2nd Computing Challenges Workshop (COMCHA),
2-4 Oct 2024

Technology trends and hardware architectures for HEP computing

Brij Kishor Jashal
Rutherford Appeleton Lab & IFIC, U.Valencia

Introduction

I enjoyed reading your slides from last 2 days

I mean seriously !!

A lot of technology trends have been covered already in great detail !!!

I mean seriously a lot !!

Computing challenges in Calorimetry

Tigger strategies at LHC LHCb, CMS, ATLAS

Computing challenges in tracking

AI algorithms for CTA

Computing challenges in Muon systems

Differential programming

HyperK reconstruction on GPU

From High Energy Physics to Industry

European strategy discussion
- Heterogeneous computing
- AI & ML
- Quantum & Neuromorphic
- Sustainability

Sustainability of RTA

Quantum computing and tracking

Allen optimization

Madgraph to FPGA/HLS

HEP computing challenge

Ref. V. Gligorov

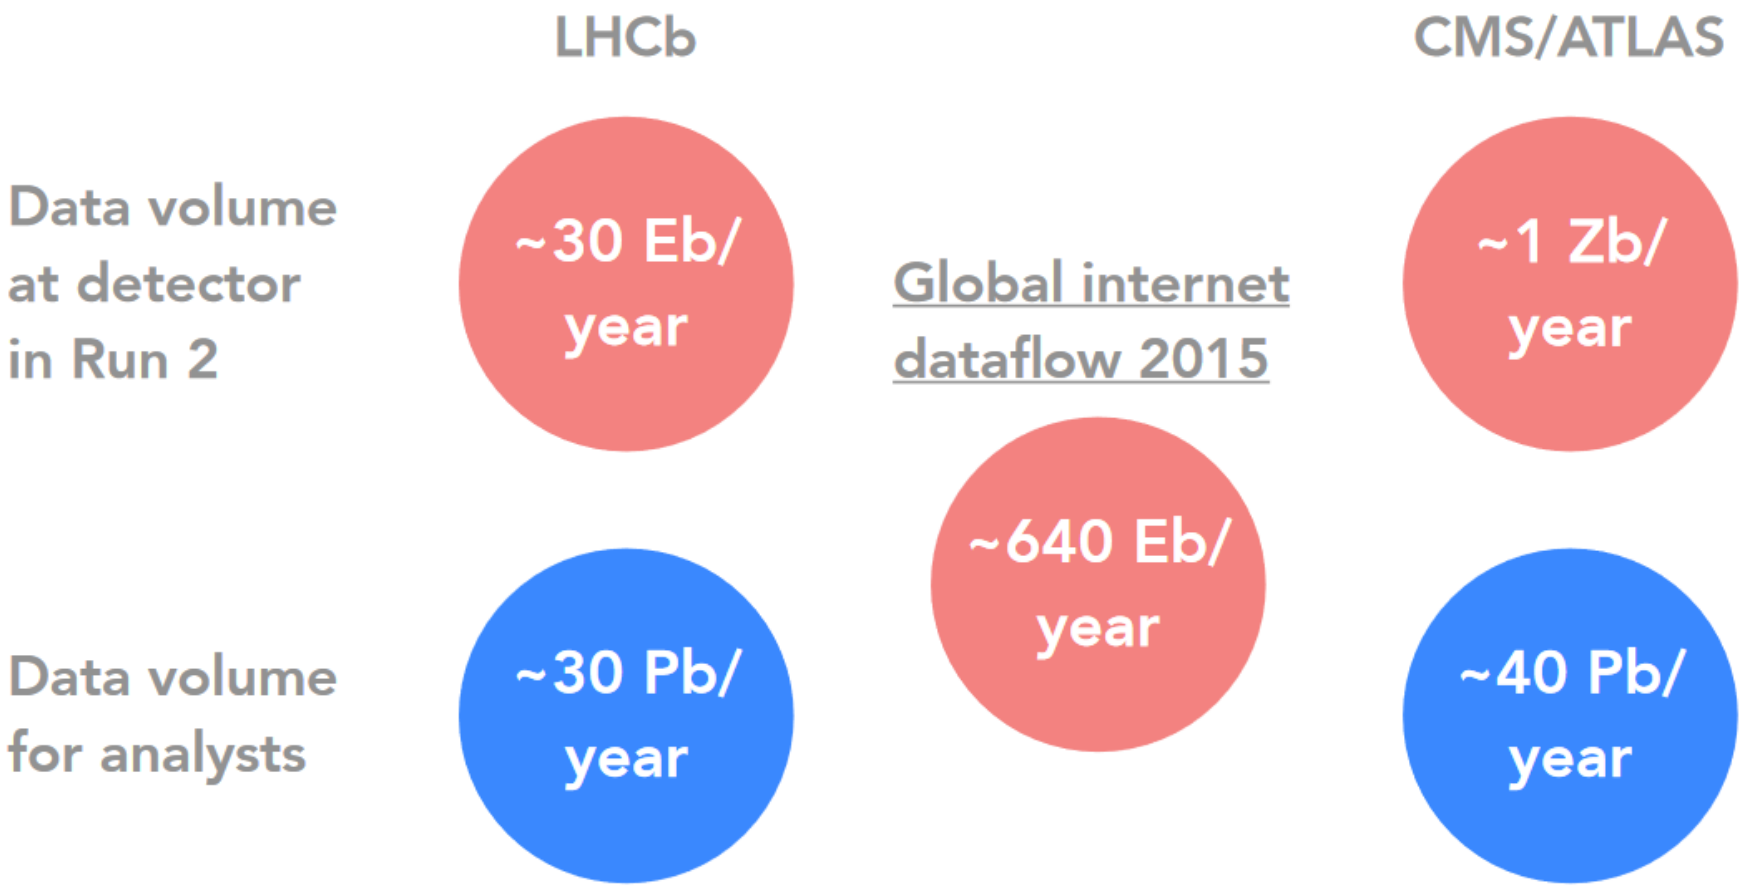

HEP computing challenge

Ref: Resource projection:

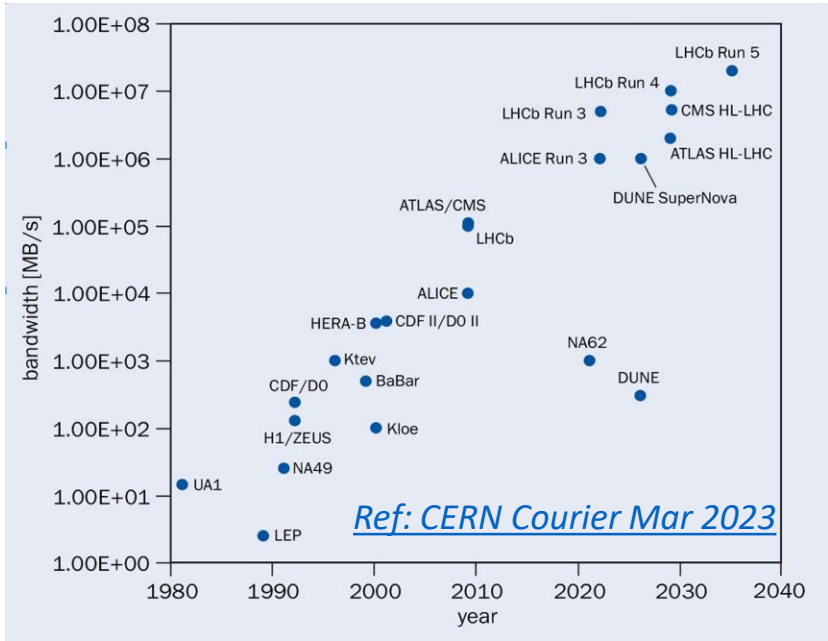

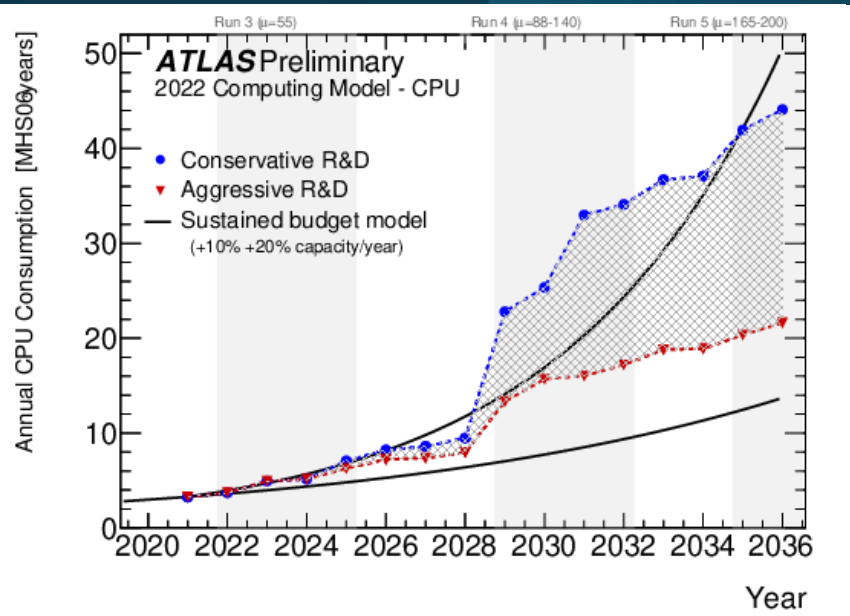

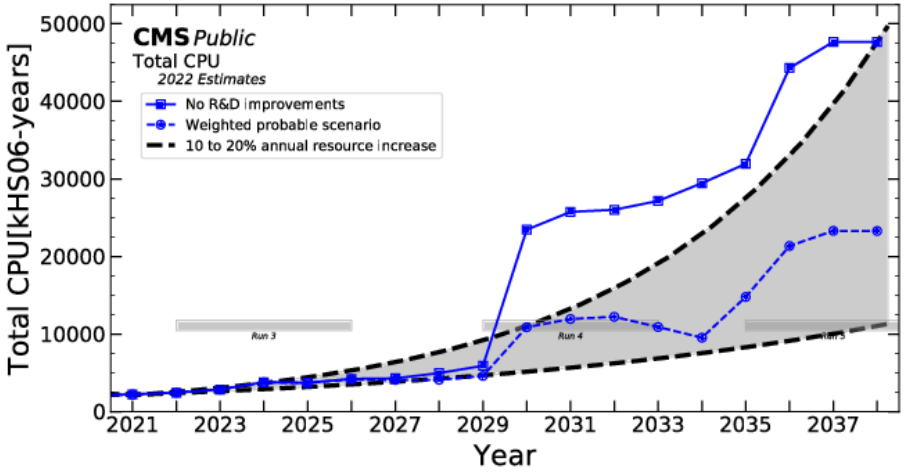

Storage challenge.

Ref: Resource projection:

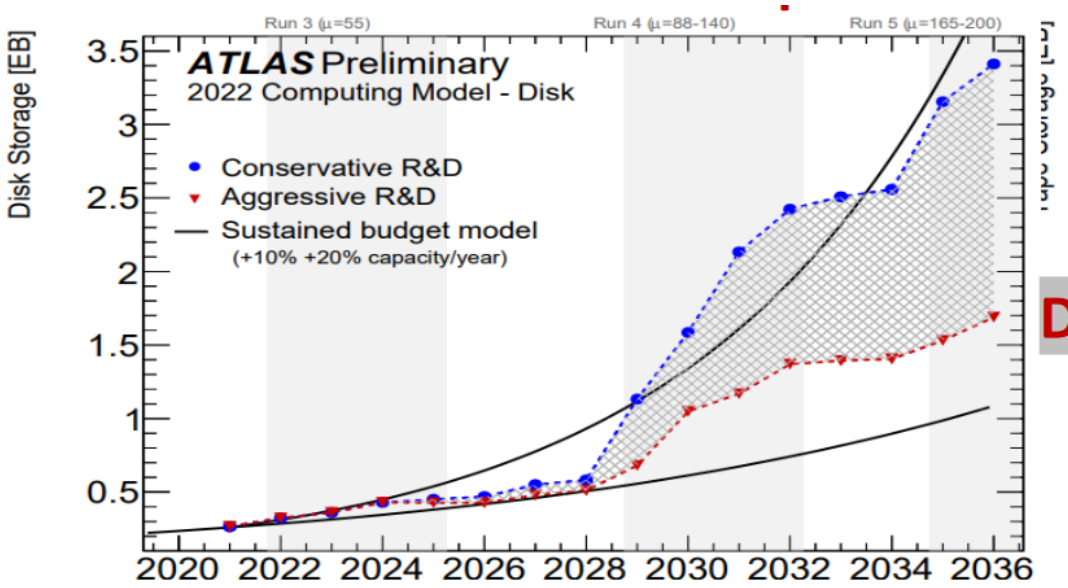

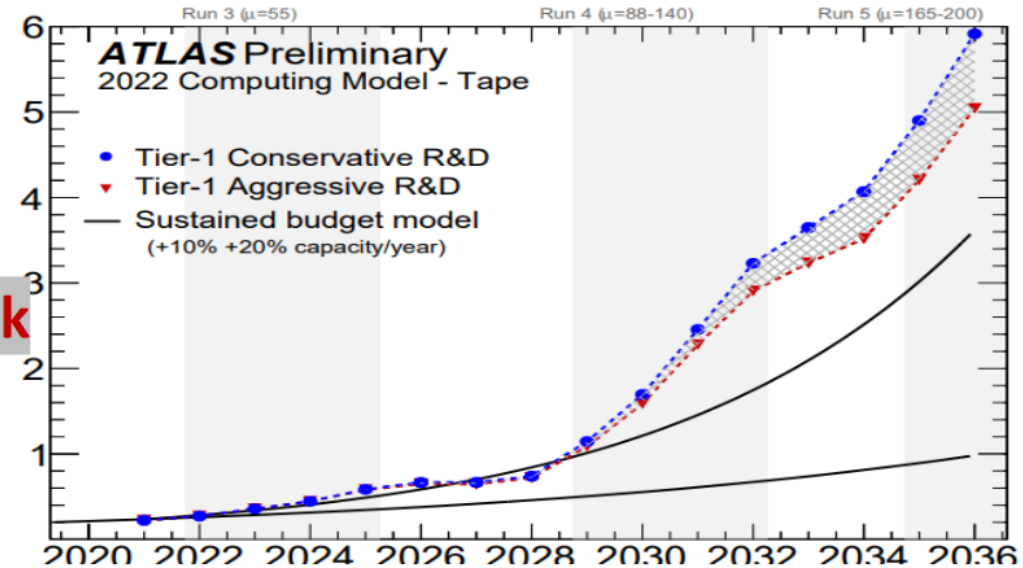

Disk

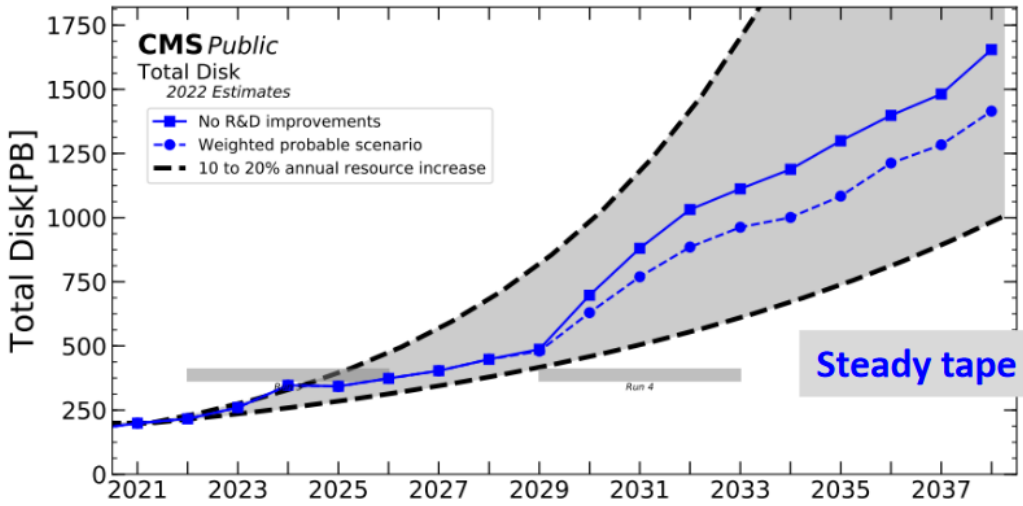

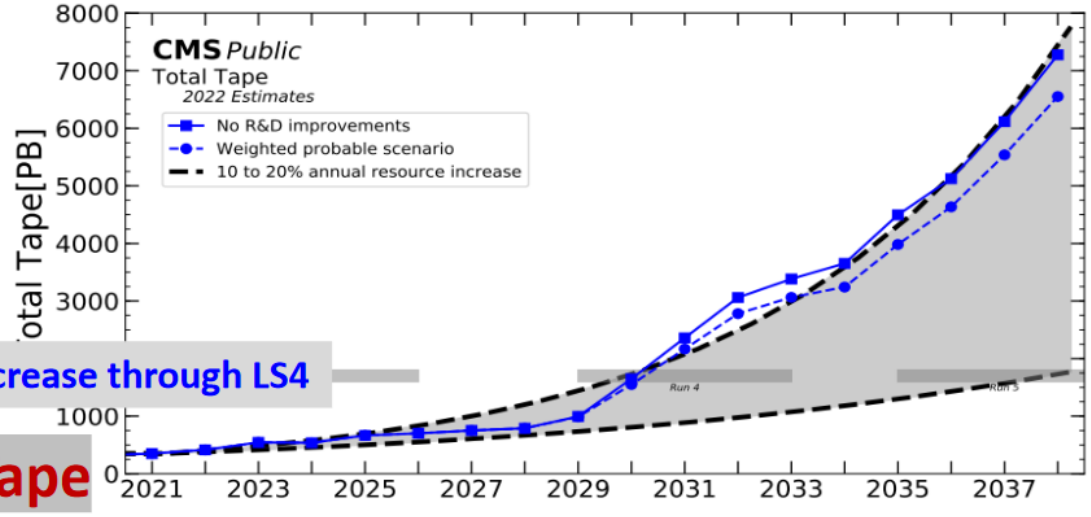

Steady tape increase through LS4

Tape

Traditional Techniques No Longer Scale

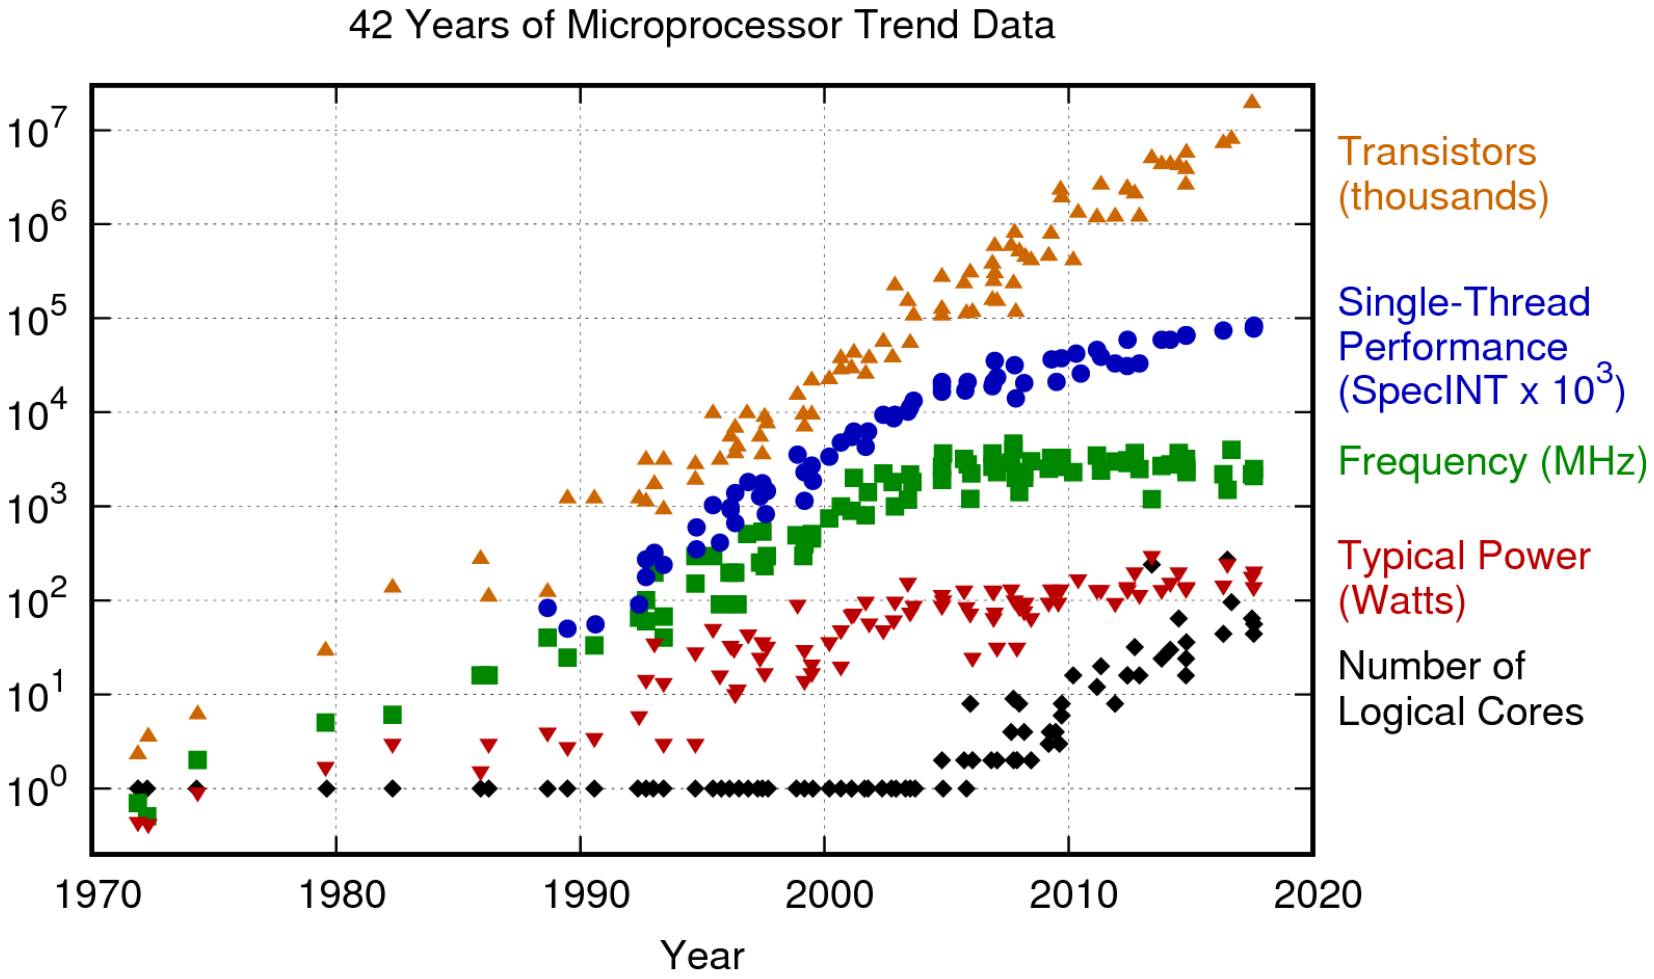

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Changed CPUs for Modern Workloads

