Tata Institute of Fundamental Research
टाटा मूलभूत अनुसंधान संस्थान

Technology trends in HEP Computing

The 8th Asia Tier Centre Forum

2nd to 4th Sep TIFR Mumbai

Brij Kishor Jashal

HEP computing challenge

	LHCb	CMS/ATLAS
Data volume at detector in Run 2	~30 Eb/year	~1 Zb/year
Data volume for analysts	~30 Pb/year	~40 Pb/year

Global internet dataflow 2015

~640 Eb/year

Ref. V. Gligorov

HEP computing challenge

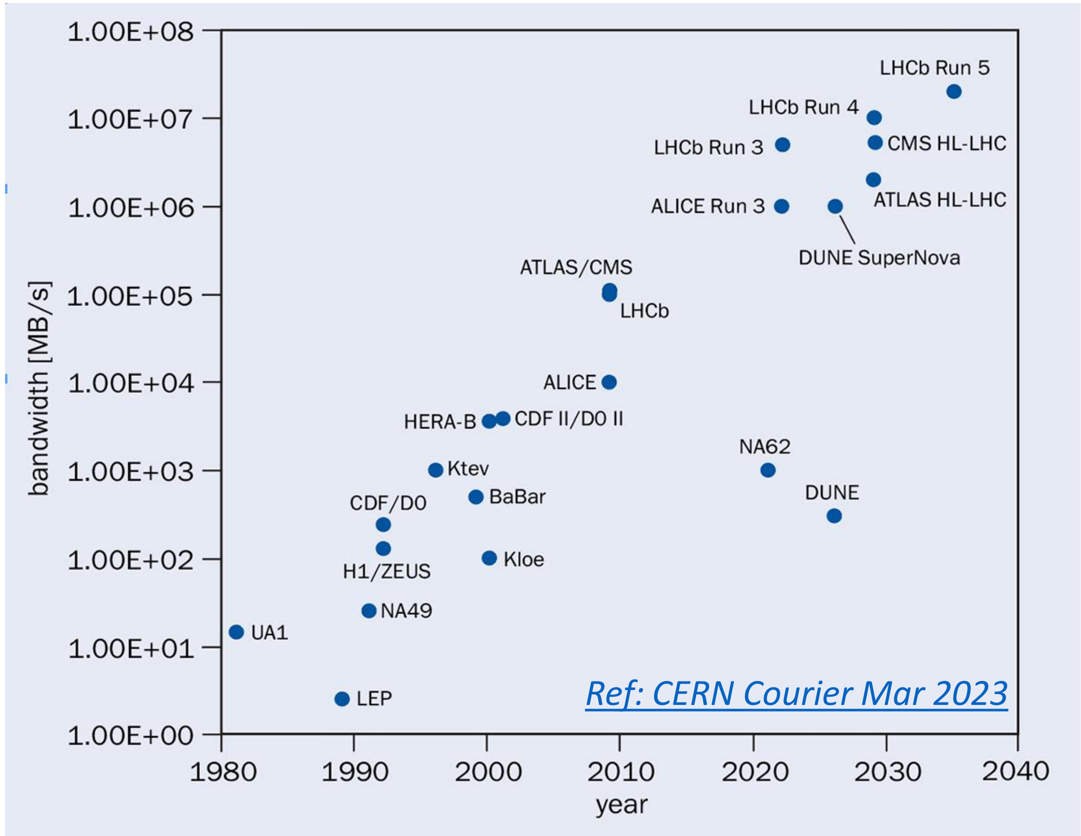

Ref: CERN Courier Mar 2023

Bandwidth of data analyzed in real-time versus the start date of various high-energy physics experiments.

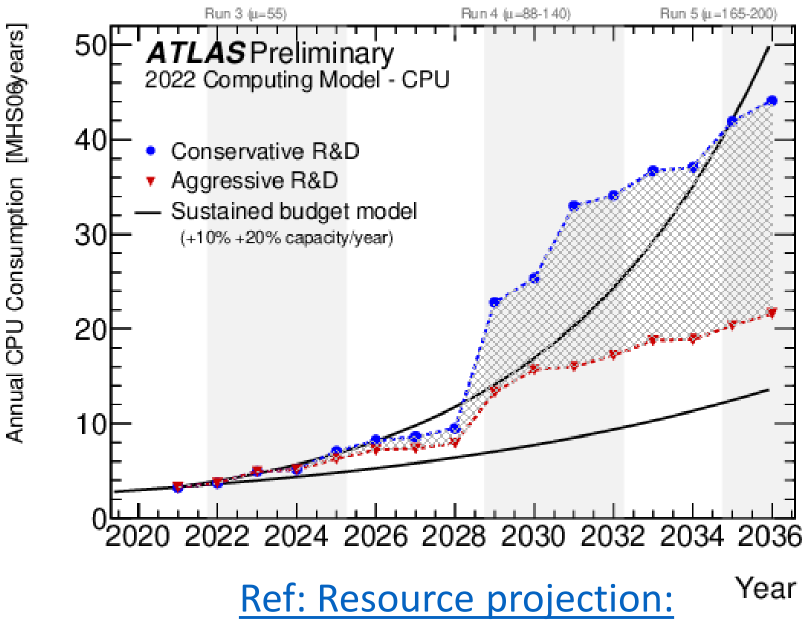

Ref: Resource projection:

