KEK
High Energy Accelerator Research Organization

GEANT4
A SIMULATION TOOLKIT

CPU Power Efficiency in Data Center Point of View

Shogo Okada

Koichi Murakami

KEK / CRC

Motivation

- We are interested in Data Center Business:
  - architecture choice, system design, procurement, system installation, operation…
  - CPU choice is an important task.
    - performance, cost efficiency, power efficiency, compatibility, …
  - Geant4 simulation occupies a big percentage of CPU usage in HEP data centers.

- Emerging CPUs other than Intel / x86
  - Intel is in very difficult situation…
  - AMD EPYC (x86)
  - ARM-based processors
    - Fujitsu A64FX, Apple M1,…

KEKCC Central Computing System

2020

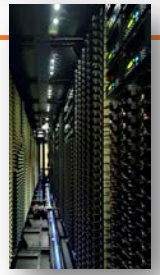

Tape / HSM

IBM

Tape Library : TS4500
Tape Drive : TS1160 x72
Max Capacity : 100 PB

GPFS Disk

Migration / Staging

IBM

ESS GL4C x4 GPFS (17PB)
ESS GL4C x2 GPFS GHI (8.5PB)
ESS GL4C x 1 GPFS (Belle II FE, 700TB)

Mellanox

IB SW : CX7500

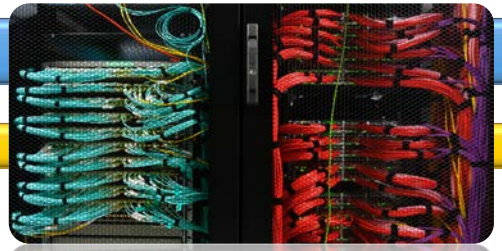

CISCO

Nexsus 9516

IB 4xEDR

10/40 GbE

40 GbE

Work / Batch Servers

Lenovo

SR630 12nodes (Work)
SD530 368nodes (Batch)
15,200 cores (Intel Xeon Gold 6230)
87TB memory

JUNIPER NETWORKS

FW SRX4100

40 GbE

SNET5

LHC

Grid UMD / Service

Lenovo

SR650 48 nodes
SR630 54 nodes
V5030E 546TB

IBM

Batch Servers

KEKCC Central Computing System

- Data analysis and MC generation 24/7
  - 15,000 CPU cores (Intel Xeon Gold 6230)
  - ~ 1M Jobs / day
  - MC generation is essential for all experiments.
    - It may take a significant fraction of CPU jobs. (guess)

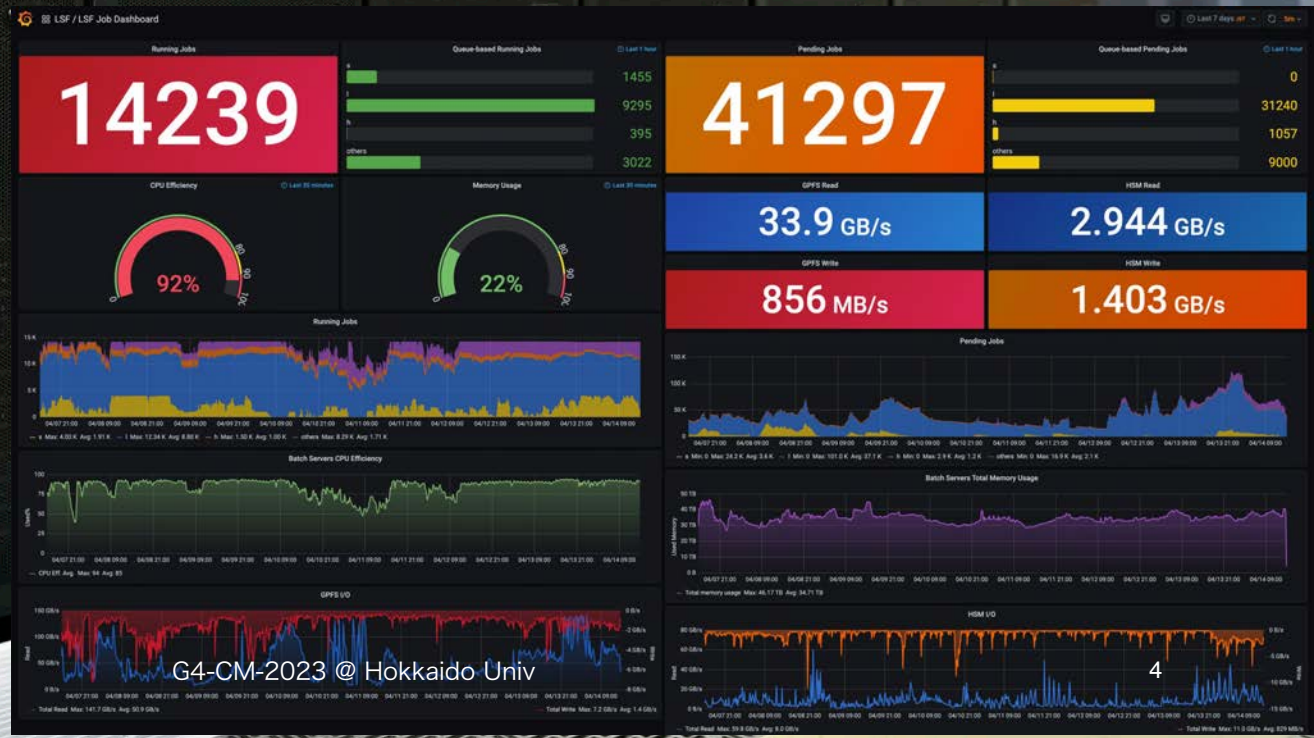

Power Consumption

- Power Consumption Breakdown:
  - Batch Servers (blue): 63%
  - Storage (green): 24%
  - Servers (mazenta): 8%
  - Network (yellow): 4%
- Batch calculation consumes 60% of the total system energy.

