

# PAPI Component: Powercap

Philip (Phil) Vaccaro  
ICL Lunch Talk 9/22/2016



THE UNIVERSITY OF  
TENNESSEE  
KNOXVILLE

# What is Powercap?

- Framework for monitoring and limiting power on Intel processors
- Exposed by Linux kernel as a Sysfs directory tree
- Available in kernel 3.13 and later
- Relies on Intel RAPL hardware counters for measurements

# Intel RAPL

- Hardware counters corresponding to power measurements
- Power measurements are accessed in Linux via MSR (machine specific registers)
- Available in Sandy Bridge and later

# PAPI Power Components

## RAPL Component

- Reads/writes power attributes
- Accesses MSRs to monitor and set power attributes
- Requires root access

## Powercap Component

- Reads/writes power attributes
- Accesses Linux powercapping framework via Sysfs
- Only requires Linux kernel 3.13 or later
- Root not needed

# Why add a Powercap component?

- Low barrier of entry
  - Only need to have kernel 3.13 or later and Sandy Bridge or later
- Don't need root access to poke at MSR's
- Intuitive Sysfs interface

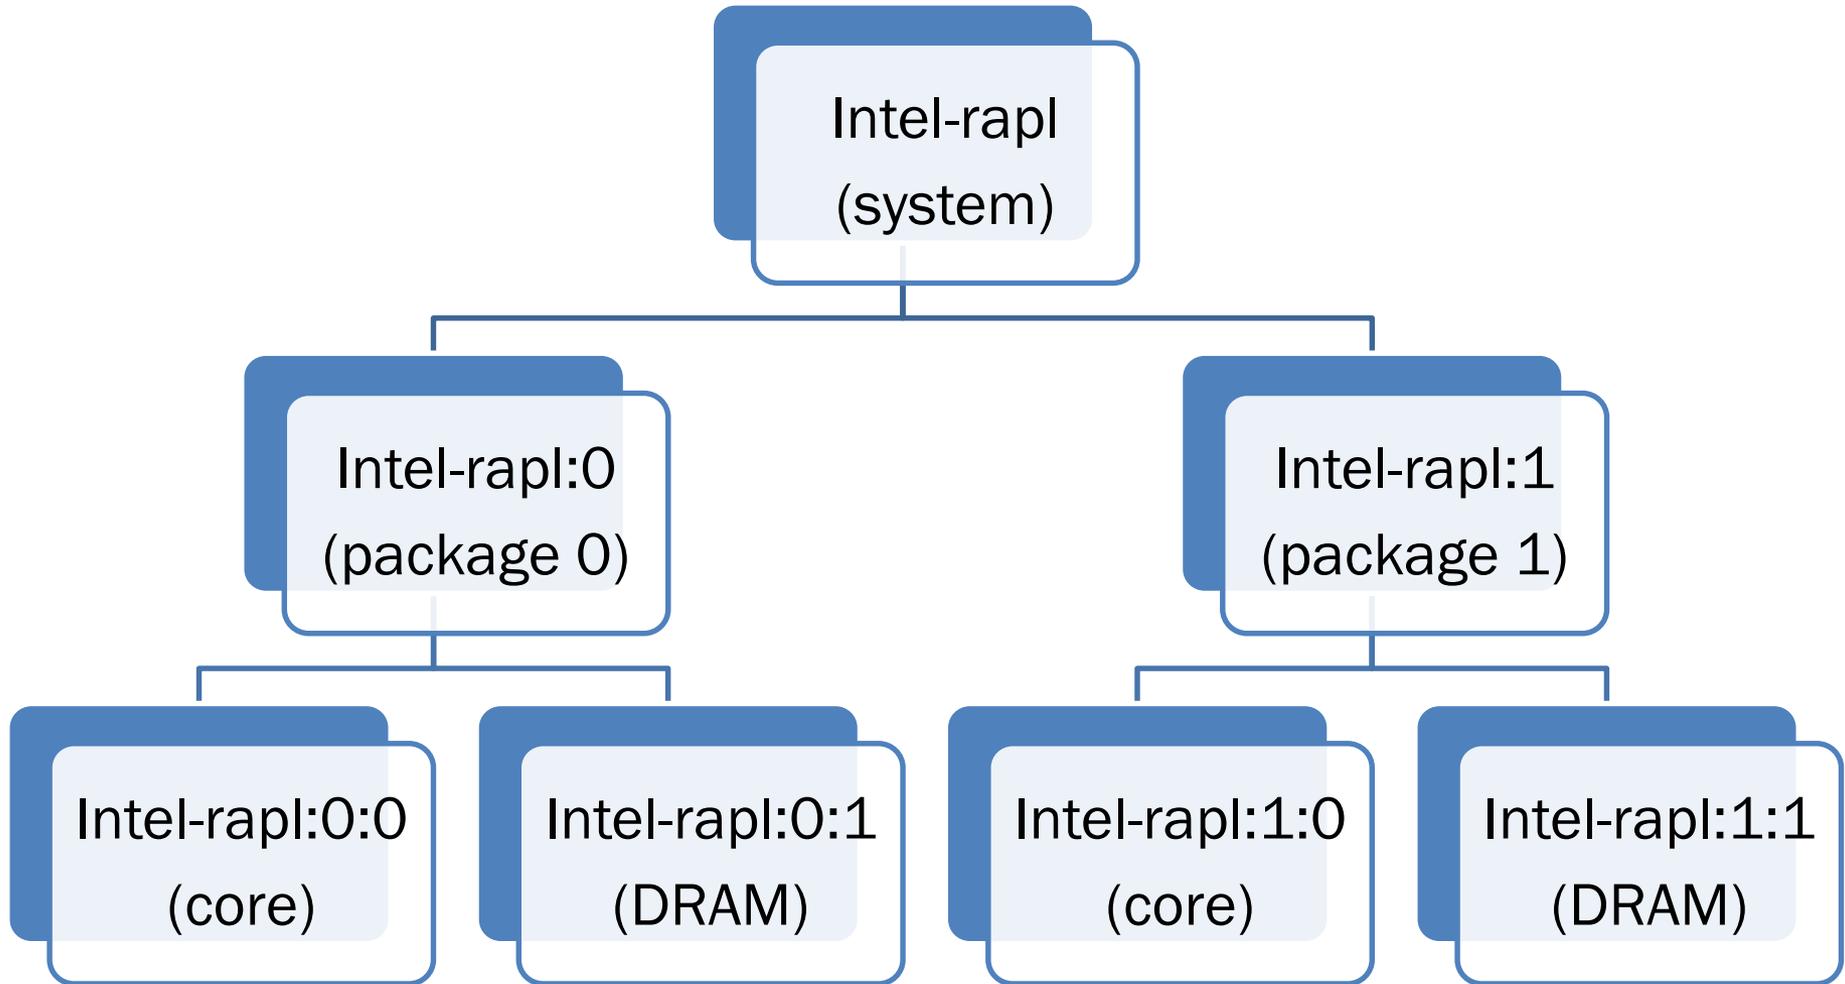
# Sysfs Structure

- Located at /sys/class/powercap
- Organized into hierarchical directories which represent power zones
- Different systems have a different subset of available power attributes
- Top level is package level
  - lower levels can correspond to core, uncore, or DRAM

