# Performance Tests
## (A64FX, Apple M1, AMD)

*Koichi Murakami (KEK)*

*Virtual Geant4 Collaboration Meeting 2021*

# 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, …
  - Geant4simulation accounts for big percentage of CPU usage in HEP data centers.

- Emerging CPUs other than Intel / x86
  - Intel is in very difficult situation…
    - https://news.in-24.com/business/144257.html
  - AMD EPYC (x86)
  - ARM-based processors
    - Fujitsu A64FX, Apple M1,…

Power consumption study

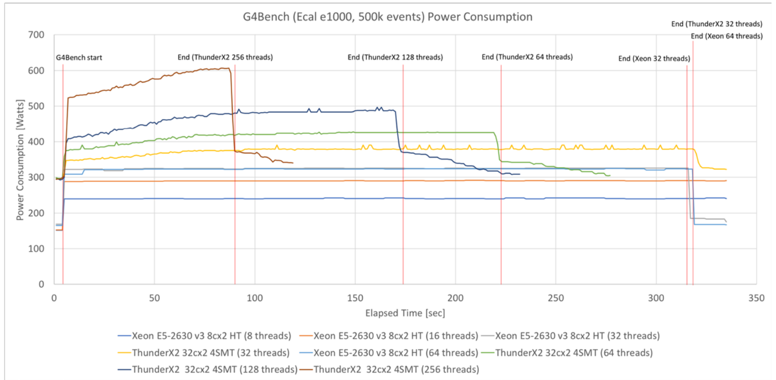

# What should know about ARM

- Arm Ltd. develops the architecture and licenses it to other companies. They utilize Arm's architectural model as a kind of template, building systems that use Arm cores as their central processors.
  - CPU characteristics differs from each other.

- ARM is a family of reduced instruction set computing (RISC) architectures for computer processors.
  - x86 is CISC (Complex Instruction Set Computer)
  - No compatibility between RICS and CISC
  - History of CISC vs. RISC

# Tested Arm Processors

### Marvell ThunderX2 Cavium

- HPE Apollo 70 : Marvell ThunderX2 CN9980 (32cores, 2.0GHz) x2 + 256GB DDR4 / 180W TDP
- ISA : ARM v8.1
- Micro Arch: Vulcan / 16nm
- unfortunately they withdrawed from HPC.

### Fujitsu A64FX

- Fujitsu FX700 : A64FX (48cores, 2.0GHz) x1 + 32GB HBM2 / 150W TDP (guess)
- ISA : ARM v8.2-A + SVE (Scalable Vector Extention)
- Micro Arch : Originally Designed / 7nm

### Apple M1

- Mac mini : Apple M1 / 8GB Unified / 39W (max)
- ISA : ARM v8.4 (No SVE)
- Micro Arch : Firestrorm + Icestorm / 5nm / big.LITTLE-based
- 4 HP cores (Firestorm) + 4 HE cores (Icestorm) / 3.2GHz (big), 2GHz (little)

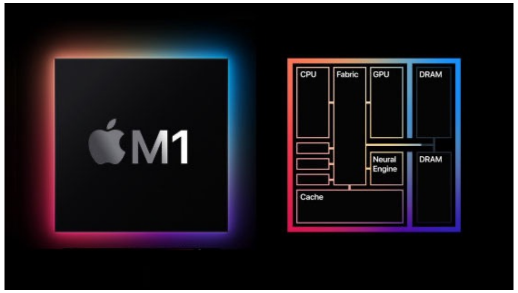

# Intel / AMD Processors

- Intel Xeon
  - Xeon Gold 6240, 18cores (2.6/3.9GHz) x2 + 192GB
  - Code : Cascade Lake, 14nm
  - Future test : Ice Lake Xeon
- Intel Core i9 (reference)
  - i9-9900K, 8cores (3.6/5.0GHz)
  - Code : Coffee Lake (9th Gen), 14nm
  - Note : single thread performance is better than Xeon series.
- AMD / Rome
  - Ryzen Threadripper PRO 3995WX, 64cores (2.7/4.2GHz) x1 + 256 GB
  - Micro Arch : Zen2 (Rome), 7nm
- AMD EPYC / Millan
  - EPYC 7313P 16cores (3.0/3.7GHz)x1
  - Mirco Arch : Zen3 (Millan), 7nm
  - EPYC 7643 48cores (2.3/3.6GHz) x2 planned

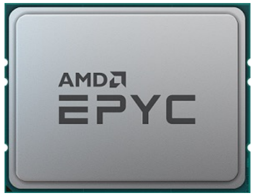

# Benchmark Programs

- G4Bench (Geant4 v.10.7.p1)
  - https://github.com/koichi-murakami/g4bench
  - 3 Tests:
    - ECAL-e1000 : EM shower for 1GeV electron
    - HCAL-p10 : Hadron sandwich calorimeter for 10GeV proton
    - VGEO-x18/e20/p200 : Water phantom voxel for
      - X-ray (18MV Linac) / 20MeV electron / 200MeV proton
  - Observations:
    - EPS : #Events per (milli)second
    - SPS : #Steps per second : same trend with EPS
    - Edep : total energy deposit for physics regression
  - Continuous testing using GitHub Actions and Docker
    - https://hub.docker.com/repository/docker/koichimurakamik6/geant4-runtime
    - https://github.com/koichi-murakami/geant4_runtime

- Chem-Bench
  - GitHub private repository (limited access)
  - G4DNA benchmark : G-value w/ IRT (physics/chemical processes)
  - Observations: EPS (#Histories per min)