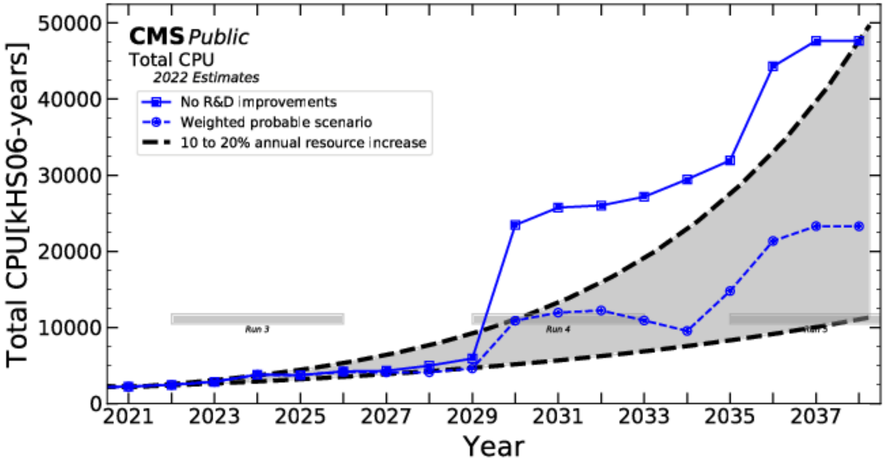

Traditional Techniques No Longer Scale

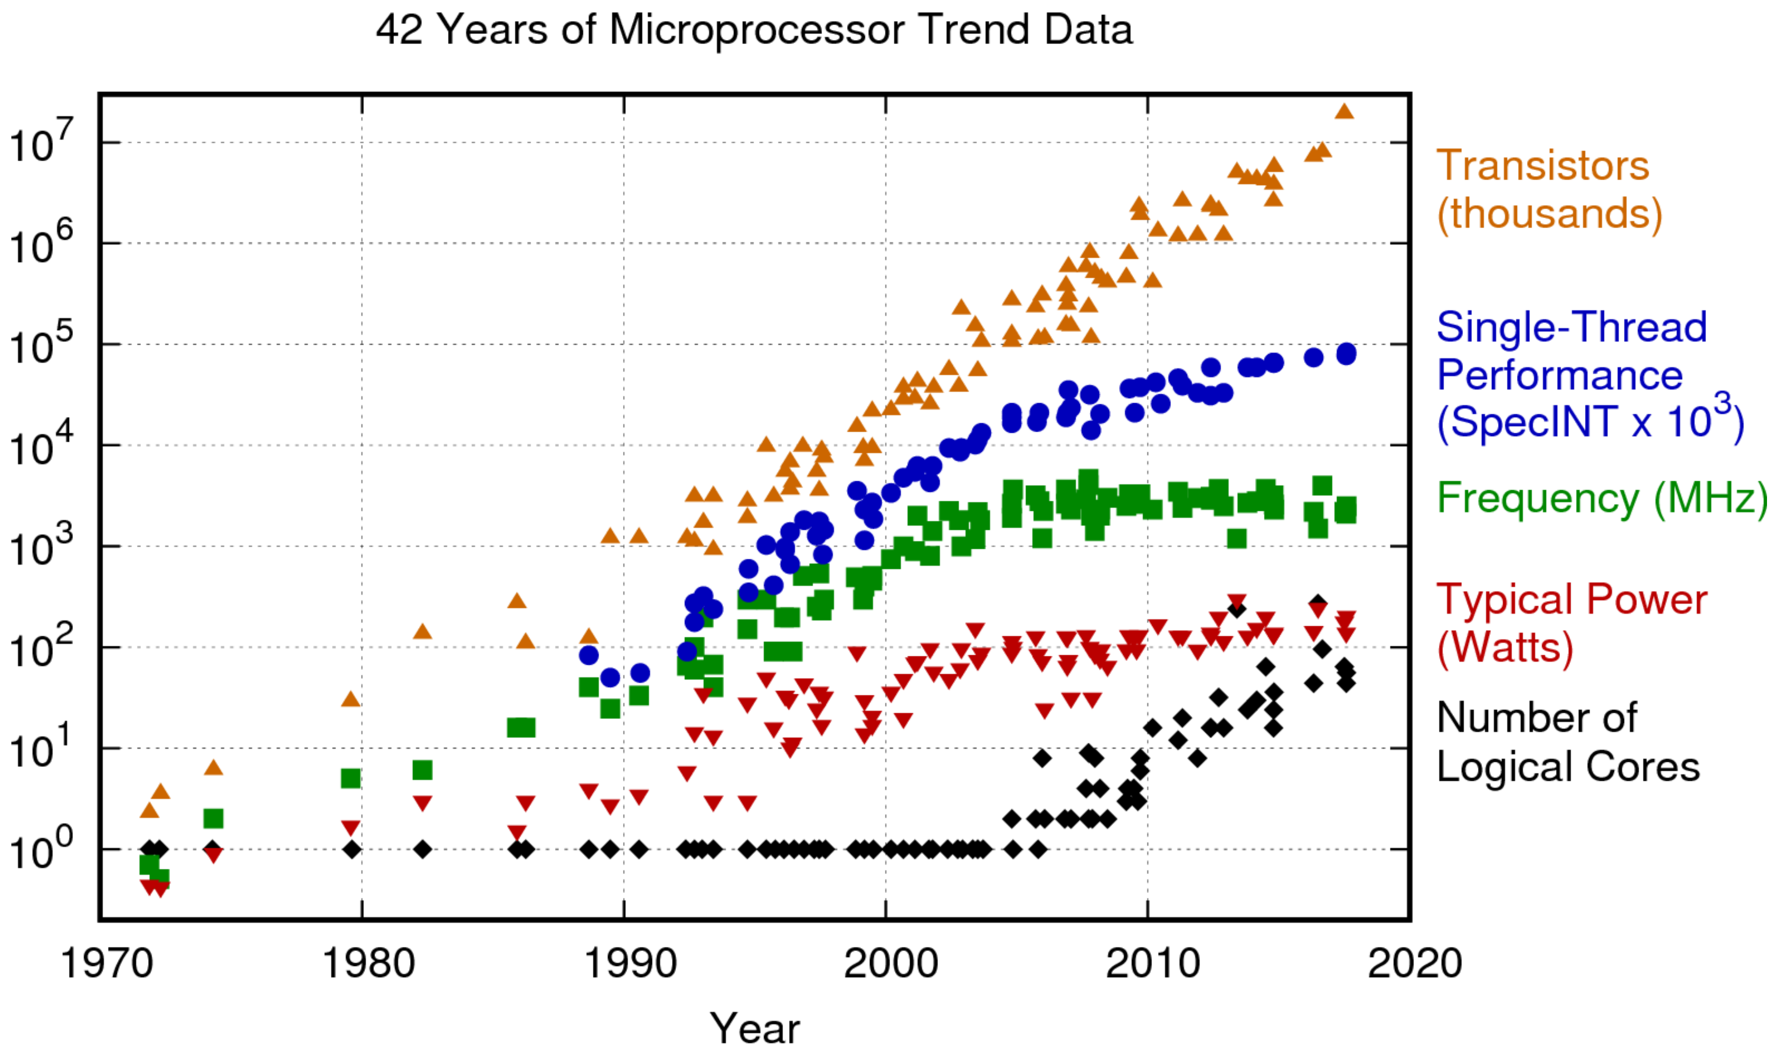

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Silicon