Prof. Simon McIntosh-Smith

PI Isambard Tier-2 HPC service

Head of the HPC research group

University of Bristol

The evolution of computer architecture and its implications for meshing

Challenge 1: the slowing of Moore's Law

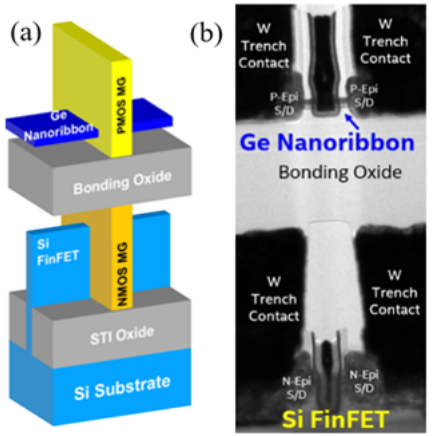

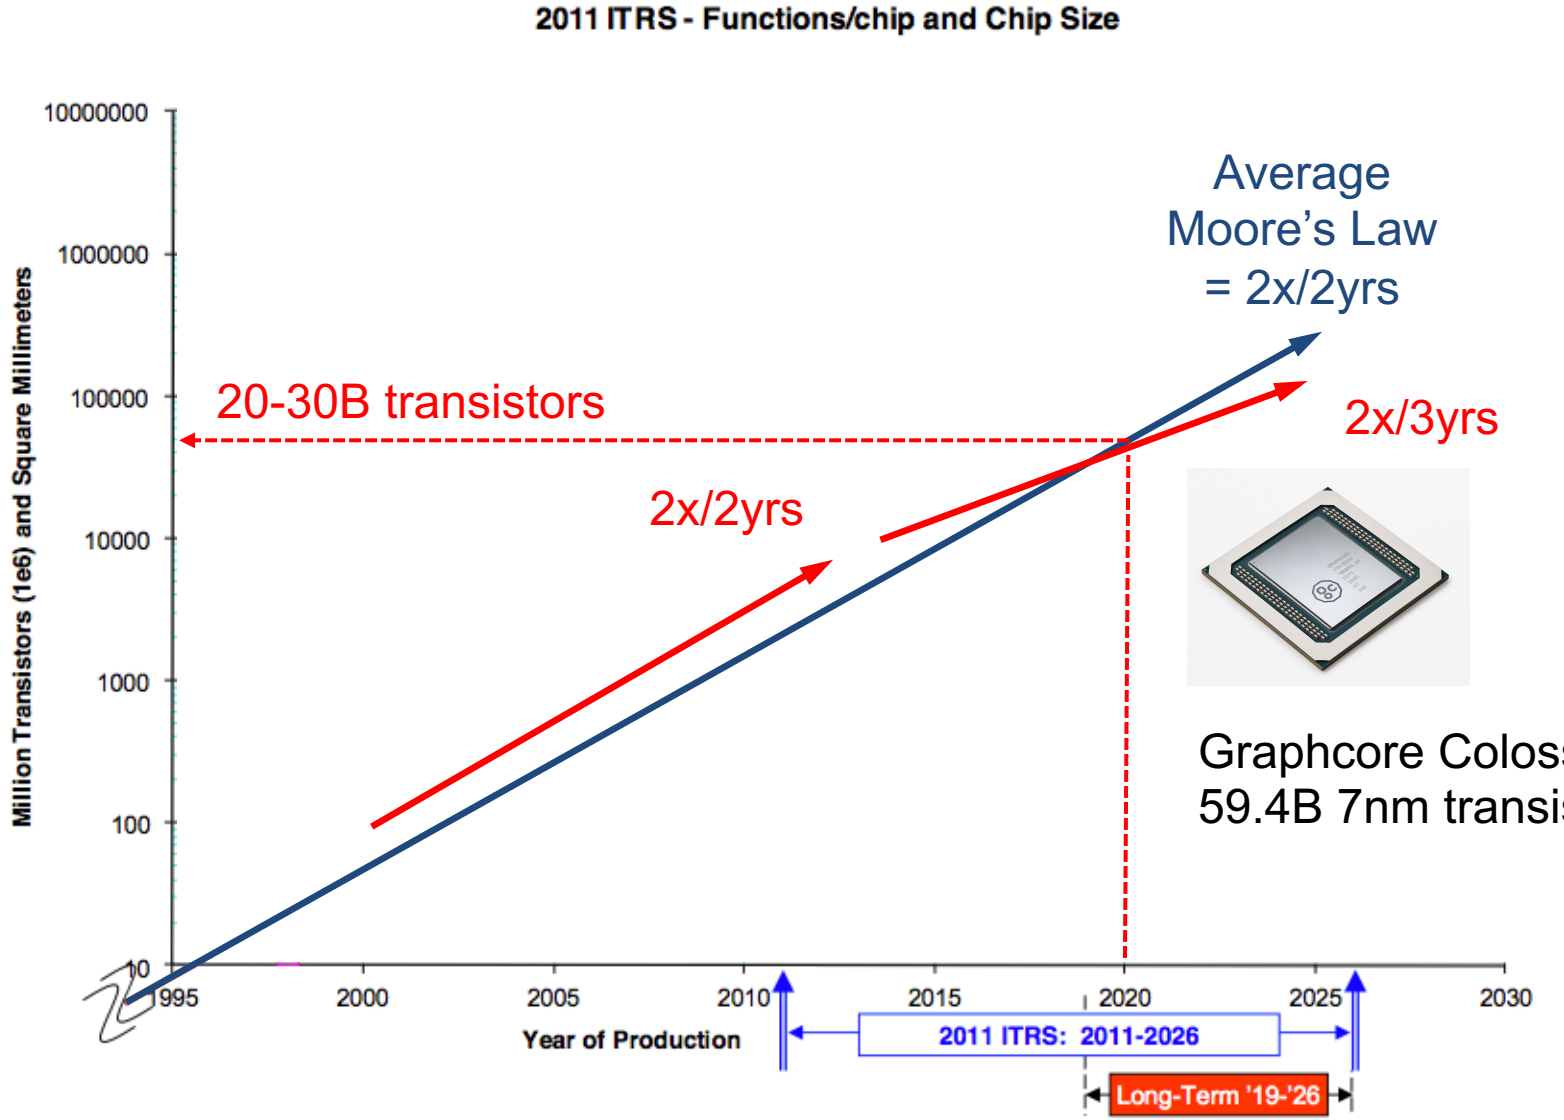

Graphcore Colossus MK2 GC200 IPU
59.4B 7nm transistors in 823mm^2

http://www.itrs.net/Links/2011ITRS/2011Chapters/2011ExecSum.pdf

Challenge 2: the end of Dennard Scaling ($P \propto V^2$)

