POWER PROFILING OF GEMTC MANY TASK COMPUTING

Sean Wallace

BACKGROUND

- Many Task Computing focuses on using many resources over a short period of time.

- Historically, hardware accelerators (GPUs, Xeon Phi) haven't been invited to the party.

- GeMTC is a framework which enables MTC workloads to run on NVIDIA GPUs.

- While still in its infancy, it also supports the Xeon Phi.

WHY POWER?

- Power consumption of MTC workloads is very understudied.

- MTC workloads are fundamentally different than traditional HPC workloads, surely they exhibit different characteristics in their power consumption?

- Do certain accelerators perform more efficiently than others?

MONEQ

- Power monitoring library originally written for IBM Blue Gene/Q supercomputers extended to provide support for hardware accelerators.
- Built upon vendor supplied APIs.
- Provides automated profiling of user applications at sub-second intervals.

EXAMPLE

```
int main(int argc, char **argv){
    MonEQ_Initialize();

    kernel_call<<<dimGrid,
dimBlock>>>(a, N);

    MonEQ_Finalize();
}
```

```
int main() {
    MonEQ_Initialize();

    #pragma omp parallel for
num_threads(num_threads)
    for (i = 0; i < N; i++) {
        for (j = 0; j < M; j++) {
            C[0] = A[0] + B[0];
        }
    }

    MonEQ_Finalize();
}
```

NVIDIA K20

- Peak Performance (Double):
  - 1.17 Tflops
- Peak Performance (Single):
  - 3.52 Tflops
- Memory Bandwidth:
  - 208 GB/sec
- CUDA Cores
  - 2496

INTEL XEON PHI

- 1.053 GHz Clock Speed
- 60 Cores
- 8 GB Memory
- 16 Memory Channels
- 512-bit wide instruction vectors