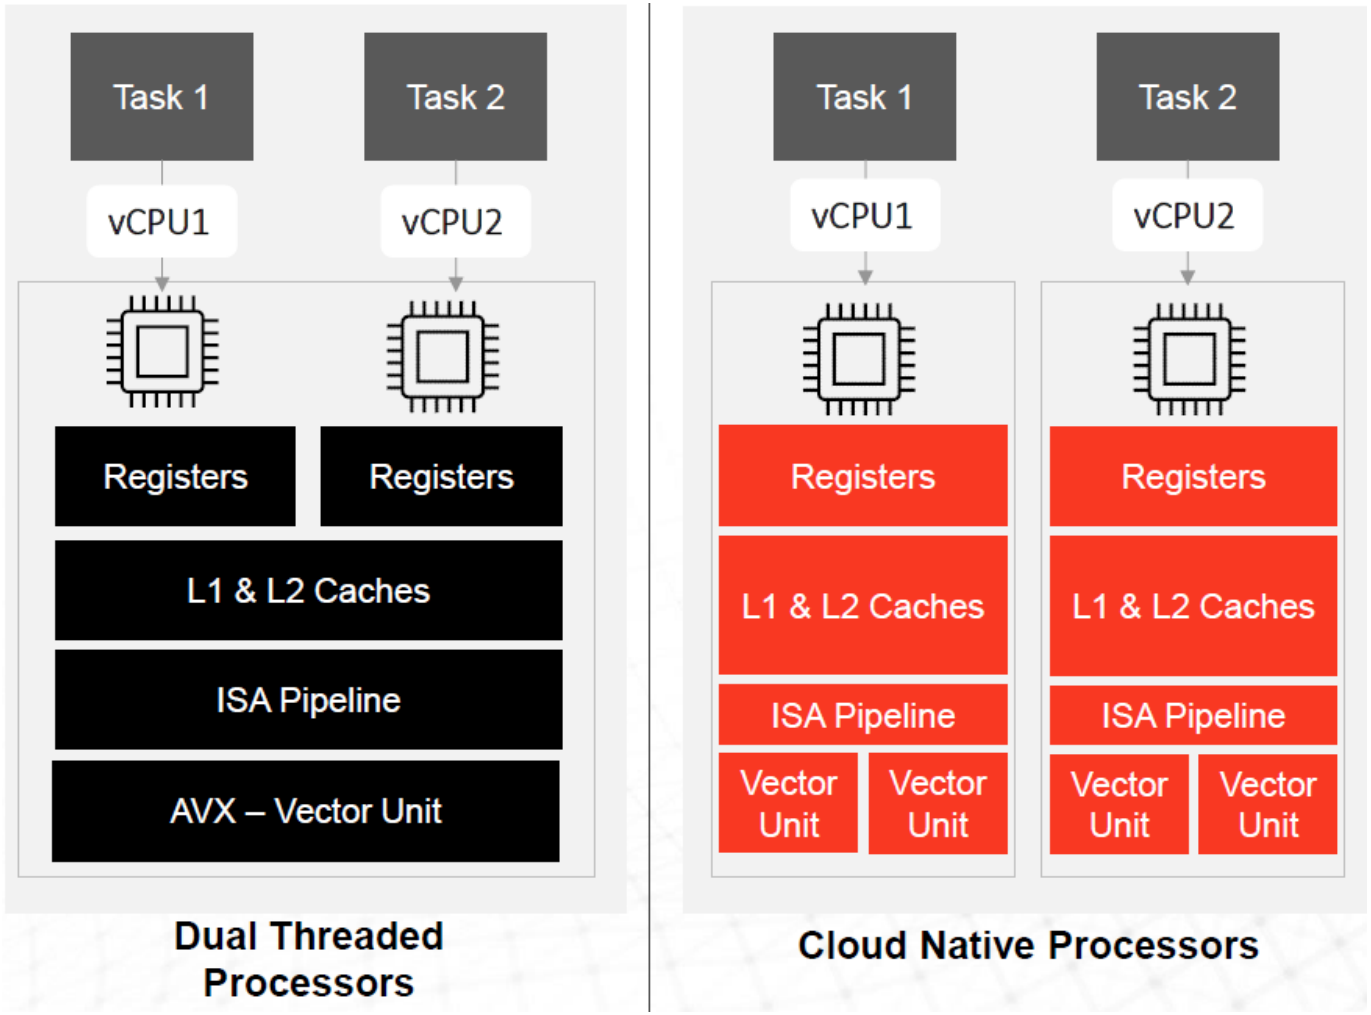

<u>For all workloads</u>

- Private resources in each core
- Resists noisy neighbour influence
- Predictable latency
- Linear Scaling
- Up to 384 Vector Engines
- Scale out requires SW optimisation
- Containerization of services and

