_n3_1777594610.sh
arguments = $(Cluster) $(Process) 3 1
should_transfer_files = IF_NEEDED
when_to_transfer_output = ON_EXIT
transfer_input_files = generated_condor/run_wirecell_worker_n3_1777594610.sh, test-subgraphs.jsonnet, drifted-cosmics_2_wct-frame.npz, toolkit/spng/cfg/, toolkit/wire-cell-data
output = logs/wirecell.$(Cluster).$(Process).n3.out
error = logs/wirecell.$(Cluster).$(Process).n3.err
log = logs/wirecell.$(Cluster).n3.log
request_cpus = 1
request_memory = 16 GB
request_disk = 32 GB
request_gpus = 1
```

# SGPU0004 Node

- 10 L40 GPUs

```
bash-5.1$ cpufreq
Python Version: 3.9.21.final.0 (64 bit)
Cpuinfo Version: 8.0.0
Vendor ID Raw: AuthenticAMD
Hardware Raw:
Brand Raw: AMD EPYC 9355 32-Core Processor
Hz Advertised Friendly: 2.1422 GHz
Hz Actual Friendly: 2.1422 GHz
Hz Advertised: (2142222000, 0)
Hz Actual: (2142222000, 0)
Arch: X86_64
Bits: 64
Count: 128
Arch String Raw: x86_64
L1 Data Cache Size: 3 MiB (64 instances)
L1 Instruction Cache Size: 2 MiB (64 instances)
L2 Cache Size: 64 MiB (64 instances)
L2 Cache Line Size:
L2 Cache Associativity:
L3 Cache Size: 1048576
Stepping: 1
Model: 2
Family: 06
```

```
bash-5.1$ nvidia-smi
Thu Apr 30 23:02:43 2026
```

NVIDIA-SMI 575.57.08		Driver Version: 575.57.08		CUDA Version: 12.9	
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile Uncorr. ECC
Fan	Temp	Pwr:Usage/Cap	Memory-Usage	GPU-Util	Compute M. MIG M.
Perf					
0	NVIDIA L40S	On	00000000:03:00.0	Off	0
N/A	36C P8	24W / 350W	3MiB / 46068MiB	0%	Default N/A
1	NVIDIA L40S	On	00000000:04:00.0	Off	0
N/A	38C P8	22W / 350W	3MiB / 46068MiB	0%	Default N/A
2	NVIDIA L40S	On	00000000:63:00.0	Off	0