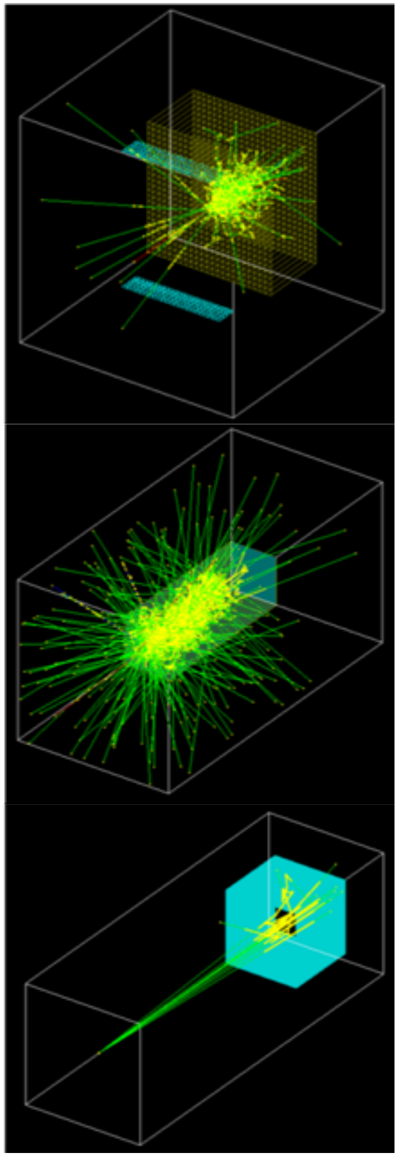

# Notes on A64FX Compiler Study

For A64FX machine, we tested with 2 compilers:

- GCC 8.3.1
- Fujitsu compiler (v4.2.0) :
  - Default options  and w/ Optimization flags

- Obseved slight improvements (gcc < FC < FC-opt, 1%$\Delta$)
- No significant differences were observed between these settings.

