Evaluation of **ARM CPUs** for IceCube available through GKE (Google Kubernetes Engine)

Igor Sfiligoi, Frank Würthwein – UCSD
David Schultz, Benedikt Riedel - University of Wisconsin Madison

UC San Diego

WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

Benchmark setup

- Google offered some credits in 2022 to IceCube to be used on the Google Kubernetes Engine (GKE)
  - A managed Kubernetes setup on the Google Compute Cloud
- Google has several generations of CPUs available
  - x86 from both Intel and AMD
    - in both Hyper-threaded and not mode
  - ARM CPUs from Ampere Computing (no hyper-threading)
- IceCube's main CPU-heavy application is CORSIKA
  - Cosmic air shower simulator
  - x86 binaries available on CVMFS
  - ARM binaries compiled in GKE

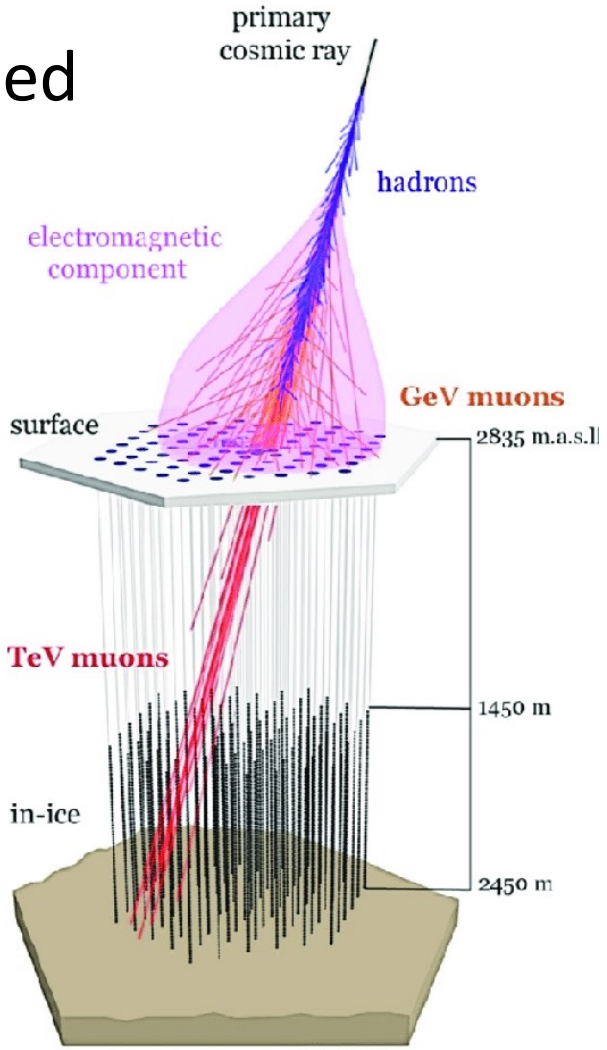

CPUs used

- GKE being a Cloud platform, a lot of details are hidden
- But this is what can be observed:
  - t2a - one vCPU per core
    - Ampere Altra Arm Q64-30 processor with all-core frequency of 3.0 GHz
  - t2d - one vCPU per core
    - AMD EPYC 7B13 - 2.45 base 2.8 effective 3.5 max boost
  - c2d - two vCPU per core (i.e. HT)
    - AMD EPYC 7B13 - 2.45 base 2.8 effective 3.5 max boost
  - c2 - two vCPUs per core
    - Intel® Xeon® Gold 6253CL - 3.1 base 3.8 all core turbo 3.9 single core