N/A	36C P8	34W / 350W	3MiB / 46068MiB	0%	Default N/A
3	NVIDIA L40S	On	00000000:64:00.0	Off	0
N/A	50C P0	87W / 350W	11999MiB / 46068MiB	0%	Default N/A
4	NVIDIA L40S	On	00000000:65:00.0	Off	0
N/A	36C P8	23W / 350W	3MiB / 46068MiB	0%	Default N/A
5	NVIDIA L40S	On	00000000:83:00.0	Off	0
N/A	38C P8	35W / 350W	3MiB / 46068MiB	0%	Default N/A
6	NVIDIA L40S	On	00000000:84:00.0	Off	0
N/A	37C P8	33W / 350W	3MiB / 46068MiB	0%	Default N/A
7	NVIDIA L40S	On	00000000:85:00.0	Off	0
N/A	33C P8	23W / 350W	3MiB / 46068MiB	0%	Default N/A
8	NVIDIA L40S	On	00000000:E3:00.0	Off	0
N/A	31C P8	24W / 350W	3MiB / 46068MiB	0%	Default N/A
9	NVIDIA L40S	On	00000000:E4:00.0	Off	0
N/A	37C P8	33W / 350W	3MiB / 46068MiB	0%	Default N/A

# Profiling information

- Nsys profile to profile the workflow
  - Produces nsys-rep and sqlite output that can be analyzed offline
  - Low level information in terms of GPU and CPU compute, memory during runtime
- CUDA NVTX to annotate nodes in the workflow to get insights on targeted parts of the code

# Full Coverage

```
// this bool indicates exception, and is probably ignored
bool Graph::execute()
{
    auto nodes = sort_kahn();
    l->debug("executing with {} nodes", nodes.size());

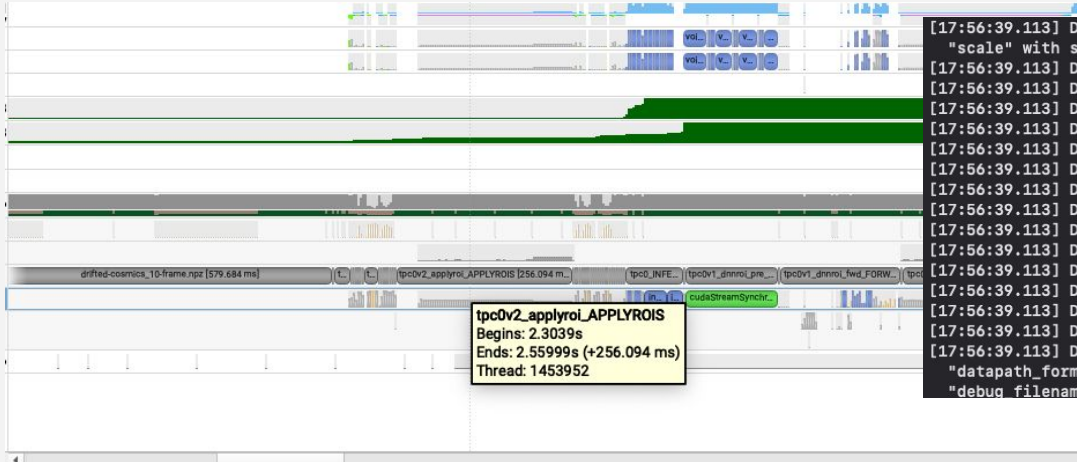
    while (true) {
        int count = 0;
        bool did_something = false;

        for (auto nit = nodes.rbegin(); nit != nodes.rend(); ++nit, ++count) {
            Node* node = *nit;
            NVTX_SCOPED_RANGE((node->instance_name()).c_str());
            auto core = std::clock();
            auto wall = std::chrono::high_resolution_clock::now();
        }
    }
}
```

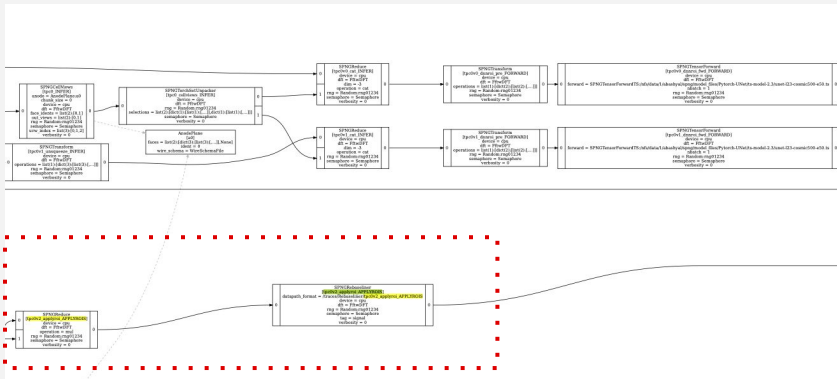