Sustainability - Embodied Energy in Silicon Manufacture

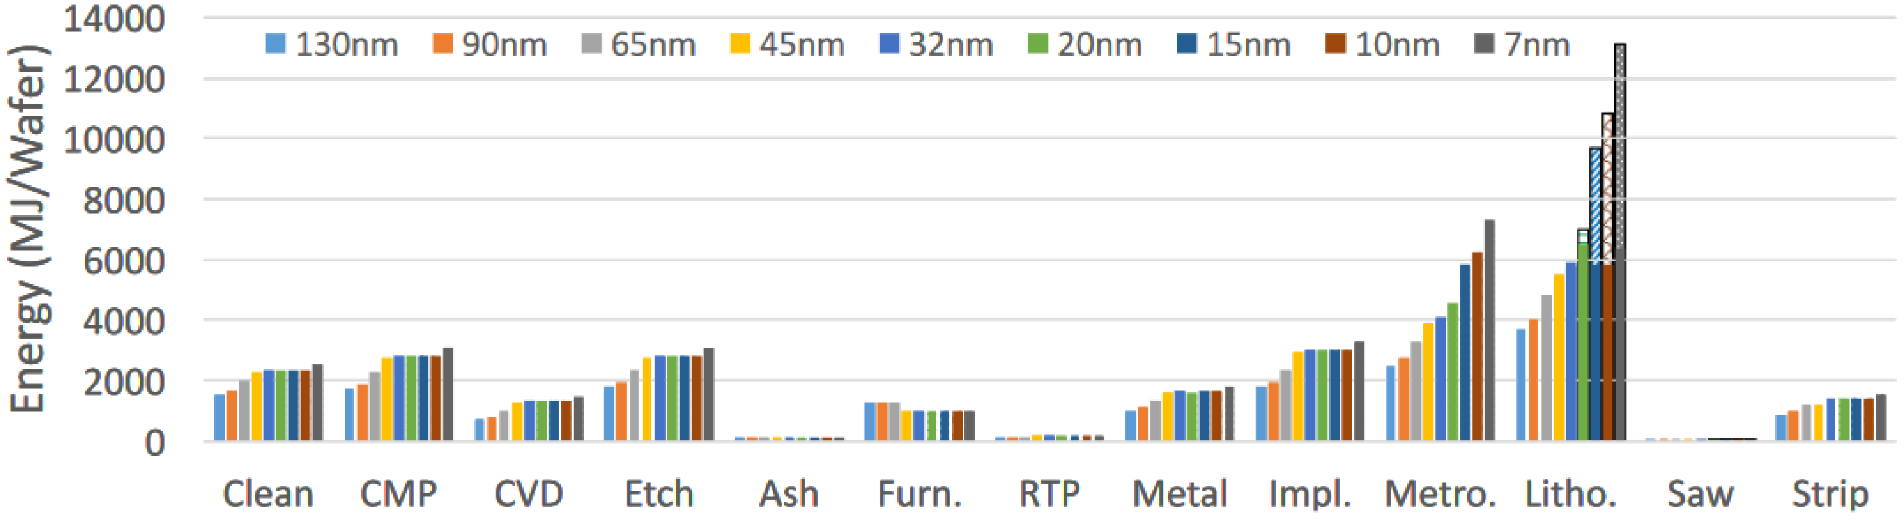

Energy used per process step calculated using the process model.