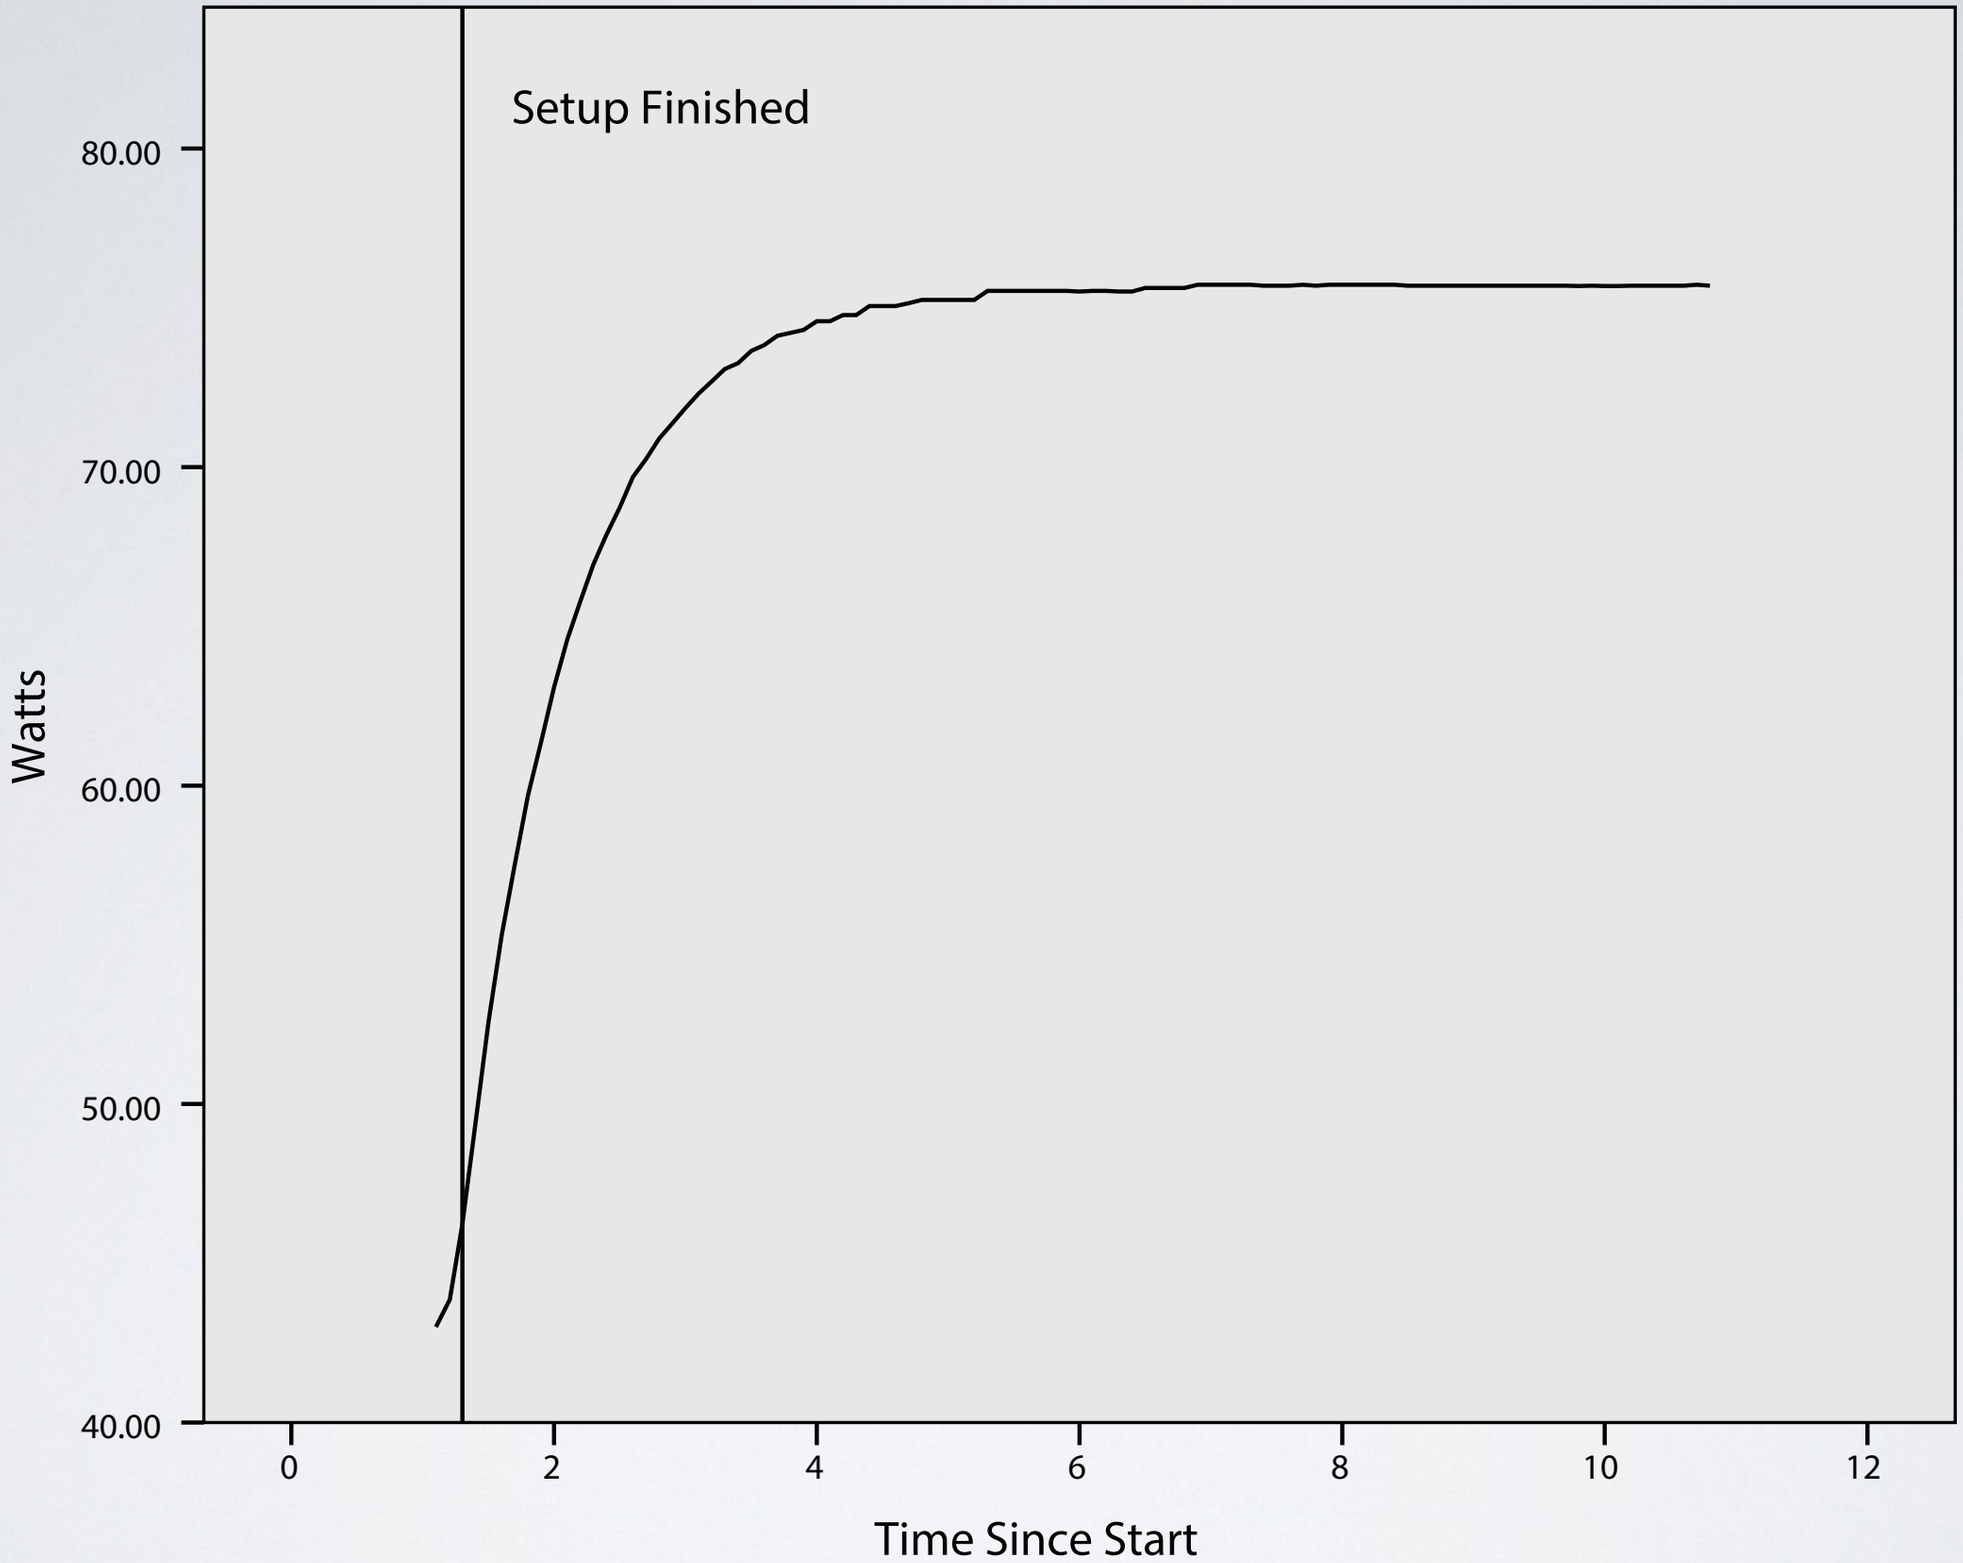

SLEEP WITH GEMTC ON NVIDIA