Added in PGrapher/Graph.cxx



# Unique names based on node instances

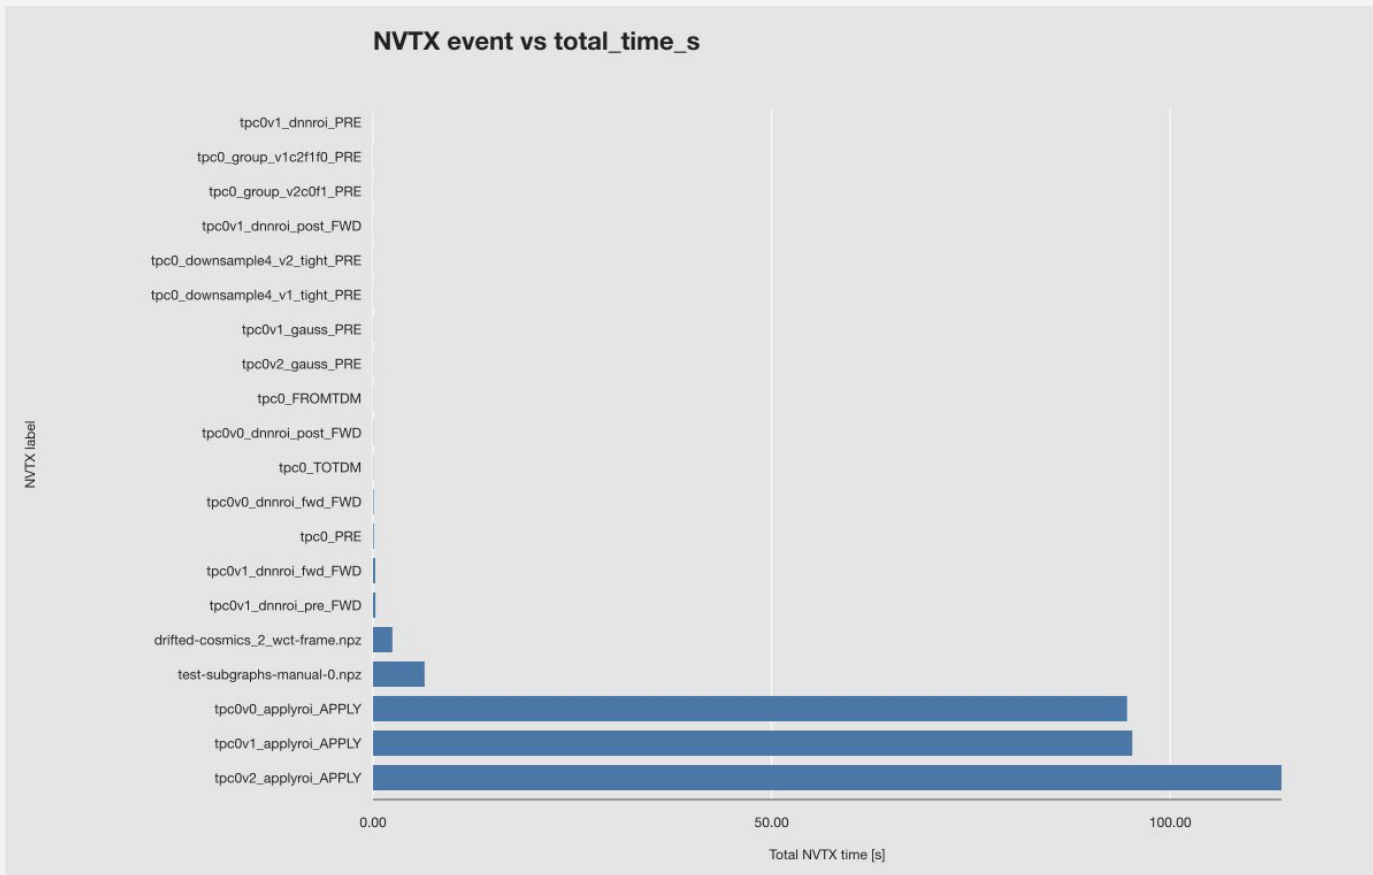


```
[17:56:39.113] D [ spng ] <Transform:tpc0v1_scale_dnnroi_dense_INFER> Applying transform operation "scale" with scalar=0.00025 and ndims=0
[17:56:39.113] D [ main ] configuring component: "SPNGReduce": "tpc0v0_applyroi APPLYROIS"
[17:56:39.113] D [ spng ] <Reduce:tpc0v0_applyroi APPLYROIS> fanin multiplicity: 2
[17:56:39.113] D [ main ] configuring component: "SPNGRebaseliner": "tpc0v0_applyroi APPLYROIS"
[17:56:39.113] D [ main ] configuring component: "SPNGReduce": "tpc0v1_applyroi APPLYROIS"
[17:56:39.113] D [ spng ] <Reduce:tpc0v1_applyroi APPLYROIS> fanin multiplicity: 2
[17:56:39.113] D [ main ] configuring component: "SPNGRebaseliner": "tpc0v1_applyroi APPLYROIS"
[17:56:39.113] D [ main ] configuring component: "SPNGReduce": "tpc0v2_applyroi APPLYROIS"
[17:56:39.113] D [ spng ] <Reduce:tpc0v2_applyroi APPLYROIS> fanin multiplicity: 2
[17:56:39.113] D [ main ] configuring component: "SPNGRebaseliner": "tpc0v2_applyroi APPLYROIS"
[17:56:39.113] D [ main ] configuring component: "SPNGReduce": "tpc0v0_cat_INFER"
[17:56:39.113] D [ spng ] <Reduce:tpc0v0_cat_INFER> fanin multiplicity: 2
[17:56:39.113] D [ main ] configuring component: "SPNGReduce": "tpc0v1_cat_INFER"
[17:56:39.113] D [ spng ] <Reduce:tpc0v1_cat_INFER> fanin multiplicity: 2
[17:56:39.113] D [ main ] configuring component: "SPNGTransform": "tpc0v0_unsqueeze_INFER"
[17:56:39.113] D [ spng ] <Transform:tpc0v0_unsqueeze_INFER> Transform config: {
  "datapath_format": "",
  "debug_filename": ""
```




Workflow that is being profiled