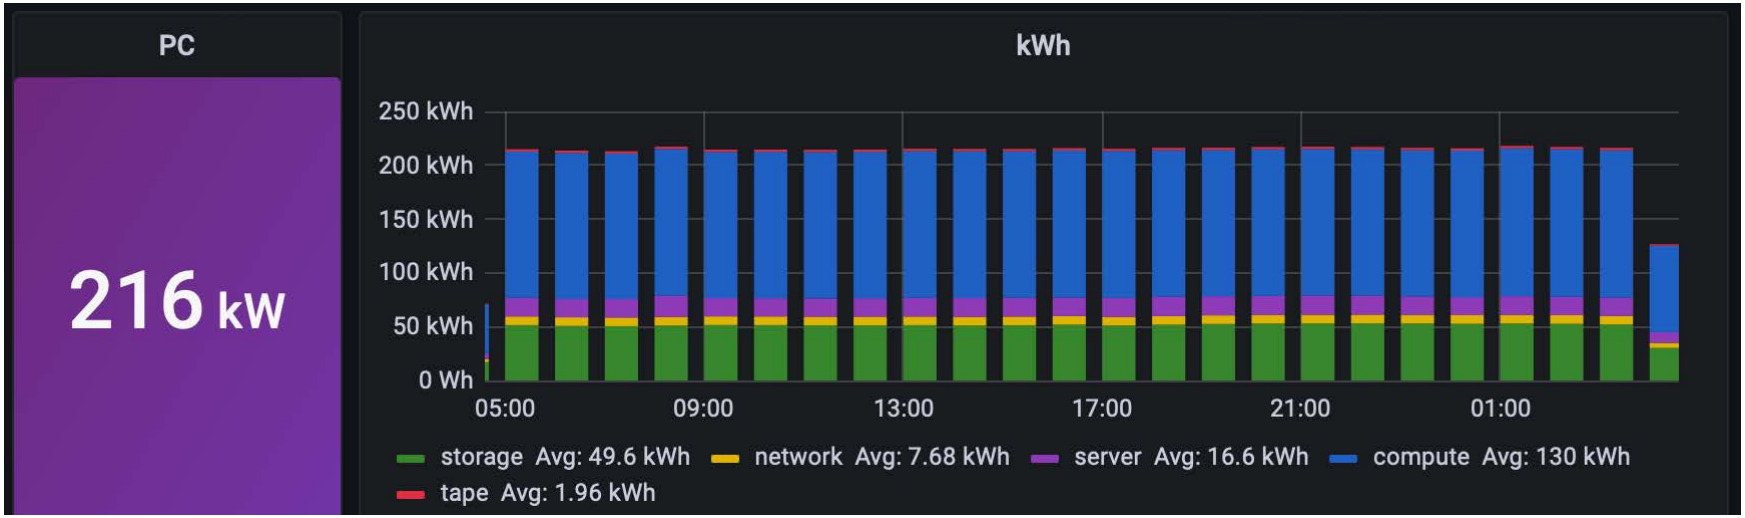

New CPU Metic: Power Efficiency (Apple)

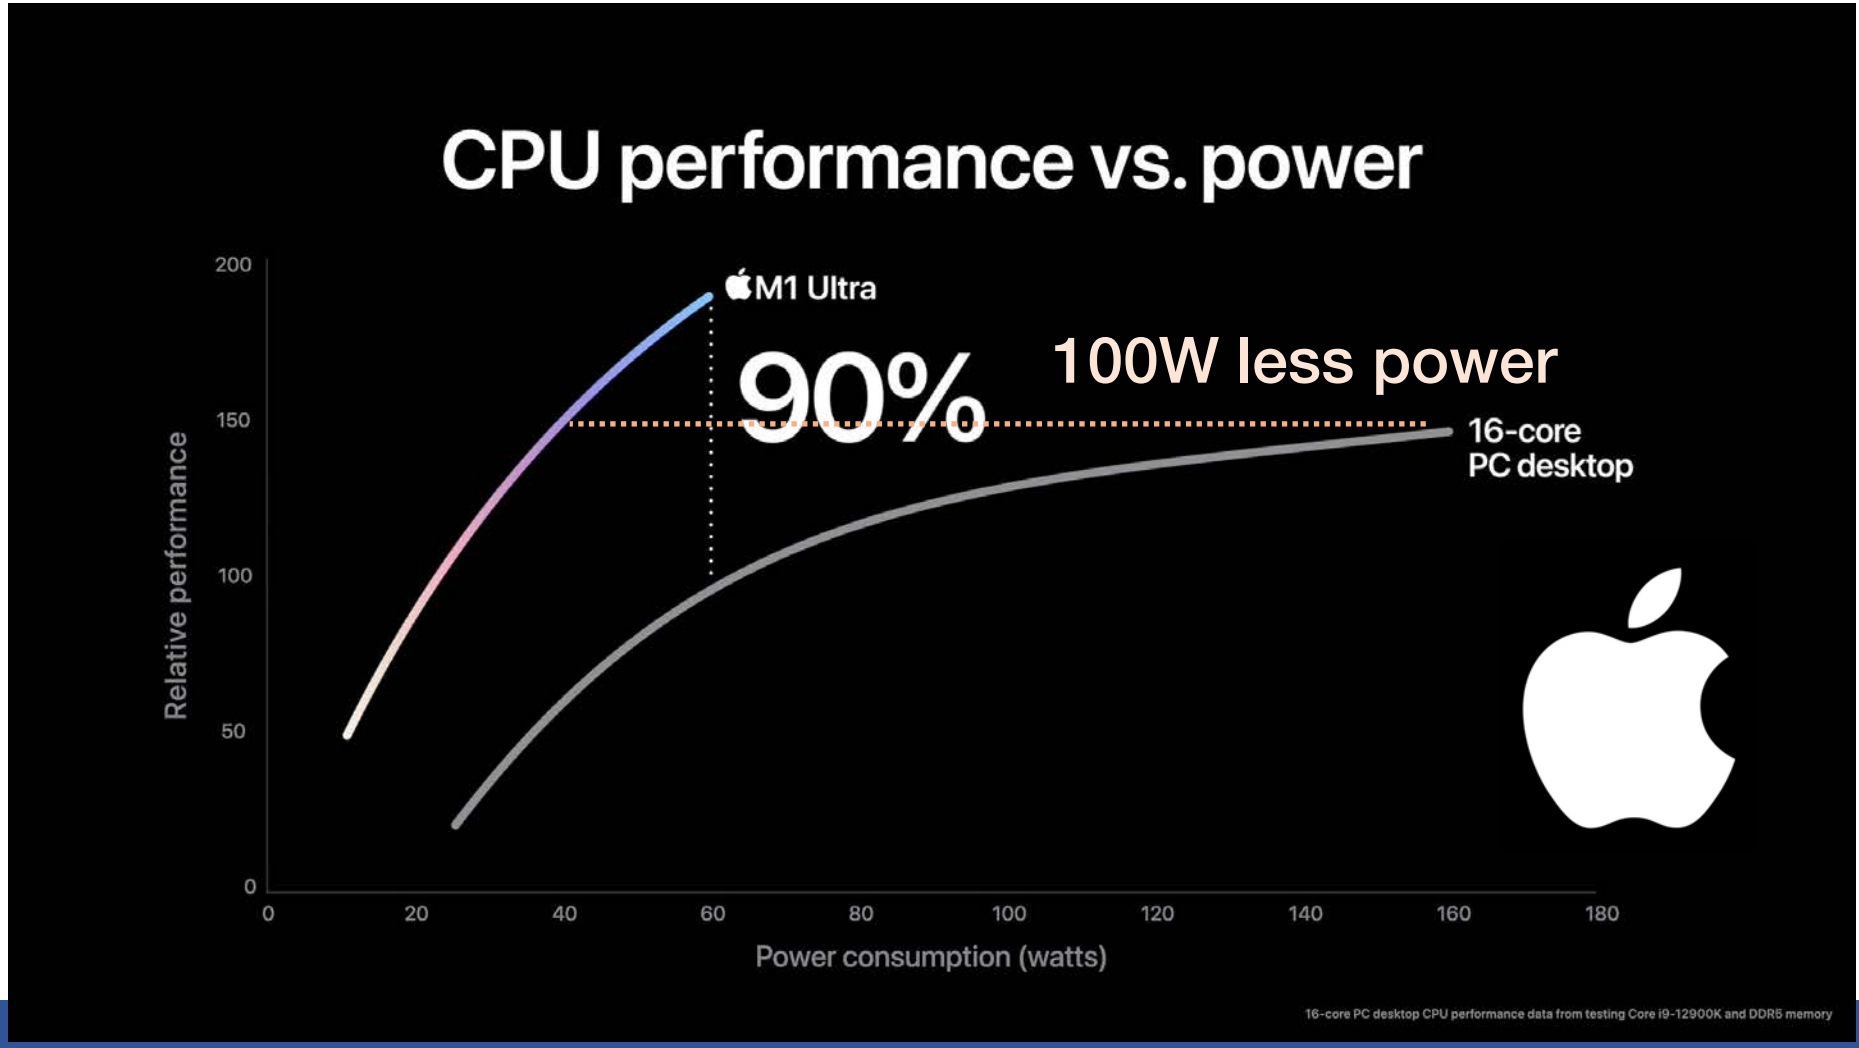

Power Performance: Intel (vs. AMD / Apple)