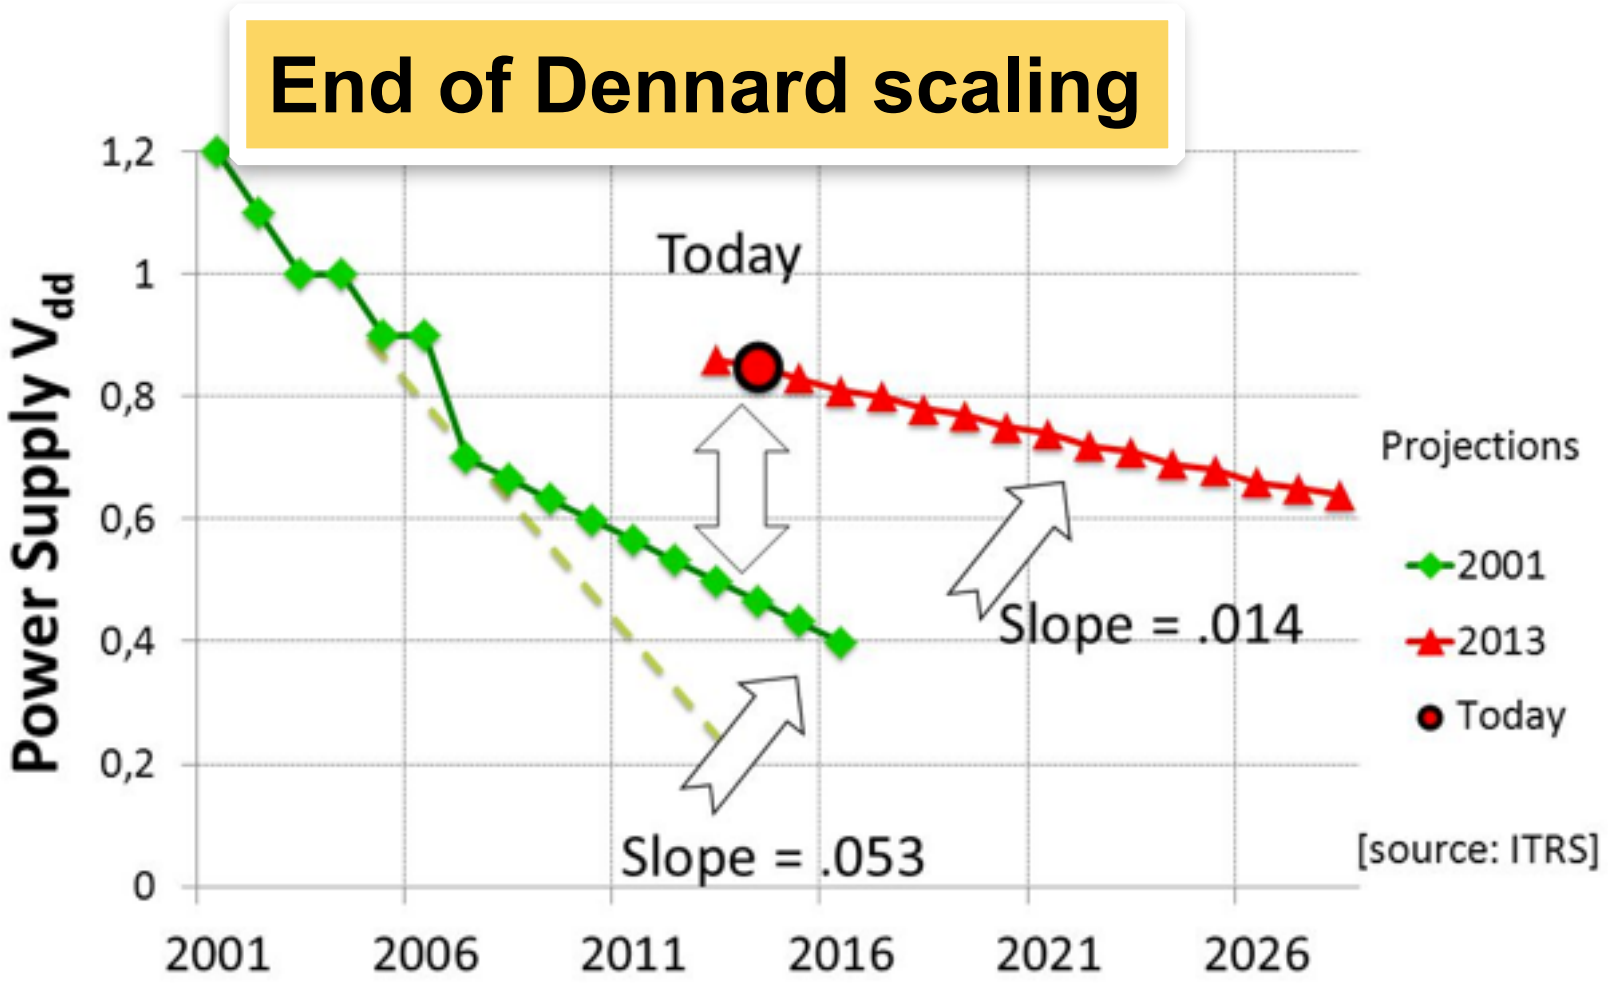

University of BRISTOL

https://sites.cs.ucsb.edu/~zhaoxia/papers/Defense_slides.pdf

ASIMOV

Challenge 2: the end of Dennard Scaling ($P \propto V^{2)}$

Electrical Properties of Transistors

YEAR OF PRODUCTION	2015	2017	2019	2021	2024	2027	2030
Logic device technology naming	P70M56	P48M36	P42M24	P32M20	P24M12G1	P24M12G2	P24M12G3
Logic industry "Node Range" Labeling (nm)	"16/14"	"11/10"	"8/7"	"6/5"	"4/3"	"3/2.5"	"2/1.5"
Logic device structure options	FinFET FDSOI	FinFET FDSOI	FinFET LGAA	FinFET LGAA VGAA	VGAA, M3D	VGAA, M3D	VGAA, M3D
DEVICE ELECTRICAL SPECS							
Power Supply Voltage - Vdd (V)	0.80	0.75	0.70	0.65	0.55	0.45	0.40

INTERNATIONAL TECHNOLOGY ROADMAP FOR SEMICONDUCTORS 2.0, 2015

Challenge 3: the exponential growth of parallelism

- The coming generation of Exascale supercomputers will contain a diverse range of architectures at massive scale
  - **Fugaku**: Fujitsu A64fx Arm CPUs
  - **Perlmutter**: AMD EYPC CPUs and NVIDIA GPUs