Domain Specific Architectures: a New Cambrian Explosion

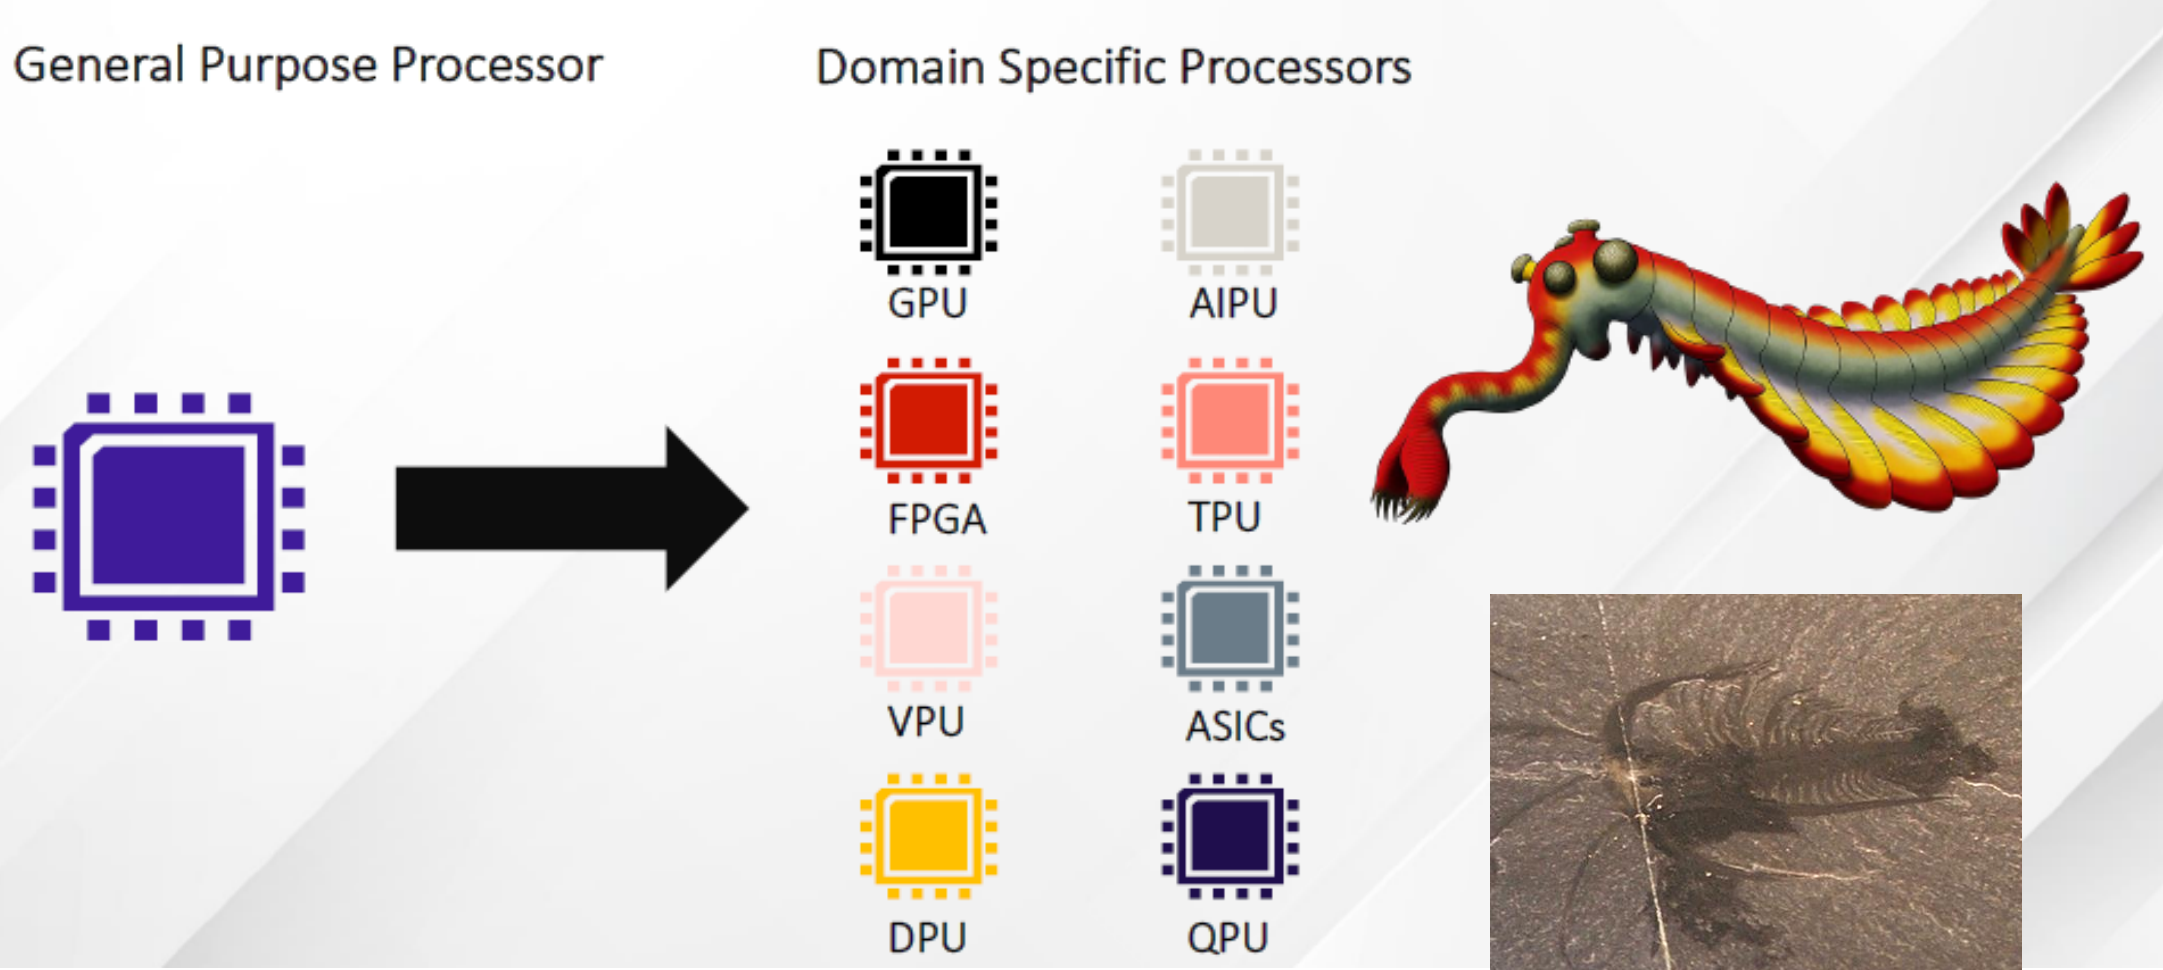

Architecture-aware optimizations in hot sections of our code can yield huge gains overall

Silicon Sustainability - Embodied Energy in Silicon Manufacture

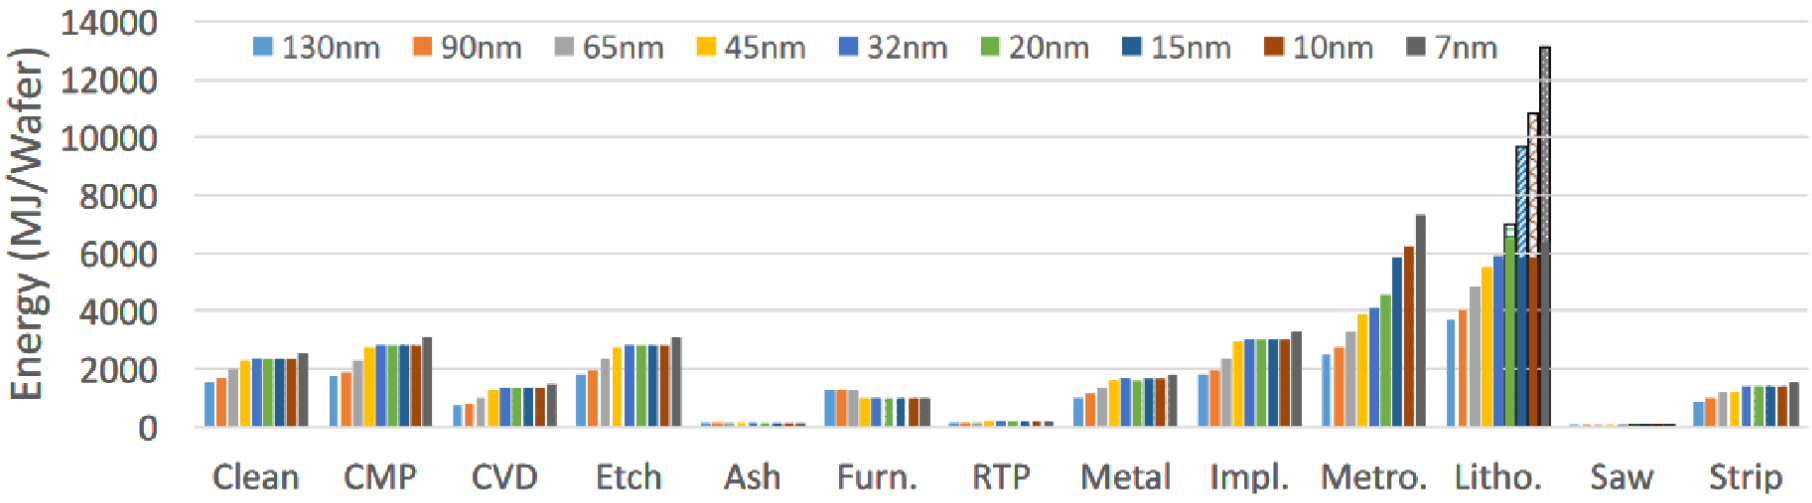

Figure 4. Energy used per process step calculated using the process model.

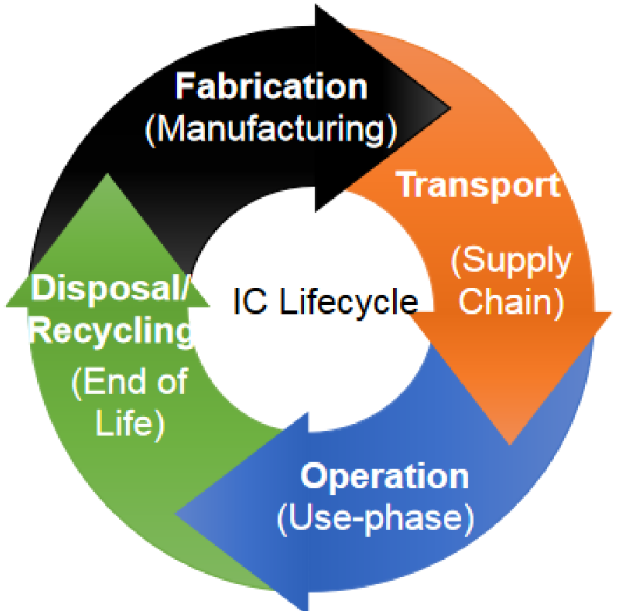

Architecture	Energy per trigger (mJ)	Gain	Total gain
E5-2630-v4 Xeon			
Before SW optimization	39.9	1.0x	
w/Physics optimizations	21.0	1.9x	1.9x
w/SIMD optimizations	8.4	2.5x	4.8x
7502 EPYC			
w/SIMD optimizations	3.2	2.6x	12.5x
Event Building Node, NR			
1 GPU	3.1	1.03x	12.9x
2 GPUs	2.4	1.29x	16.6x
3 GPUs	2.1	1.15x	19.0x
Dedicated GPU machine			
4 x 2080 Ti + 2 Network Cards	2.8	1.14x	14.3x
5 x 2080 Ti + 3 Network Cards	2.5	1.12x	16.0x
Pure GPU machine			
8 x 2080 Ti + Onboard Network	2.1	1.15x	19.0x