# Single Thread Peformance (EM)

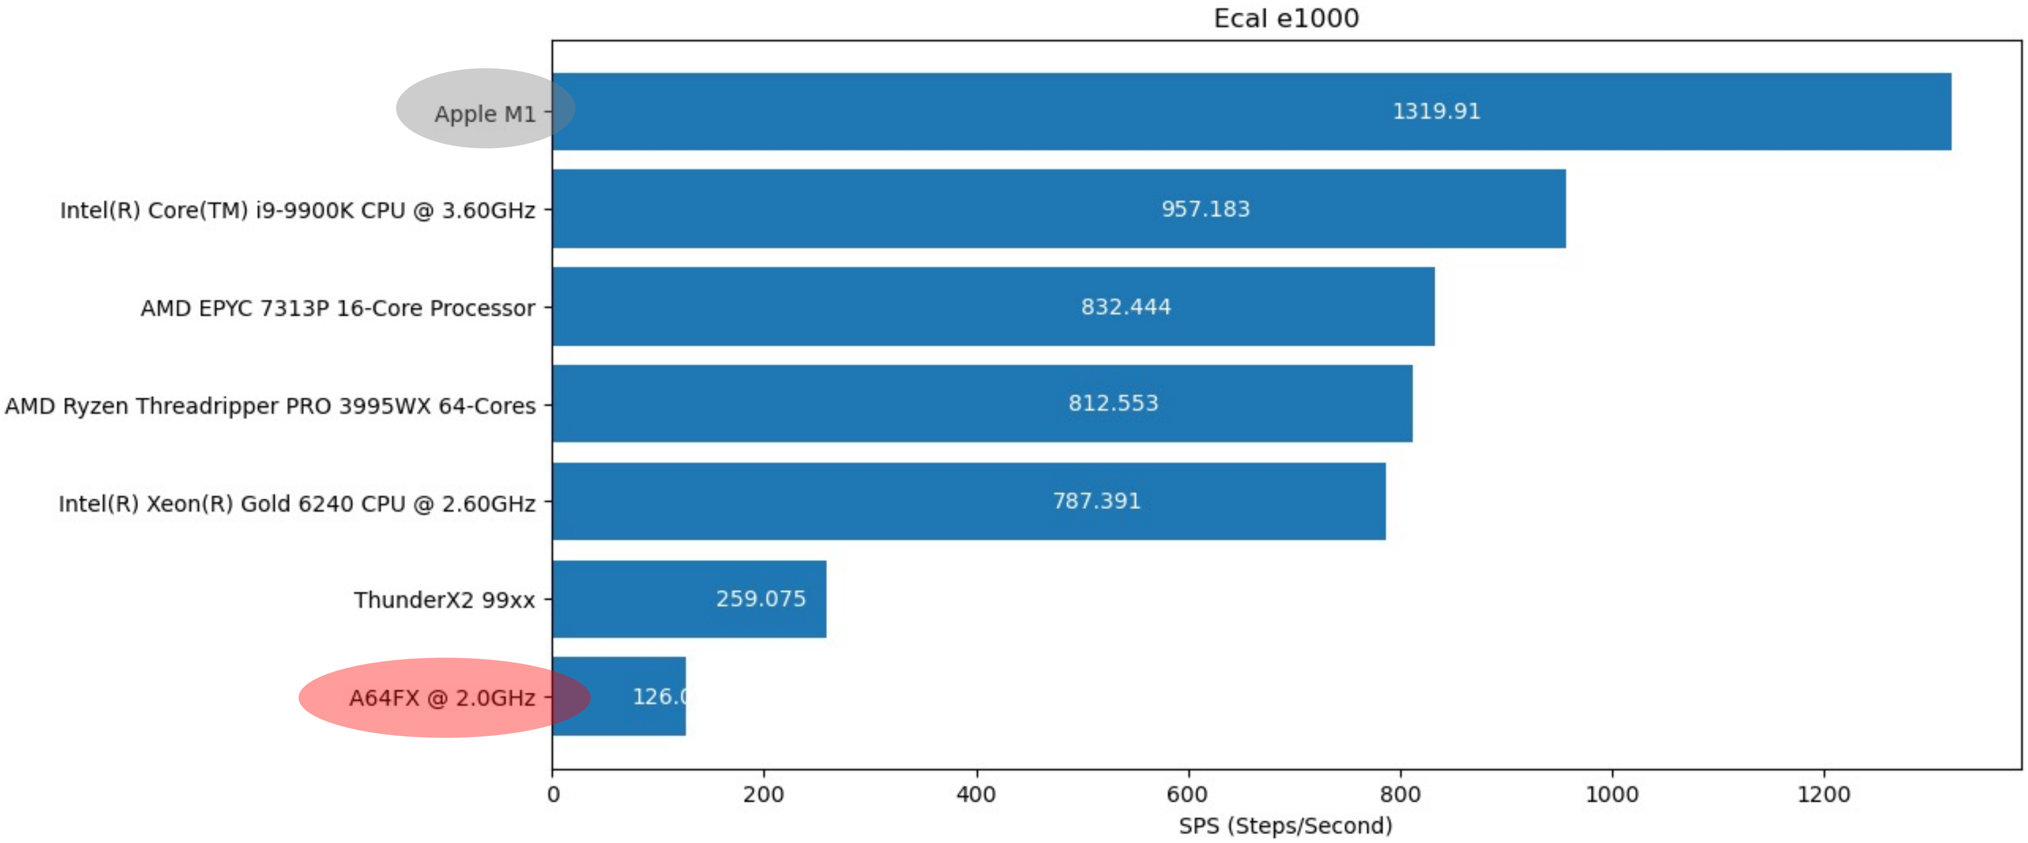

Edep : 967.048 (x86) / 967.125 (M1) / 967.374 (A64FX)

# Single Thread Peformance (Hadron)

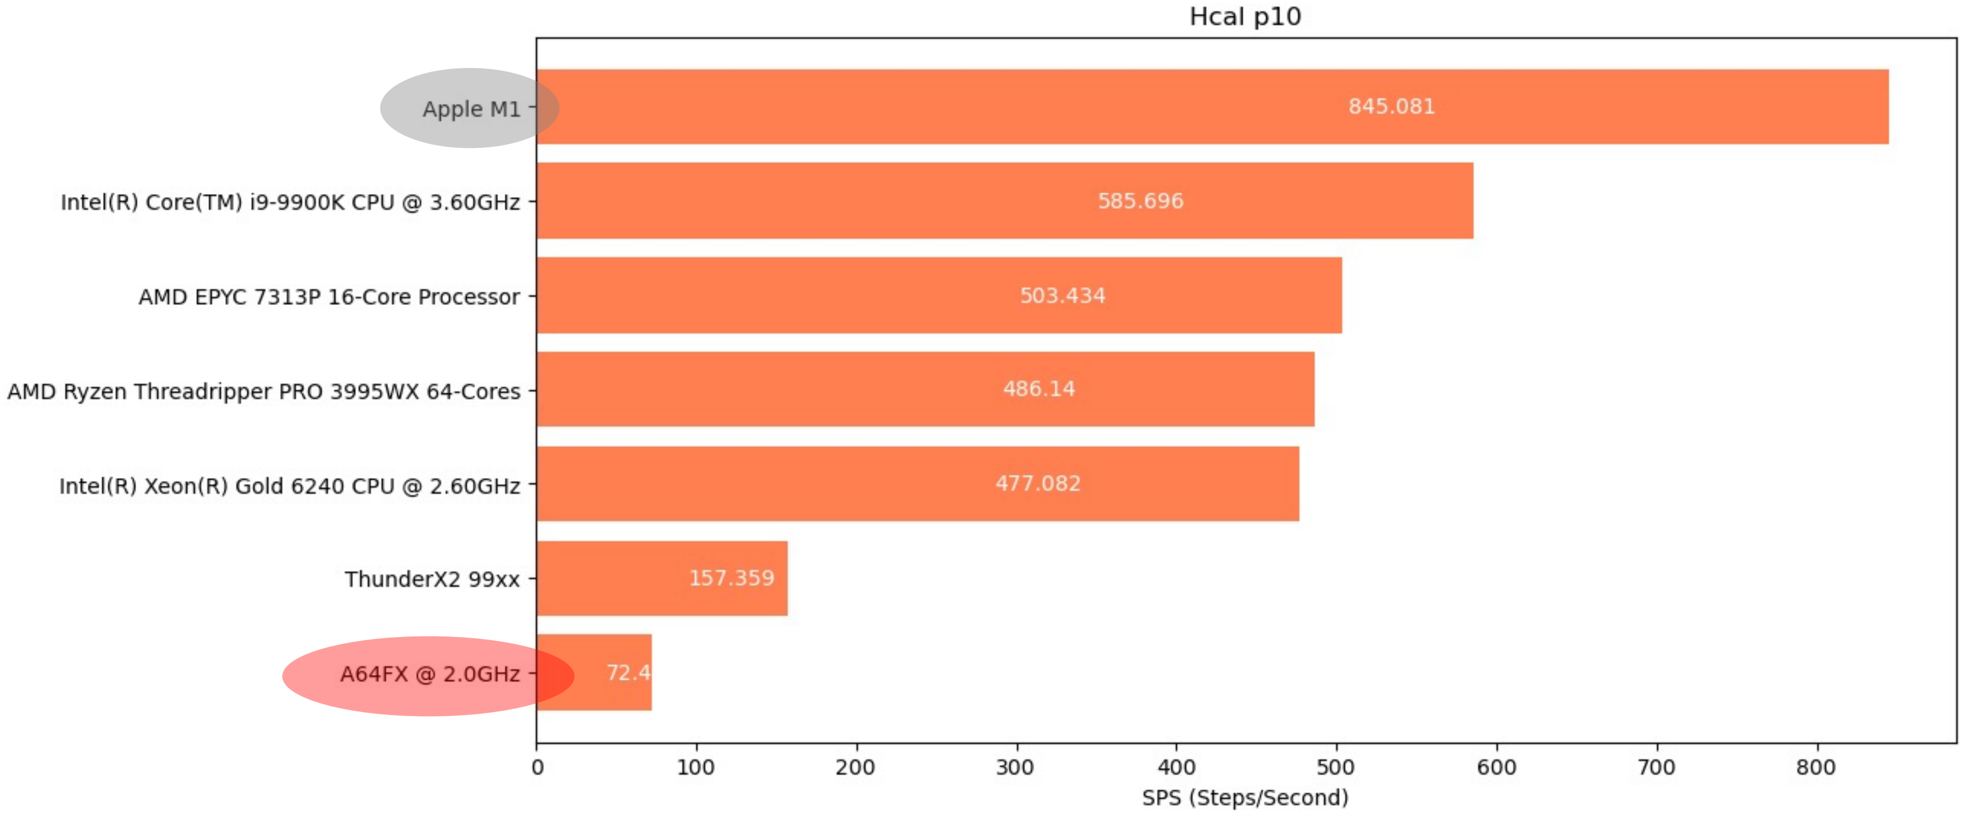

- No siginificant difference is obserbed from EM.