Same HW, different setup, different pricing

Benchmark results

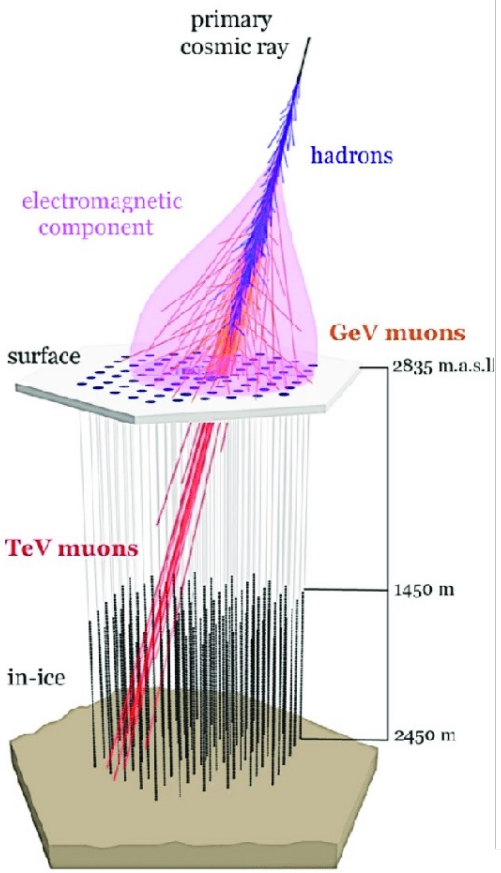

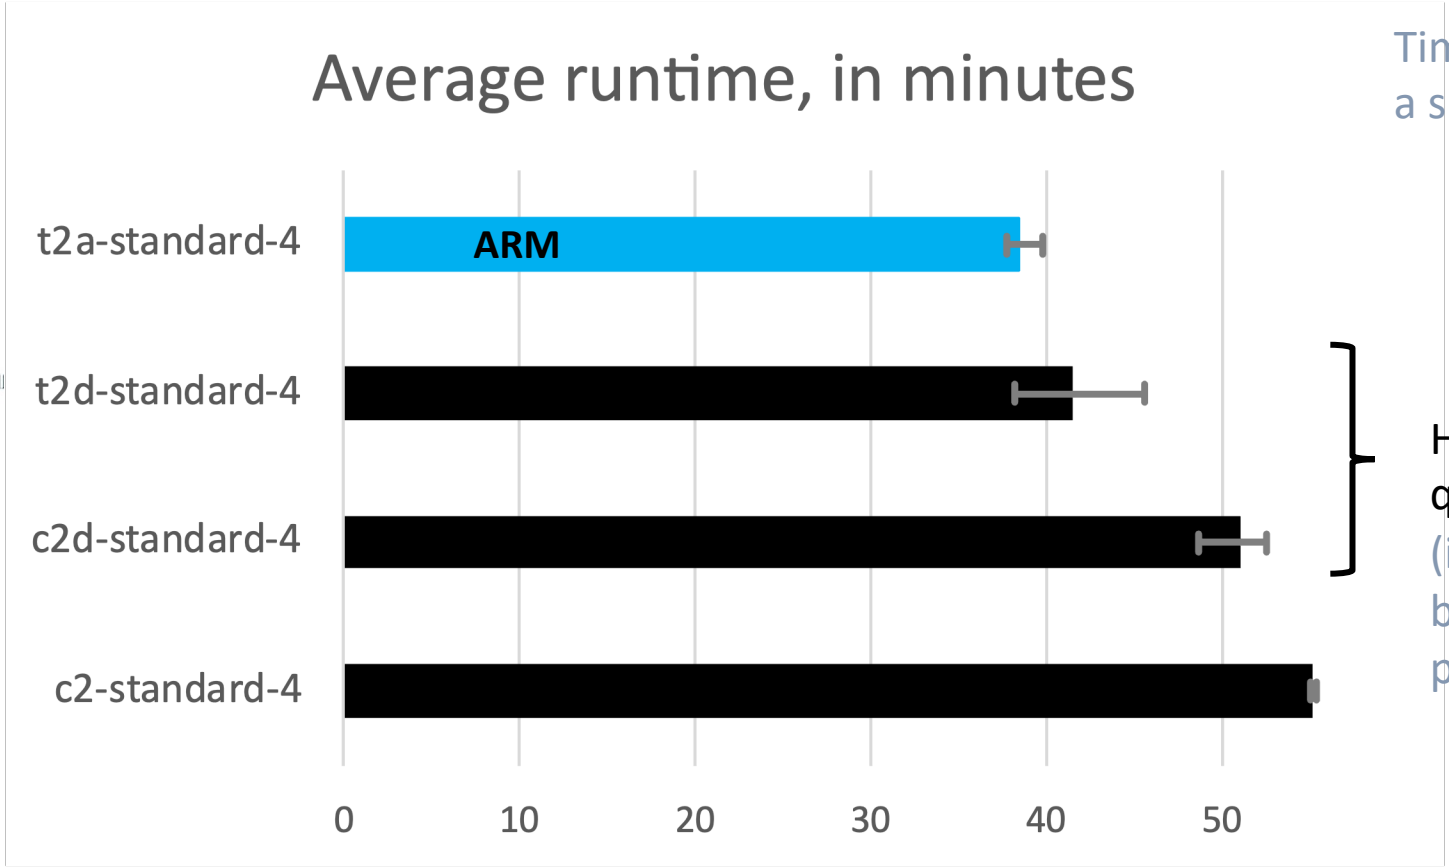

Time to completion using a single vCPU per process

Hyper-threading actually quite effective
(it does take longer but using half the physical cores)