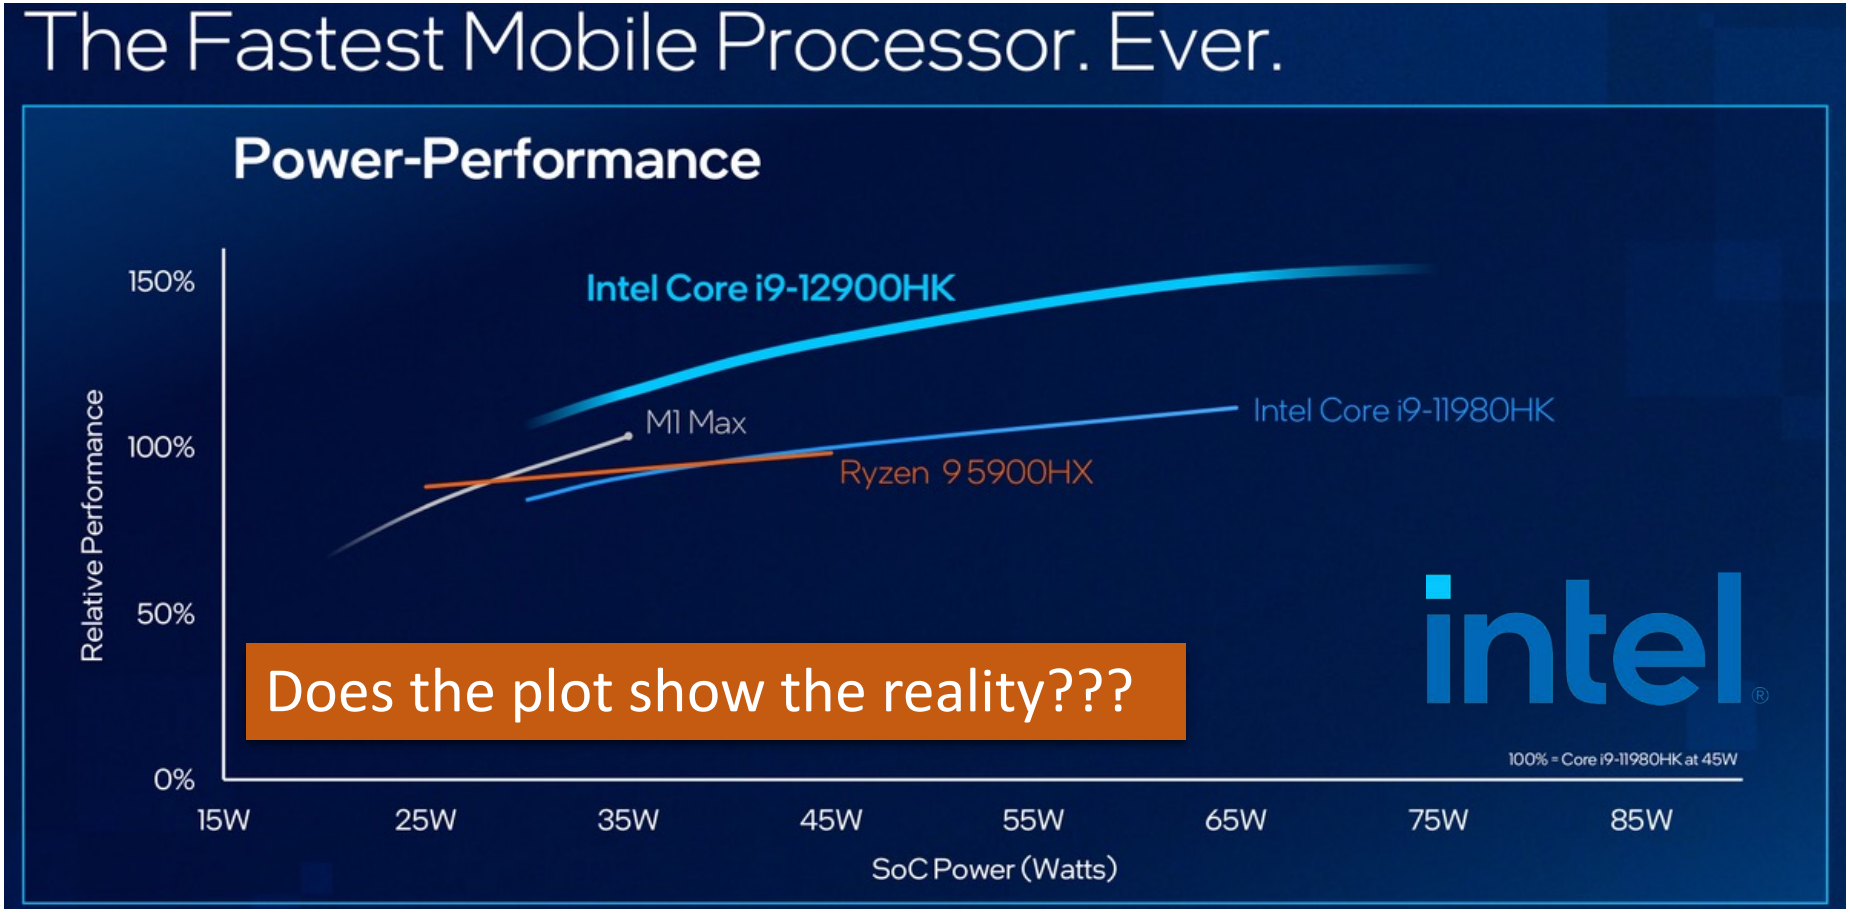

Power Performance: AMD vs. Intel

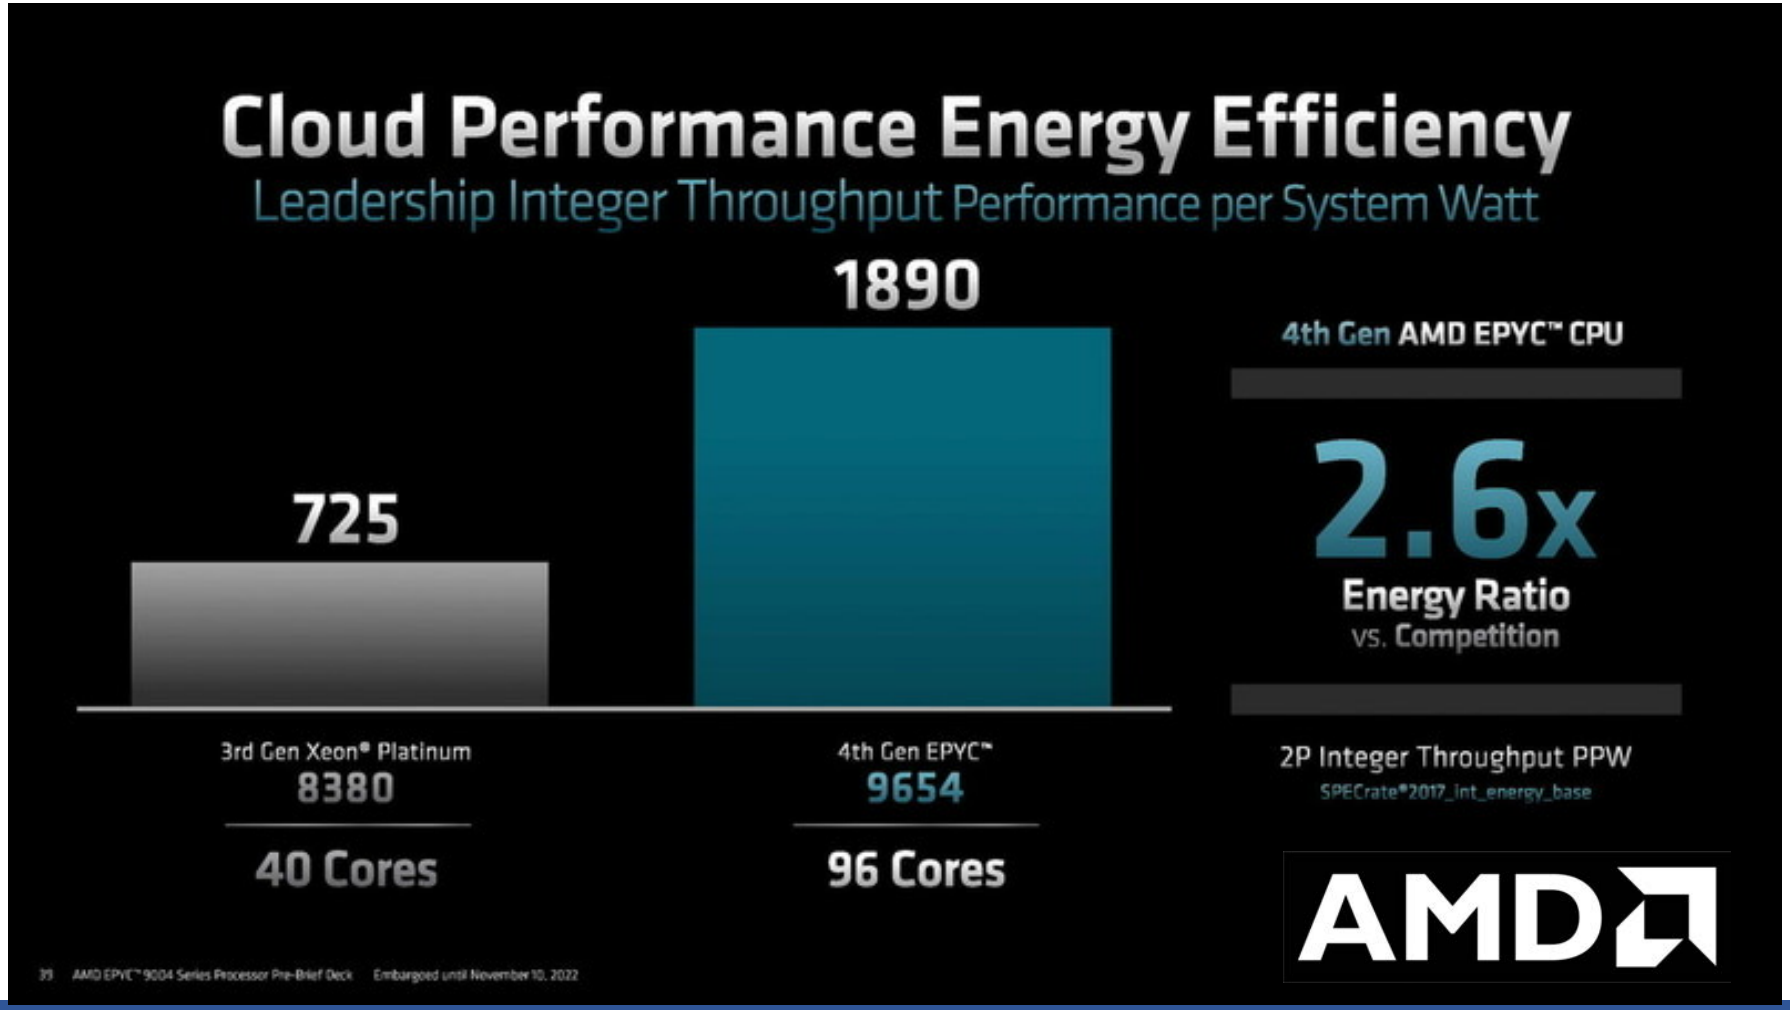

Thoughts on Power Performance

- CPU power efficiency is becoming one of the important metrics.
- Plots of “performance” vs. ”power consumption” are always biased by each vendor.
  - The actual power performance depends on applications.
  - We have to be very careful to understand these plots.
- Apple silicon shows excellent power performance in Desktop/Laptop.
- Intel is seeking peak performance; it’s never-ending…,
  - But they cannot neglect power performance.
  - Intel seems to be behind AMD (EPYC).
- AMD EPYC processors are increasing the share of HPC systems.

Situation of ARM CPU

- ARM Architecture
  - Designed to be power efficient for devices like embedded systems, smartphones, etc.
  - Apple Silicon, A64FX, …
  - But, each CPU seems to have a different direction.

- Fujitsu A64FX (48cores)
  - ARMv8.2-A + SVE (Scalable Vector Extention)
  - In reality, it is designed and optimized for vectorized applications
  - Not suitable for generic applications: much lower performance

- New emerging ARM for HPC
  - Ampere Alta (ARMv8.2+) : 128 cores max
  - NVIDIA Grace Hopper (Neoverse, Armv9.0-A): 72 cores max

Benchmarks

Benchmark Applications

- Using Geant4 ver. 10.7.4
  - G4 v11 series cannot be built in the A64FX compiler.
- G4Bench: https://github.com/koichi-murakami/g4bench
  - Electromagnetic shower simulations
  - For the v11 series, g4bench2 is available on GitHub.
- Microdosimetry
  - Simulating physics and chemical interactions using Geant4-DNA
- Evaluations:
  - Computing performance in Geant4 sequential and multi-threading mode
  - Power efficiency

Results of Sequential Mode (1)