Edep : 844.723 (x86) / 844.986 (M1) / 845.345 (A64FX)

# Multi-Thread Peformance (EM)

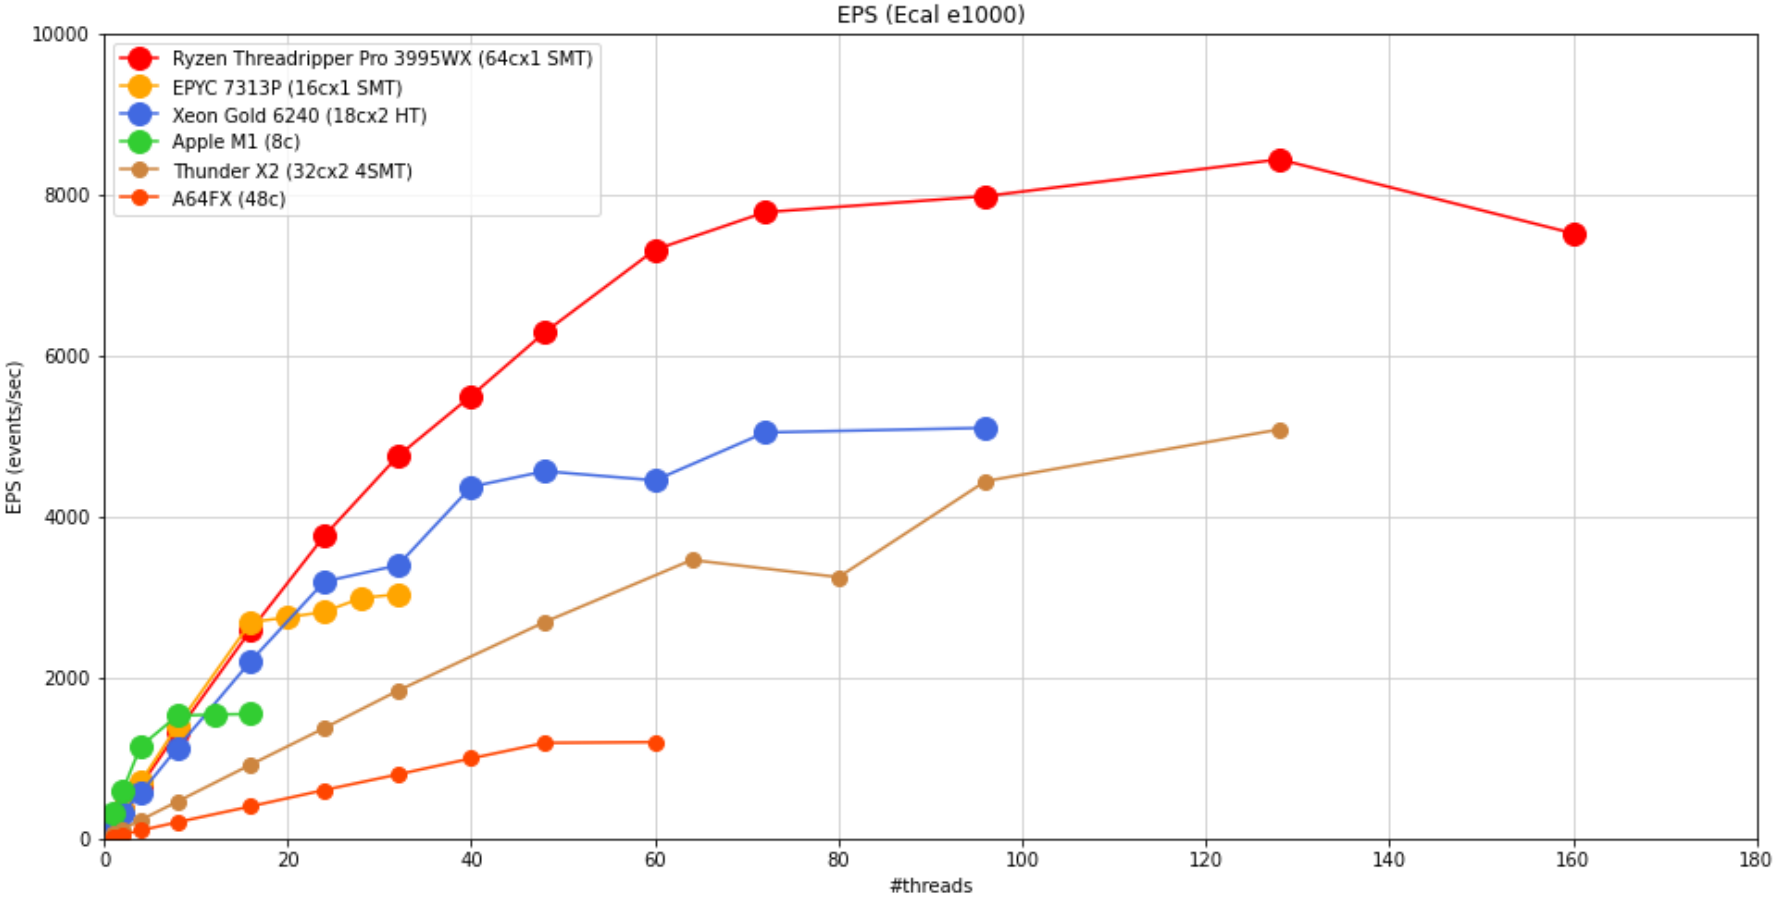

- *Apple M1 has 2 types cores (4HP/4HE)*
- *ARM shows good linearty. (RISC characteristics)*

# Chem-Bench results of Apple M1

- Single-thread EPS (histories/min) :

| | |
|---|---|
| Apple M1 | 366 |
| Intel Xeon Gold 6132 | 150 |

- MT :
  - same tendency with G4Bench
  - slight vCPU overhead for M1 (Parallles + Ubuntu)