Benchmark results

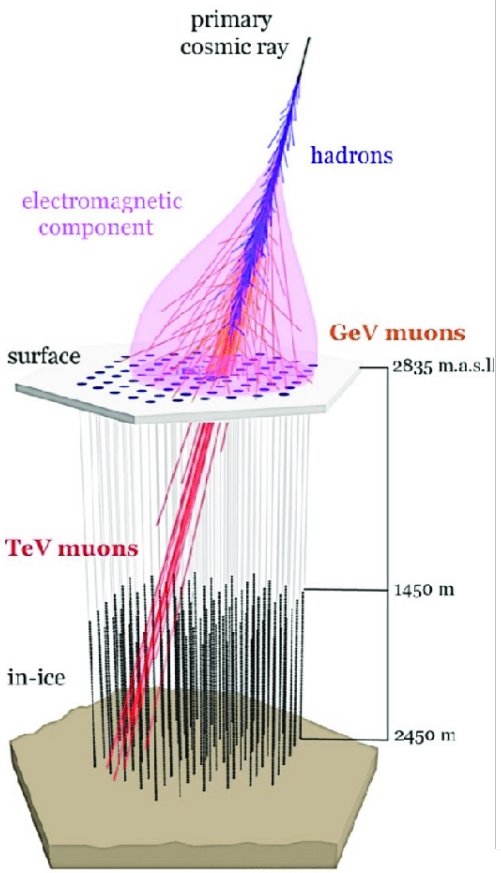

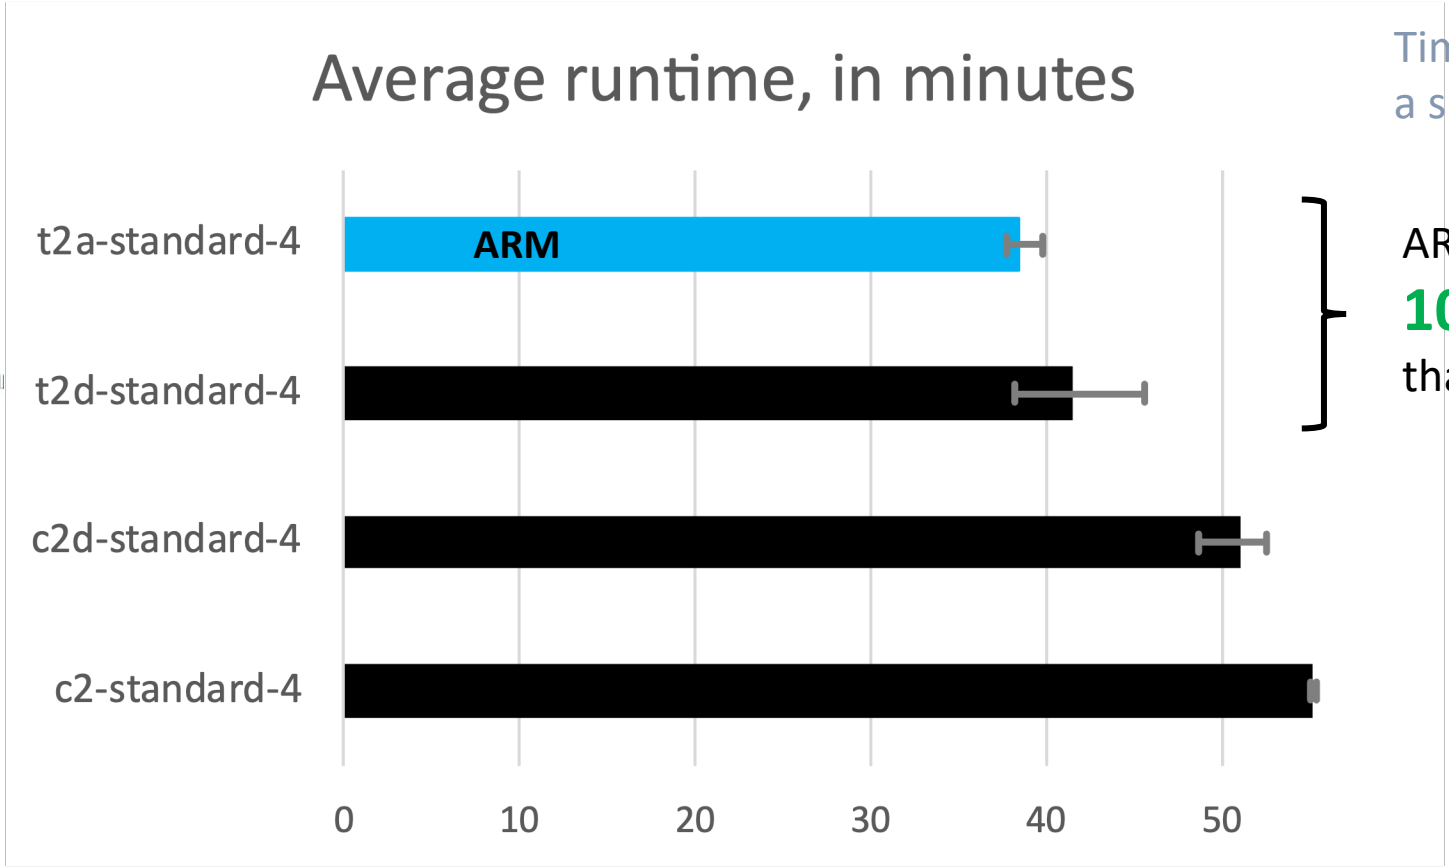

Time to completion using a single vCPU per process

ARM about **10% faster** that the closest x86

ARM CPUs faster than all tested x86 CPUs (at least in early 2022)

Let’s consider cost in GKE

Machine Type	CPU	vCPUs	Physical cores	Cost ($/h)
t2a-standard-4	**ARM**	**4**	**4**	**0.154**
t2d-standard-4	x86-64	4	4	0.169
c2d-standard-4	x86-64	4	2	0.1816
c2-standard-4	x86-64	4	2	0.2088

(Was accurate in early 2022)