Processor industry

Ref: Resource projection:

Only 3 companies in the world capable of fabricating leading-edge chips ("5nm node" or less)

	Samsung	TSMC	Intel
Revenue :	44 B$/y	67B$/y	54 B$/y
Fabs (leading edge) :	6	10 (5)	15 (7)
Sites :	South Korea	Taiwan +(40B$ investment Arizona)	US, Israel, Europe
Customer (main) :	Smartphones: ARM	AMD : all CPU and GPU; Apple : ARM Nvidia: GPU Sony, Microsoft: game console CPU+GPU some Intel processors (2024)	Intel: CPU and integrated GPU

Very complex fabrication process
A wafer stays 3 month in a fab and runs through ~1000 processing steps

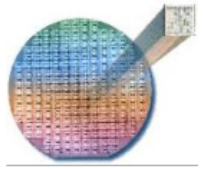

Very few companies can provide:

- Ultra-pure silicon wafers
- Special photoresist
- Precise photolithography masks
- Ultra-pure chemicals

- A new fab requires investments of >10B$
- All 3 companies want to invest each ~100B$ during the next years in new fabrication units

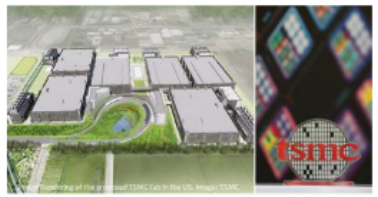

Only one company (ASML) provides EUV Extreme Ultraviolet Lithography equipment

Monopoly, single source suppliers…….

Disk storage trends

Ref: Resource projection:

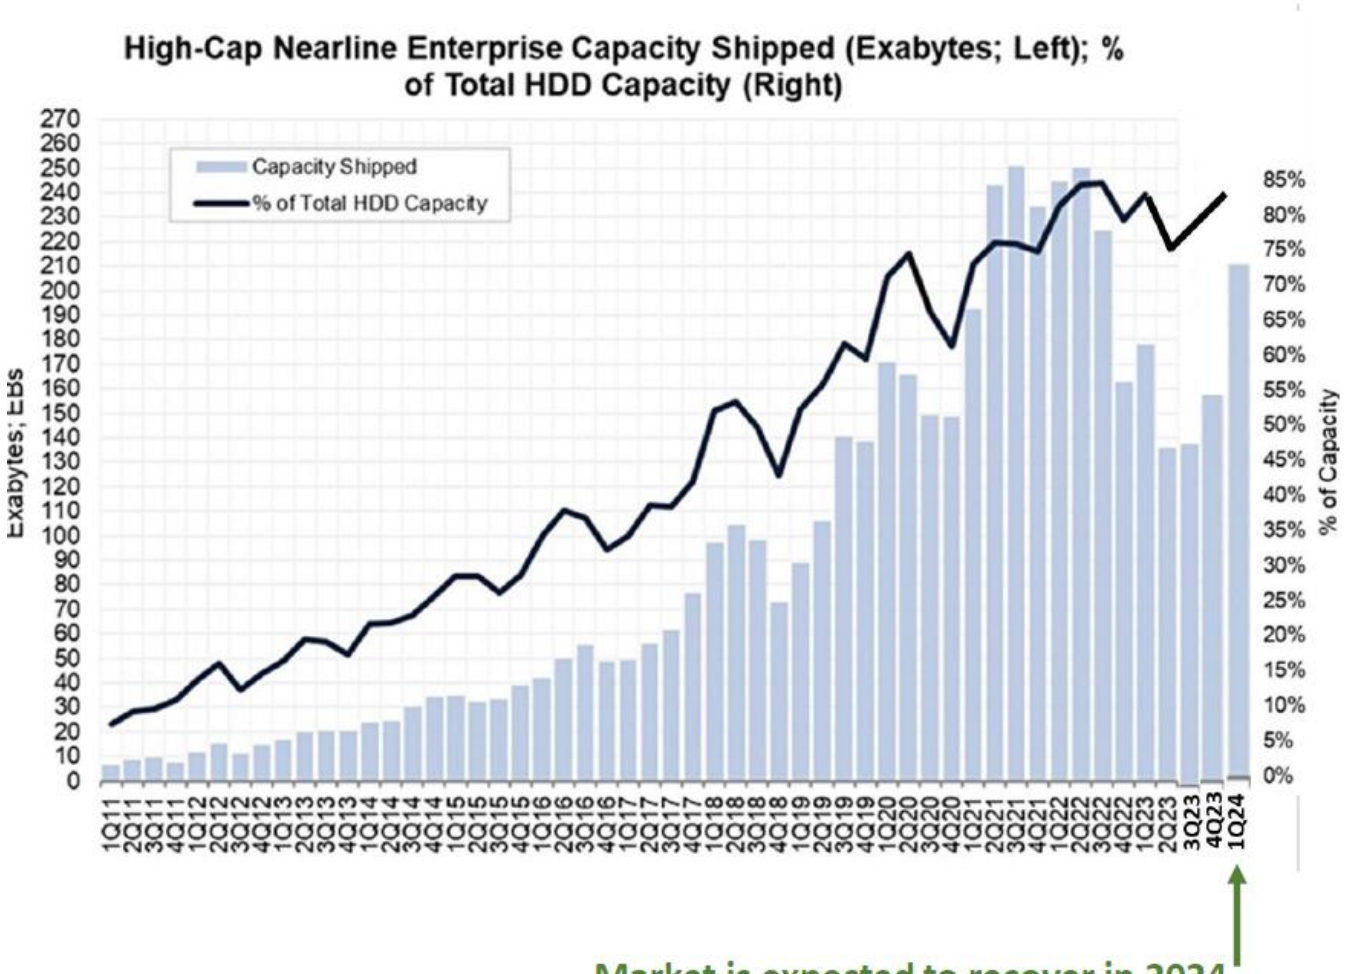

Market is expected to recover in 2024

Capacity drives (>8TB) dominating the shipments
Strong trend towards SMR drives

Notebooks are already 100% equipped with SSDs PCs in the near future

Economic turbulences and stockpiling during COVID-19 → revenues and sales dropped in 2023

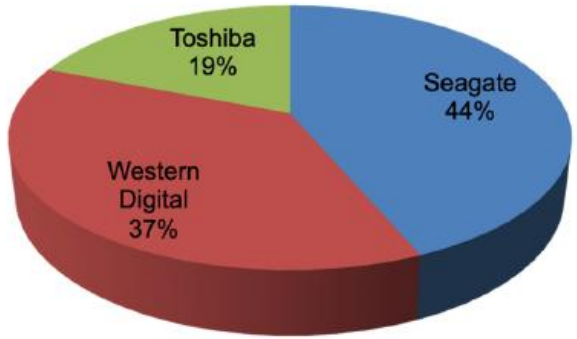

2023 HDD Market Share to Date
COUGHLIN ASSOCIATES IMAGE

Only 3 companies dominating the market

Processor industry

Ref: Resource projection:

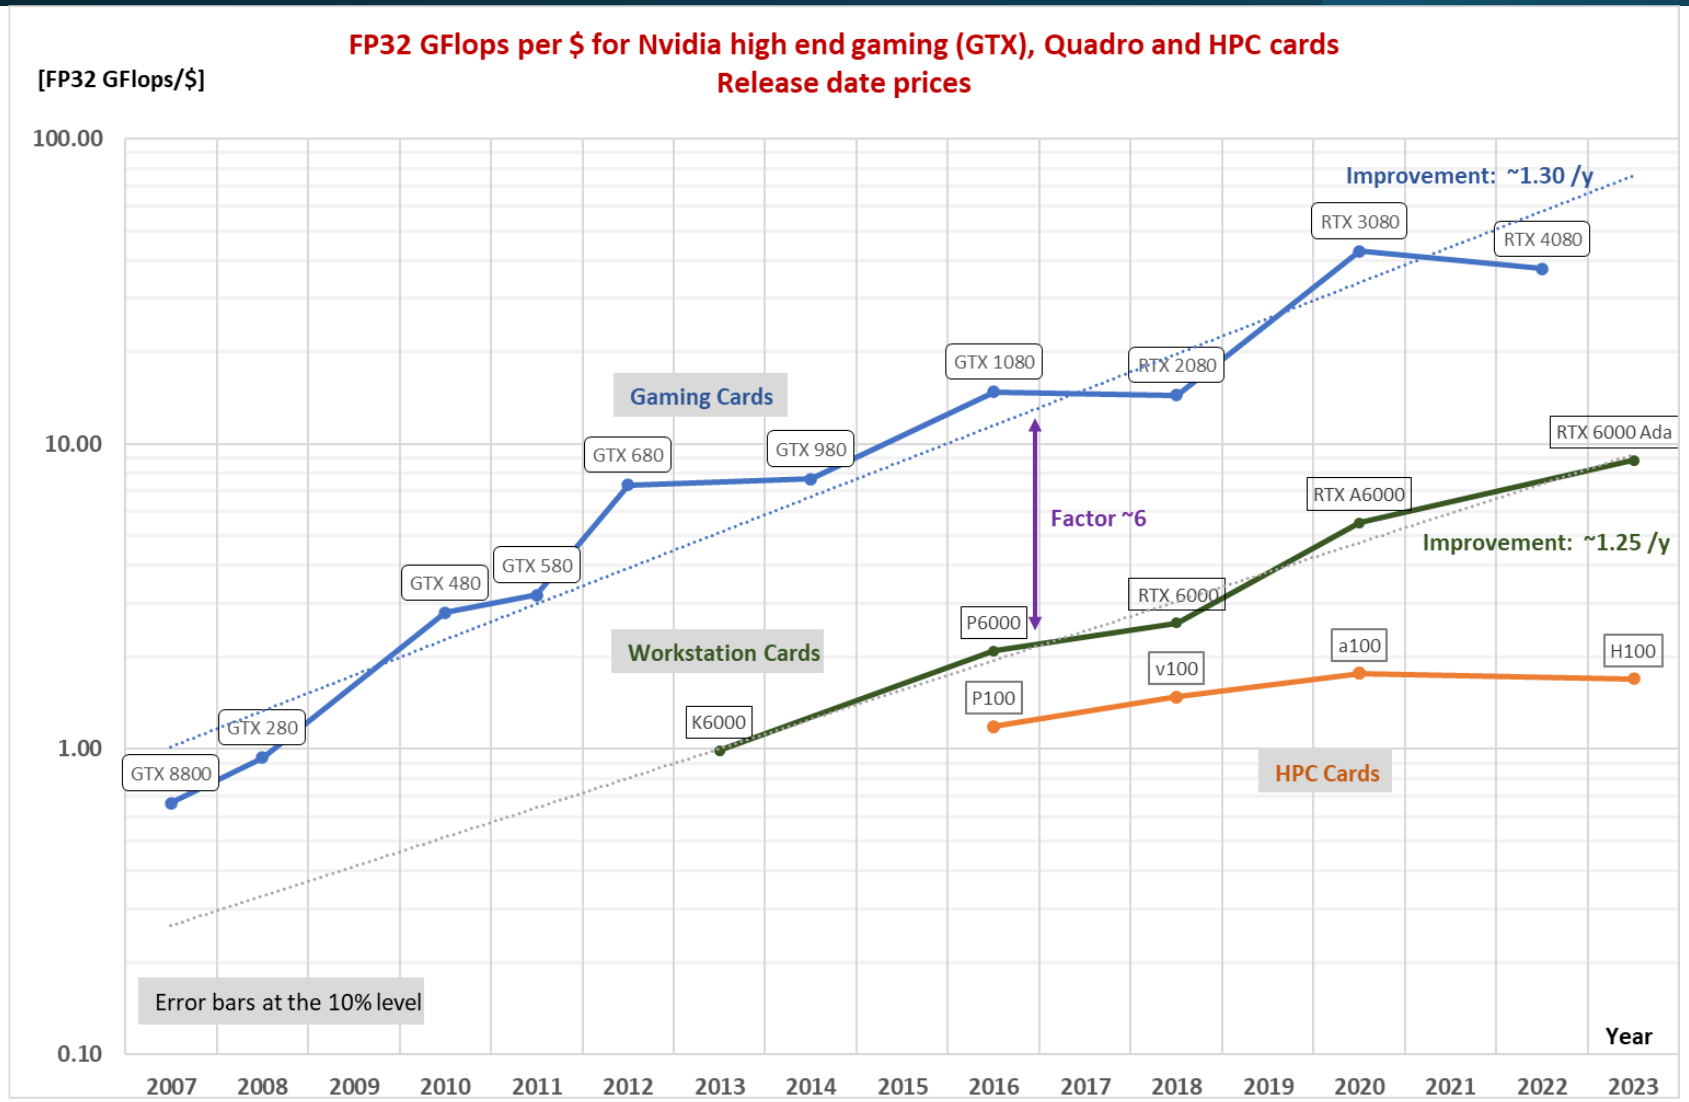

Software approaches beyond evolutionary baseline: HLT

Ref: Looking-into-trigger-future-U2 D.camporaa

Complexity challenges ahead: Despite clever simplifications, the complexity bounds highlight significant challenges for future Runs.

	Theoretical problem	Simplification	
Data sorting	$O(n^2)$	$O(n \cdot log(n))$	**Quicksort or merge sort**
Track seeding	$O(2^n)$	$O(n \cdot log(n)^2)$	**Geometry or physical constrains,**
Track following	$O(2^n)$	$O(n \cdot log(n))$	**Kalman filter to most likelihood path**
Likelihood minimisation	$O(2^n)$	$O(n^6)$	**Gradient descent from exp to high-deg pol**
Clustering	$O(n^2)$	$O(\|V\| + \|E\|)$	**Graph based clustering**
Selections	$O(2^n)$	$O(n^2)$	**Exp to Quad**

Need for advancements in algorithms:

- Similar challenges for MC simulations and offline processing.
- Development of more advanced and efficient data traversal algorithms is essential to manage exponentially growing data throughput,

Hybrid architectures in offline computing ??

For a 30 MHz HLT1 at LHCb

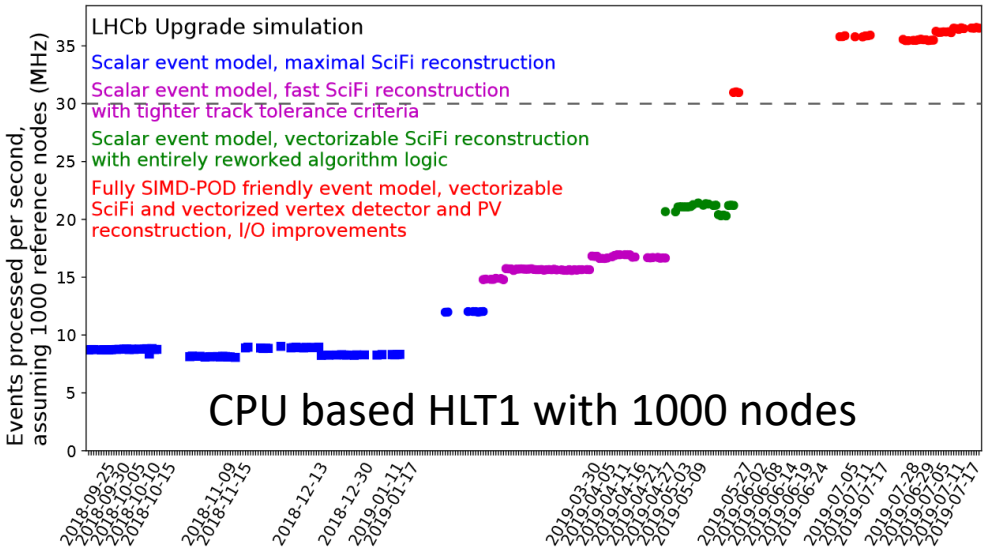

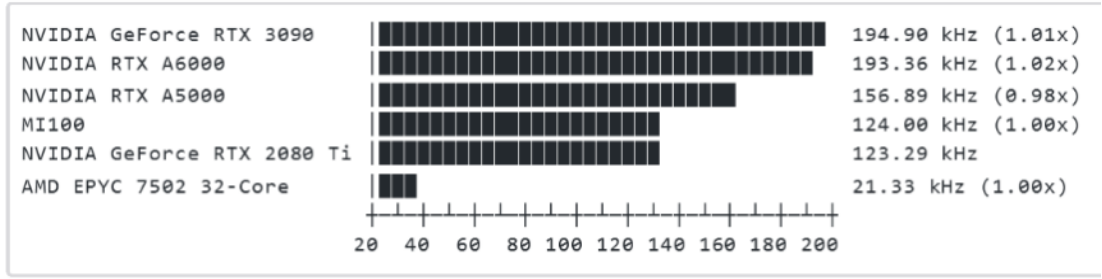

GPU based HLT1 for Run3 with O(200) GPUs

Almost negligible GPU utilization in offline workloads

- WLCG does not take GPUs as standard pledge yet
- Because almost negligible workloads
- We must be able to utilize the non CPU resources for offline computing