- AMD EPYC Millan
  - x1.5 faster than Xeon Gold (Sky lake)
  - Ubuntu Docker is 30% faster??? (under investigation)
    - Ubuntu 20.04 is faster than RH8.4 in bare metal
    - only for Chem-Bench. no observation in G4Bench

*by Shogo*

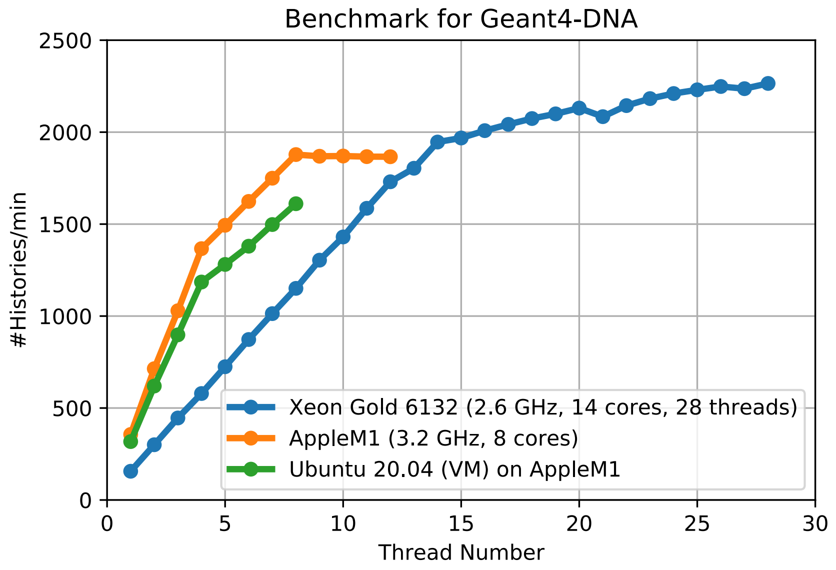

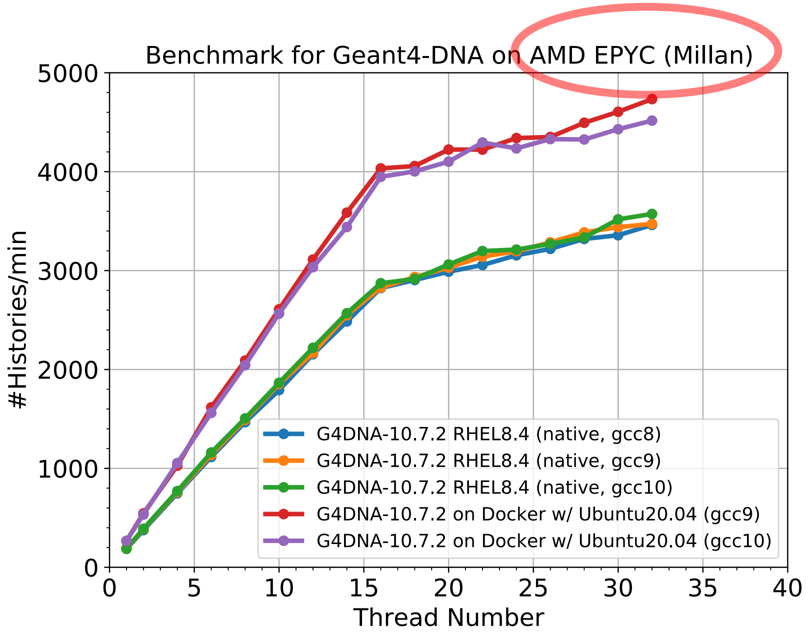

# Chem-Bench results (2)

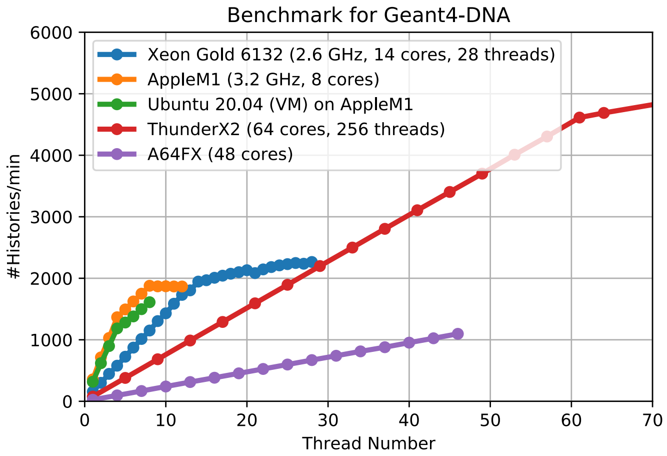

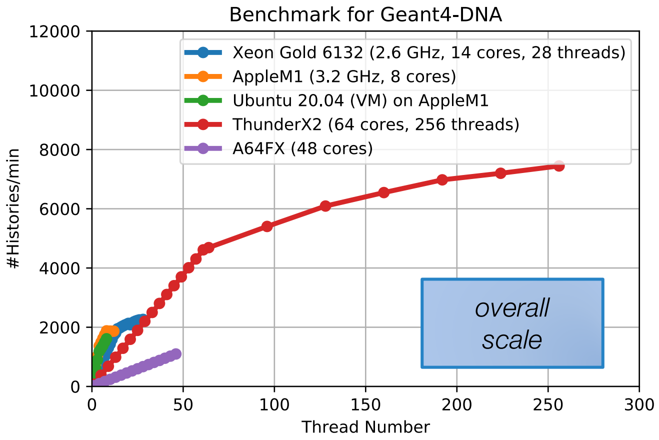

*overall scale*

Same trend with G4Bench

- Apple M1 fastest