  - **Frontier**: AMD EPYC CPUs and Radeon GPUs
  - **Aurora**: Intel Xeon CPUs and Xe GPUs
  - **El Capitan**: AMD EPYC CPUs and Radeon GPUs

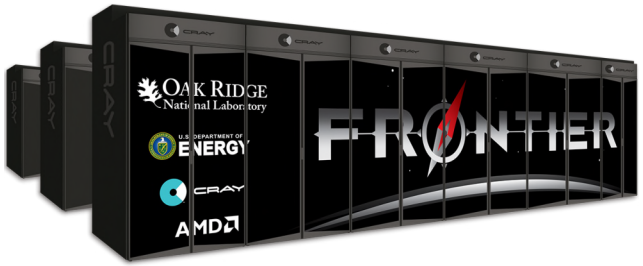

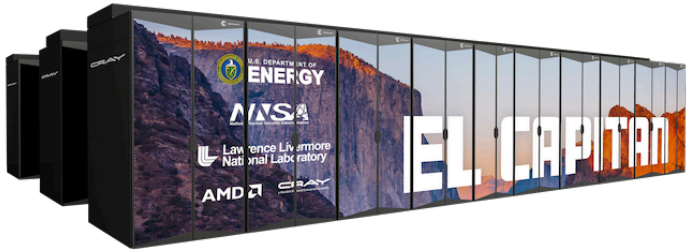

University of BRISTOL

The Next Platform, Jan 13th 2020: "HPC in 2020: compute engine diversity gets real"
https://www.nextplatform.com/2020/01/13/hpc-in-2020-compute-engine-diversity-gets-real/

2017 Turing Award lectures from Hennessy and Patterson

John L. Hennessy and David A. Patterson won the 2017 A.M.Turing Award at the 45th International Symposium on Computer Architecture (ISCA) in Los Angeles:

"A New Golden Age for Computer Architecture: Domain-Specific Hardware/Software Co-Design, Enhanced Security, Open Instruction Sets, and Agile Chip Development"

University of BRISTOL

https://iscaconf.org/isca2018/turing_lecture.html

ASIMOV

Opportunity: post-Dennard improvements via specialisation

- Processors in the mobile space have been doing this for years, e.g. Apple A12 in iPhone

