# An Open-Source Power Monitoring Framework for Real-Time Energy-Aware GPU Scheduling Research

Mohsen Karimi, Yidi Wang, and Hyoseung Kim
University of California, Riverside
mkari007@ucr.edu, ywang665@ucr.edu, hyoseung@ucr.edu

***Abstract*—General-purpose graphics processing units (GPUs) have received substantial interest from the real-time systems community as they can be much more powerful than CPUs on massively-parallelizable, data-intensive workloads. However, GPUs operations are usually power-intensive, and when used in systems with stringent power constraints such as automobiles and battery-powered devices, the use of GPUs makes power analysis difficult and can lead to power/energy budget violations. Although many of today's GPUs have onboard sensors to report their power consumption, they are usually slow and imprecise to be used for real-time GPU research, especially when observing the power consumption behavior of GPUs and developing scheduling techniques in dynamic workload scenarios. In this work, we present an open-source GPU power monitoring framework that measures the actual power consumption of GPUs at runtime with fast and precise responses. The framework can measure the voltage and current of multiple GPUs using an array of INA260 sensors and detect current and voltage changes as little as 1.5 mA and 1.5 mV, respectively, with sampling rates of up to 5 KHz per GPU. We believe this framework will help researchers no longer rely on inaccurate readings from onboard sensors to conduct real-time energy-aware GPU scheduling research.**

## I. Introduction

GPUs, known for their parallel processing capabilities, are widely used in many applications ranging from image processing and machine learning tasks to complex simulations in computational physics. The high performance of GPUs introduces several additional considerations in system design, one of which is their high power consumption. However, it is challenging to analyze the exact power consumption of GPU-using tasks since the detailed architecture of commercial off-the-shelf (COTS) GPUs is not publicly open. While some manufacturers provide tools for power monitoring, such as onboard sensors and APIs for NVIDIA GPUs, they are not suitable for use under real-time scheduling scenarios with multiple GPU-using tasks, due to their low sampling rate and poor accuracy In this work, we design and implement an open-source GPU monitoring framework that is capable of measuring the precise power consumption of COTS GPUs. Our framework is non-intrusive and transparent to GPU workloads as it does not affect the performance of GPU tasks during power monitoring and does not require any modification to the application code.

## II. Demo Description

High-end GPUs usually receive power from PCI Express (PCI-E) as well as directly from the power supply unit (PSU). To measure the entire power consumption of each GPU, both of these powers should be measured when the

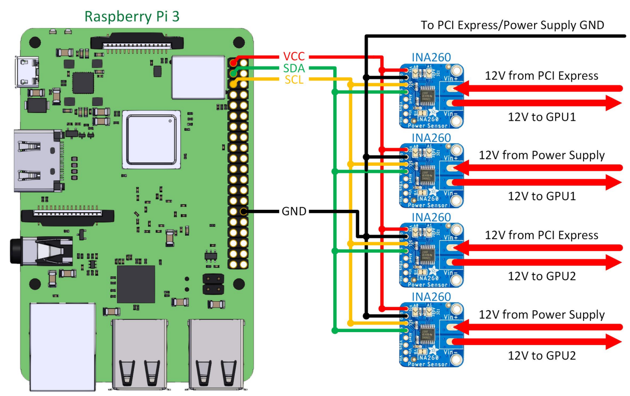


Figure 1: The connections of the sensors

GPU is operating. To measure the current consumption from each power line, a sensor should be installed in series with the power line, which makes the implementation somewhat complicated. Specifically, for PCI-E power, we use a PCI-E extension cable [2] to access the power lines, and then cut the 12V line from the PCI-E extension cable and installed a sensor in series with the line. We chose INA260 [3] as our sensor to measure voltage and current as it provides a sufficient level of sampling rate and accuracy. A similar approach is used for the separate power coming from the PSU. Consequently, we can measure the current and voltage of each GPU using two sensors, the one for energy provided by PCI-E and the other for energy provided by the PSU. Raspberry Pi 3 (RPI3) is used to collect all the data from sensors via I2C and store them in the memory SD card as csv files. The data is then transferred to the PC for further processing. The connections for the sensors are shown in Fig. 1.

We provide an open-source software library that is publicly available at [1]. We also provide an experimental setup to test the monitoring framework, which is also used in [4] to measure the power consumption of two heterogeneous GPUs, NVIDIA RTX 3070 and T400, under various real-time task execution scenarios. The setup for the experiment is shown if Fig. 2. In this setup, we are able to monitor the power consumption of the two GPUs at the same time by using four INA260 sensors. The RPI3 and PC are connected over WiFi and the Network Time Protocol (NTP) is used to synchronize time between RPI3 and PC. The data is stored in RPI3 as csv files containing the voltage and current of the sensors as well as the