- A64FX worst

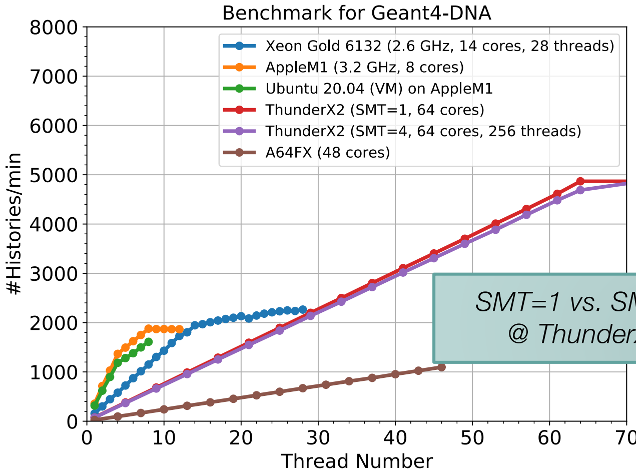

*SMT=1 vs. SMT=4 @ ThunderX2*

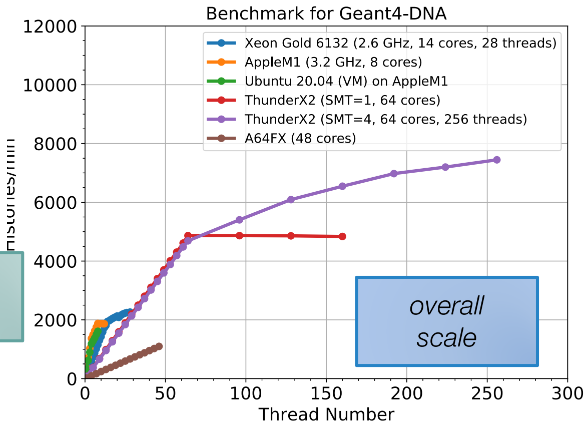

*overall scale*

SMT effect

- SMT : Simulteneous Mutithreading
- SMT=1 is slightly better upto 64 threads.
- leveling off at 64threads for SMT=1

# MT issues (Intel)

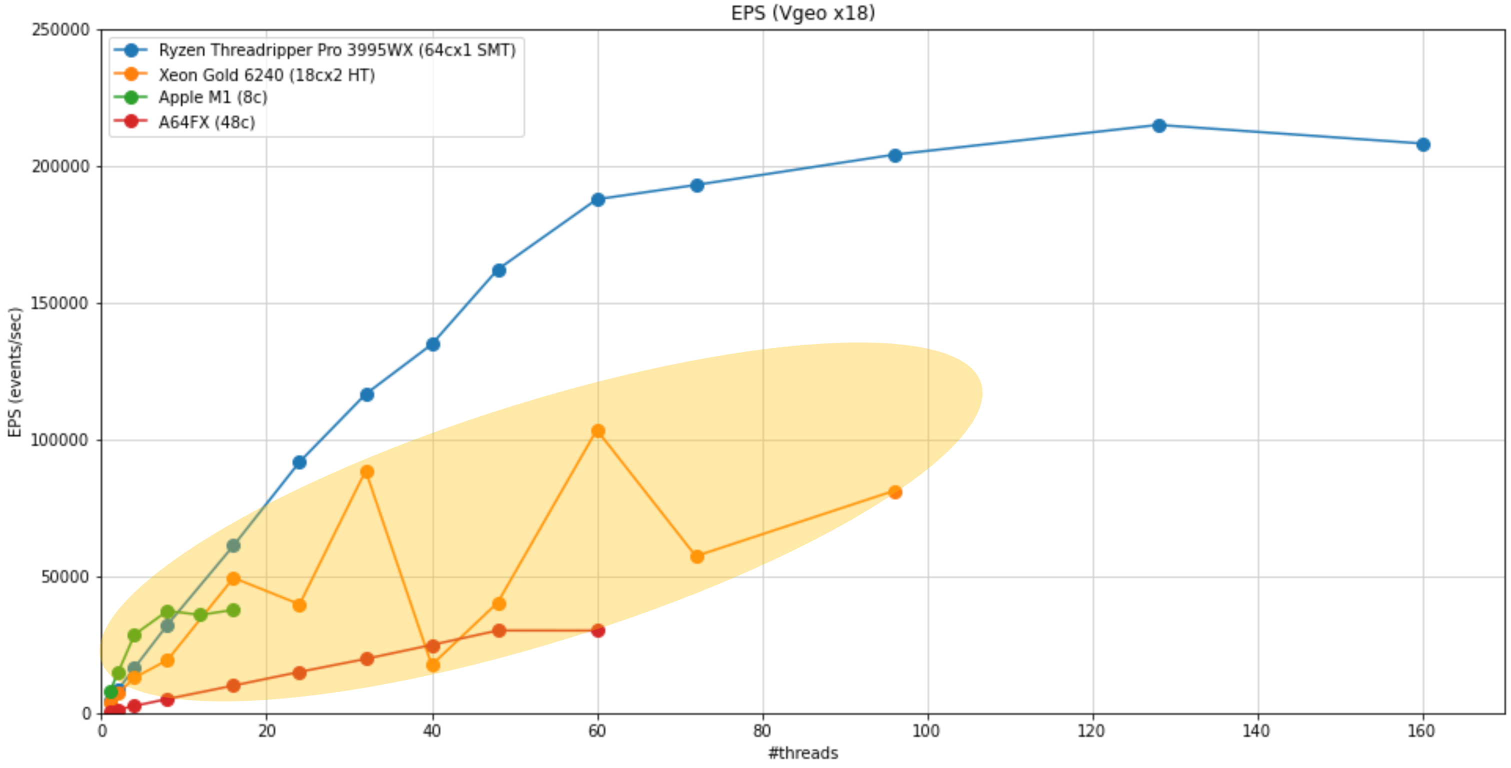

- In case of short event cycle, MT performance is quite unstable for Intel processors. (~5 events/msec)
- Other processors are ok. (AMD, ARM, Power)