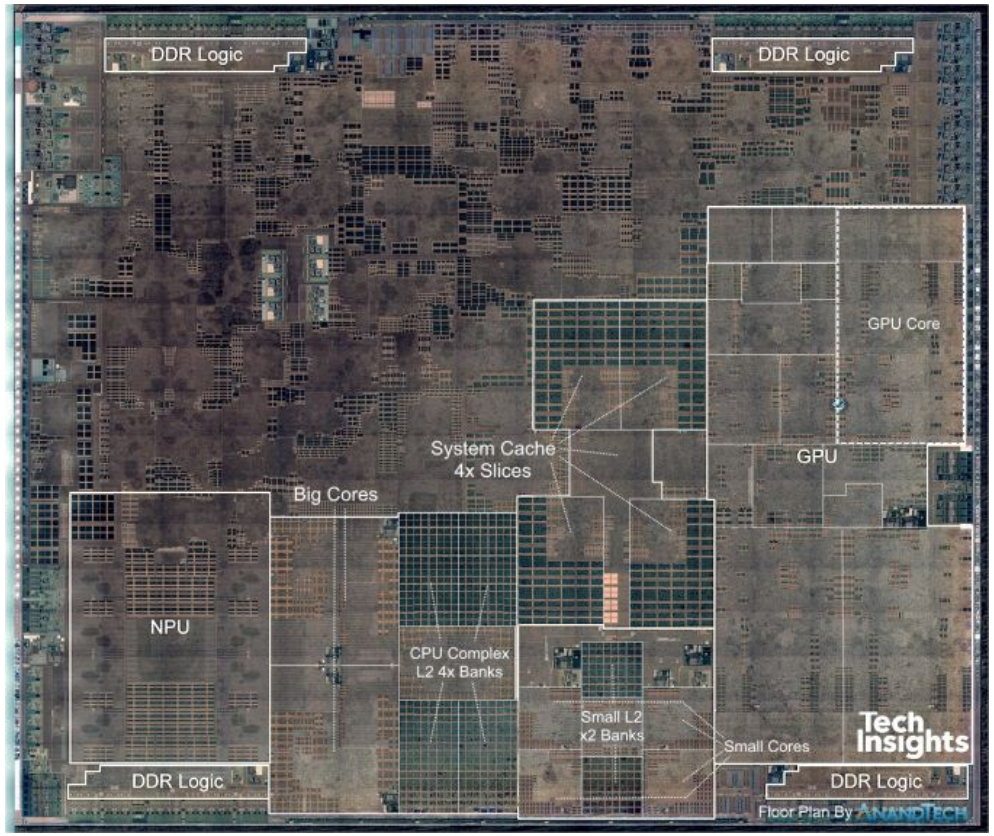

Source: Tech Insights

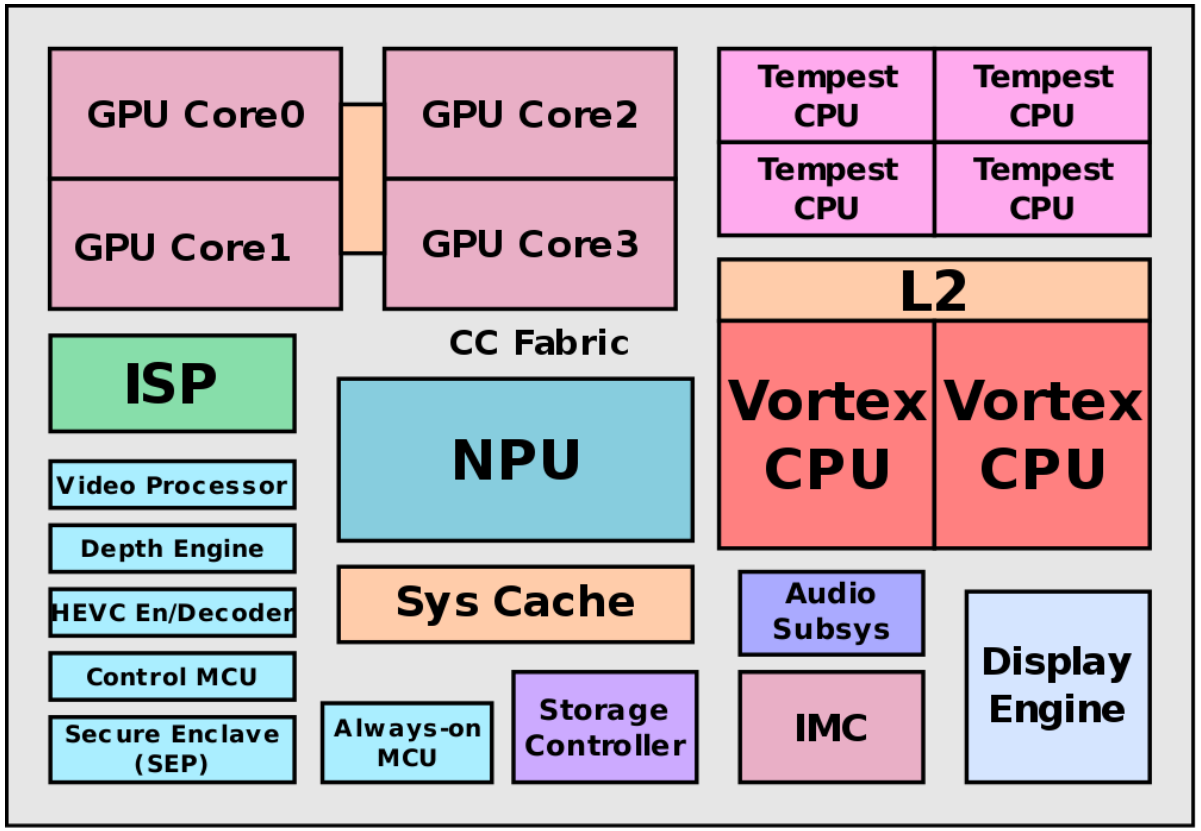

https://en.wikichip.org/wiki/apple/ax/a12

Increased heterogeneity is an important response to the slowing of Moore's Law

- Recent example: TensorCores

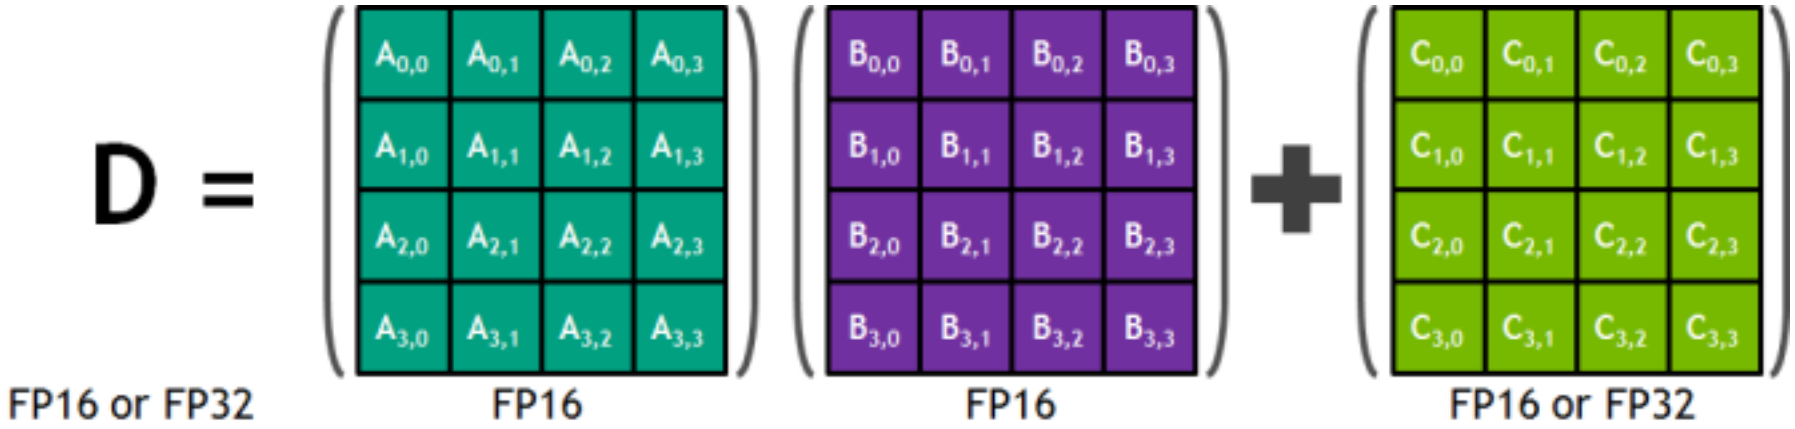

- Expect to see more application-oriented optimisations
  - Matrix multiply units in SIMD instruction sets (Arm's SME, Intel's AMX)
  - Floating point formats optimized for deep learning (BFLOAT16)

University of BRISTOL

ASIMOV

Ray Tracing cores

- NVIDIA's latest architectural innovation (Turing-class GPUs)
- Designed to accelerate the ray tracing algorithms used in graphical rendering in games, rather than for HPC
- Potential speedups of up to 10X vs CUDA code on same GPU
- Accelerates ray / surface intersection calculations
  - 10 GigaRays/s on RT cores vs 1-2 GigaRays/s in CUDA on the same GPU