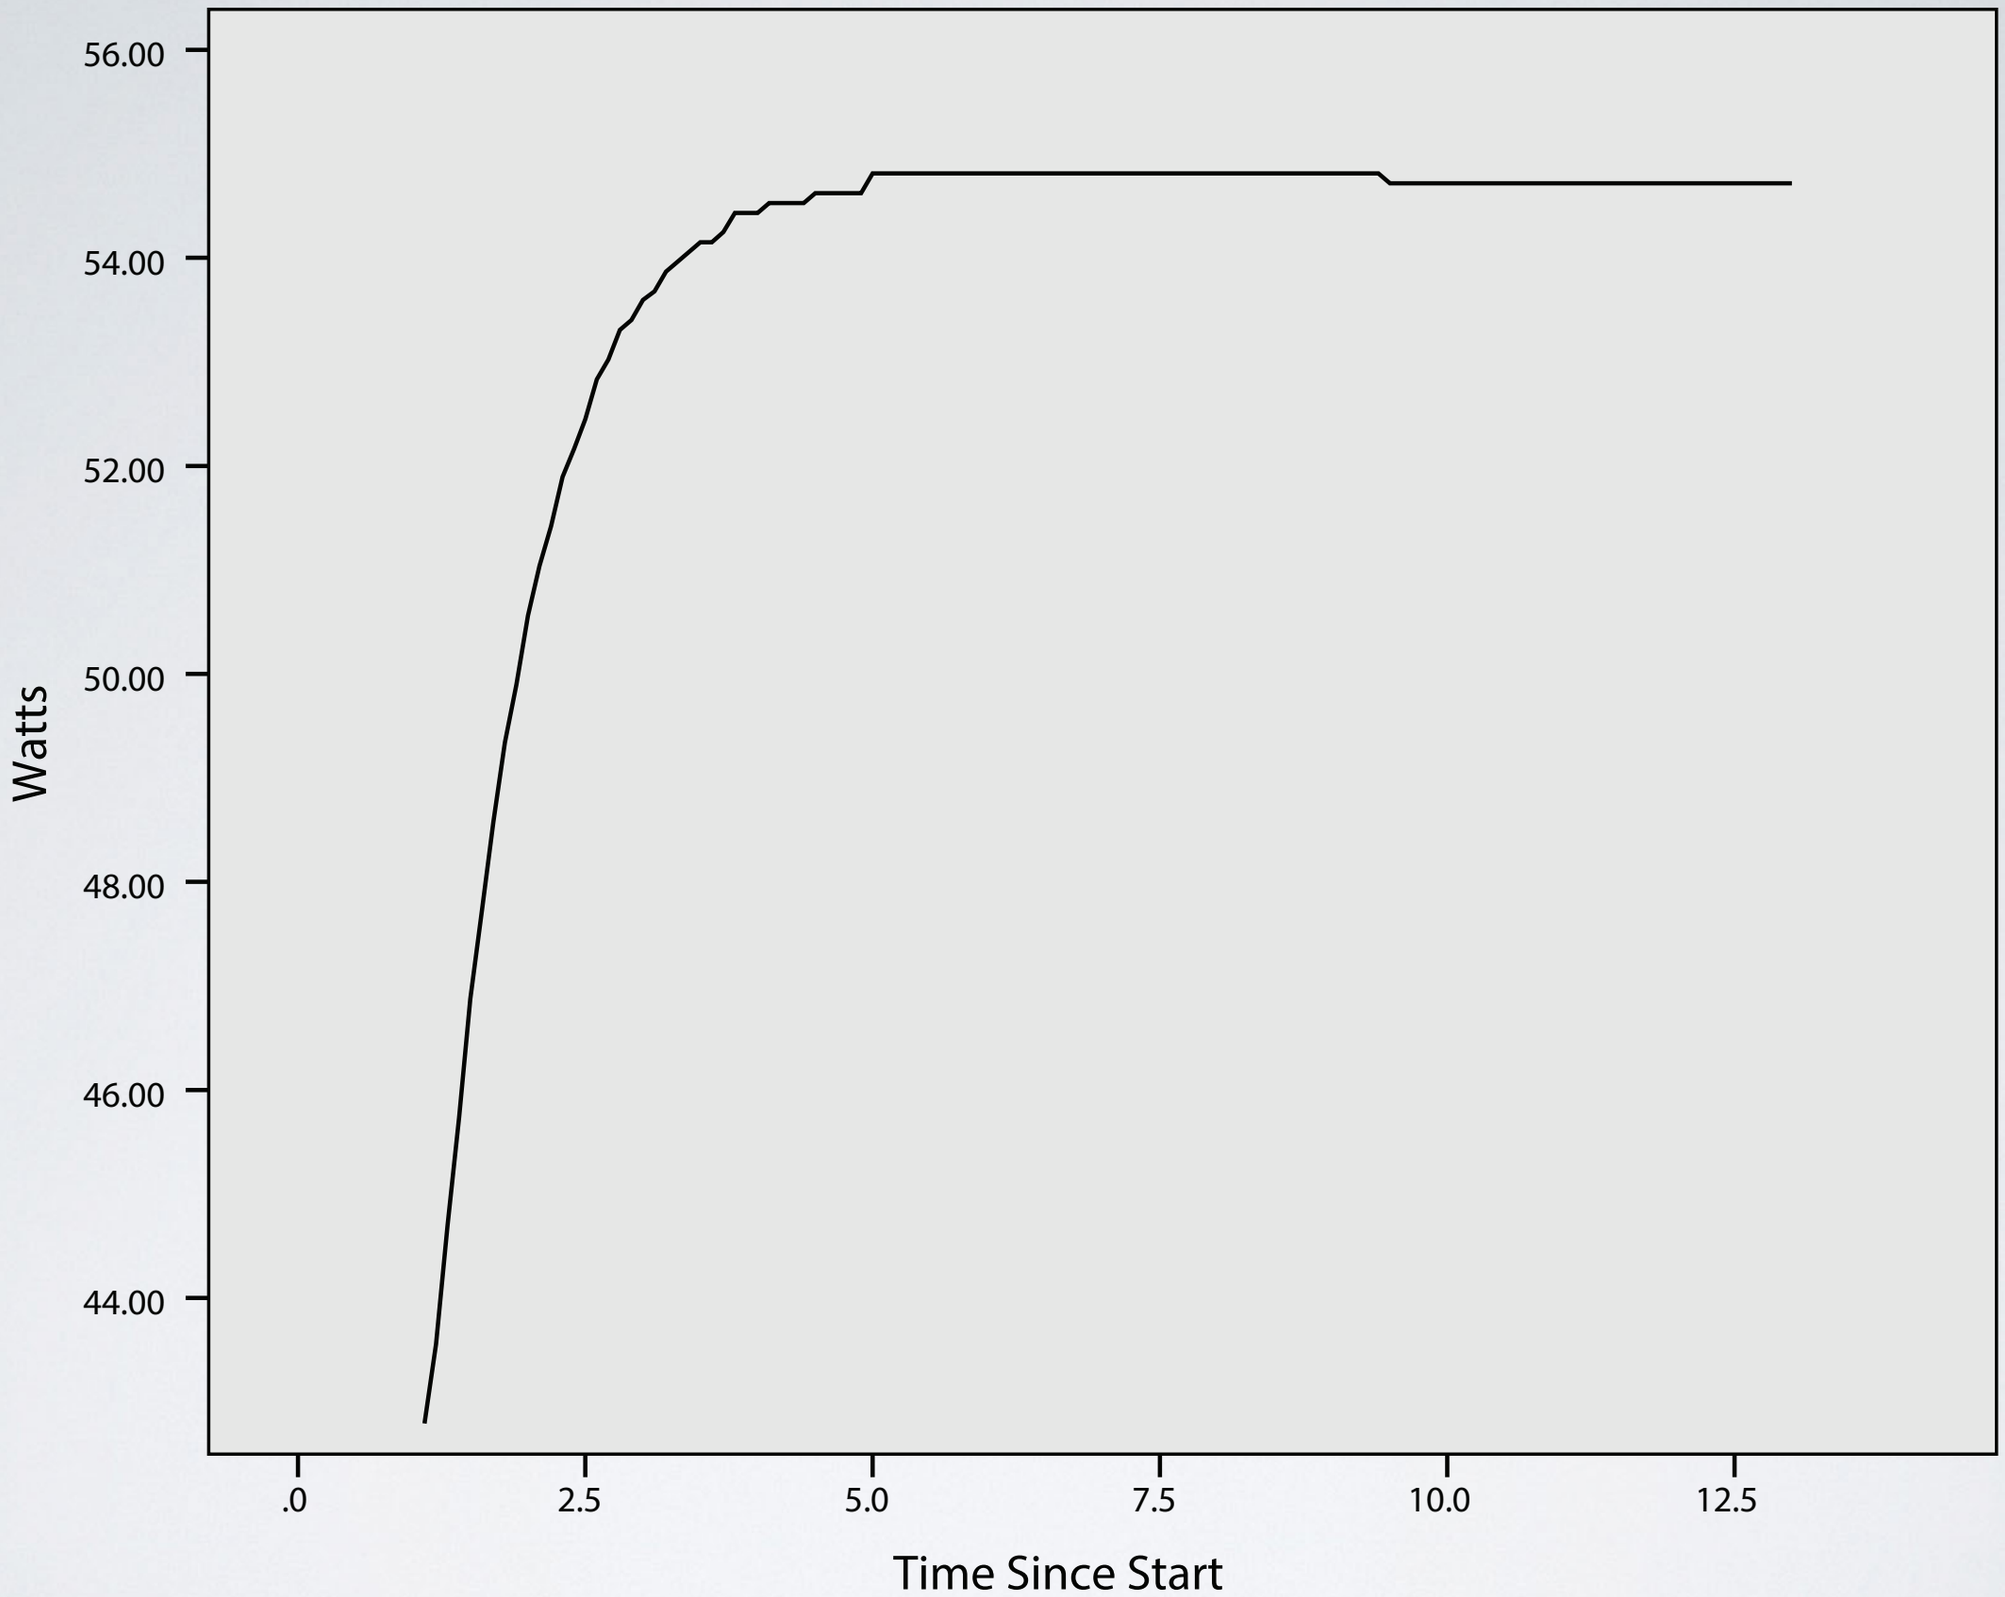

TRADITIONAL SLEEP ON NVIDIA