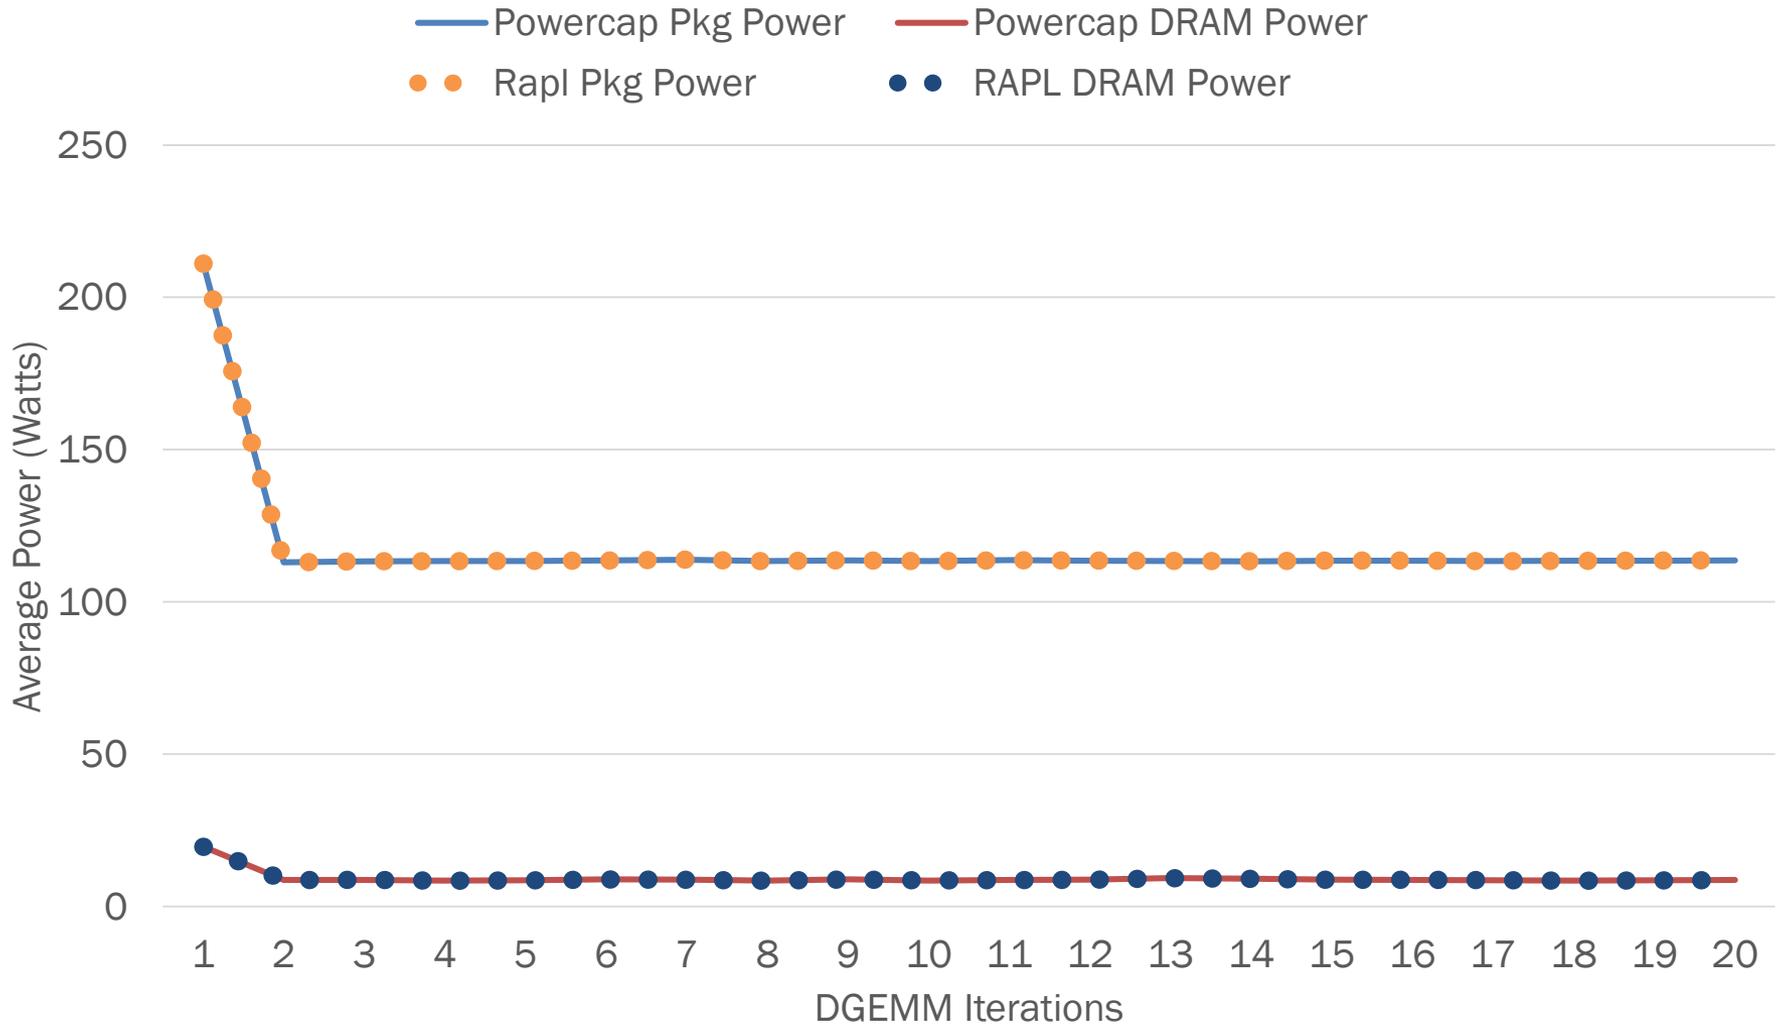
# Power Attributes Example: Scooter

- Two power zones each with core and DRAM attributes
- Each child zone has read only and some writable attributes
- For the top level zones users can set short/long term max power limit
- However, for core and DRAM there is only long term available