Pythonic ecosystem for data analysis, well advanced in HEP

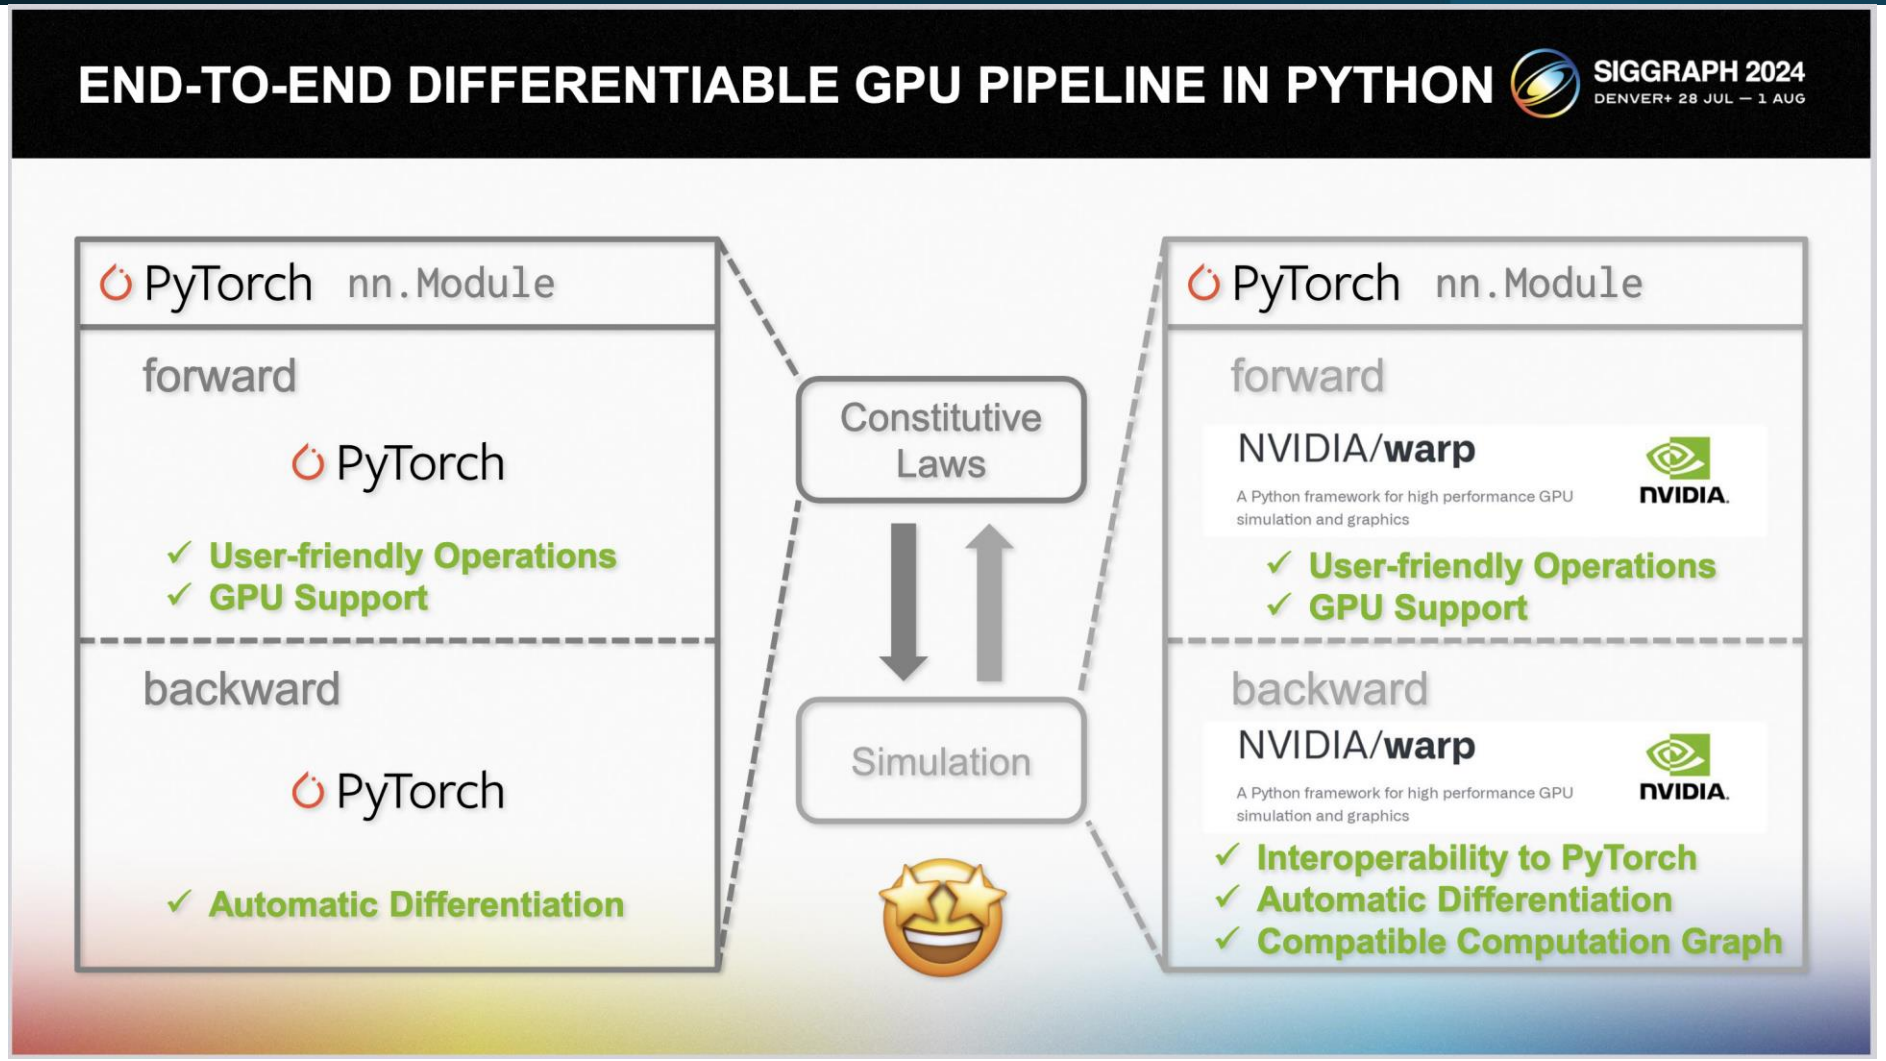

Summary

- Architecture-aware programming is a must.
- Architecture choice will be driven by throughput.
  - At High Level Trigger level, Performance over portability.
  - For offline computing, portability becomes crucial.
- Measuring Power Usage Effectiveness (PUE) for our Tier datacentres is the first step for making WLCG greener.
- To effectively utilize GPUs for offline computing, simulation chain needs significant development.
- Investment in software (and in person) will continue to pay huge dividends.

Thank you