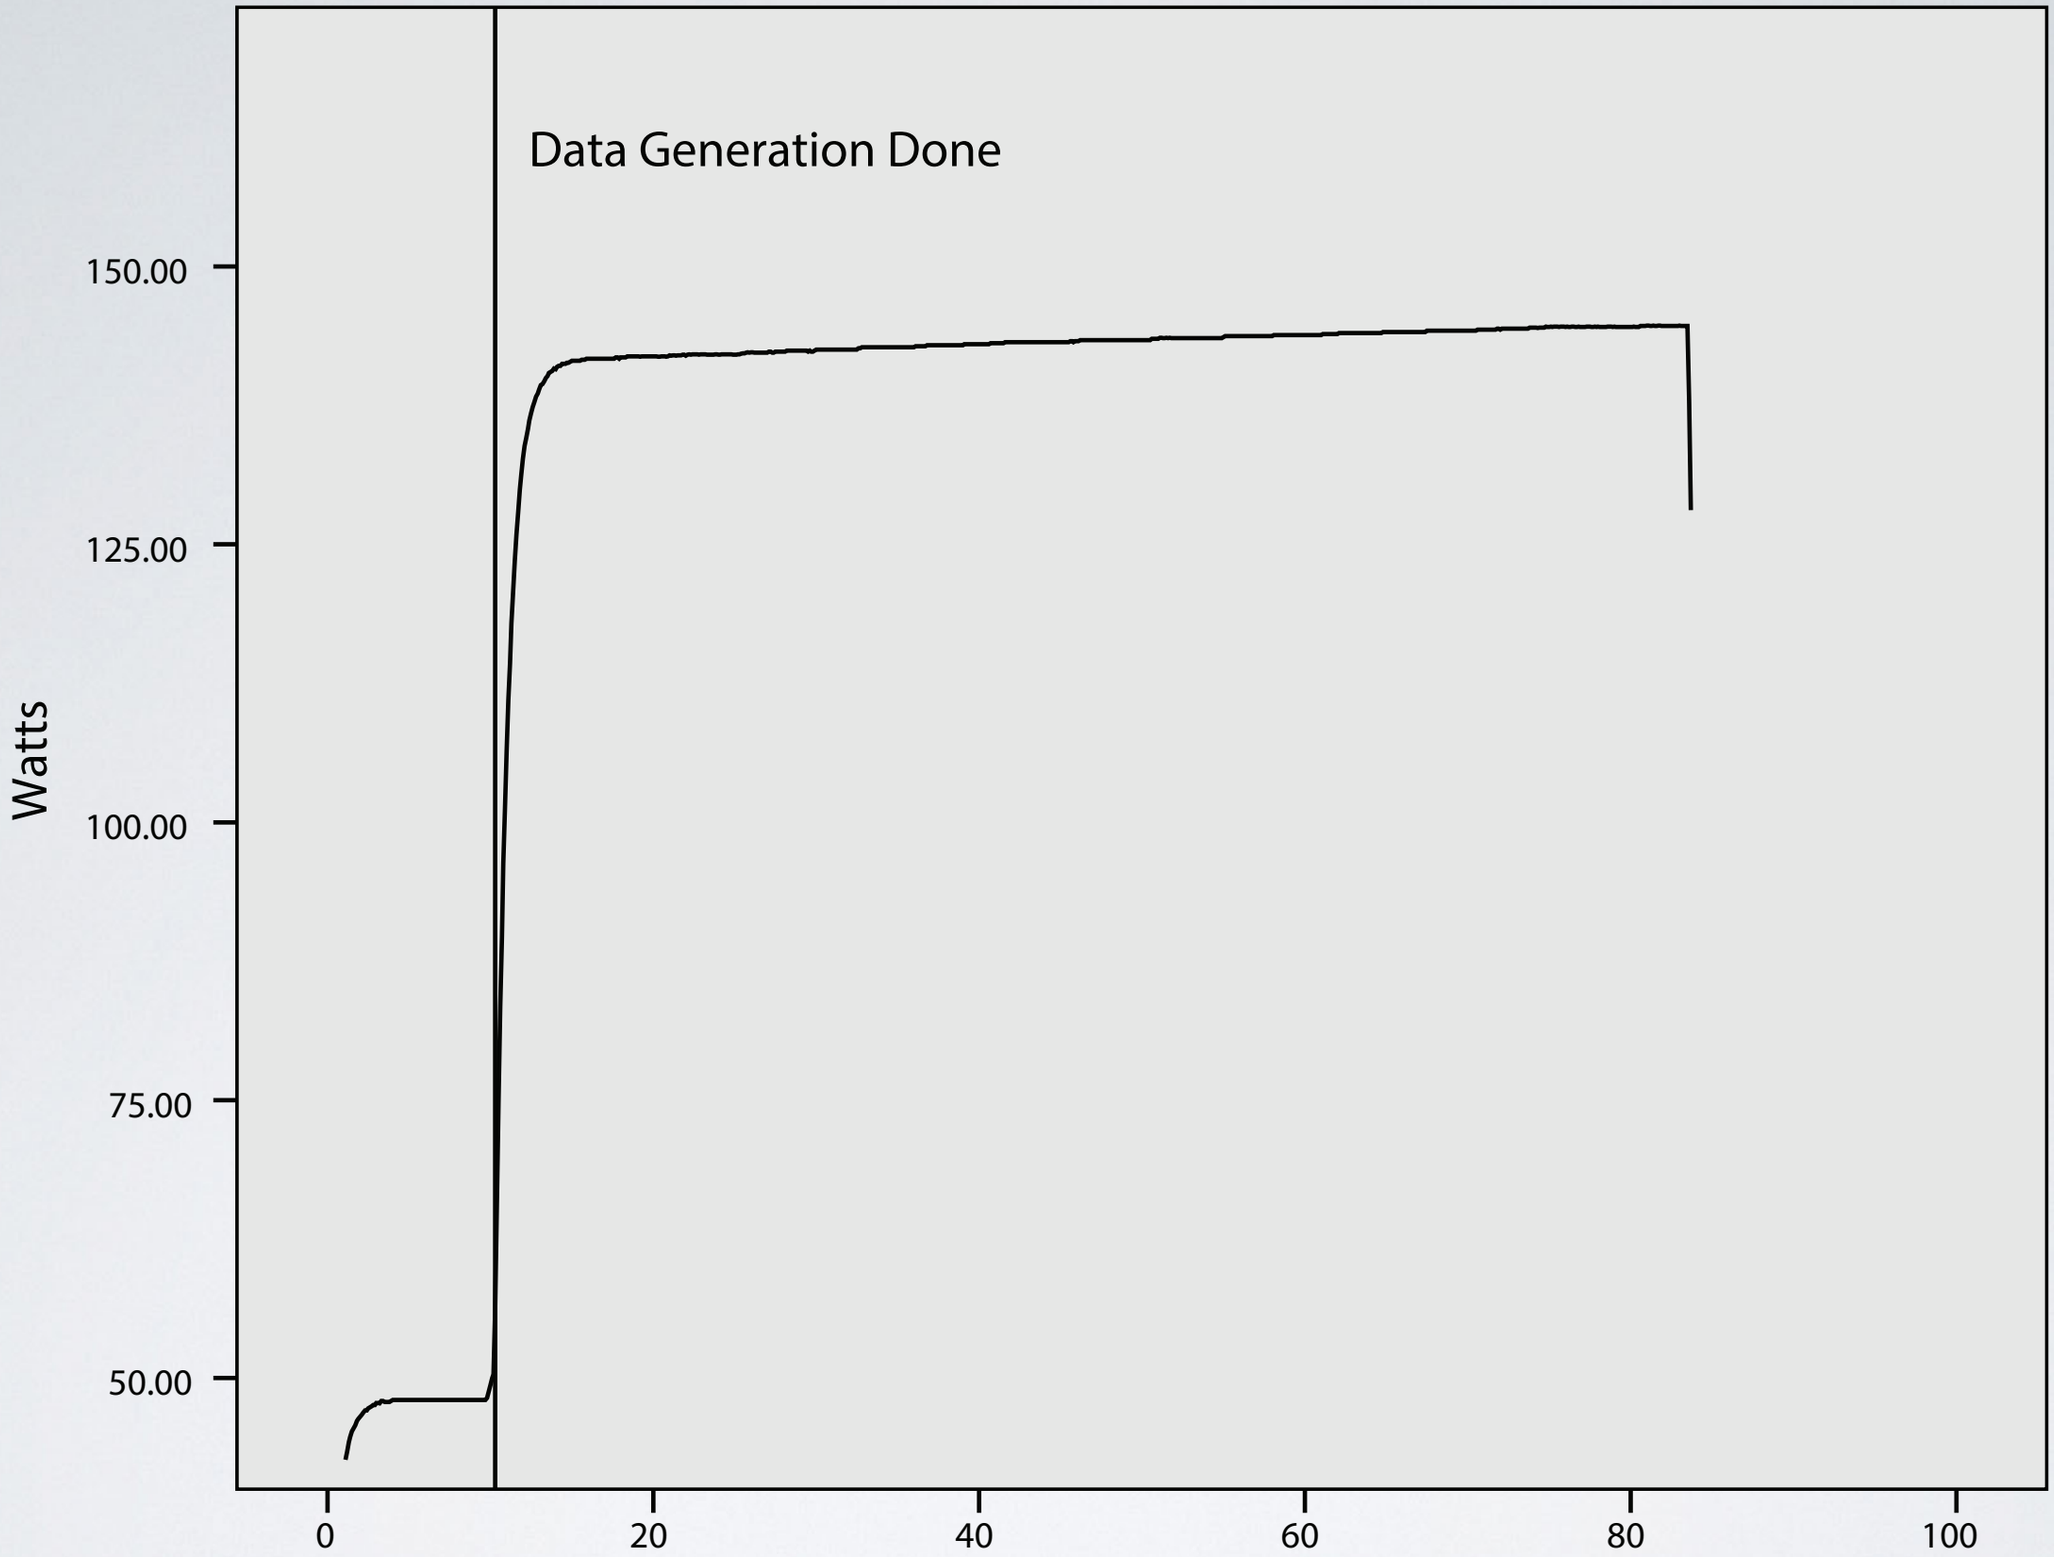

VECTOR ADD ON NVIDIA

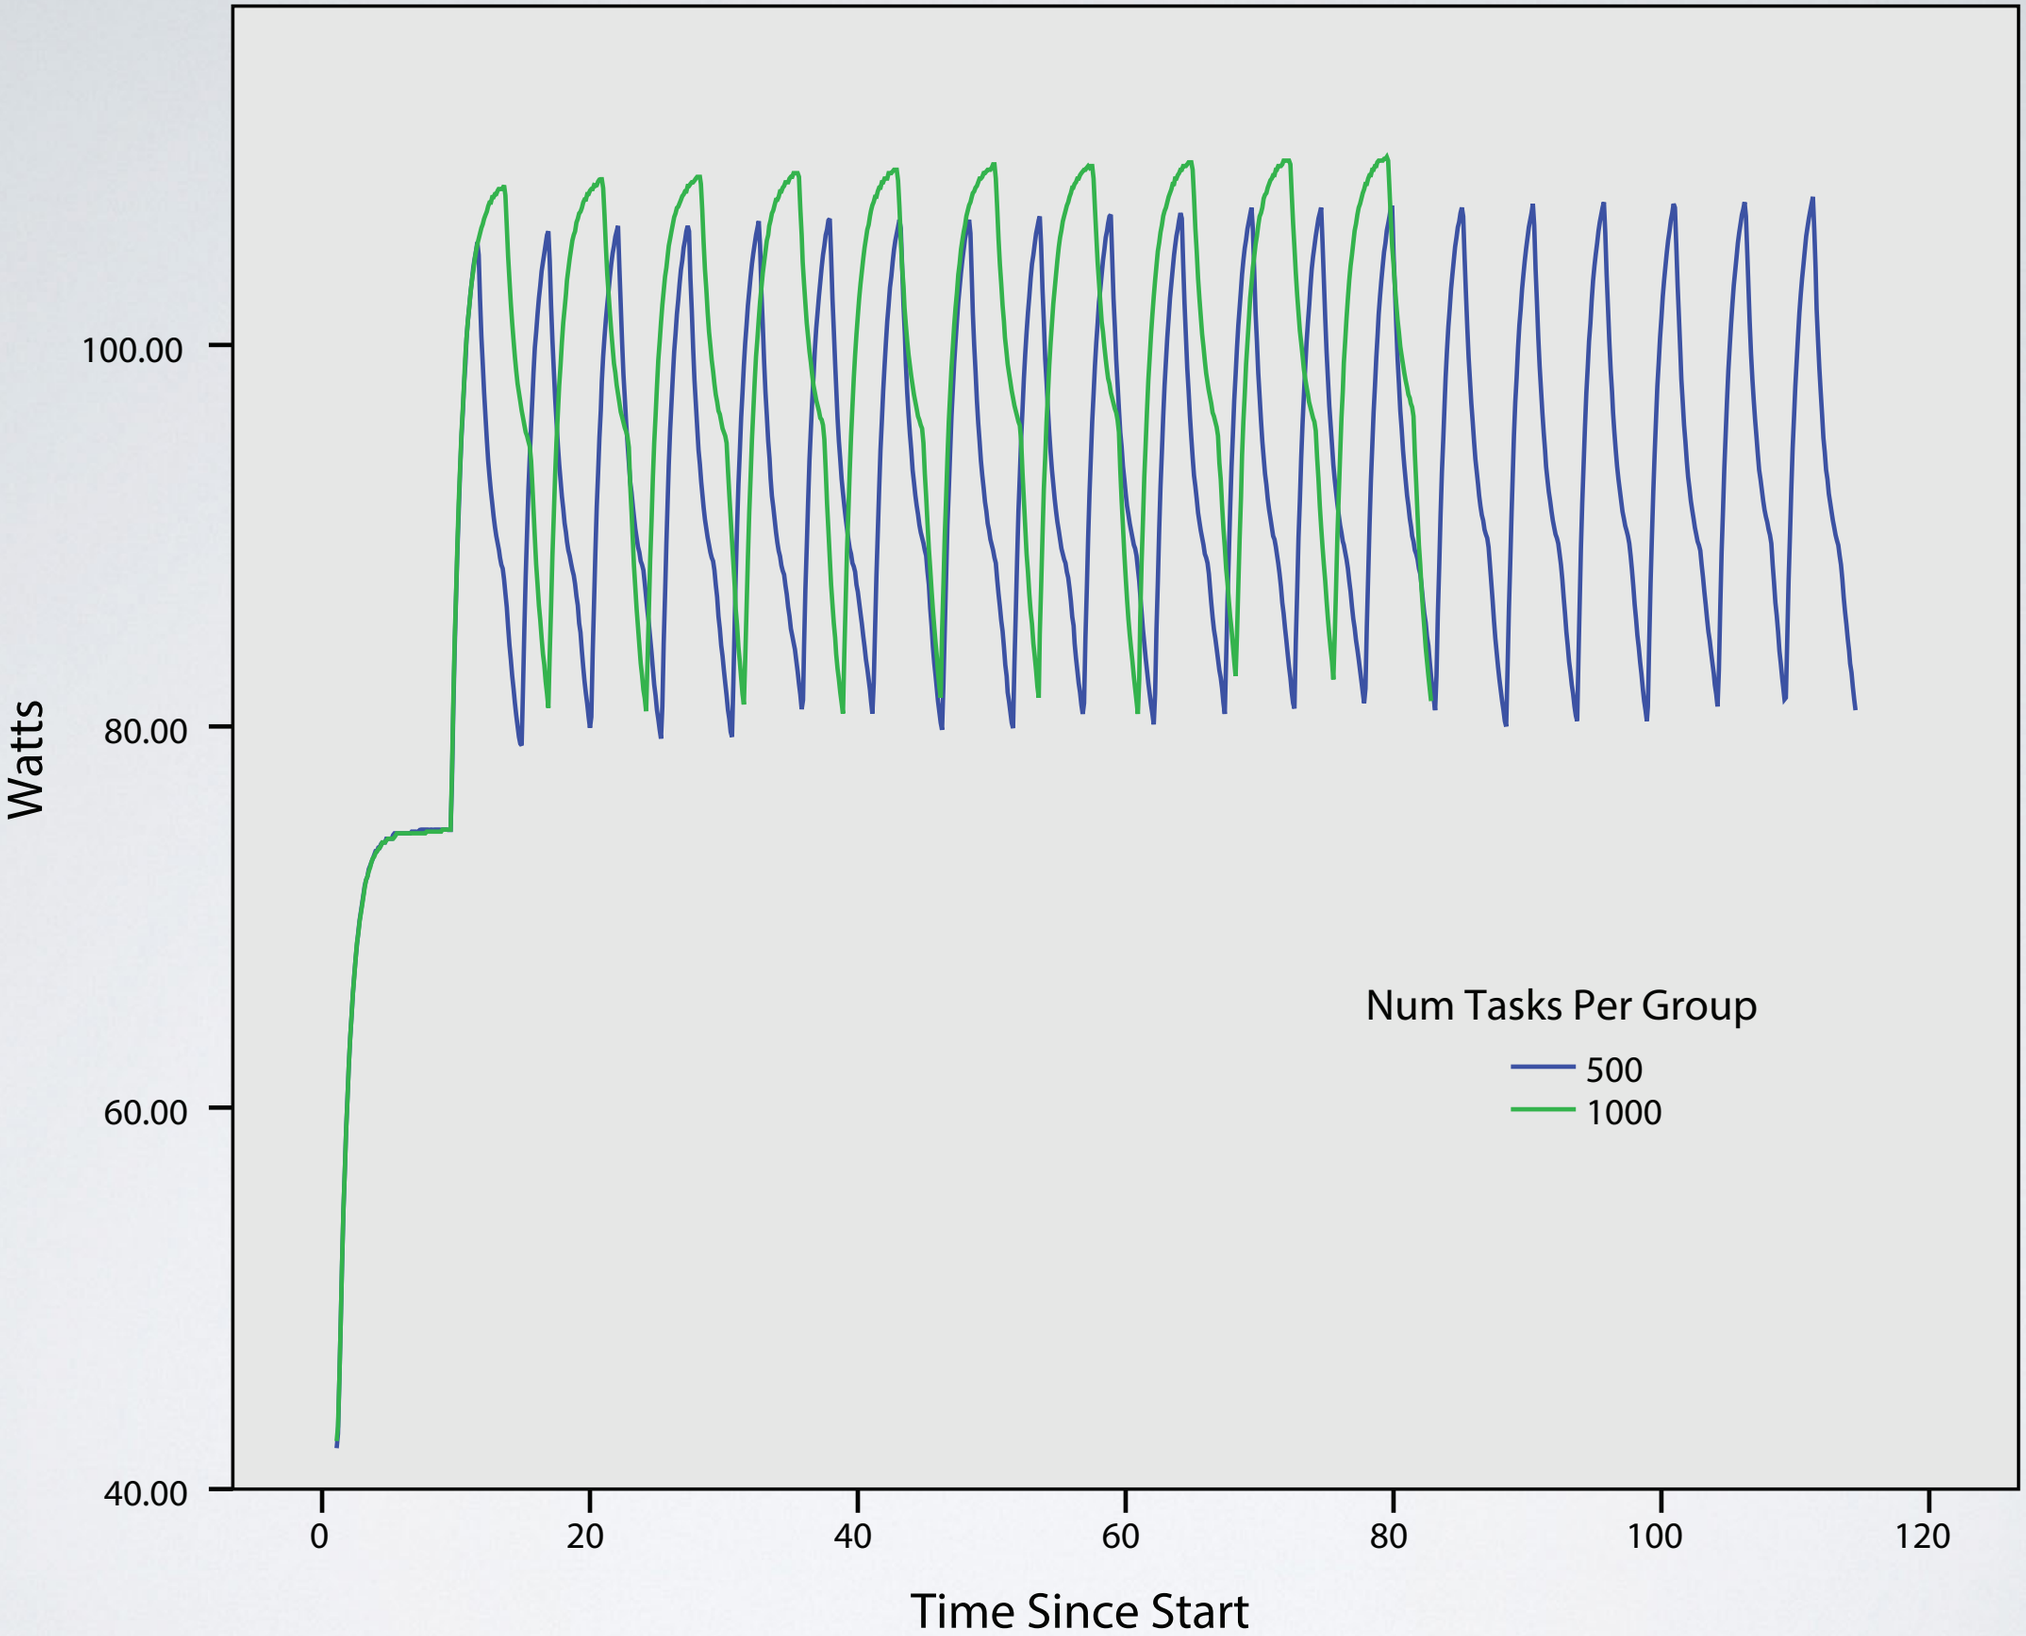

VECTOR ADD WITH GEMTC ON NVIDIA

WHAT ABOUT XEON PHI?