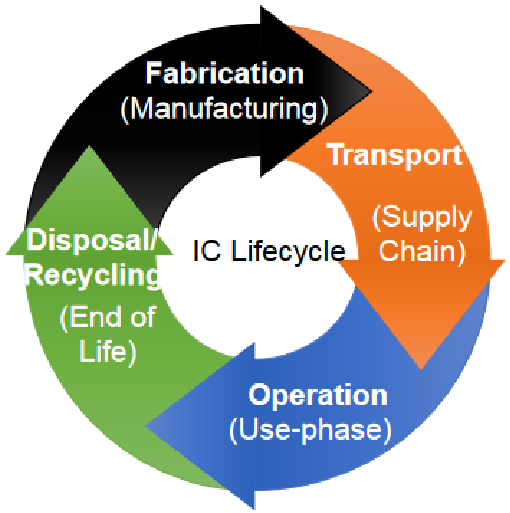

Changed CPUs for Modern Workloads

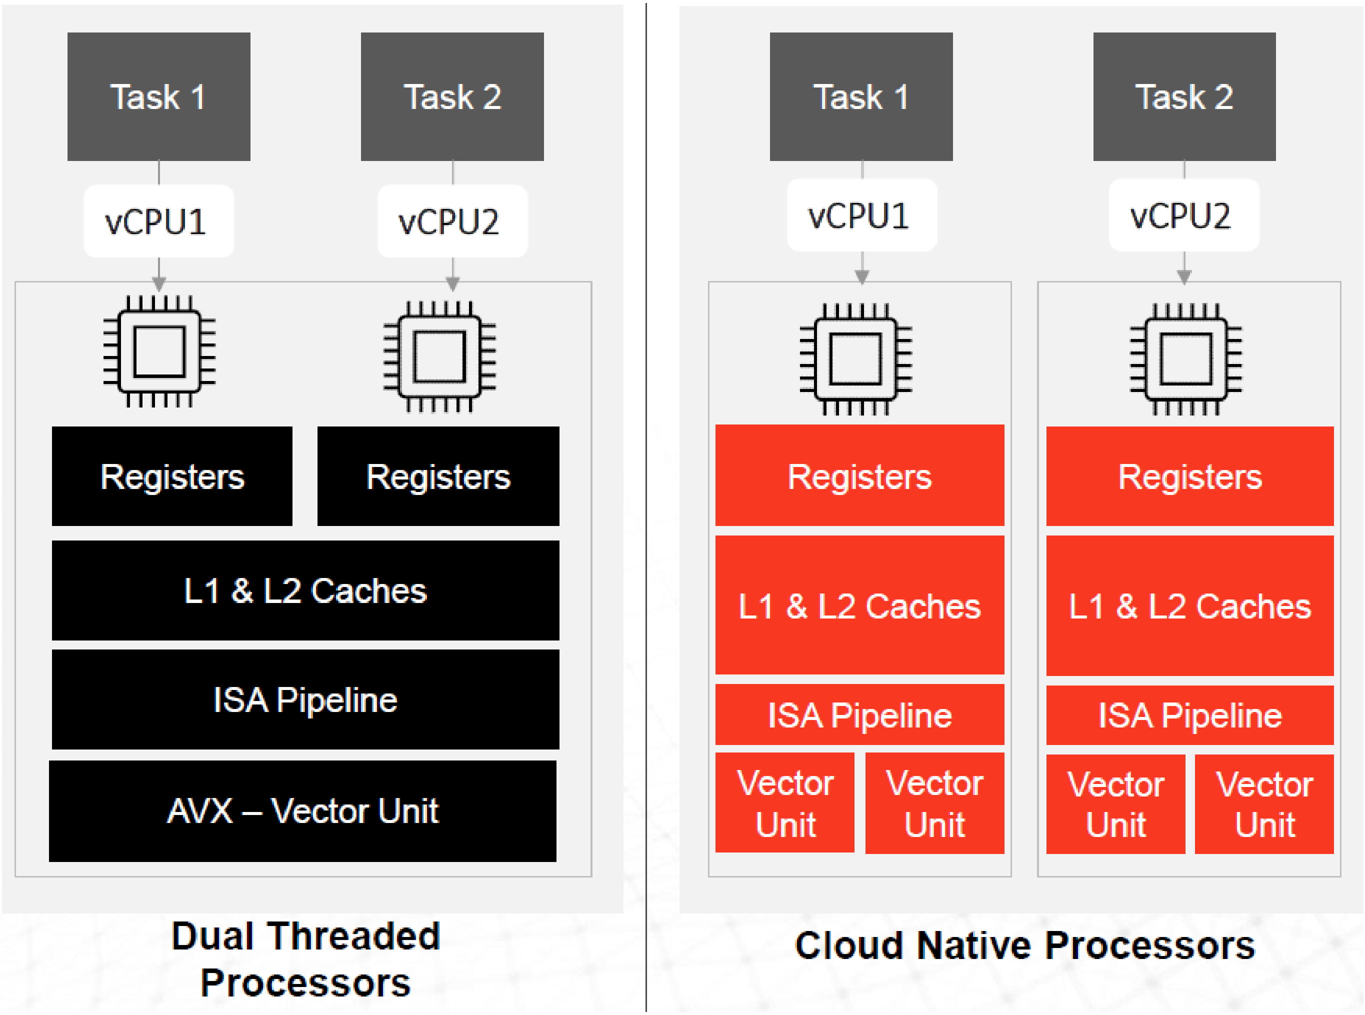

<u>For all workloads</u>

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

Domain Specific Architectures: a New Cambrian Explosion

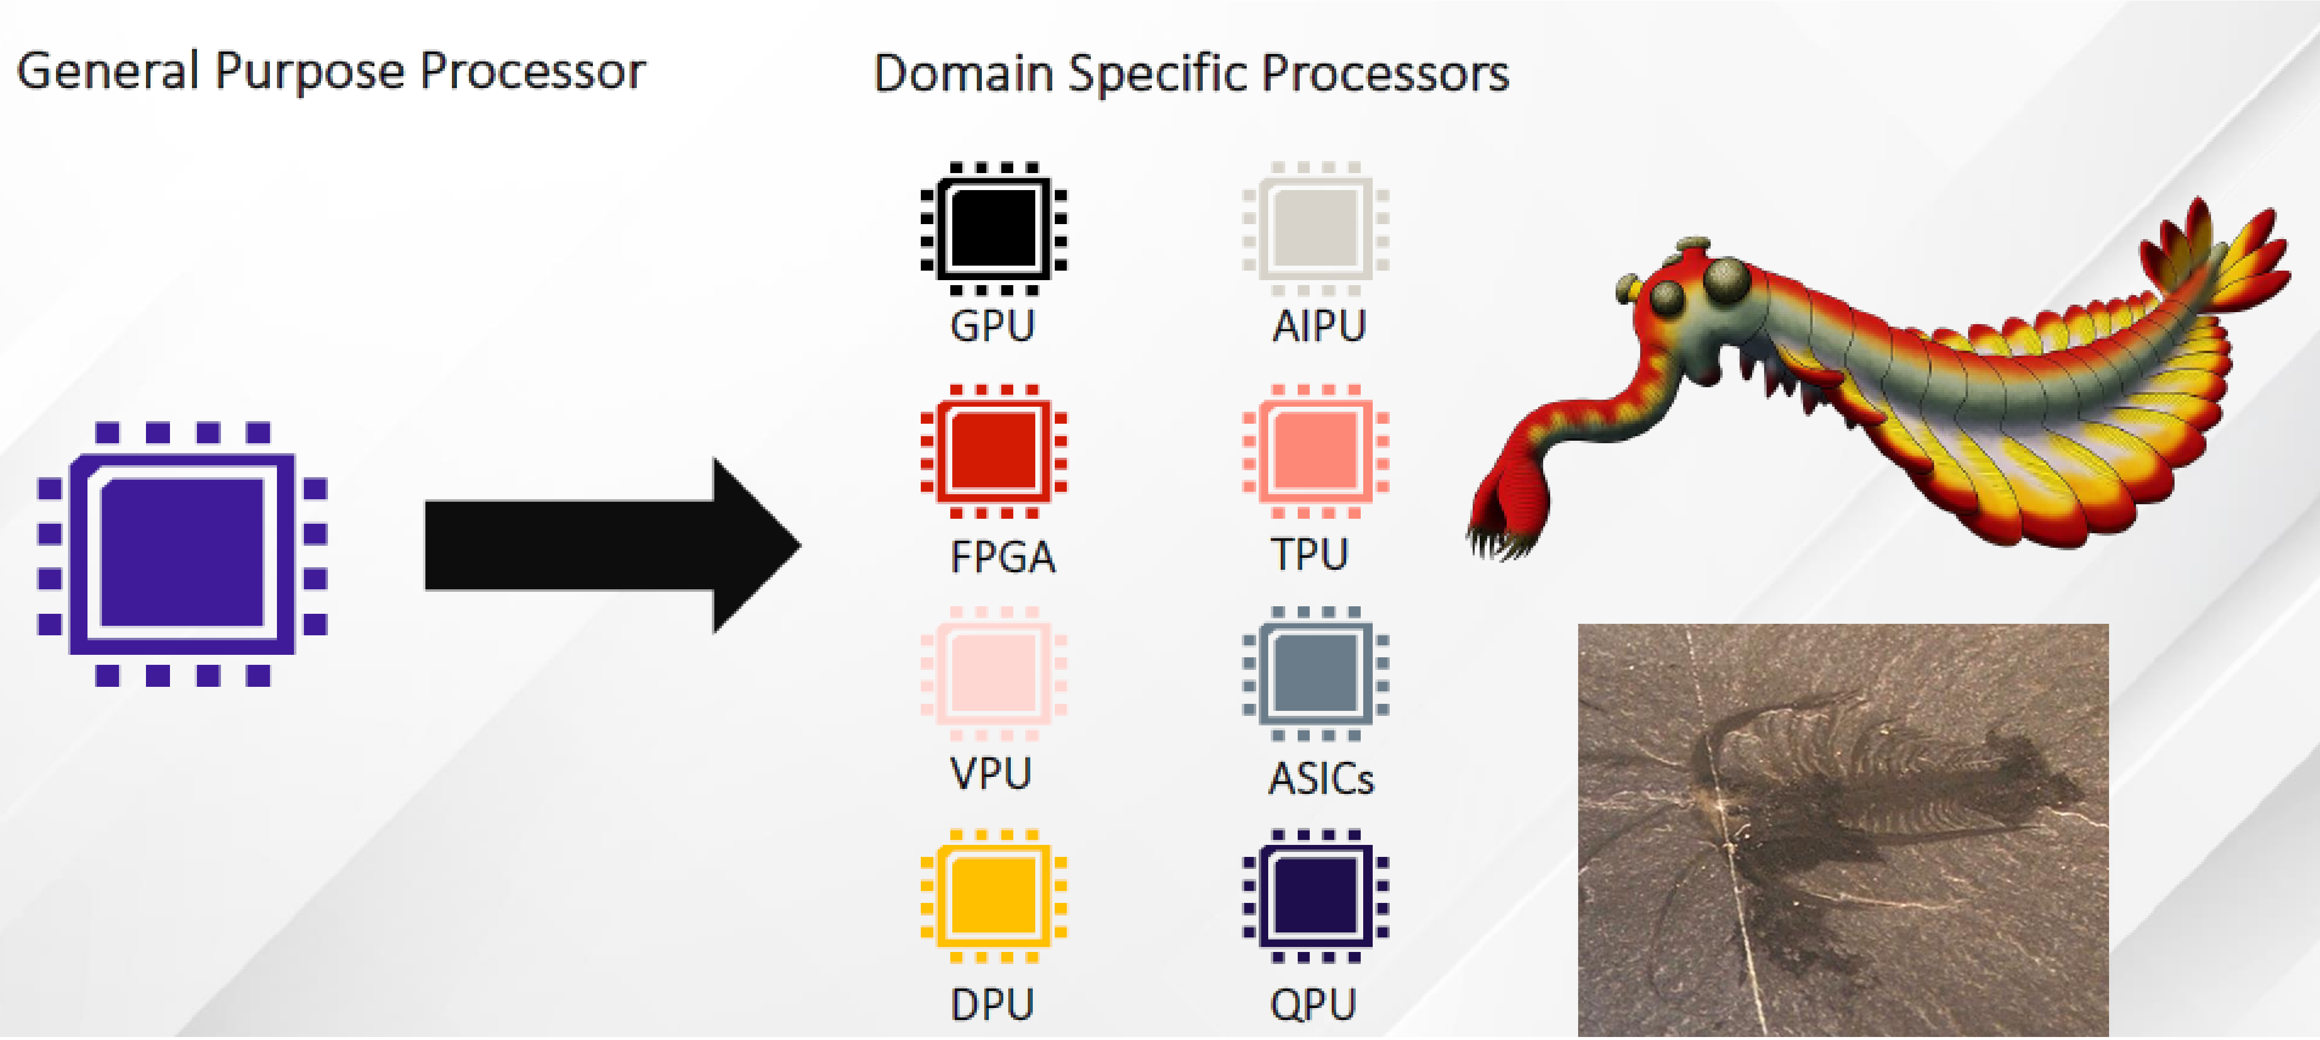