# 20 Longest Processes in terms of run time.

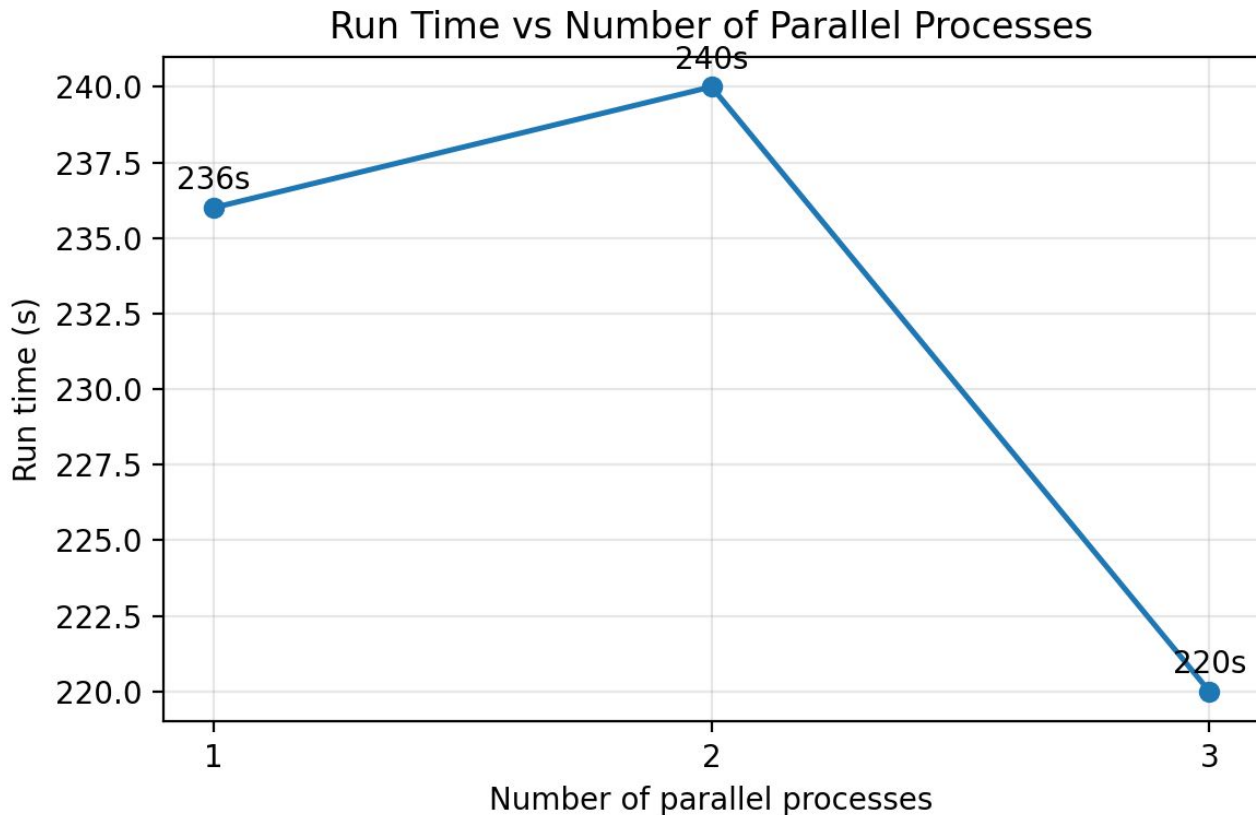