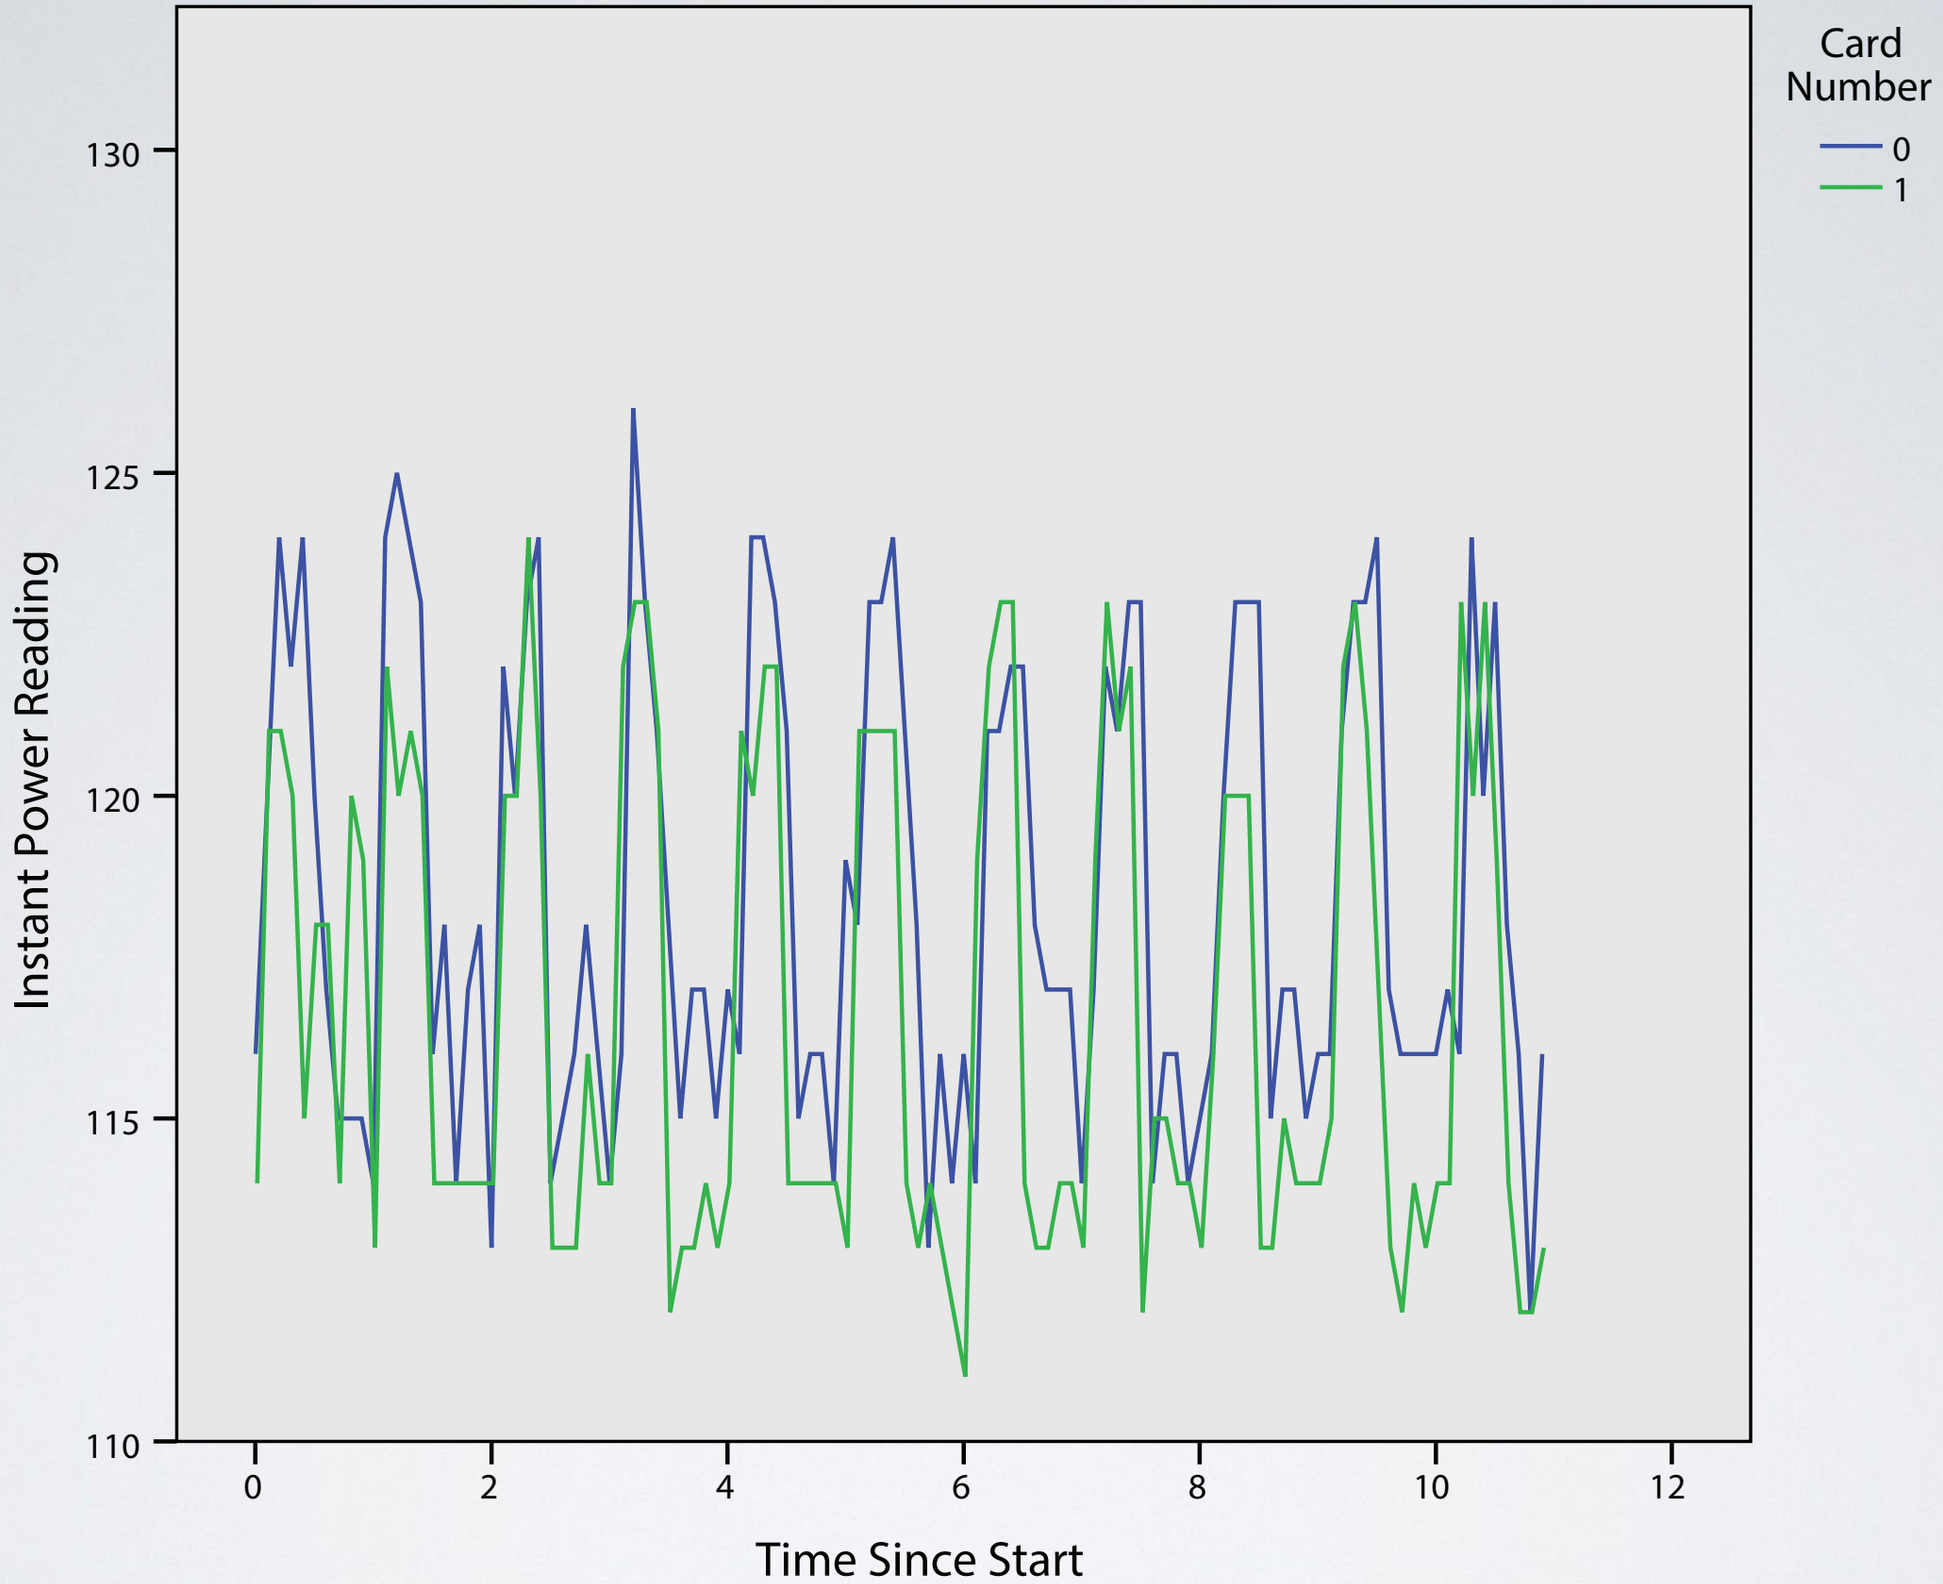

SLEEP ON XEON PHI

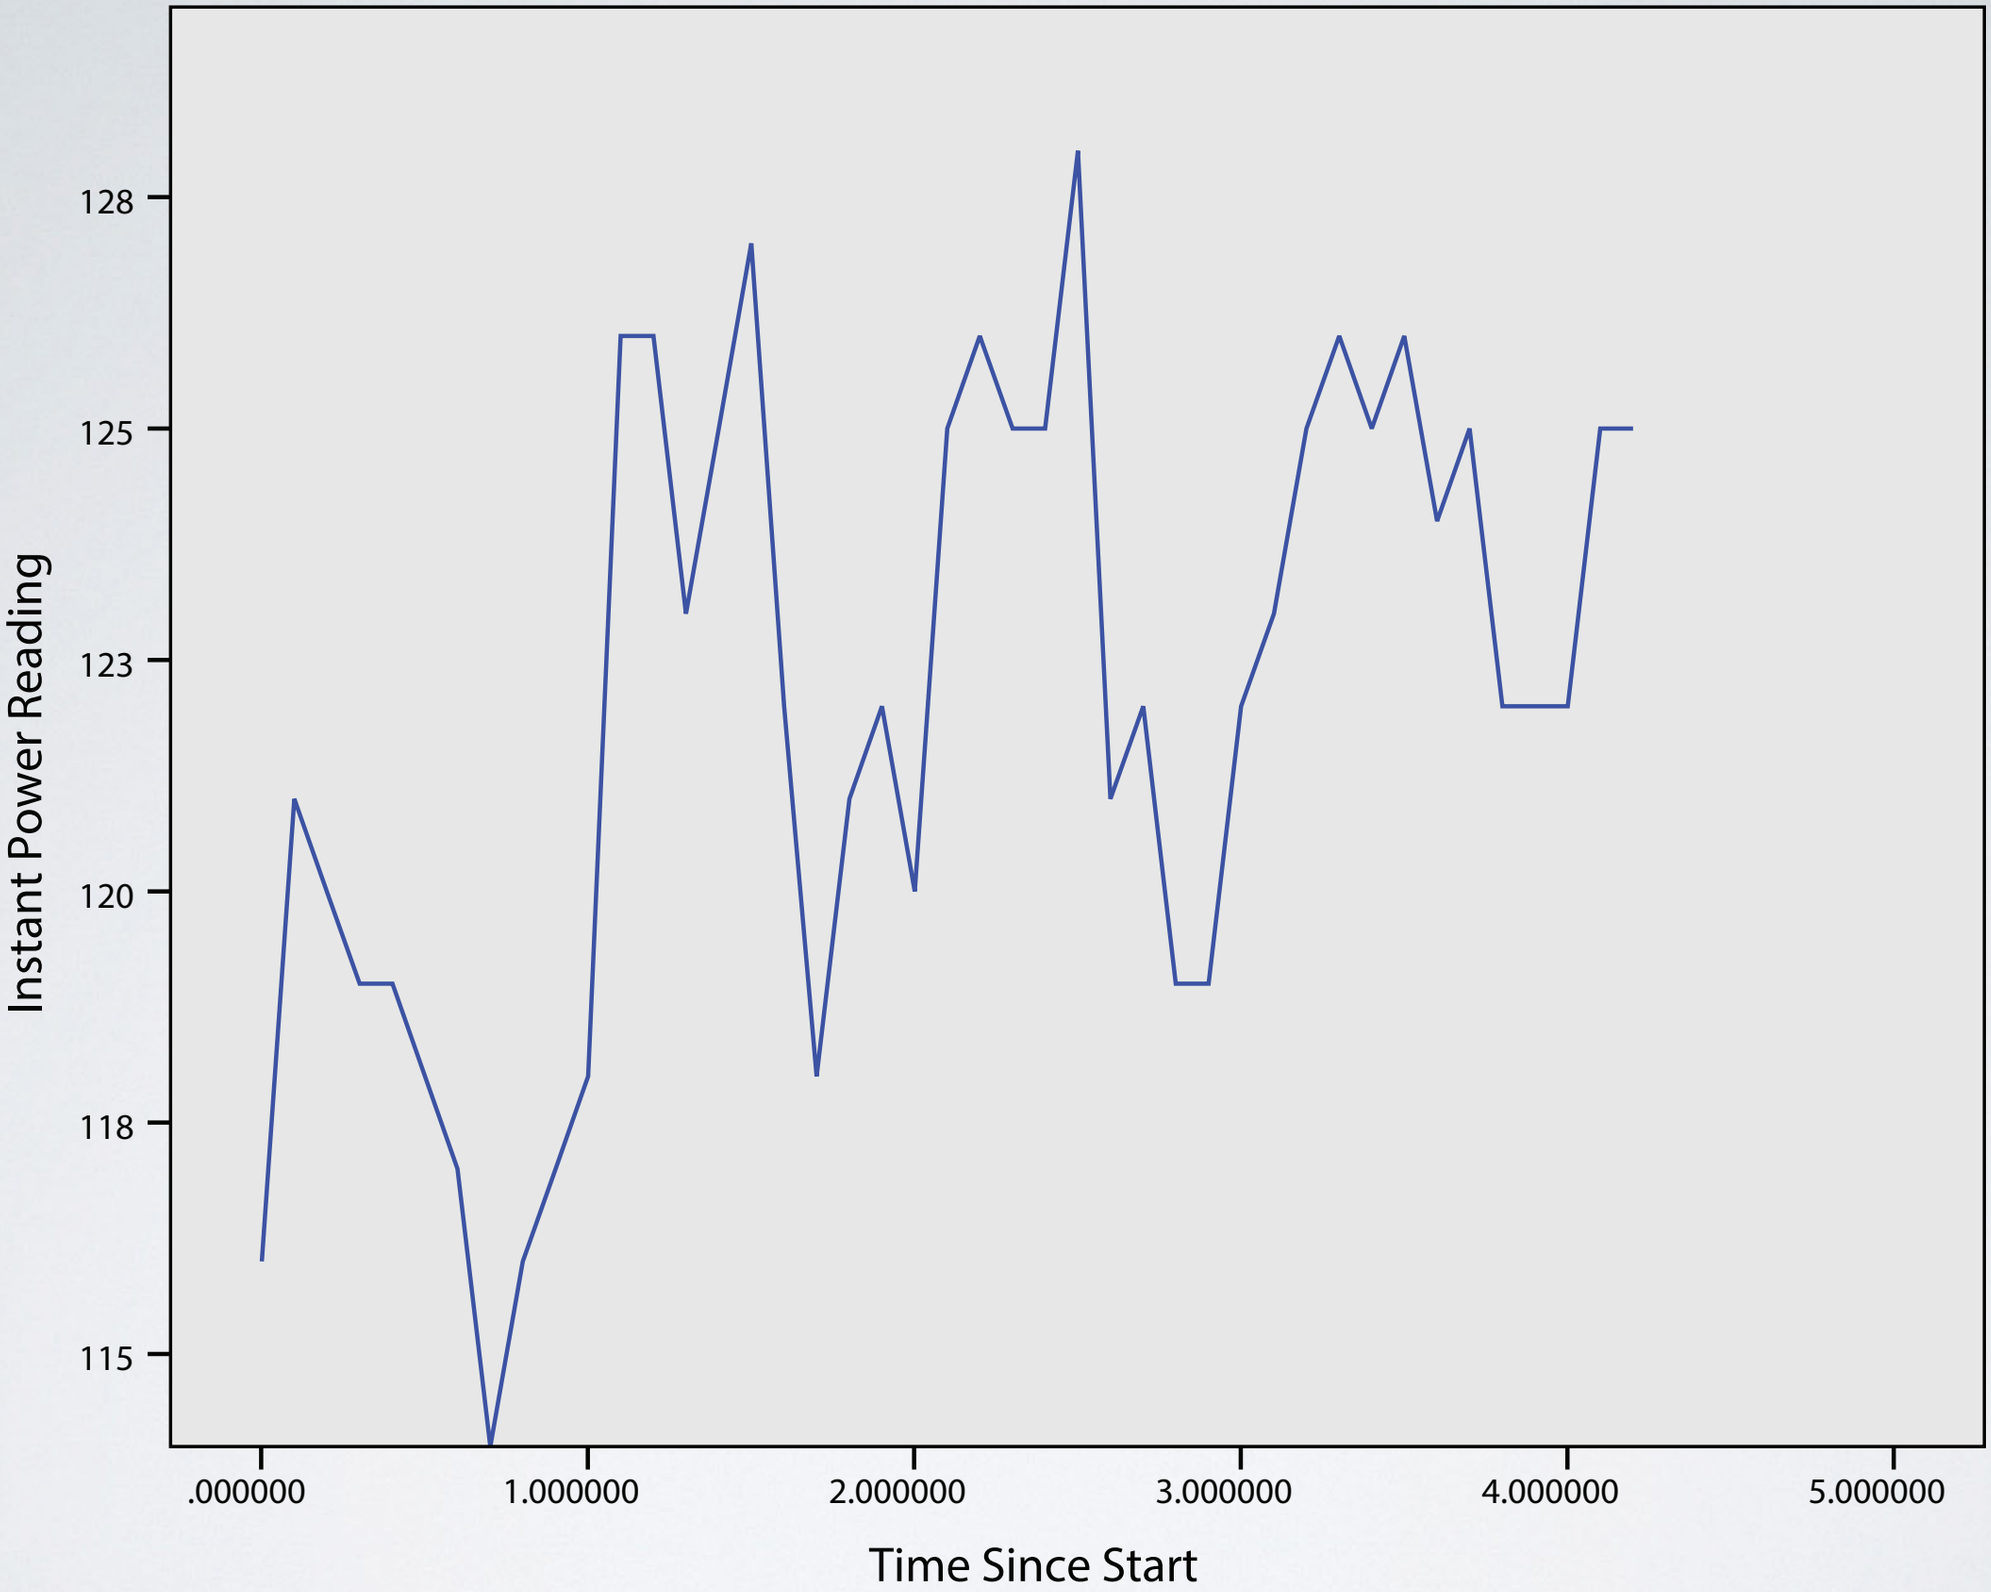

SLEEP ON XEON PHI WITH GEMTC

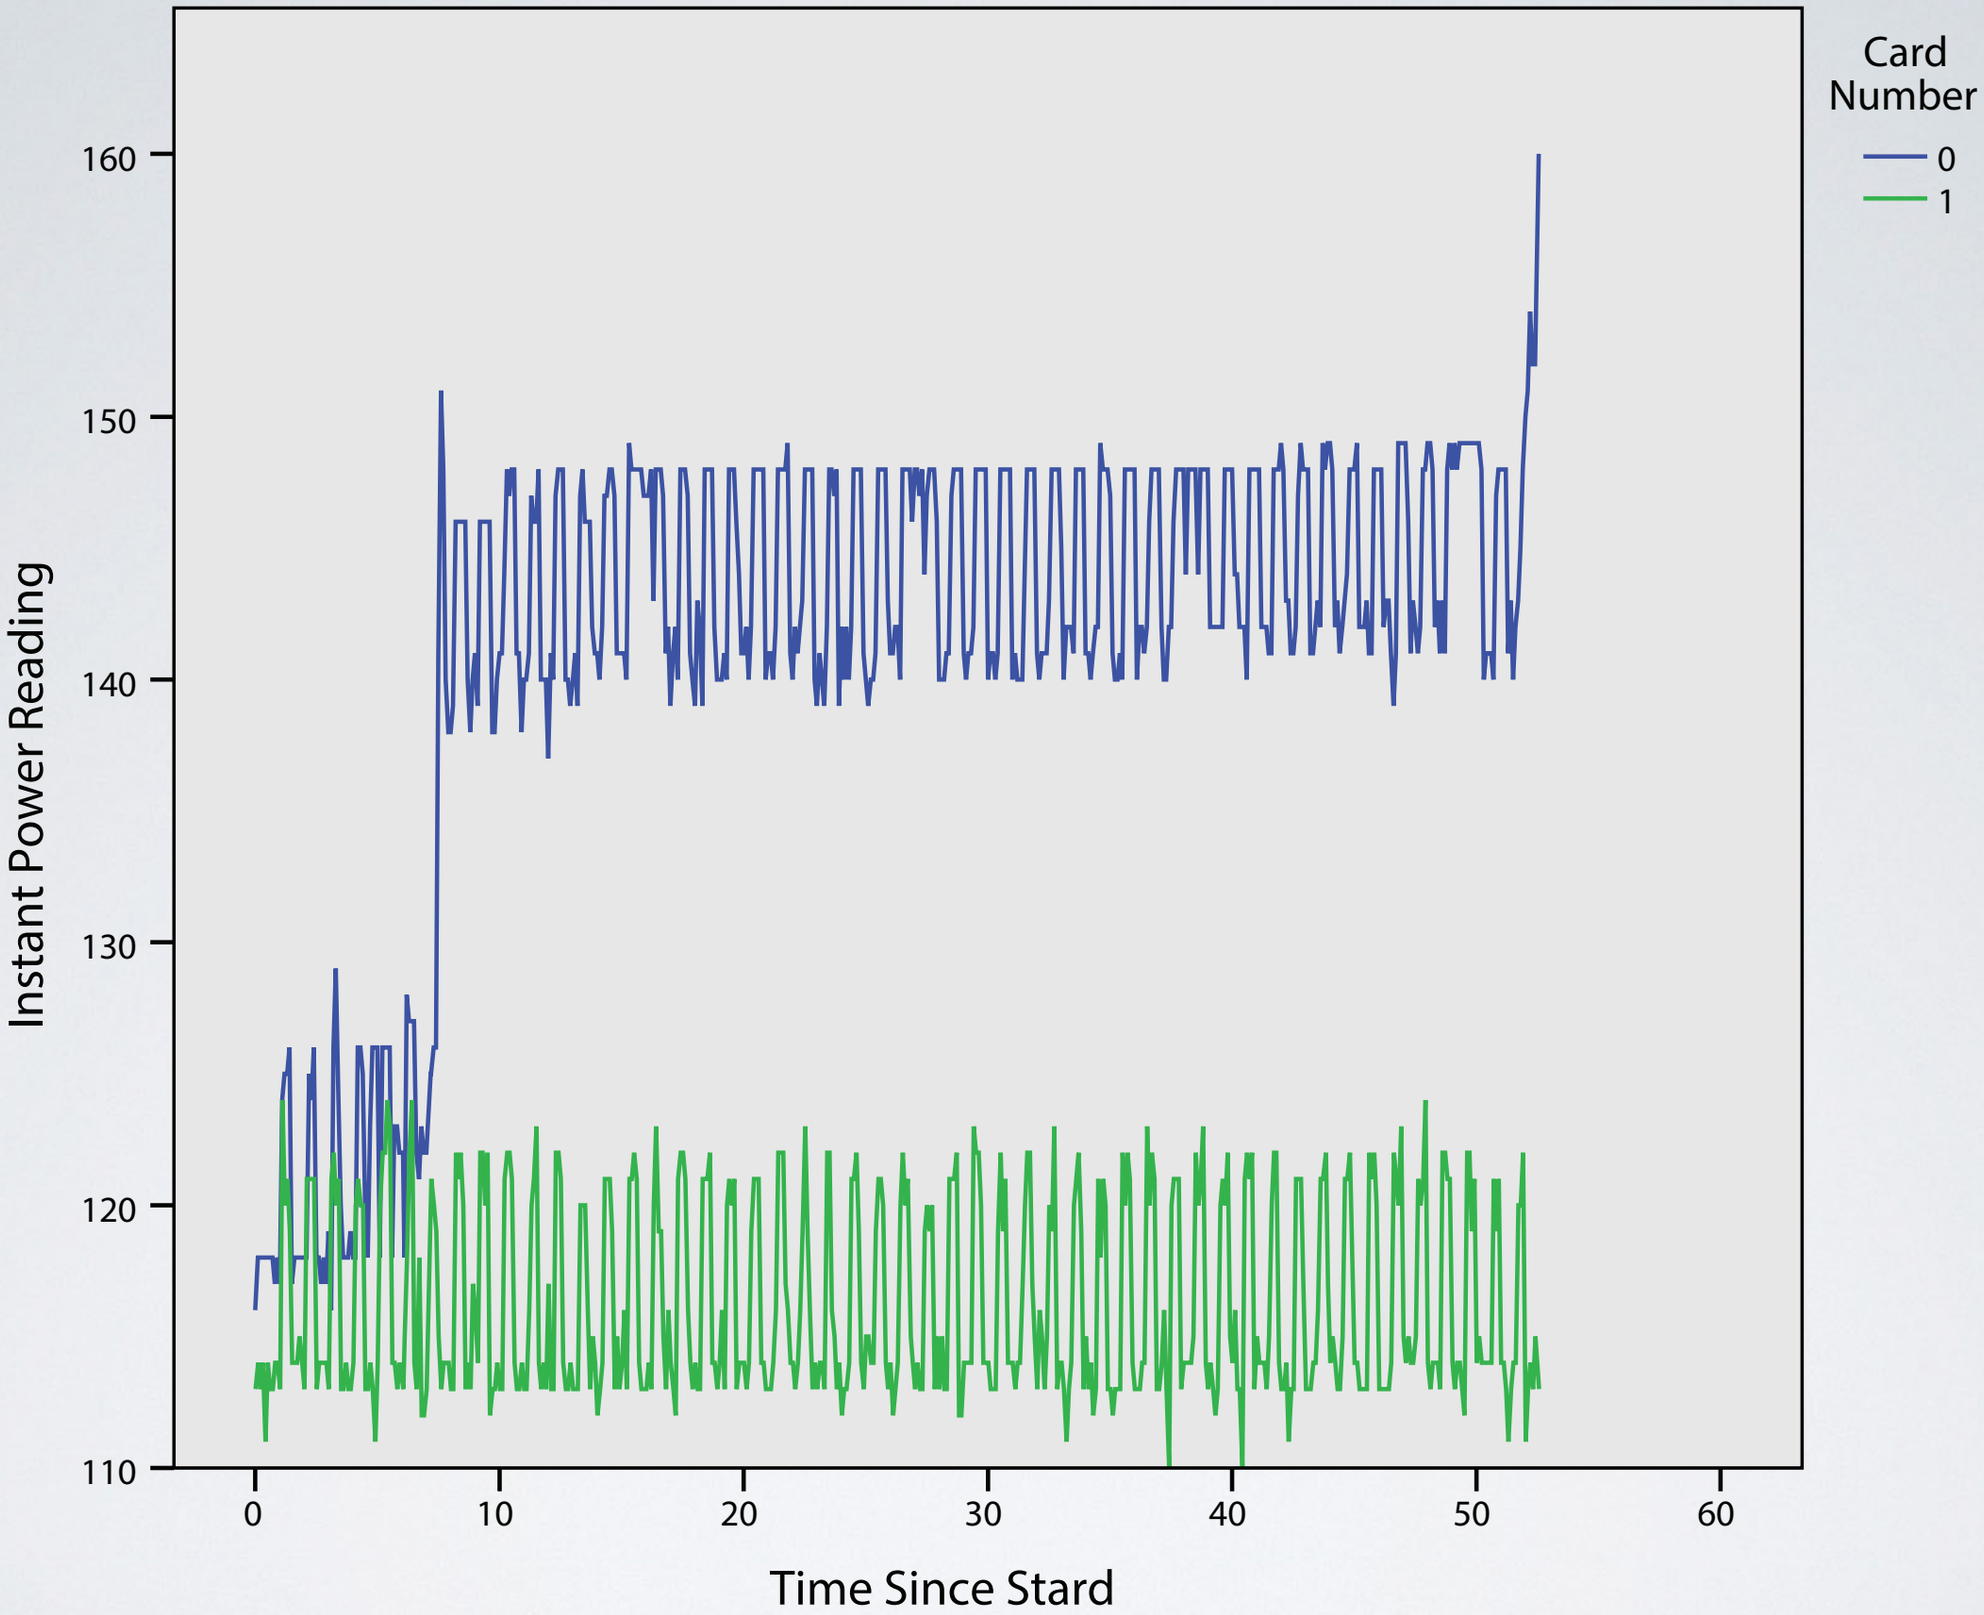

VECTOR ADD ON XEON PHI

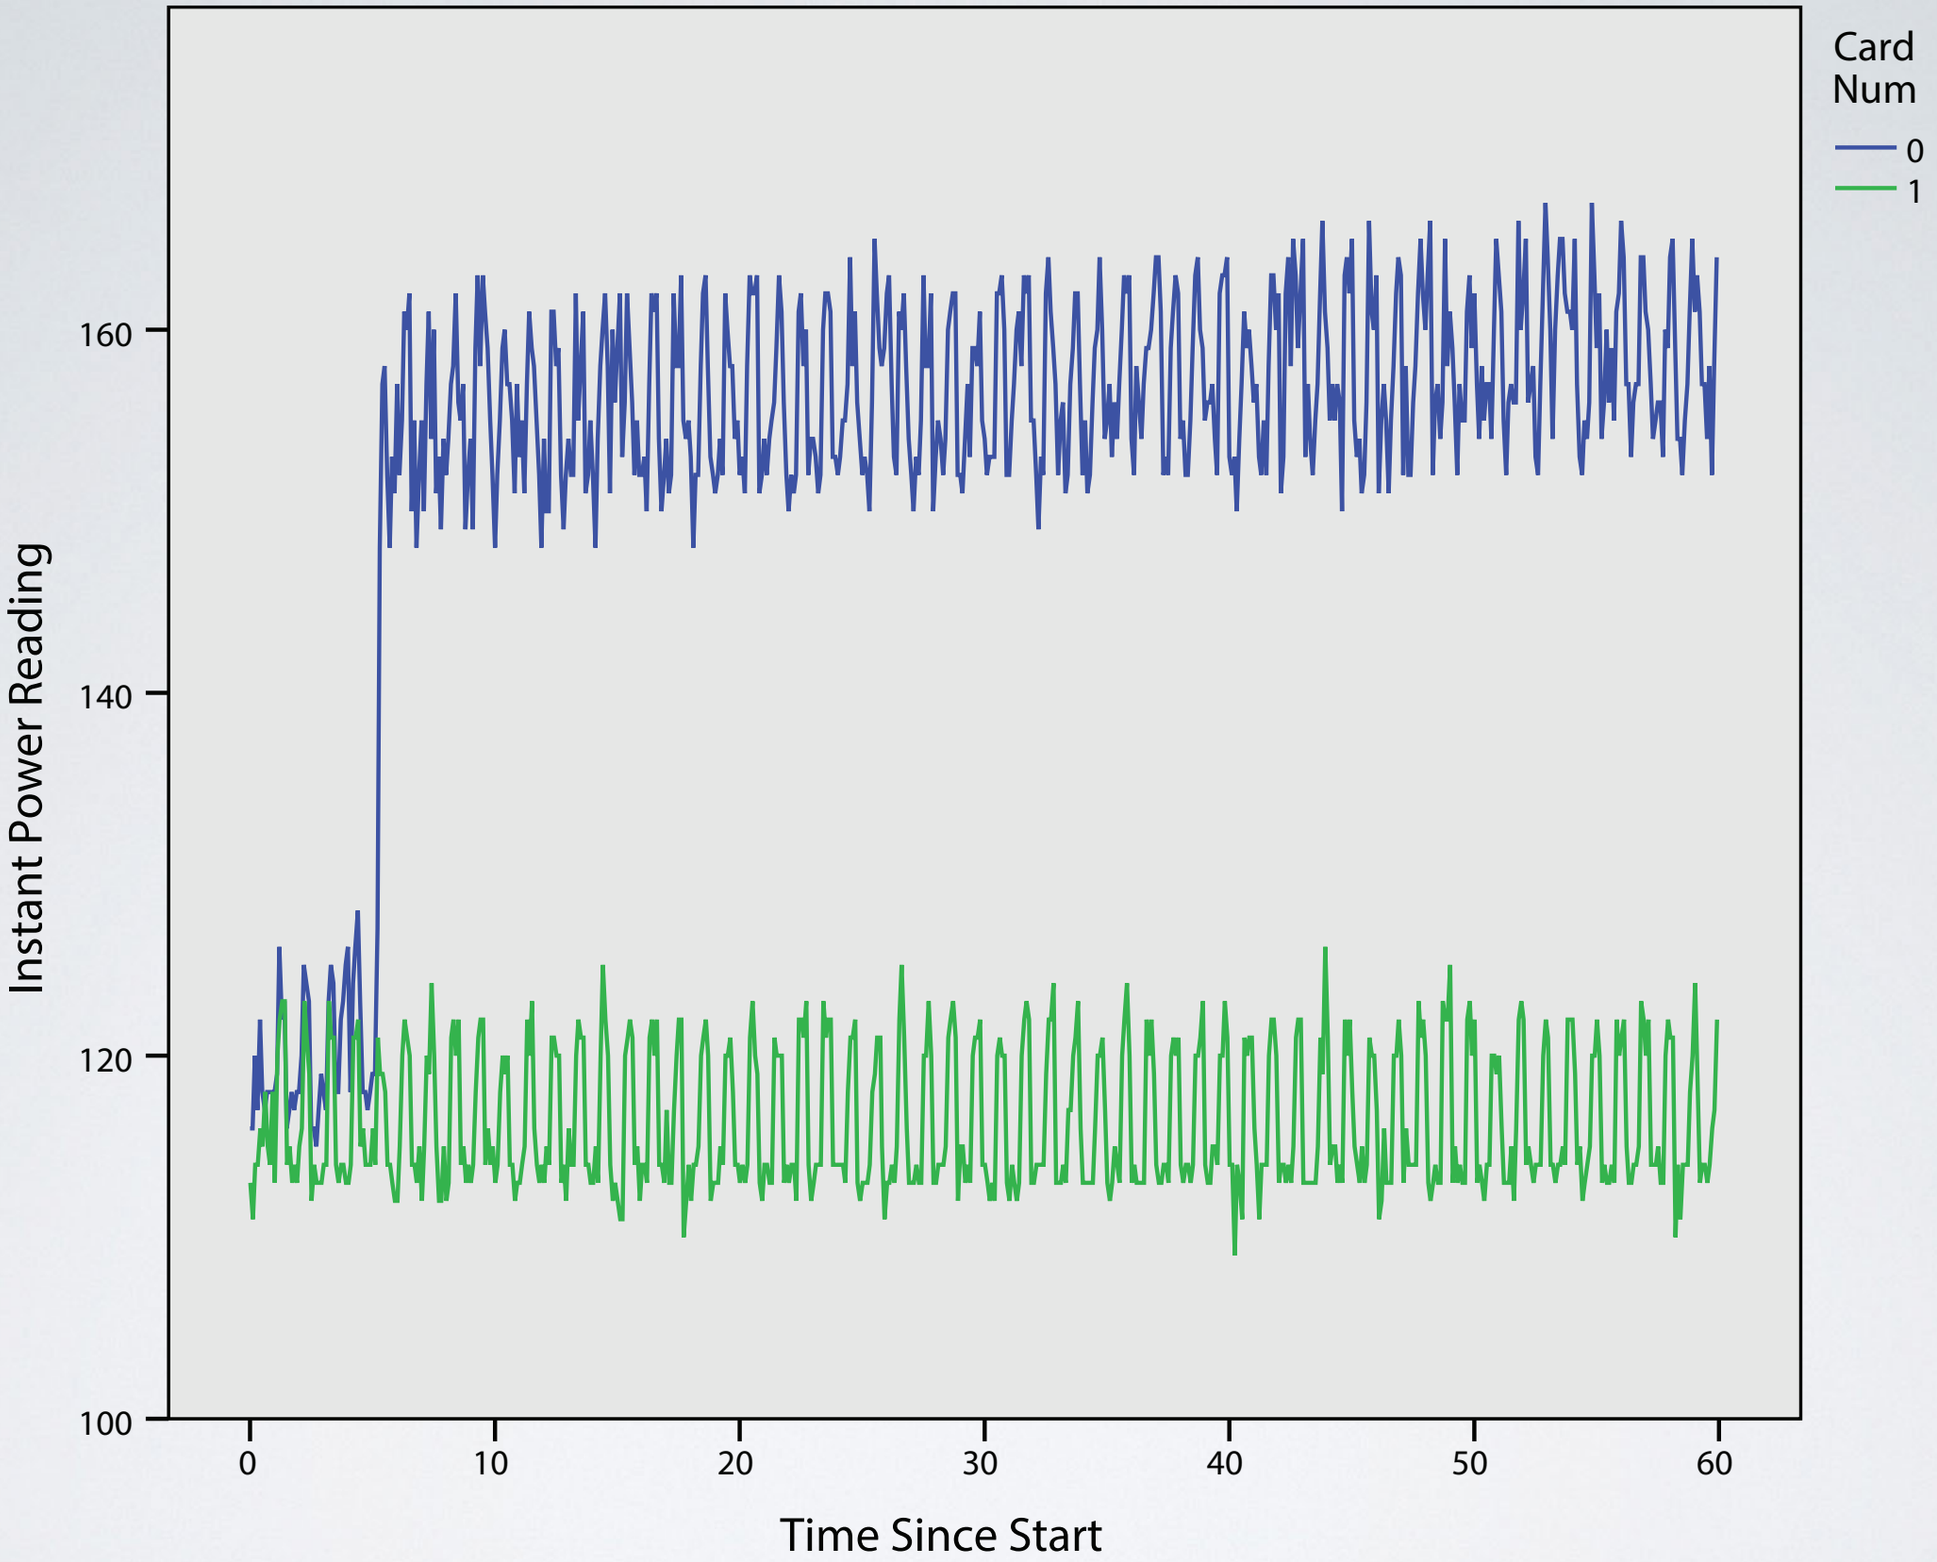

VECTOR ADD ON XEON PHI WITH GEMTC

FUTURE WORK

- Profile additional software both with GeMTC enhancement and without for direct comparison.

- Enhance support for scalability. NVIDIA component scales well to tens of thousands of tasks, however, Xeon Phi component starts having trouble at thousands of tasks.

MONEQ IS (ALMOST) OPEN SOURCE

- https://repo.anl-external.org/repos/PowerMonitoring
- Once the licensing comes through the repository will be open to the public.

QUESTIONS