Cost effectiveness

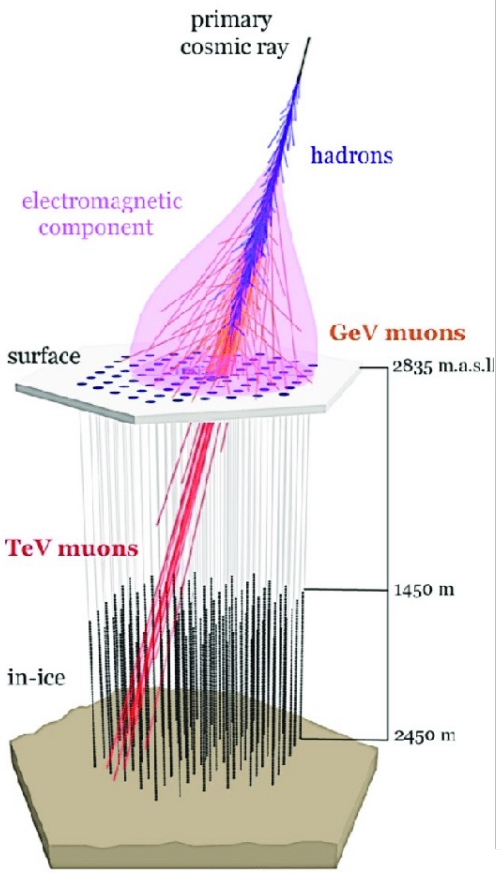

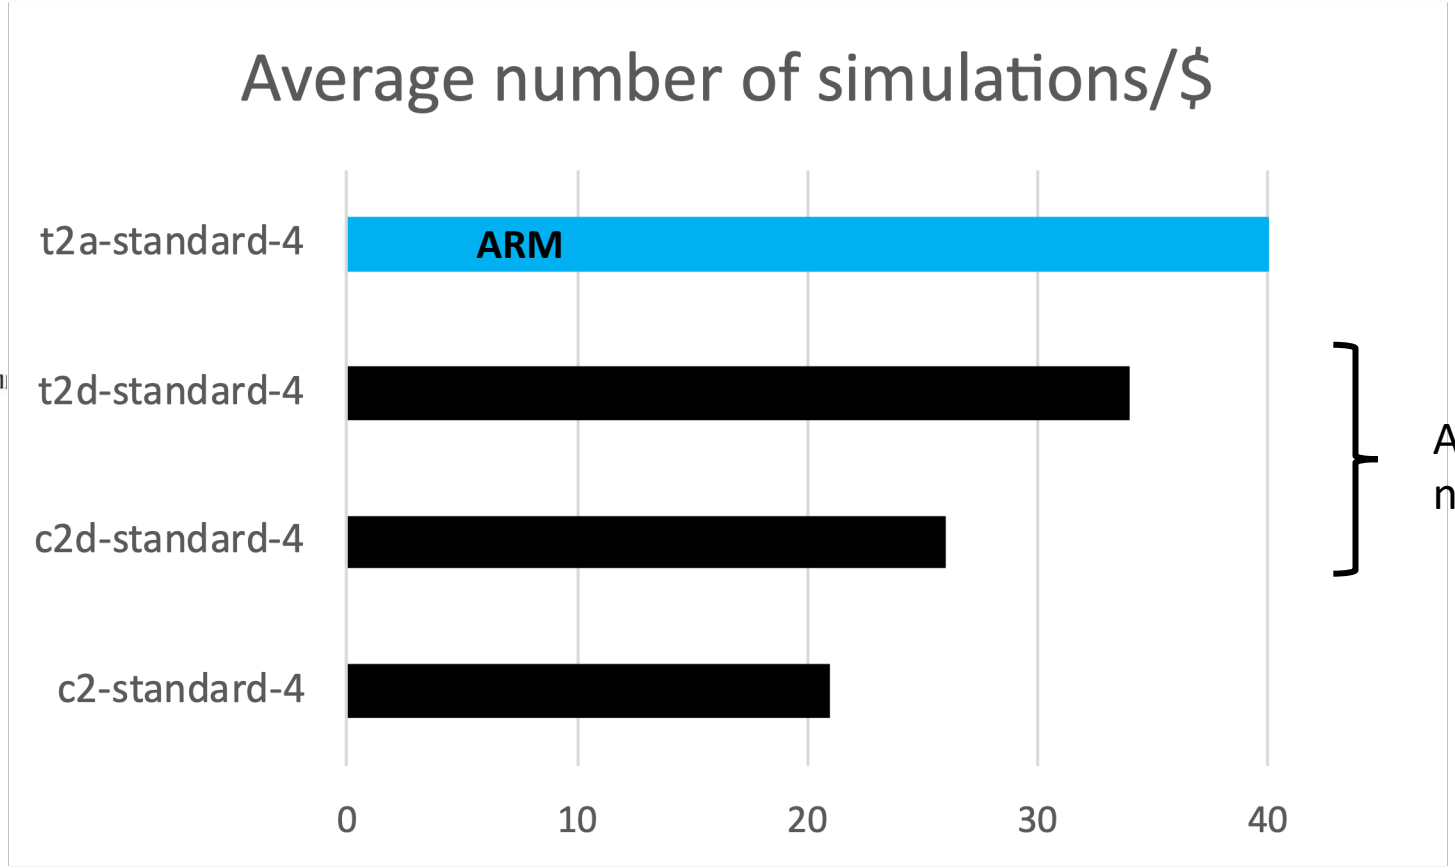

AMD HT pricing definitely not in line with performance

(at least in early 2022)