Architecture-aware optimizations in hot sections of our code can yield huge gains overall

Software approaches beyond evolutionary baseline: HLT

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 mergesort**
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, 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 advanced 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,

LHCb: GPUs at HLT1 enabled the inclusion of many new algorithms previously impossible at this level.

For a 30 MHz HLT1 at LHCb

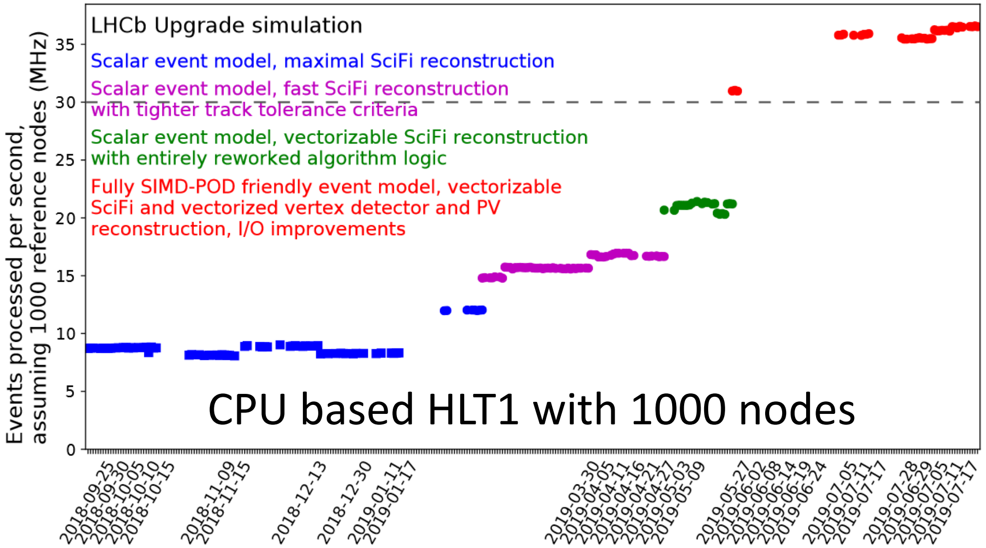

- In practice mounted server's CPUs:

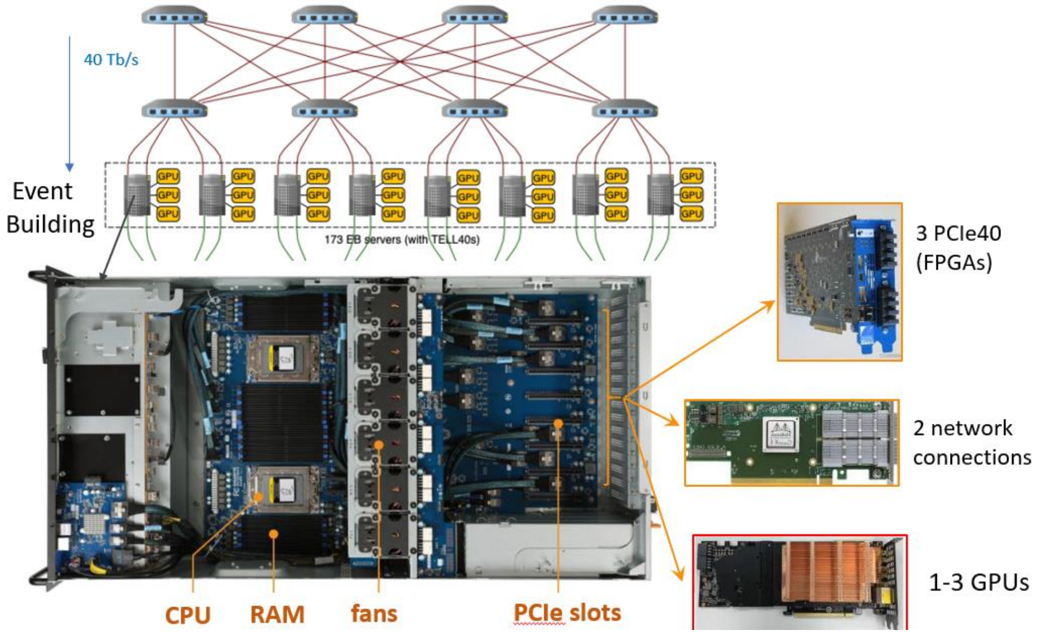

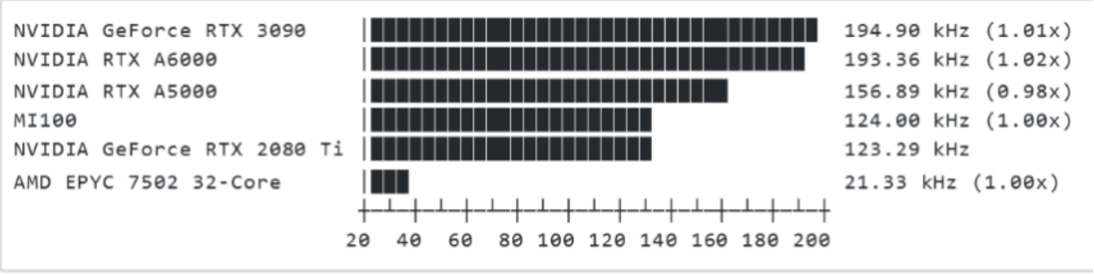

- Run3+4 architecture (Hyper-converged i.e DAQ, EB and HLT1 all together) successfully kept costs at minimum