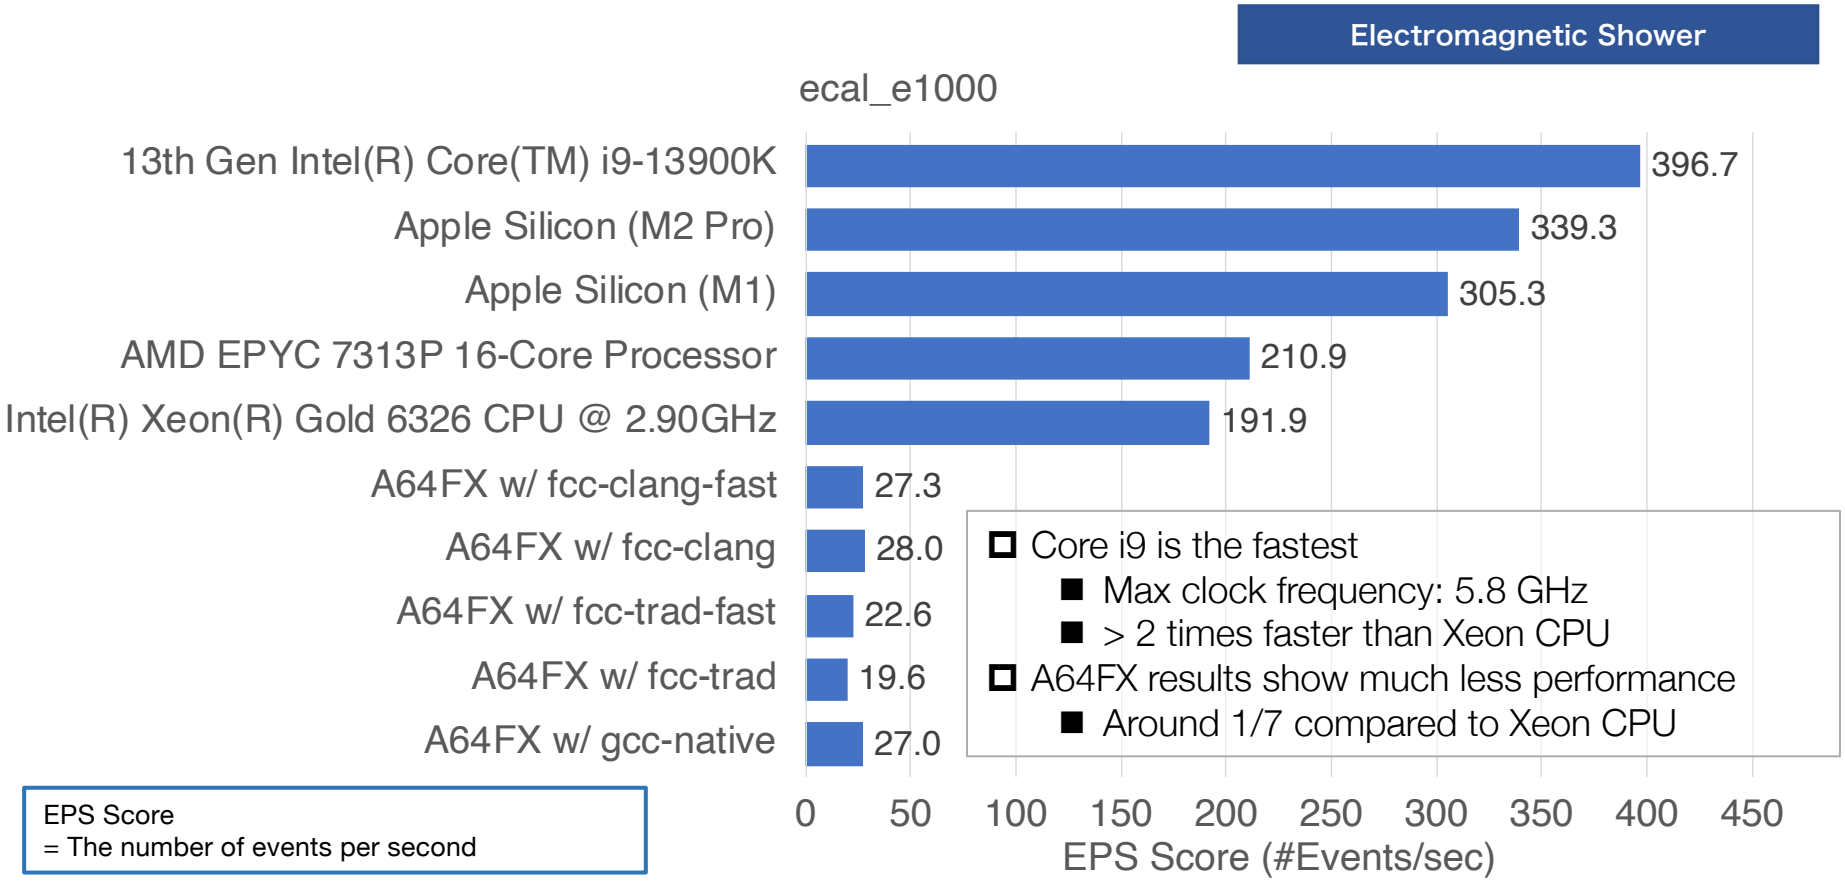

Results of Sequential Mode (2)