Cost effectiveness

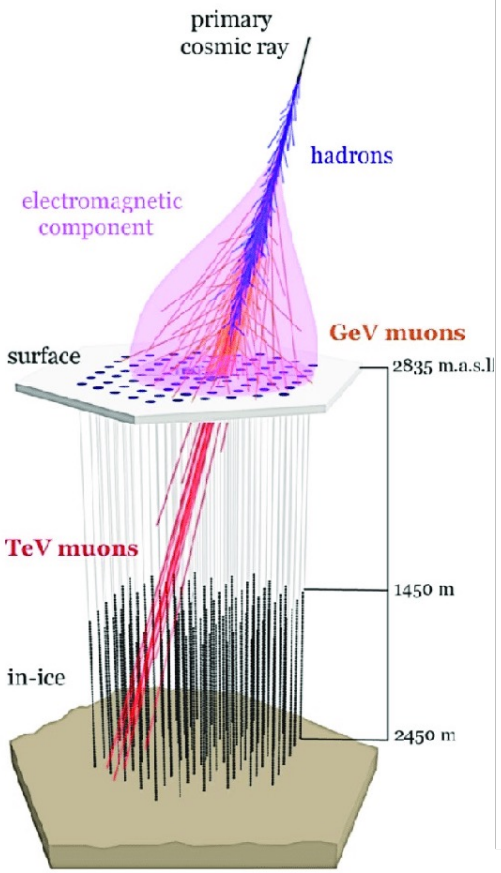

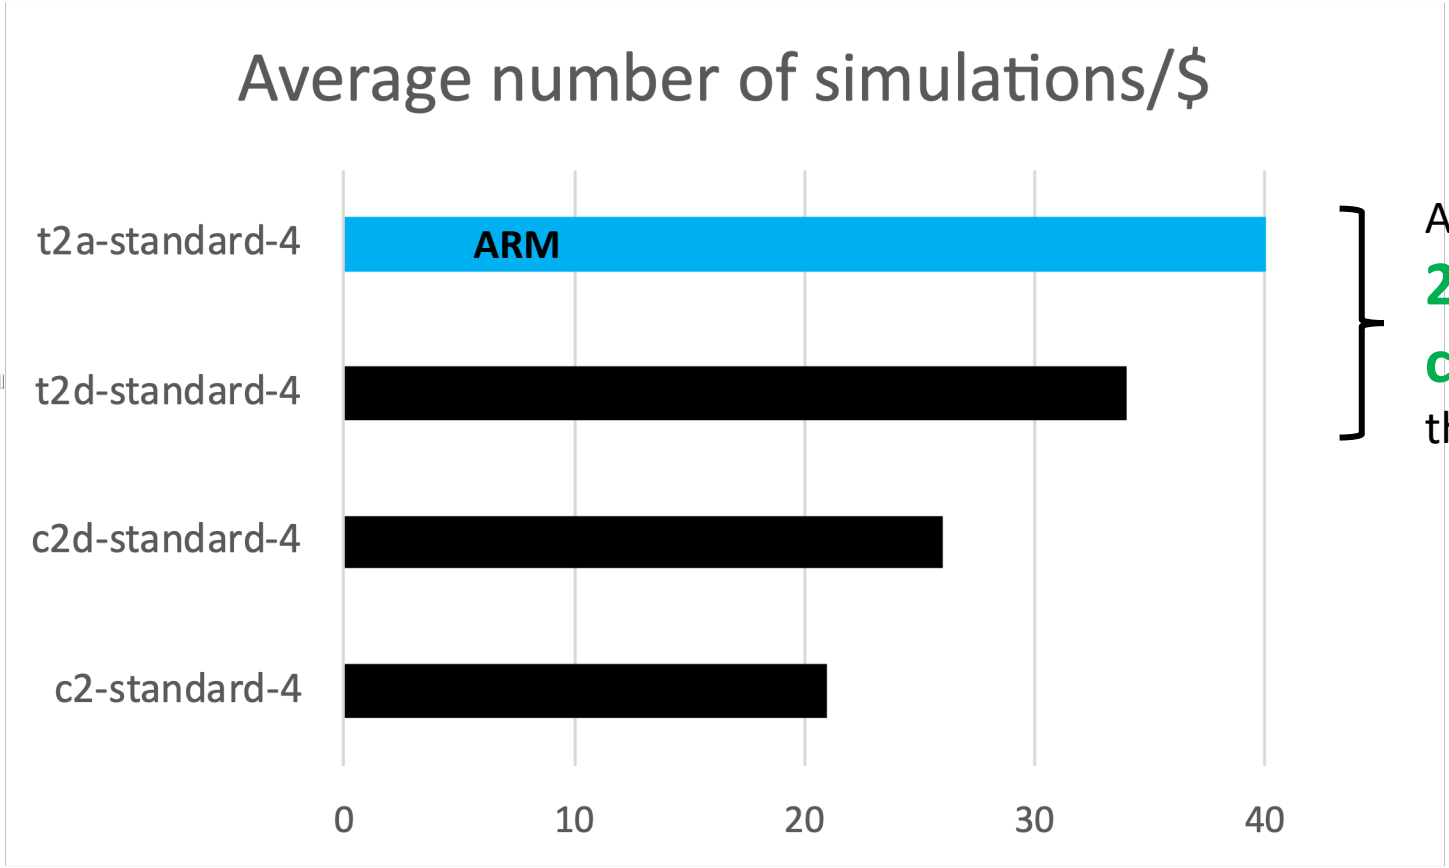

ARM about
20% more
cost effective
that the closest x86

ARM CPUs more cost-effective than x86 CPUs (at least in early 2022)

Porting experience (1/2)