GPU based HLT1 for Run3 with O(200) GPUs

Physics impact of *Downstream: LHCb as lifetime frontier experiment*

2312.14016

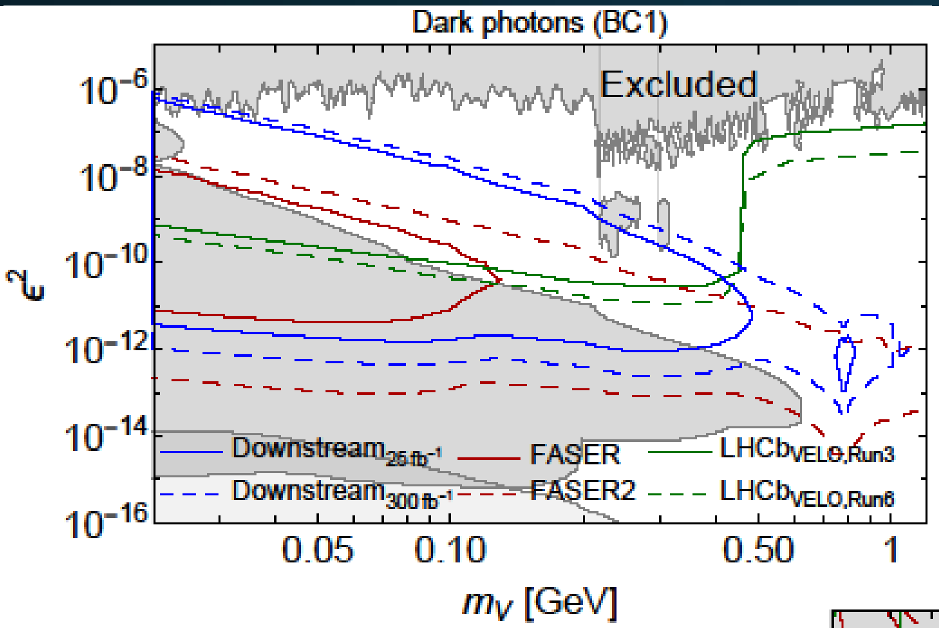

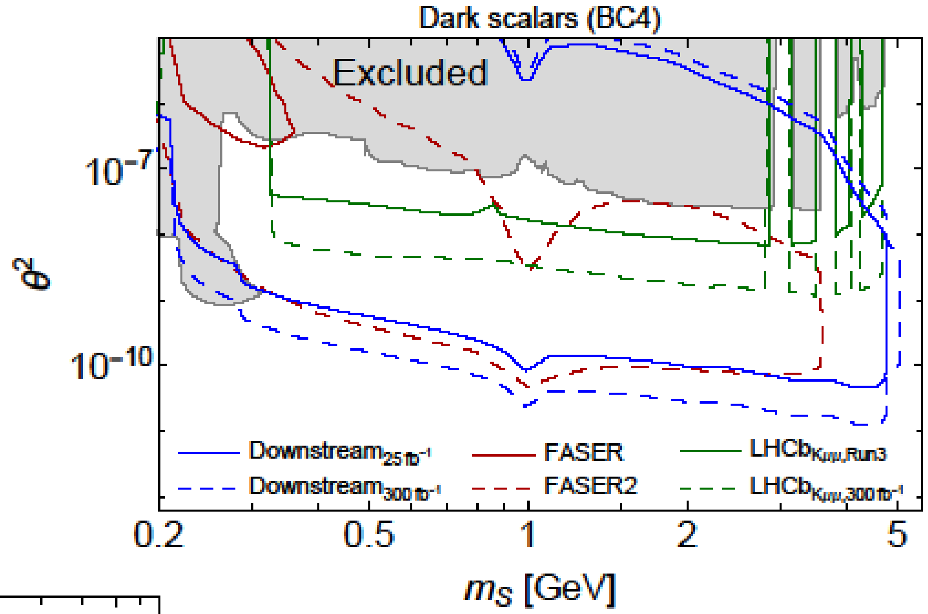

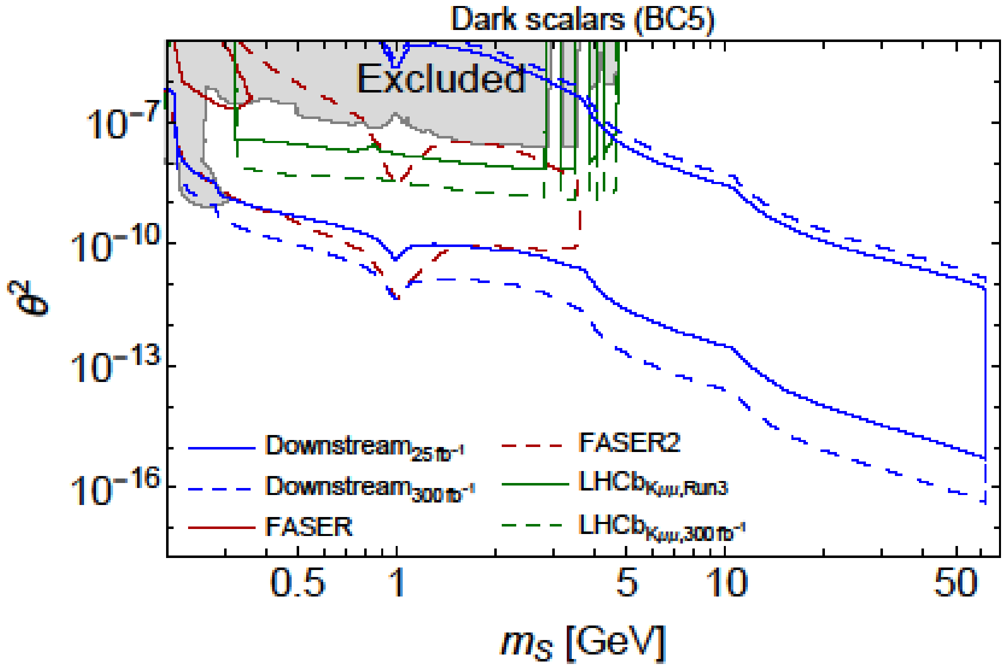

Summary

- Architecture-aware programming is a must.
- In particular, memory will play a major role to designing our future algorithms
- Architecture choice will be driven by throughput.
  - At High Level Trigger level, Performance over portability.
  - For offline computing, portability becomes crucial.
- FPGAs are key players close to the interaction point for trigger primitives.
- Measuring Power Usage Effectiveness (PUE) for our Tier centre datacentres is the first step for making WLCG greener.
- For HL-LHC and beyond, will it still be bare metal or will it all be Kubernetes orchestrated containers and services ?

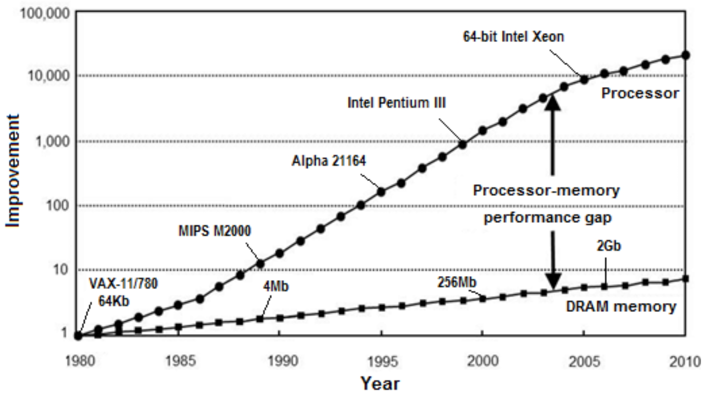

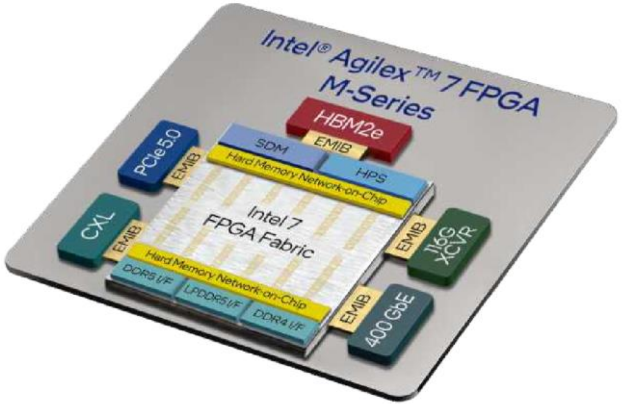

Thank you