# Top Level Sysfs Example: Scooter



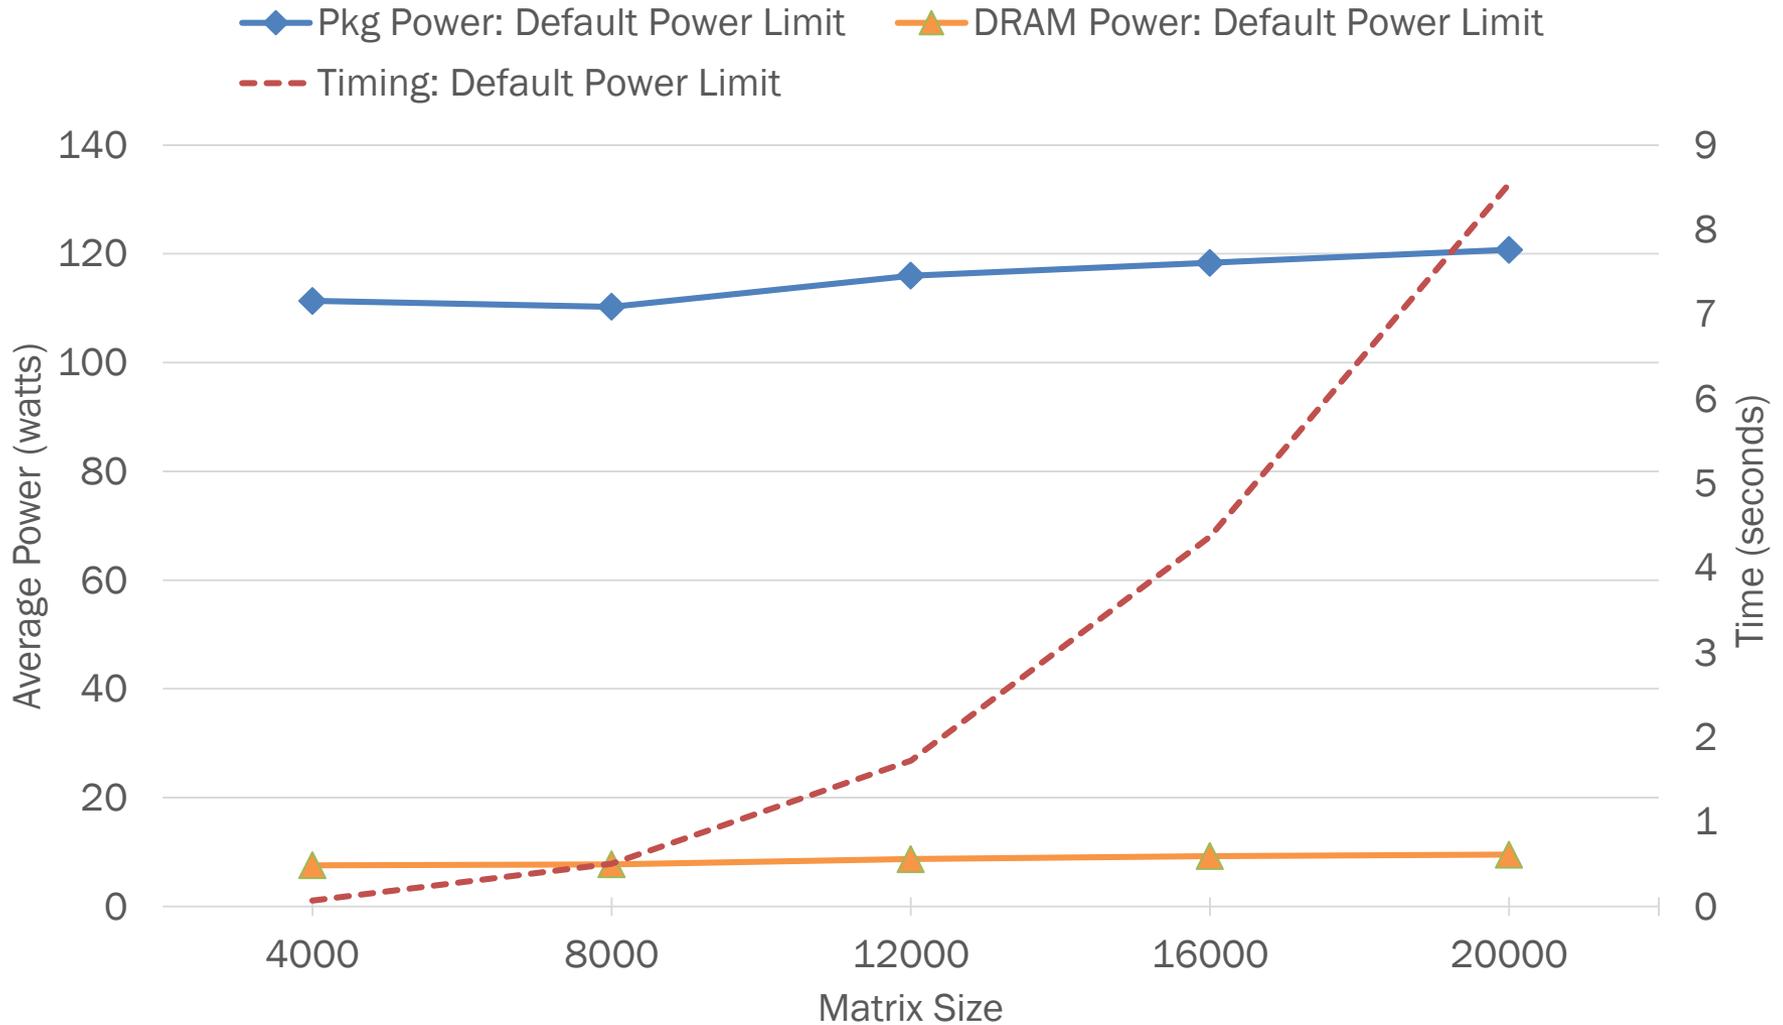
# RAPL/Powercap Power Usage



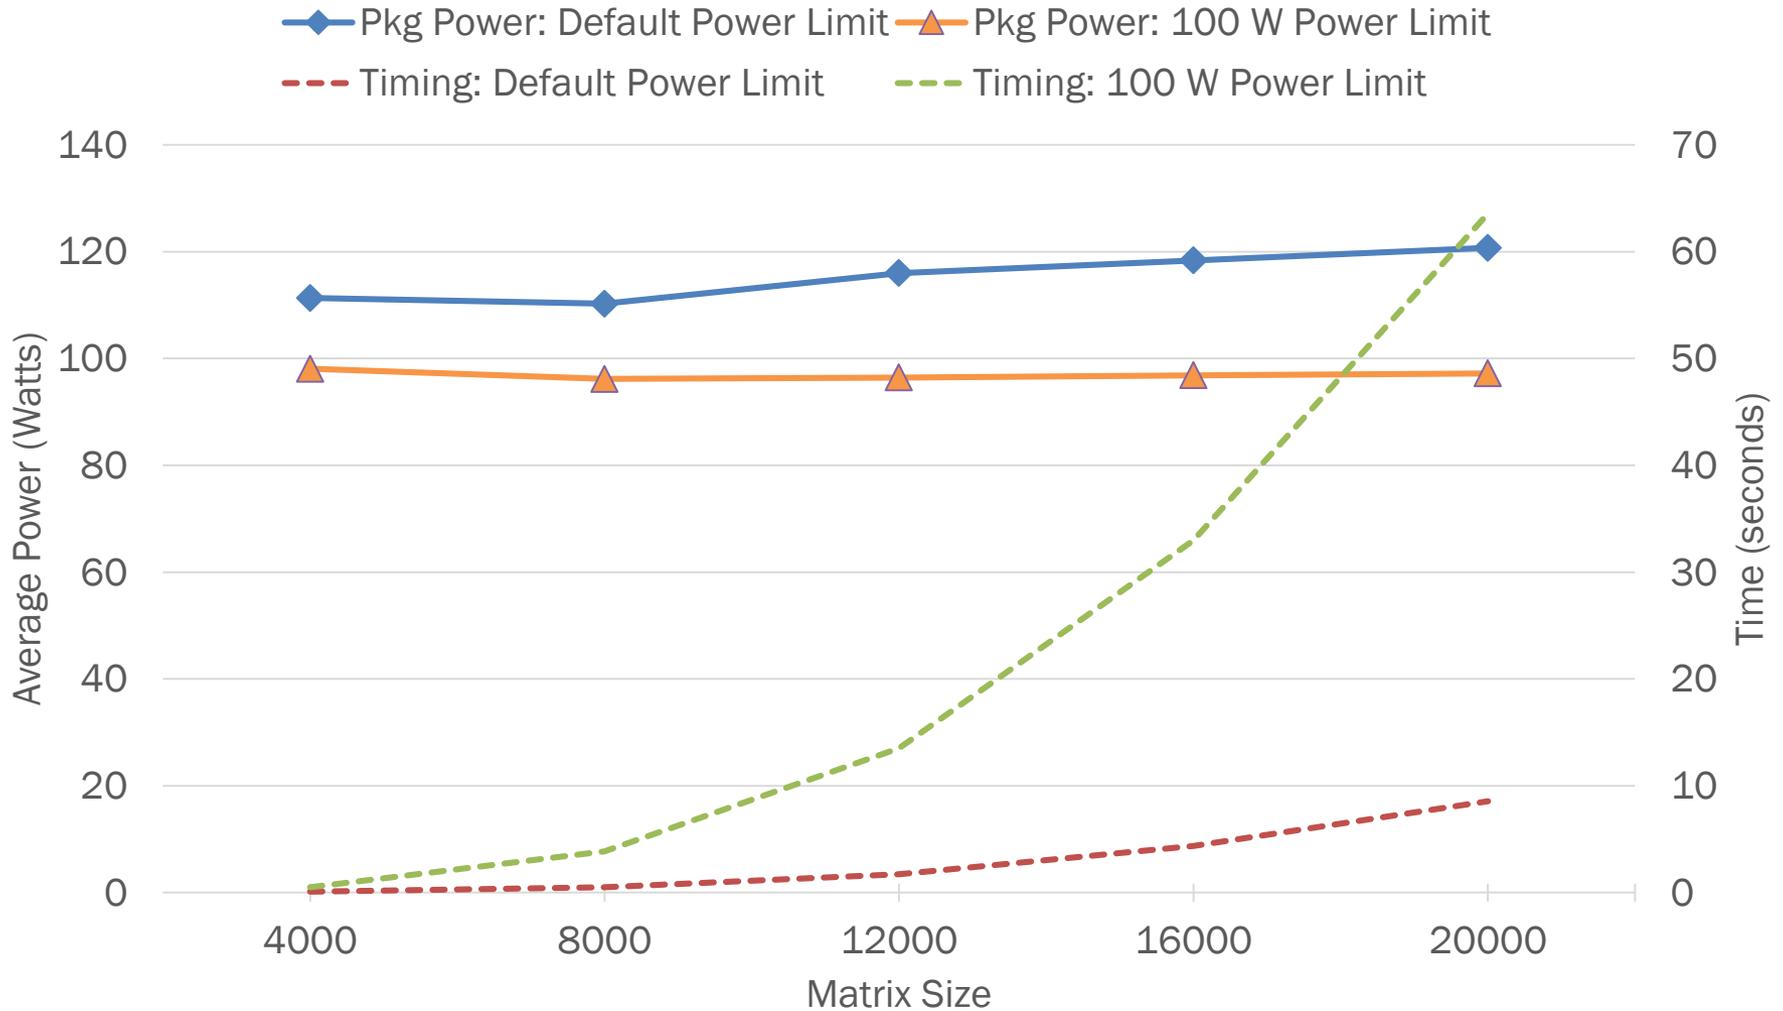
# Applying Power Limit

- As mentioned before, we can write power limits via the Powercap Sysfs.
- Want to show behavior when power limit is enacted
- Next few slides have data pertaining to this