# Memory Transfer Volume



# Processes vs Run Time

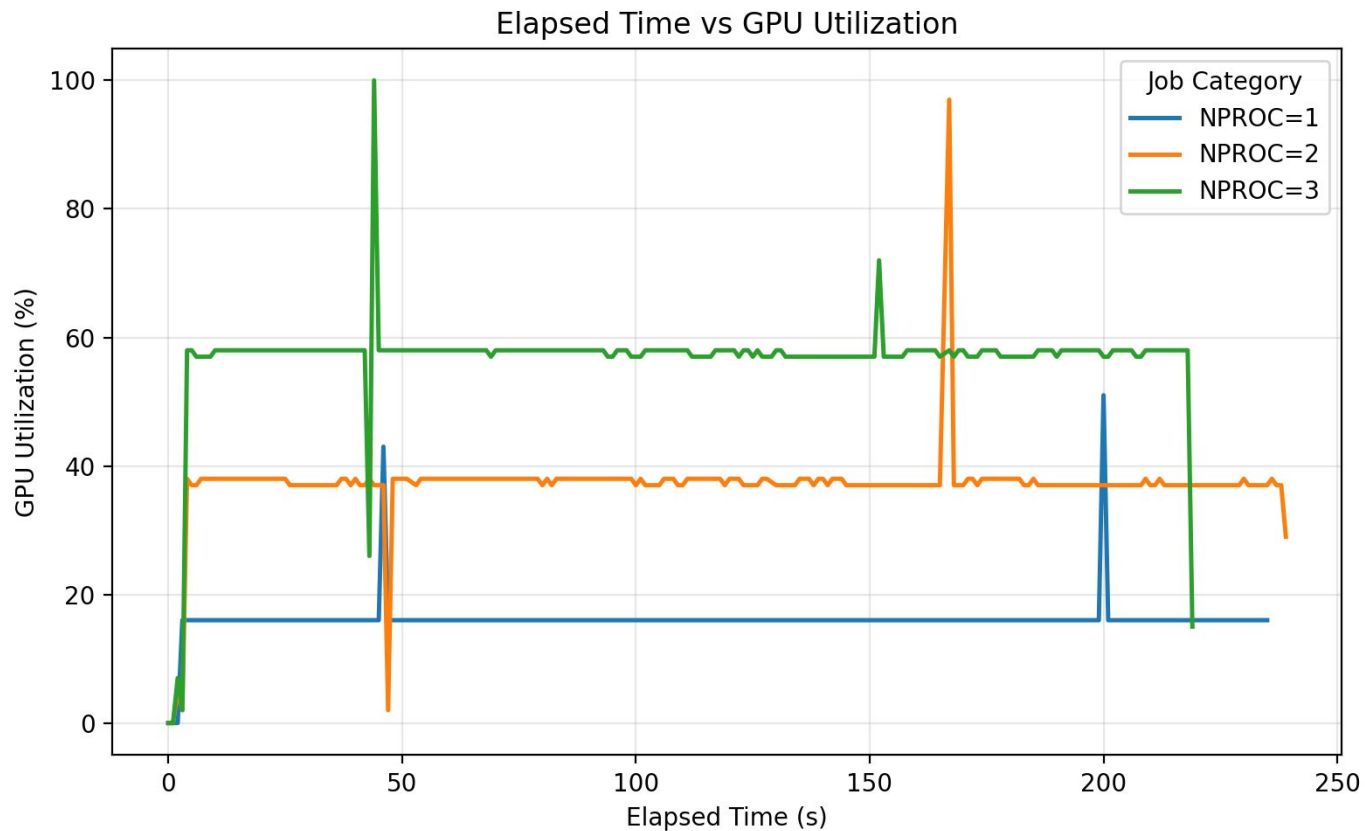


~15 GB for each job.