- OS and the like came out of the box for both platforms
  - As standard docker images (we used k8s, remember?)
- x86 binaries already available, so a no brainer
  - Loaded straight from CVMFS
- We did not have ARM binaries, so we had to compile from source
  - Likely the first time it was done for most of the IceCube SW stack

Porting experience (2/2)

- We did not have ARM binaries
  - So, we had to compile from source
- CORSIKA is a FORTRAN application
  - Built without a hitch on ARM
  - But we also needed a bunch of helper tools for the jobs to run
- Many small helper tools involved
  - Still, the process was mostly painless
  - Found a few hardcoded x86-specific compiler parameters
    - Easy to replace with ARM-equivalents
    - But took a few passes to get them all

This is it for concrete data

- The previous slides contain pretty much all that we have done

- In the next few slides I will show a few reference data about power consumption I pulled from the Web
  - Take with a grain of salt!

CPU comparison

- Ampere Altra Q64-30
  - 64 cores at 3.0 GHz
  - 180W TDP
  - 8x 3200 DDR4 RAM

https://www.anandtech.com/show/16979/the-ampere-altra-max-review-pushing-it-to-128-cores-per-socket

- AMD EPYC 7B13
  - CPU family 25, Stepping 1
  - 60 cores, base freq 2.45 Ghz
  - 8x 3200 DDR4 RAM
- Closest public TDP data is for EPYC 7713 at 225W TDP

https://www.amd.com/en/products/cpu/amd-epyc-7713

https://cloud.google.com/compute/docs/general-purpose-machines#t2d_machines

Intel Xeon 6253CL supposedly has 18 cores at 200W TDP

CPU comparison

- Ampere Altra Q64-30
  - 64 cores at 3.0 GHz
  - 180W TDP
  - 8x 3200 DDR4 RAM

- AMD EPYC 7B13
  - CPU family 25, Stepping 1
  - 60 cores, base freq 2.45 Ghz
  - 8x 3200 DDR4 RAM
  - Closest public TDP data is for EPYC 7713 at 225W TDP

CPU power consumption is just a fraction of the total node power usage, of course!

https://www.anandtech.com/show/16979/the-ampere-altra-max-review-pushing-it-to-128-cores-per-socket