# MT Performance Stability Study

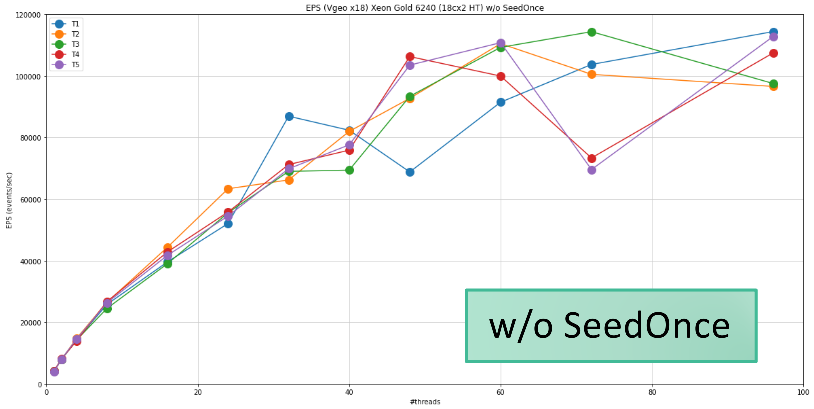

w/o SeedOnce

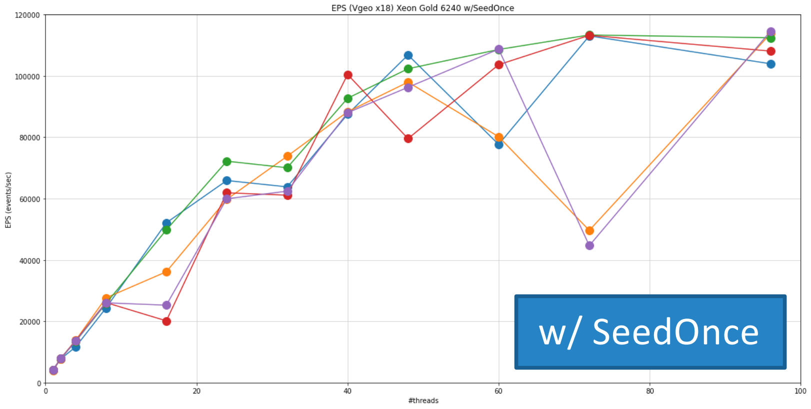

w/ SeedOnce

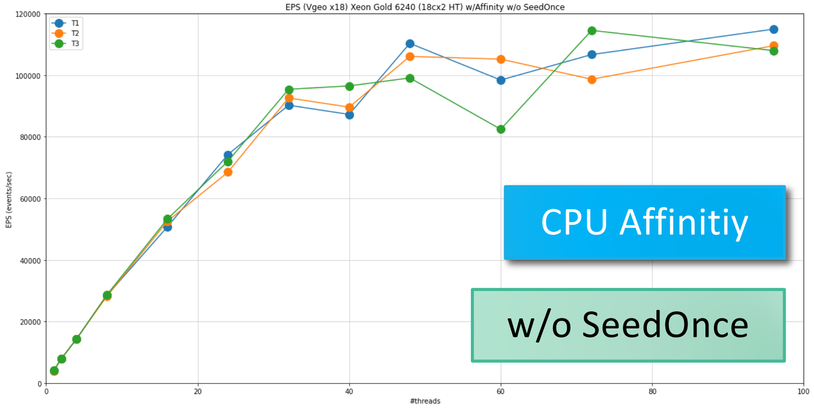

CPU Affinitiy

w/o SeedOnce

- 5-sequential measurements
- much better results than previous??
- Some confusing : needs more investigation
  - kernel update might improve?

36 physical cores

- SeedOnce = 1 is 30% faster for lower #threads<20. Still unstable for higher #threads.
- No significant difference in ECAL (w/ vs. w/o SeedOnce)
- CPU affinity setting can help.
- Static library linking do not affect much.