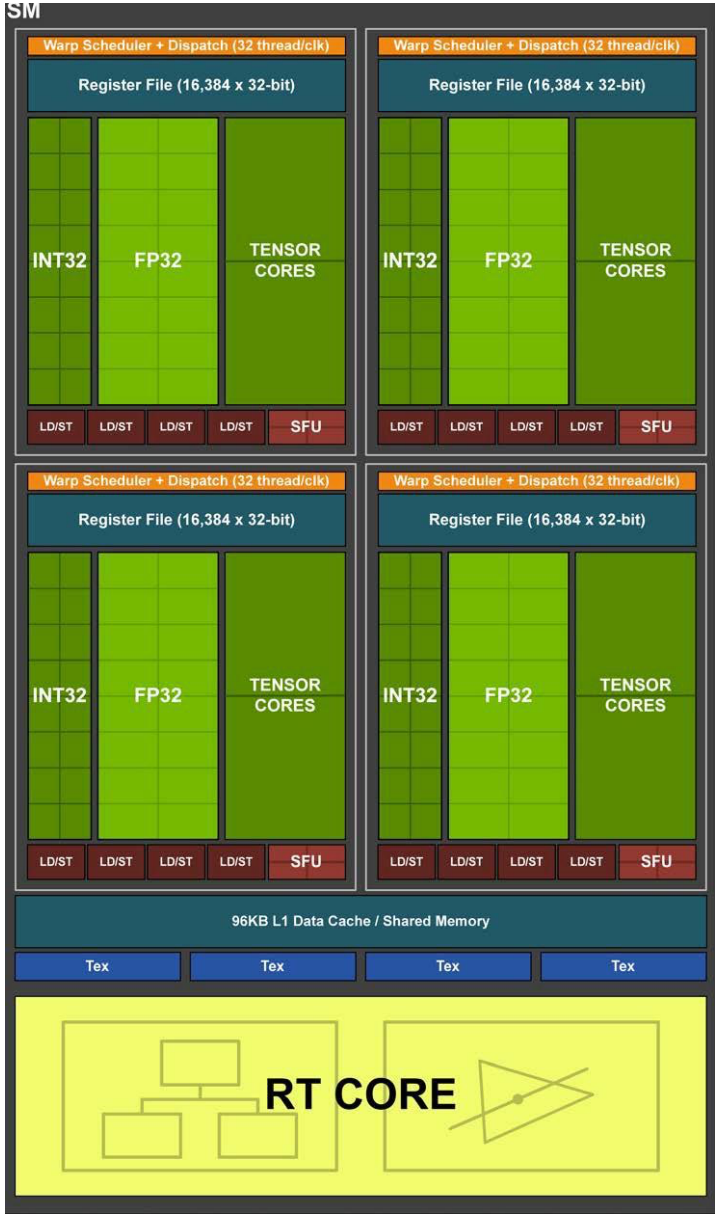

Observation: Monte Carlo particle transport has similarities to RT

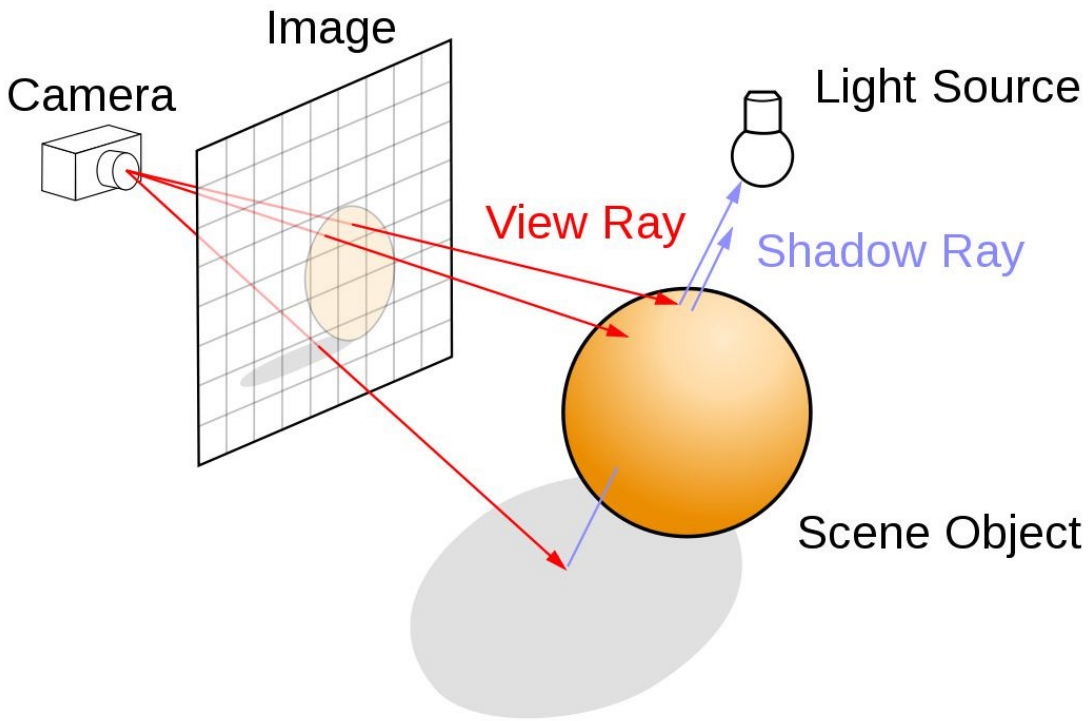

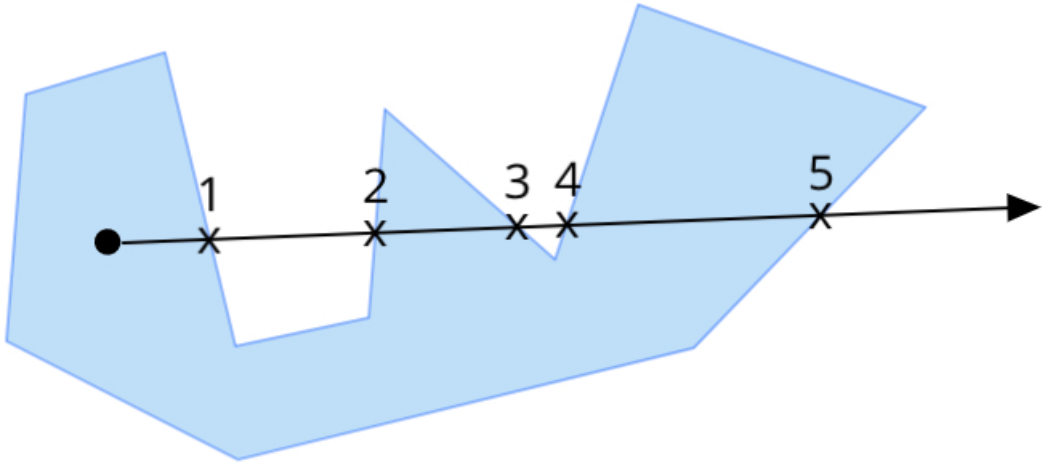

Both require large numbers of linear geometric queries
to be executed over complex 3D geometric models