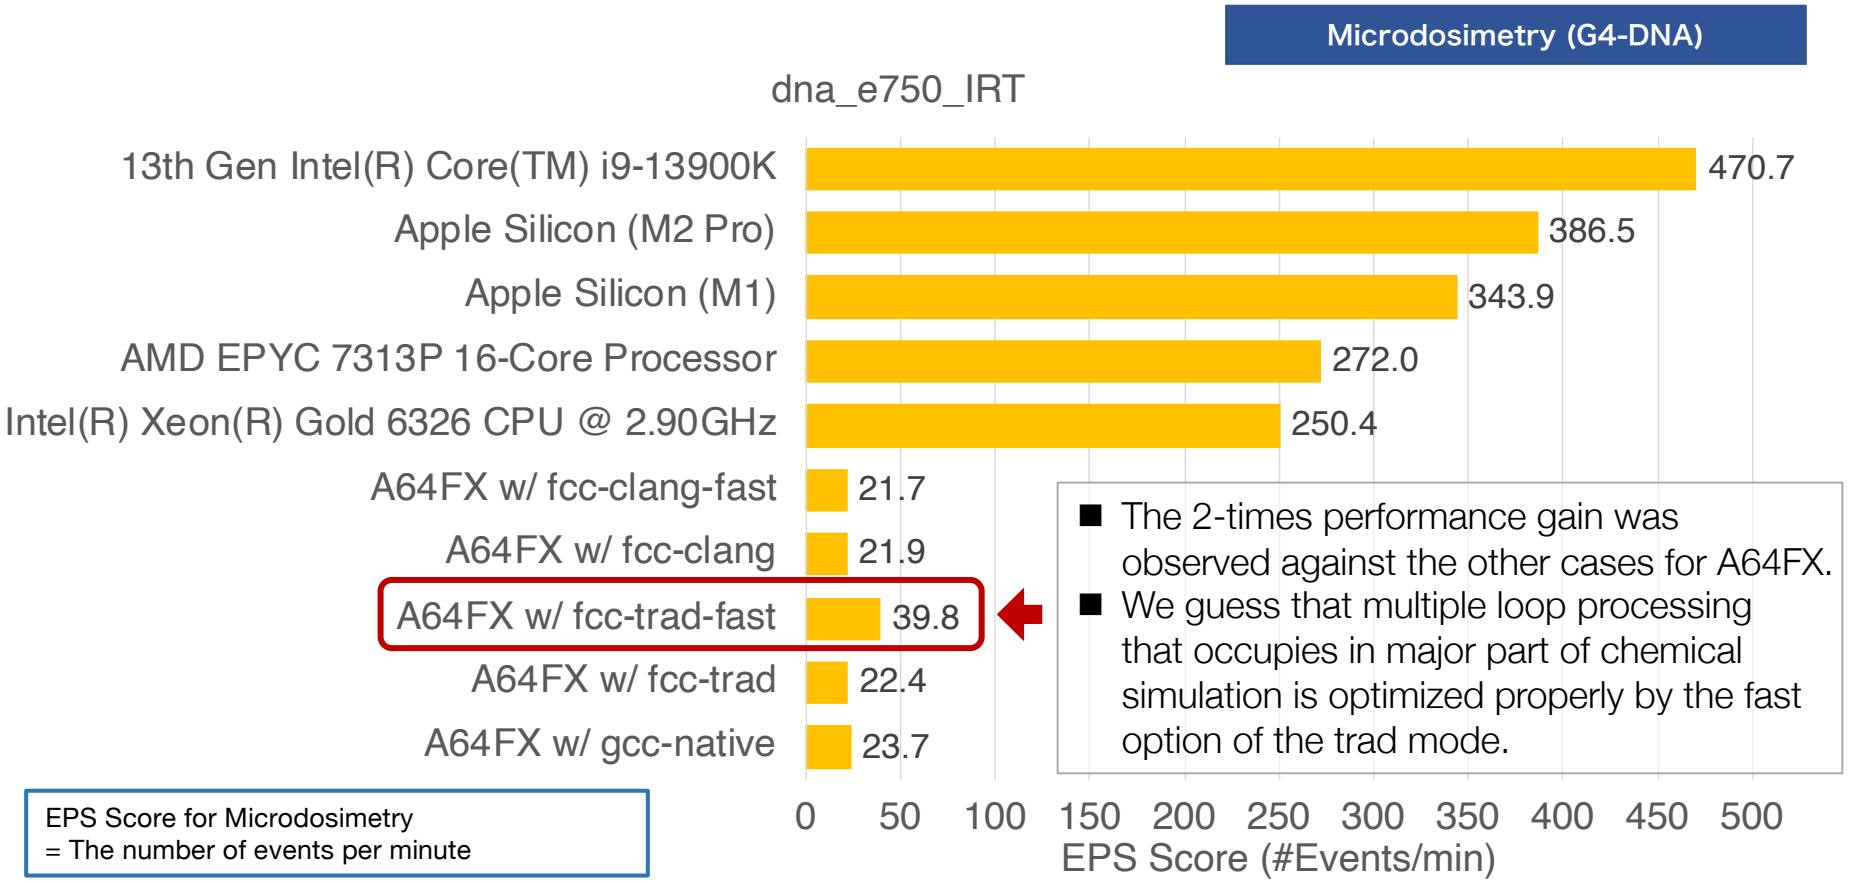

Results of Multi-Threading Mode (1)