# MT Overhead (MT/Seq for single thread)

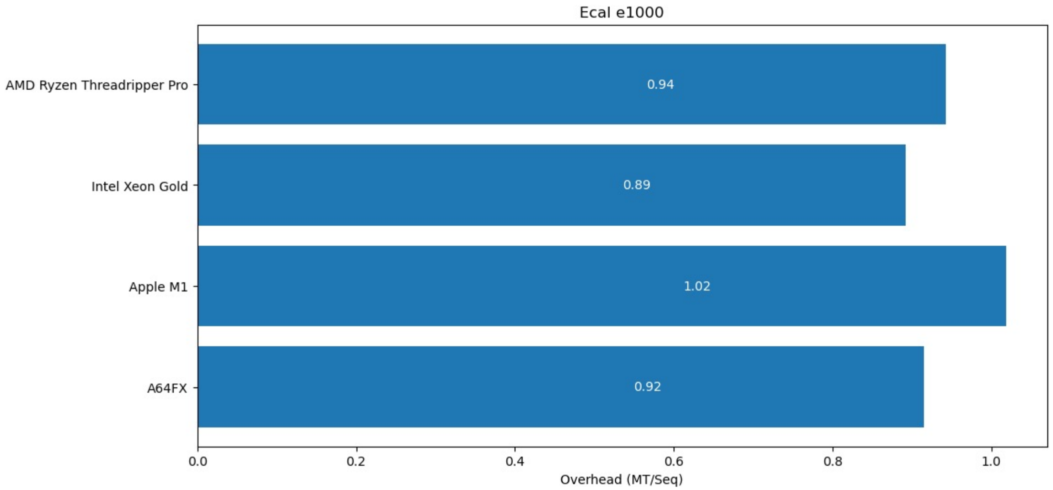

- 5-10% overhead is observed except Apple M1 for Ecal

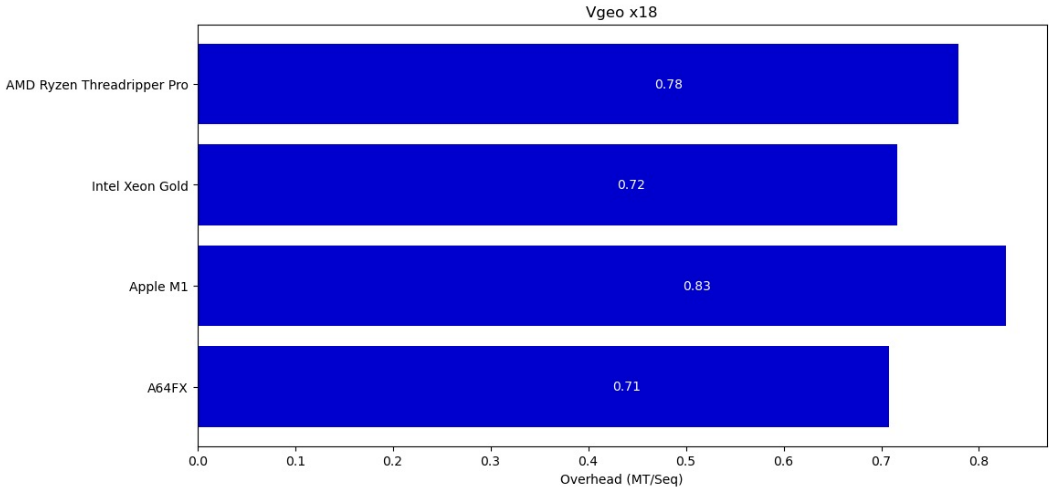

- 20-30% overhead for short-cycle events (Vgeo)
- Apple M1 shows better numbers.

# Shared vs. Static Library Linking

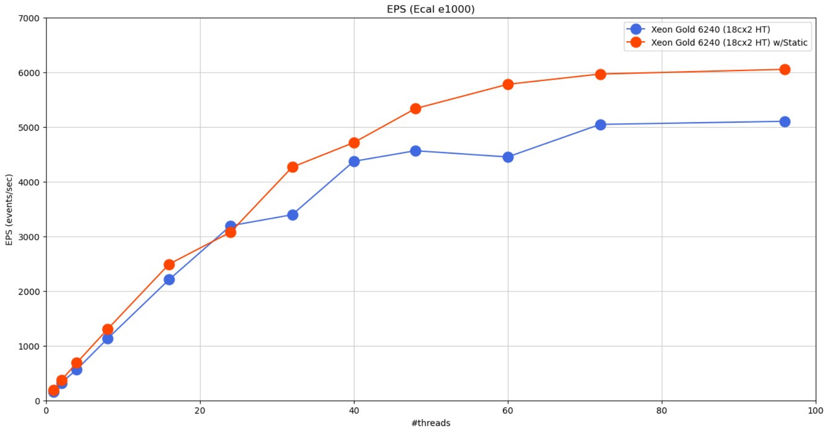

- Static library :10-20% better performance in general

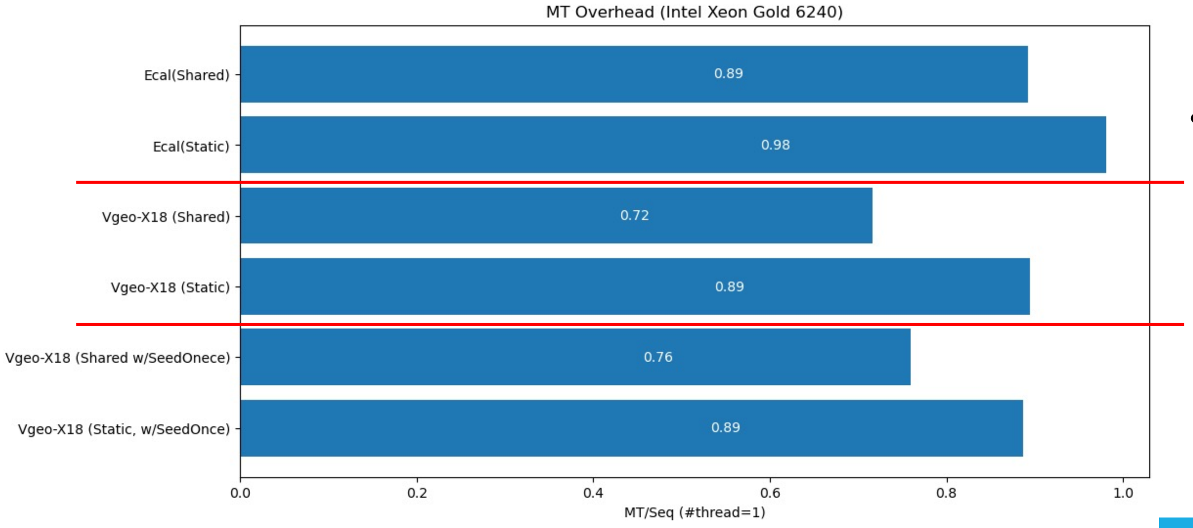

- MT event dispatch overhead is improved by using static library

# Thoughts

- Apple M1 shows good performance especially on power efficiency.
  - RISC architecture w/ more decoders enables high performance.
    - https://bit.ly/3yz5040
  - But, Apple M1 is consumer processor, not for HPC.
- A64FX performance is quite poor w/o SVE.
  - The concept of chip design is based on SVE usage.
  - Complier optimization does not help automatic SVE usage.
  - SVE offloading -> Amdarhl's law
- Neoverse seems to rely on SVE from revealed information. (NVIDIA future plan for HPC)
- ARM in HPC in the future?

# Thoughts (2)

- Using Geant4 static library linking improves perforance by 10-20%.
  - Overhead of event (thread) dispatch : 5-30% -> static linking can improve well.

- Instability of Intel processors (Xeon) for short-event cycle (5events/msec) simulation
  - SeedOnece=1option shows better performance, but does not improve.
  - CPU affinitiy setting can help some.
  - Need more investigation
  - Latest generation CPU improves the situation? : should check for Cascade lake Xeon (Planned)
  - Meltdown/Spectre patches might affect???
  - Spin-wait loops on HT-enabled system??" -> should check for HT-disabled
    - https://software.intel.com/content/www/us/en/develop/articles/long-duration-spin-wait-loops-on-hyper-threading-technology-enabled-intel-processors.html

# Backup

# Notes on Static Library Linking

- Shared vs. Static libraries
  - Static linkage gets slight better performance.
    - for EPS and MT overhead

- TLS mode (MT)
  - “initial-exec” (default) is used both for share and static cases.
  - Notes:
    - “initial-exec” results in run-time seg-fault for full static link (use *–static, static linking with system libraries*)
    - “global-dynamic” is not necessary if static linking only with G4 libs.
      - not using `-static` option and link directly with libG4xxxx.a