Total memory of each L40 GPU ~45 GB.


~3 jobs to saturate the memory (too close)

# GPU Utilization as a function of run time

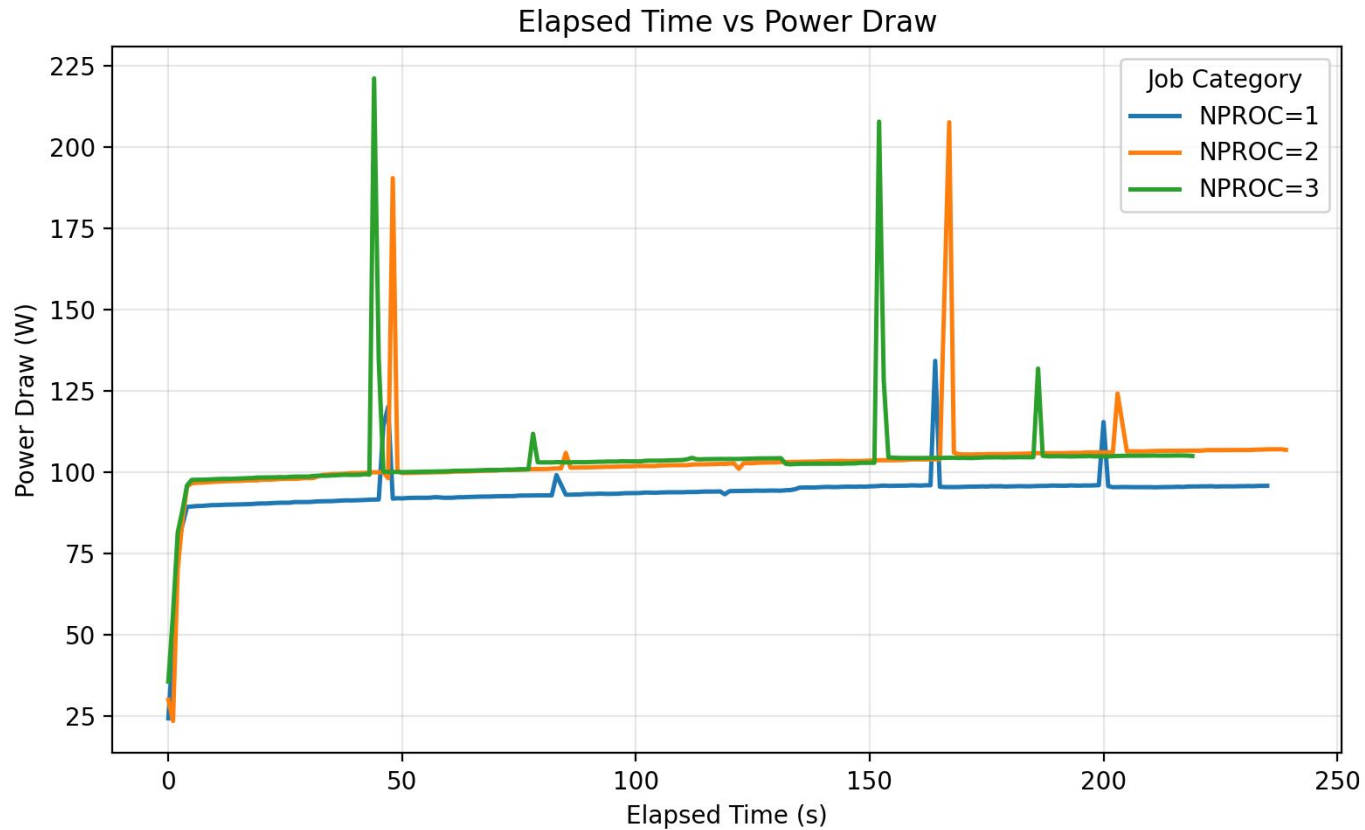


# GPU Memory allocated

Elapsed Time vs GPU Memory Used




# Power Consumption



# Temperature

Elapsed Time vs GPU Temperature



## To Do

Running with multiple gpus (4,5,all 10) with concurrent processes

Elapsed time, power, temperature

Communication with the SCDF people for the tests