https://www.amd.com/en/products/cpu/amd-epyc-7713

https://cloud.google.com/compute/docs/general-purpose-machines#t2d_machines

Intel Xeon 6253CL supposedly has 18 cores at 200W TDP

CPU power effectiveness

Disclaimer: Power usage based on estimated TDP

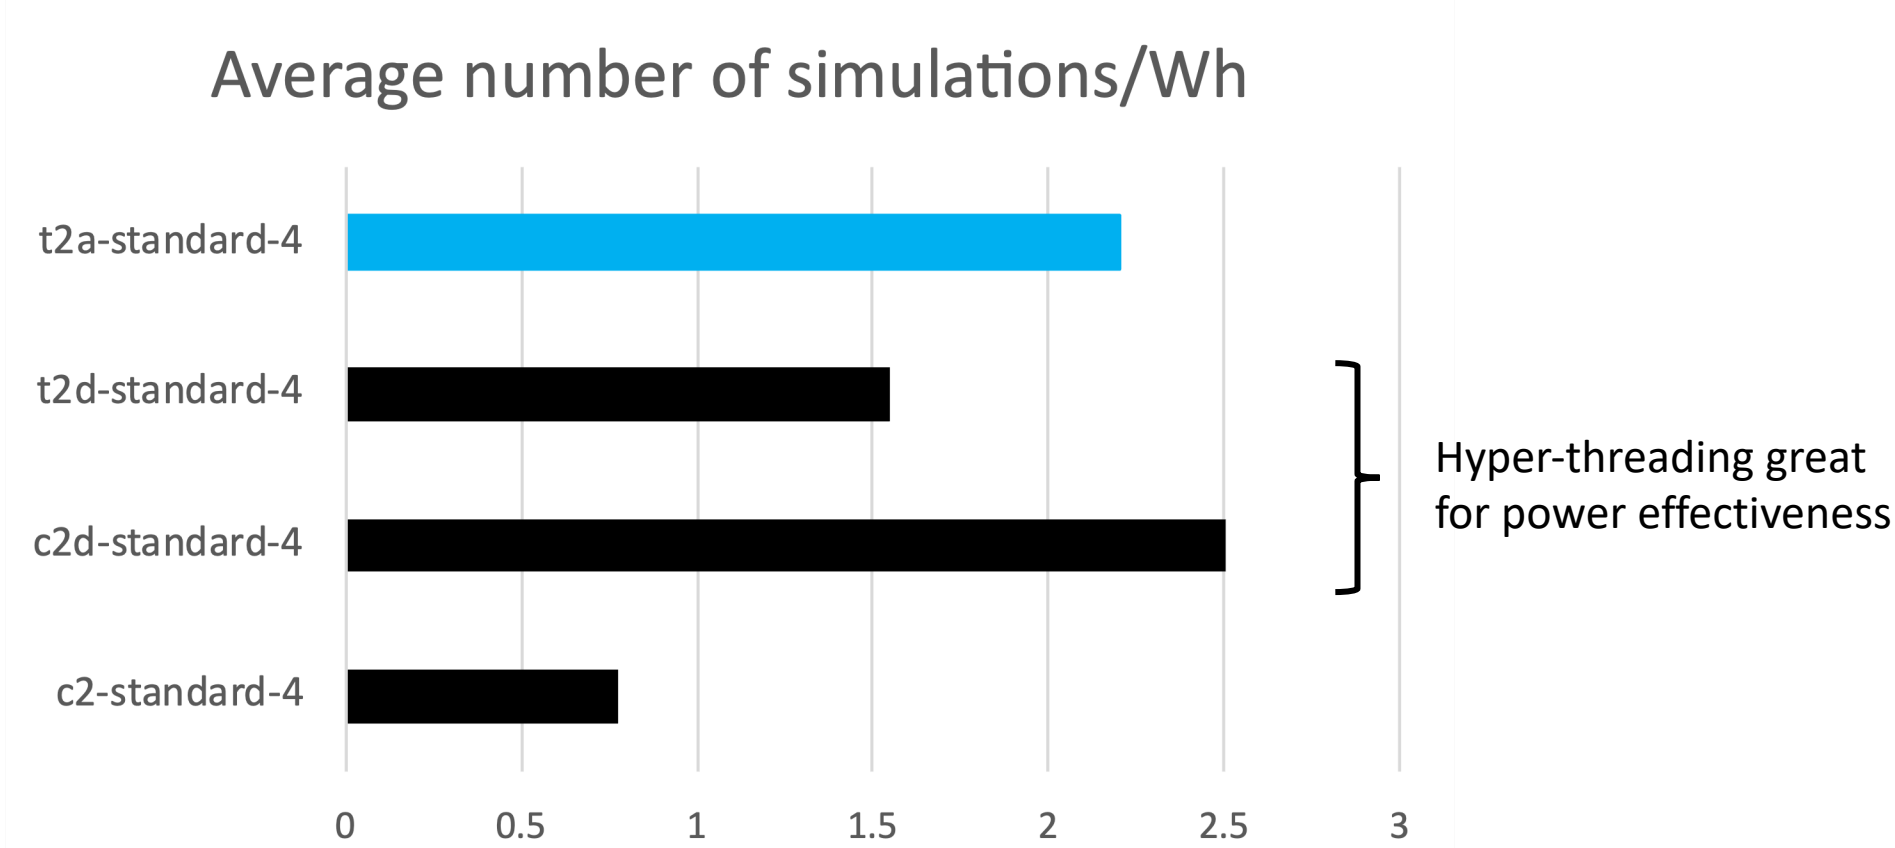

AMD EPYC in HT mode likely beats ARM CPU in power effectiveness.

But ARM CPU more power efficient when non-HT mode is desired.

In summary

- ARM CPUs now easily available in the commercial Clouds
  - Porting code to ARM not hard if you have the source code
- In GKE, using IceCube CORSIKA simulation
  - ARM cores slightly faster than x86 cores
  - And priced significantly cheaper
- No hard data on power usage, but extrapolated data imply
  - ARM CPUs only modestly more power efficient that x86 competition

Other considerations

- Google Cloud keeps adding new CPUs as time goes by
  - E.g., c3 series, based on Intel Xeon Platinum 8481C (Saphire Rapids)
  - Did not repeat the tests on these to see how they compare
- ARM-based servers are available on-prem
  - In addition to x86-based ones
  - We did not systematically check those as a mean for comparison

Acknowledgement

- This work was partially funded by the U.S. National Science Foundation (NSF) under grants OAC-1826967, OAC-1541349, OPP- 2042807, CNS-1925001 and OAC- 2103963.

- All Google Kubernetes Engine costs have been covered by Google-issued credits.

UC San Diego