Opportunity: application optimized architectures

- **AI / deep learning** (Graphcore, Google et al)
- **BLAS / linear algebra** (tensor core-like accelerators)
- **Encryption / compression**
- **Ray-tracing** (ray-surface intersections, particle tracking)

- What about accelerating esoteric memory access patterns? Sparse matrix operations? Random number generation? …?
- **Challenge: hard to target via software!**

Vector gather / scatter

```
const int loc  = some_strange_func(id);

float val = mem[loc];
```

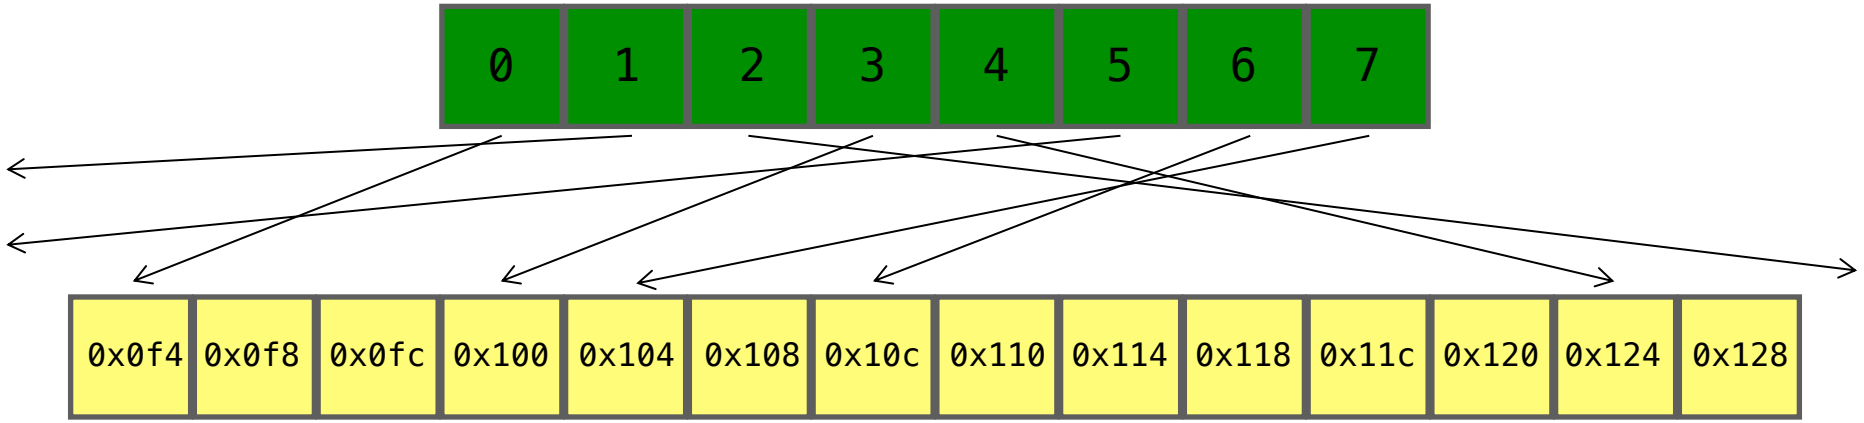

University of BRISTOL

ASIMOV

Some key issues for meshing

1. **Deep, complex memory hierarchies**
   - Meshing typically heavily memory bandwidth bound
   - HBM can offer an order of magnitude more bandwidth (TB/s vs 100s of GB/s, but for order of magnitude smaller memories (10s of GB vs 100s of GB)
2. **Massive parallelism**
   - Instruction, Data, Thread, Core, Socket, Node, …
3. **Performance portability**
   - Want to be able to mesh on diverse systems: CPU only, CPU+GPU, dense GPU systems, …