Electromagnetic Shower

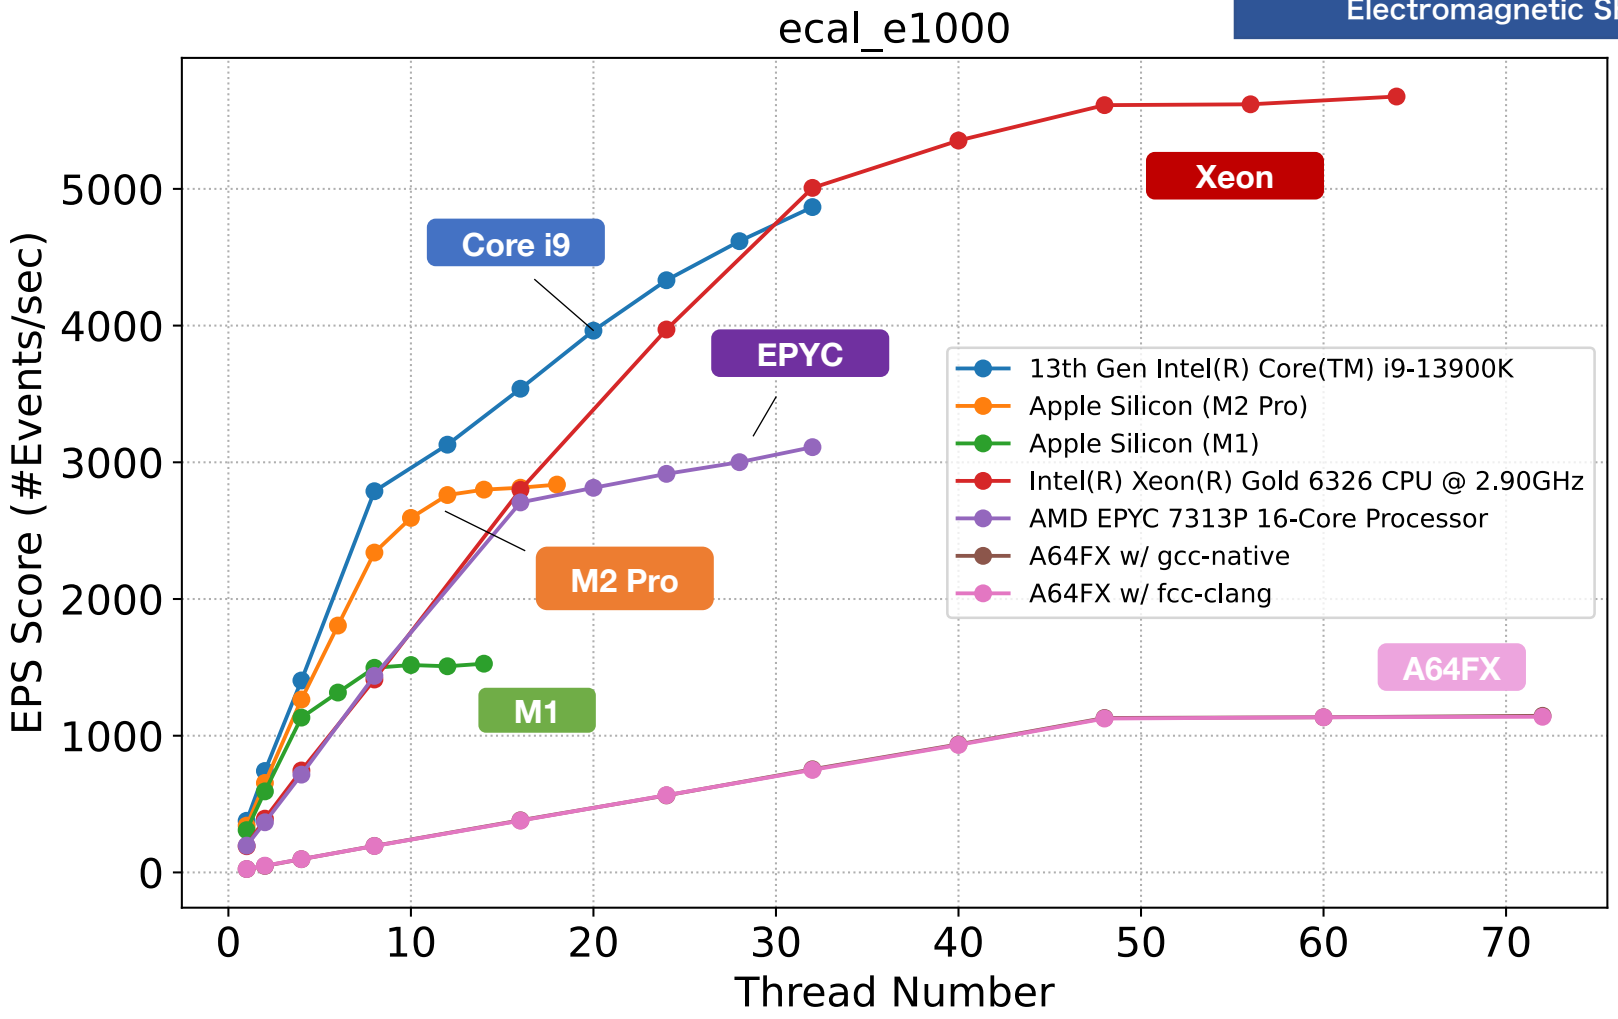

Results of Multi-Threading Mode (2)

Microdosimetry (G4-DNA)