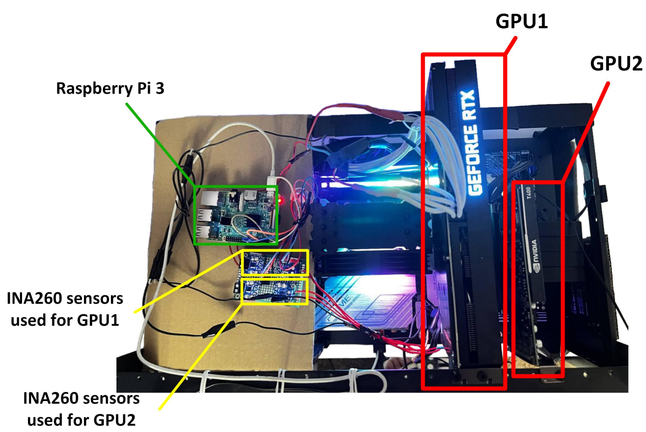


Figure 2: The hardware setup for experiment

timestamp of each sample in microseconds. The program written in C language and tested on the Linux Kernel 5.15.56-v7+ running on RPI3. Each of the INA260 sensors is capable of sampling current and voltage at the maximum speed of 140 $\mu s$ per sample. However, the bandwidth limitation of the I2C bus would make it challenging to receive data from multiple sensors at that speed. The operating system scheduler can also cause delays due to other tasks running on the system. Therefore, we set the priority of the task to *real-time* in Linux operating system and disabled unnecessary services such as GUI to reduce the delay caused by other tasks. We could achieve a sampling rate of more than 5K samples/s when recording voltage and current of two GPUs simultaneously, which corresponds to one sample every 200 $\mu s$. The block diagram of the connections for the setup with two GPUs is shown in Fig. 3.

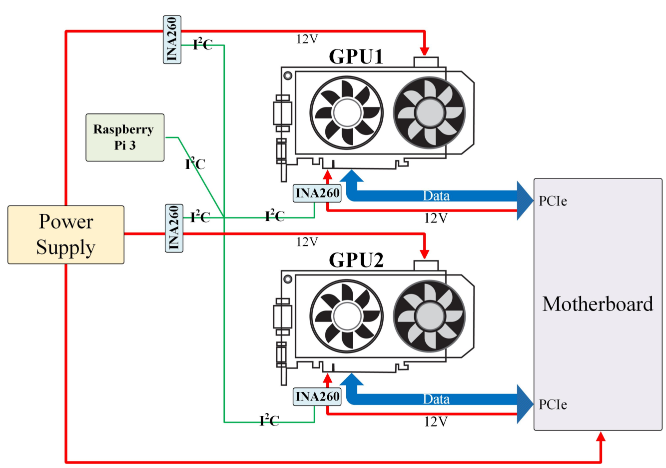


Figure 3: Block diagram of the connections for two GPU setup

To test the performance of the system, we sampled 1 million current and voltage samples and recorded the measurement time for each of the samples. Fig. 4 shows the histogram of measurement times for one million samples. The minimum, maximum, average, and 99th percentile of the measurement times are recorded as 168 $\mu s$, 224 $\mu s$, 177 $\mu s$, and 181 $\mu s$, respectively.

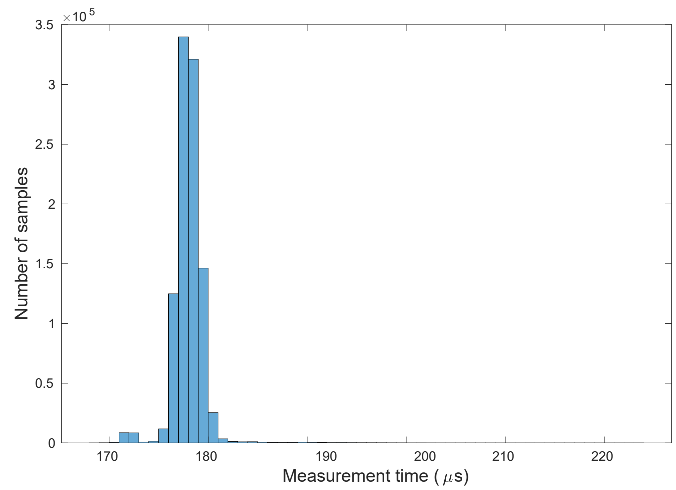


Figure 4: Histogram of measurement times

Although the current setup uses 4 sensors, our framework itself can be extended to up to 16 sensors (8 GPUs) for a single i2c bus; this limitation is imposed by the fact that the INA260 sensor can be configured to have up to 16 different I2C addresses by modifying the A0 and A1 connections of the sensor [3]. However, this limitation can be easily overcome by using multiple I2C buses and we expect our framework can be used for monitoring eight or more GPUs.

## References


[1] GPU Power Monitoring System. https://github.com/rtenlab/gpu_power_monitoring.

[2] Corsair. Pcie 3.0 x16 extension cable. https://www.corsair.com/us/en/Products/p/CC-8900419.

[3] Texas Instrument. Ina260 36v, 16-bit, precision i2c output current/voltage/power monitor. https://www.ti.com/product/INA260.

[4] Y. Wang, M. Karimi, and H. Kim. Towards energy-efficient real-time scheduling of heterogeneous multi-gpu systems. In *2022 IEEE Real-Time Systems Symposium (RTSS)*, 2022.