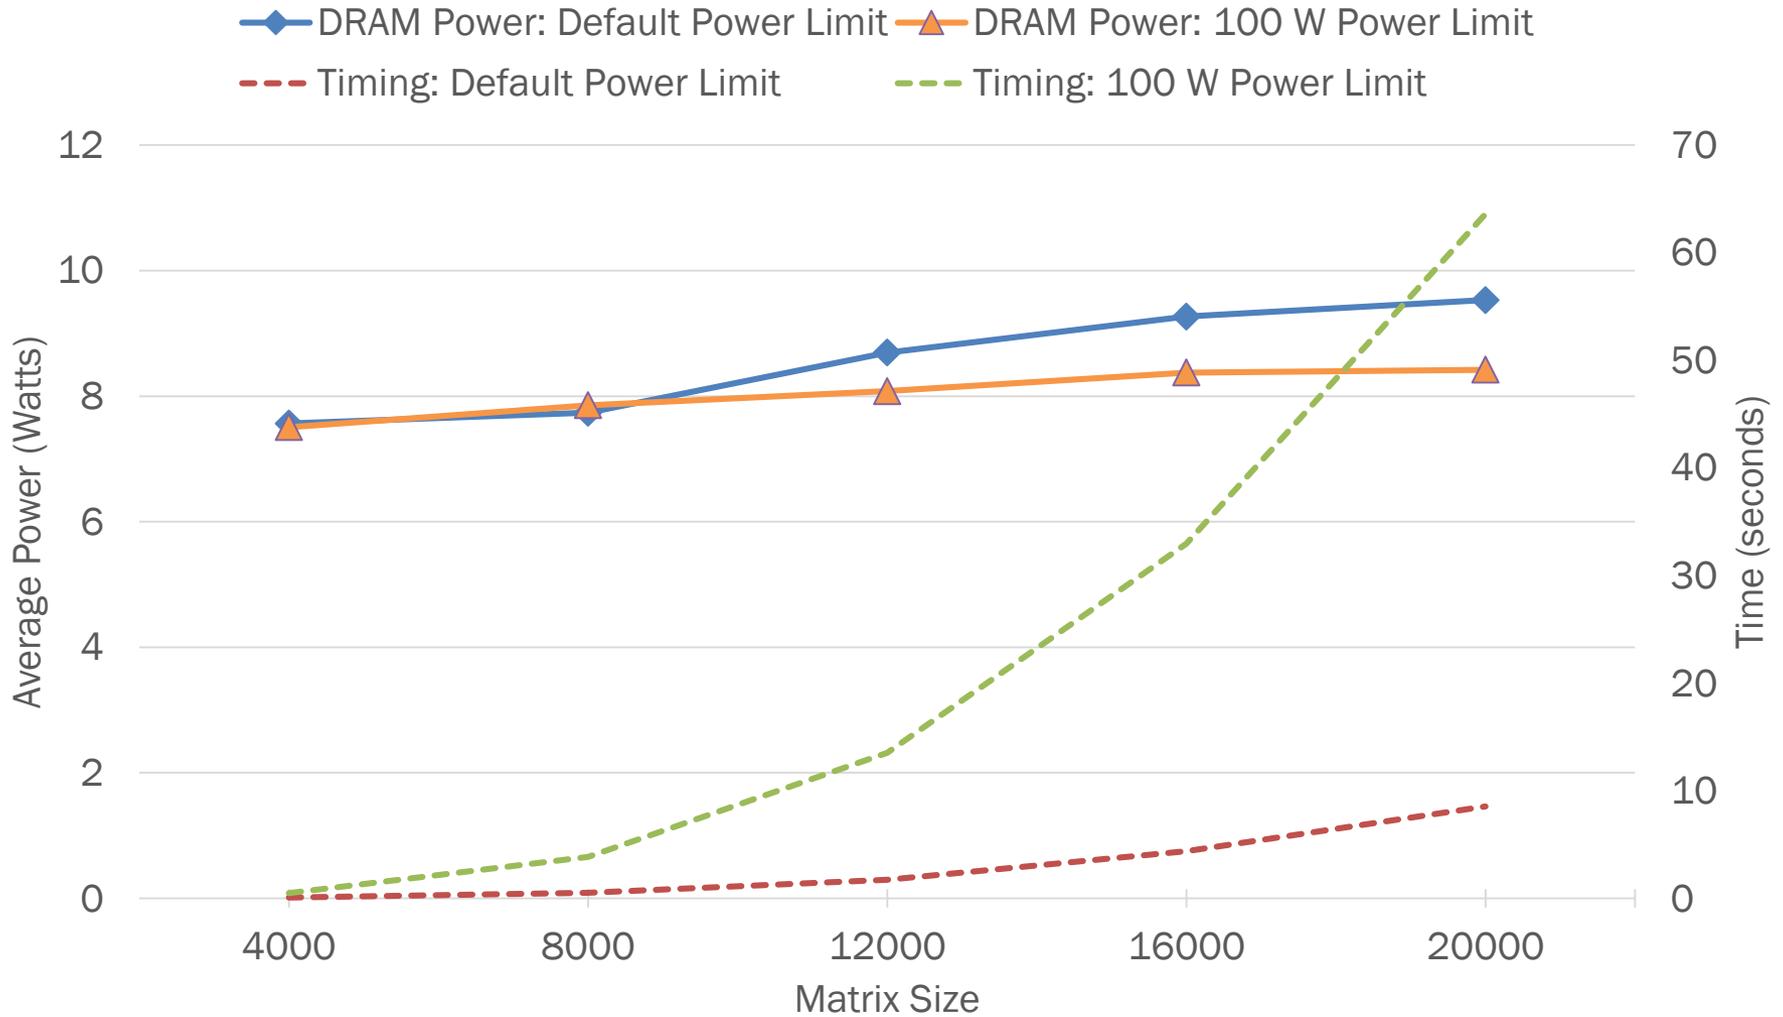
# DGEMM Powercap Data No Limit



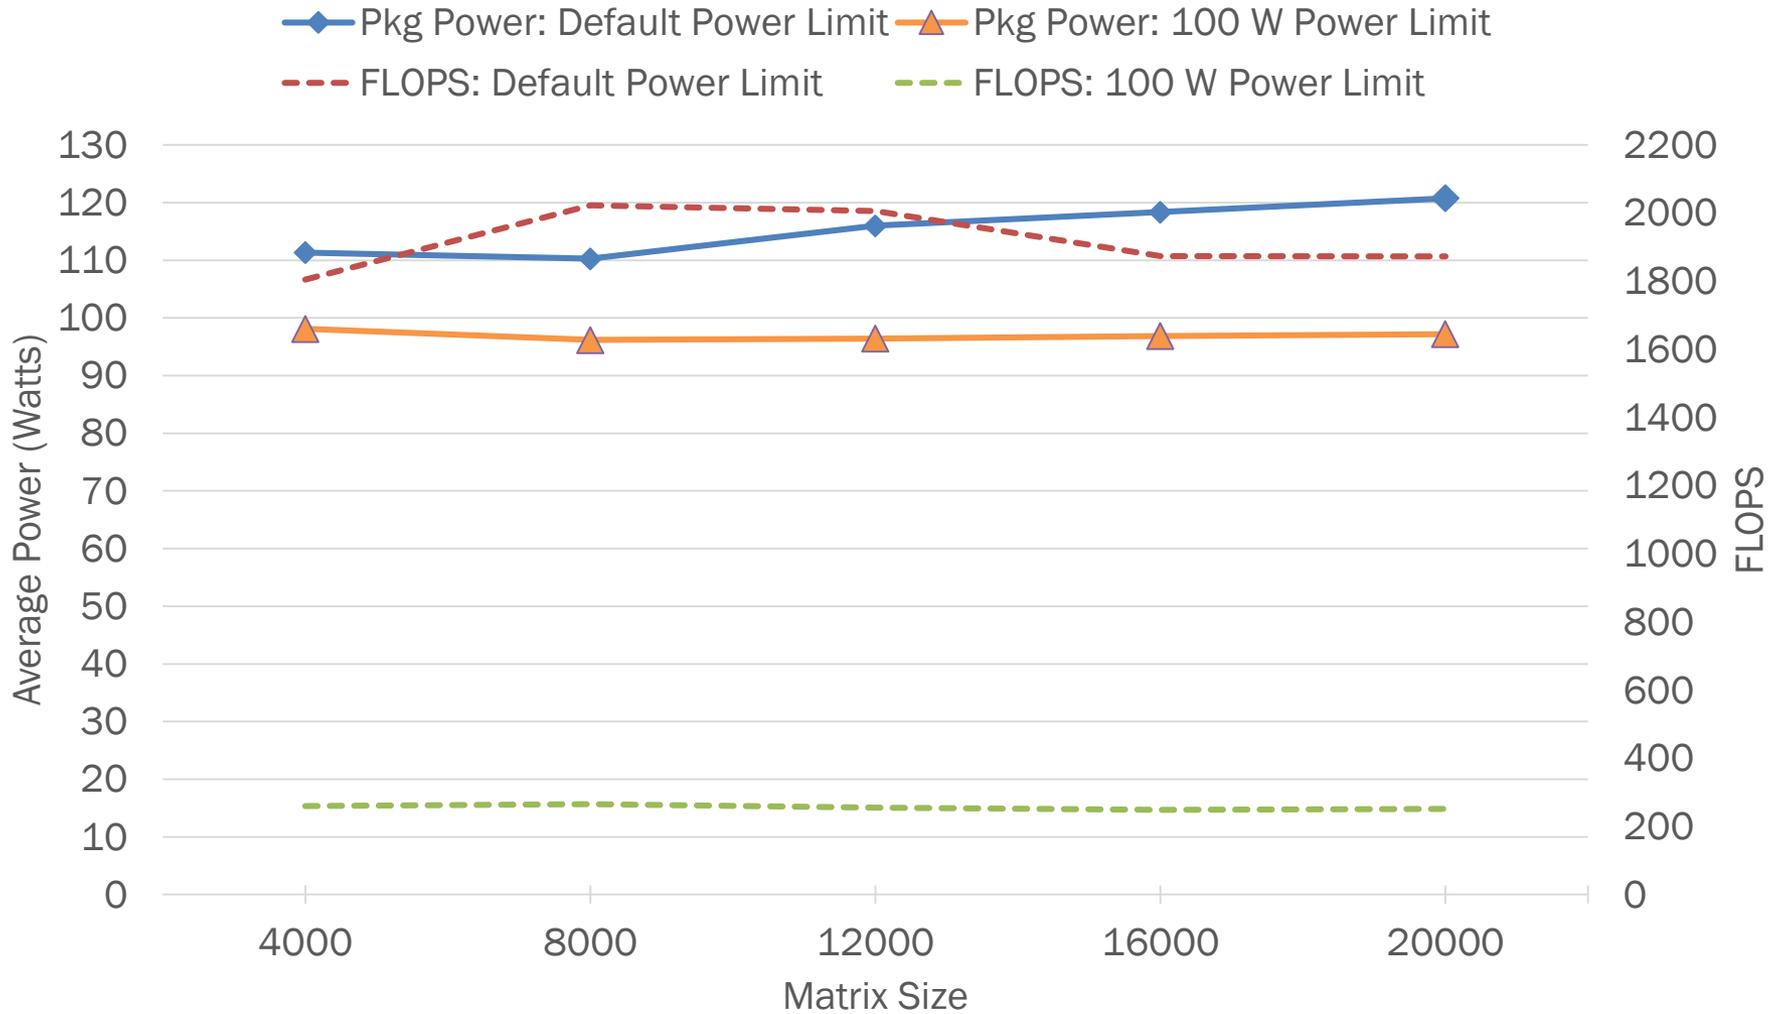
# DGEMM Powercap Data With Limit



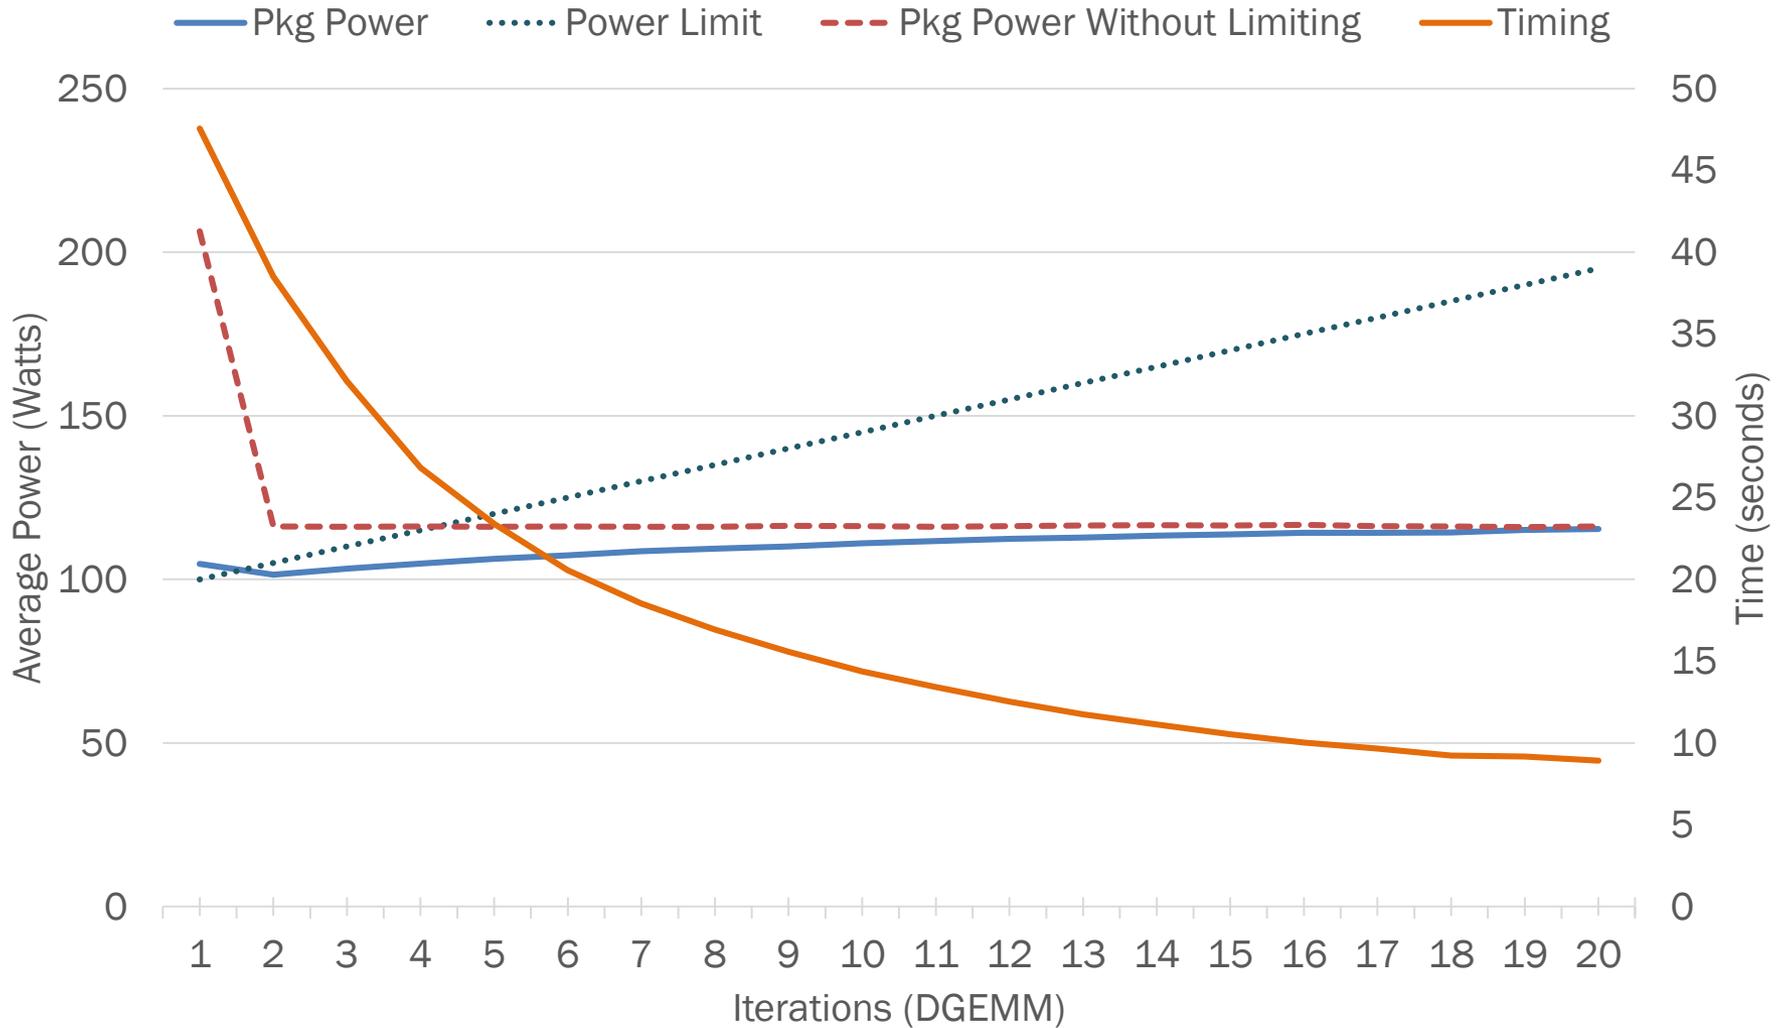
# DGEMM Powercap Data With Limit



# DGEMM Powercap Data With Limit



# Incrementing Power Limit



# Wrap up

- We can use power monitoring and limiting to help find an optimal balance between performance and power usage
- Powercap component is in the latest release of PAPI (5.5.0)
- Current work involves consolidating event names for all power related components in PAPI