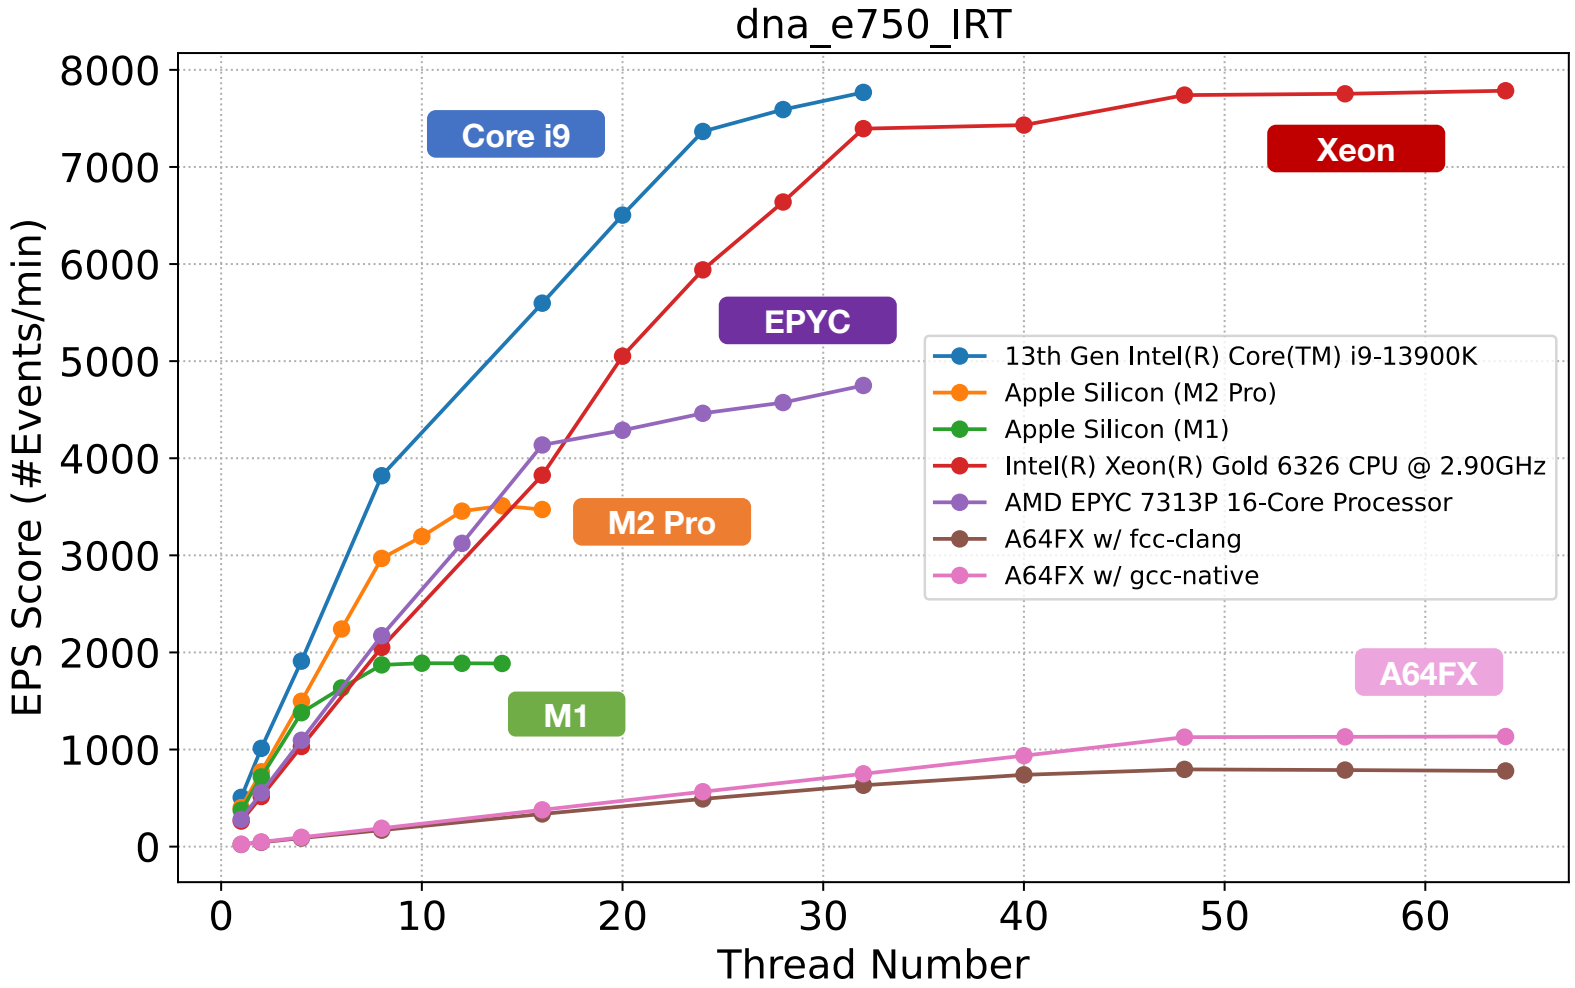

Power Performance

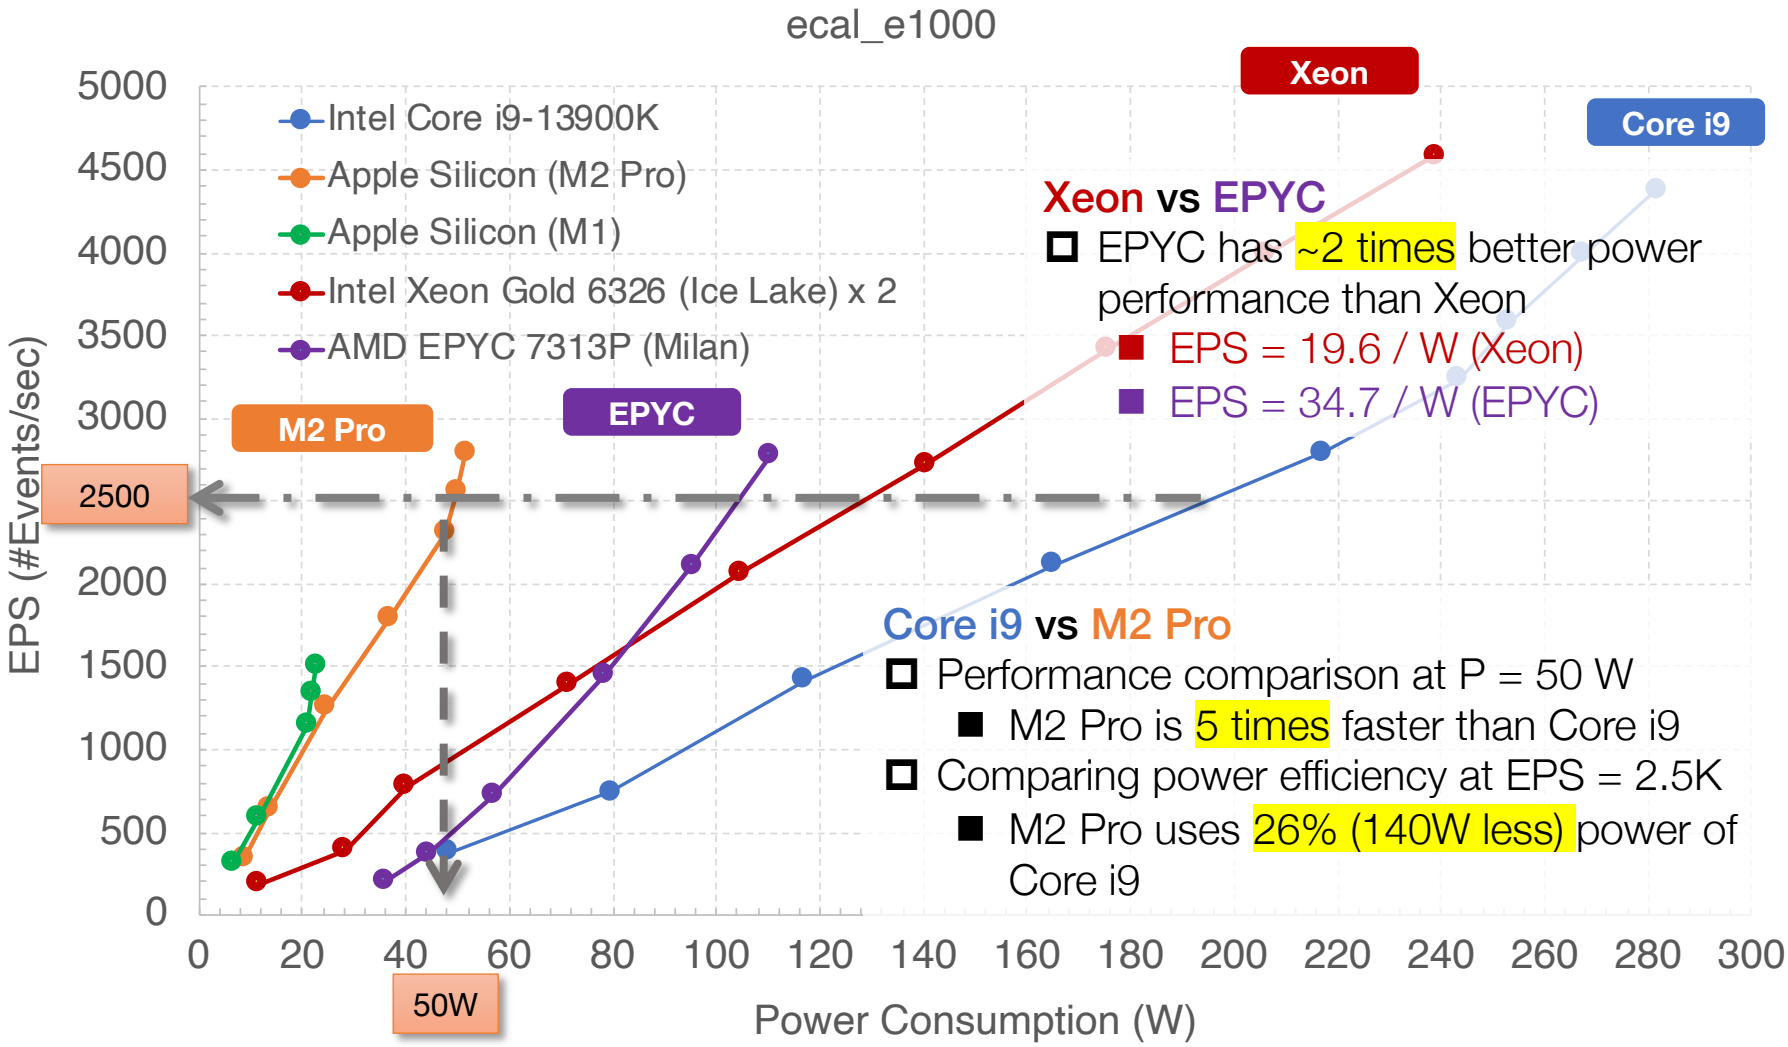

Summary

- The excellent power performance of Apple Silicon
- Unfortunately, A64FX cannot perform fully in Geant4 applications.
- Intel CPU shows the best peak performance.
  - Higher power consumption: It reaches almost 300 W.
  - The power performance of Xeon is double that of Core-i9.
- AMD EPYC shows good power performance from a data center point of view.
  - 2-times better than Intel Xeon: AMD leads ahead of Intel.
- Future plan:
  - AMD 4th Generation EPYC (Genoa, Zen4)
  - Intel 4th Generation Xeon (Sapphire Rapids)
  - New ARM processors

Benchmark Environments

CPU	Intel Core i9-13900K	Intel Xeon Gold 6326 (Ice Lake)	AMD EPYC 7313P (Milan)	Apple M2 Pro	Apple M1	A64FX
Architecture	x86	x86	x86	ARM	ARM	ARM
# of cores	P-core: 8 E-core: 16	32 cores (=16 x 2CPUs)	16 cores	P-core: 8 E-core: 4	P-core: 4 E-core: 4	48 cores
Max Clock Frequency	5.8 GHz	3.5 GHz	3.7 GHz	unknown	unknown	2.2 GHz
Memory	125 GiB	503 GiB	62 GiB	32 GB	16 GB	31 GiB
OS	Ubuntu 22.04 LTS	Alma Linux 9.2	Alma Linux 9.2	mac OS 13.5 Ventura	mac OS 13.5 Ventura	Cent OS 8.1
Compiler	GCC 11.4.0	GCC 11.3.1	GCC 11.3.1	Apple Clang 14.0.3	Apple Clang 14.0.3	GCC 8.5 / Fujistu Compiler