A quick peek into the future of memory

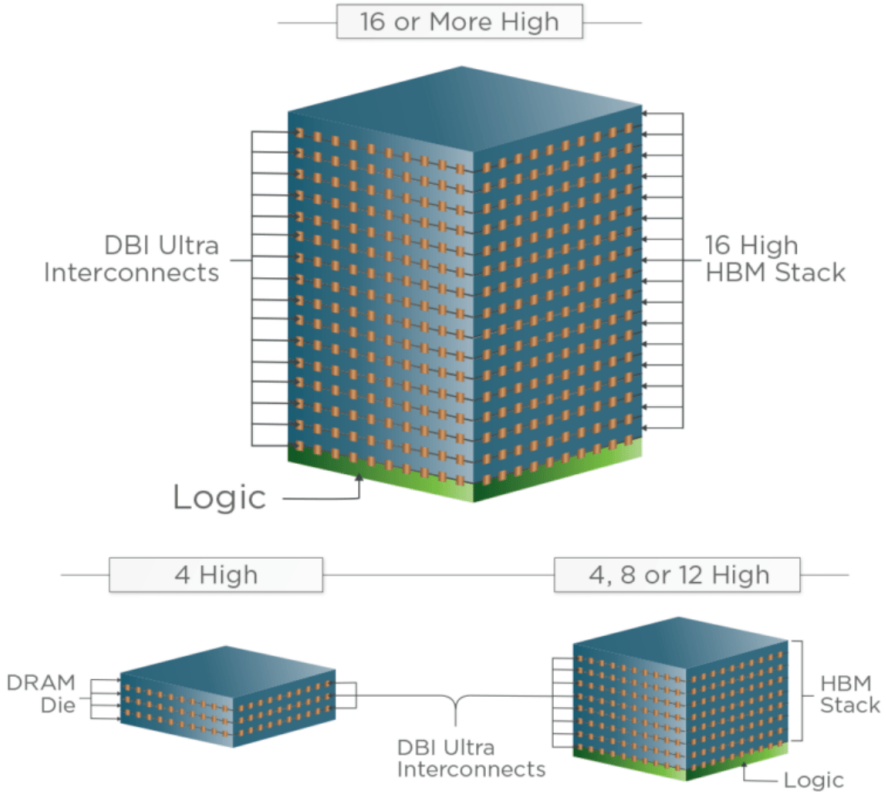

https://semiengineering.com/whats-next-for-high-bandwidth-memory/

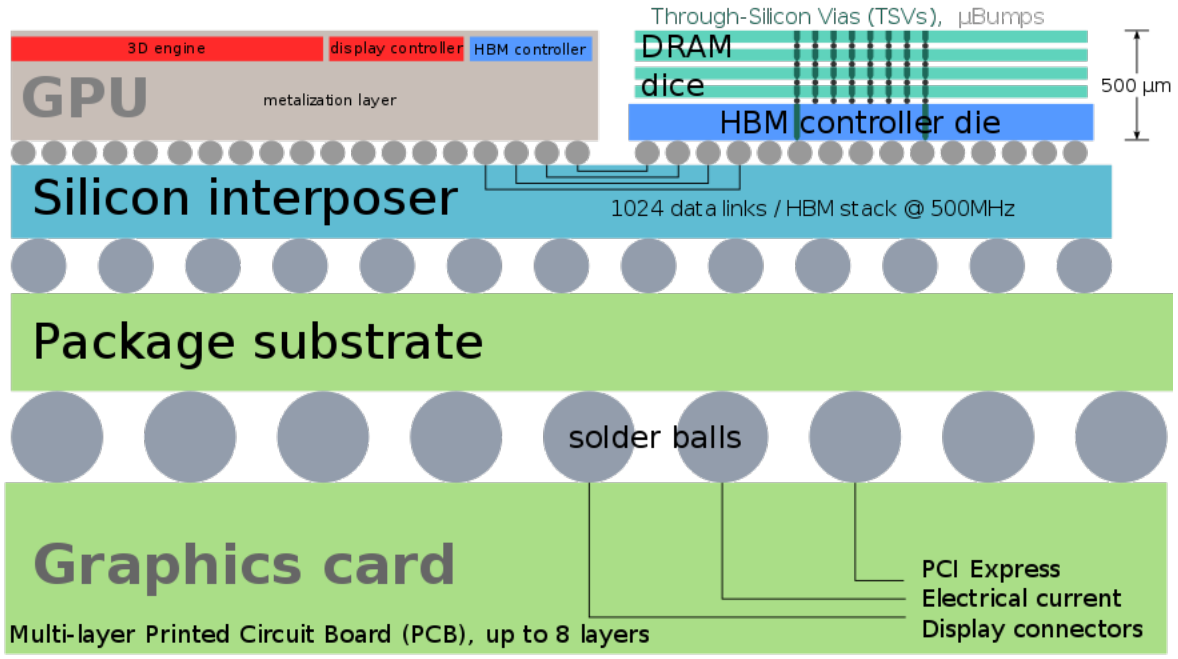

https://en.wikipedia.org/wiki/High_Bandwidth_Memory

High Bandwidth Memory (HBM)

- Already widely used in high-end graphics cards, and now in A64fx
- 90% of energy consumed by memory is used to transfer data
  - So moving things closer together can save a lot of energy
- Today's **HBM2** comes in 4-8GB capacities. 1,024 I/O pins at 2.4Gbps, equating to **307 GB/s** of bandwidth.
  - A64fx has 4 of these, providing 32GB at 1.2TB/s peak theoretical
- **HBM2E** coming soon, 8-16GB capacities. 3.2Gbps transfer rates, → **410 GB/s** bandwidth
- **HBM3** brings 4Gbps transfer rates with **512 GB/s** bandwidth. "HBM3 will be released in 2H of 2020. After HBM3, there is no concrete roadmap yet." – Jeongdong Choe, TechInsights, Dec 2019.
- **HBM *doesn't* improve latency, *only* bandwidth**
  - Might be a crucial issue for meshing

https://semiengineering.com/whats-next-for-high-bandwidth-memory

Hypothetical meshing processor

Extrapolated from various processors today (e.g. A64fx)

- **64 fast cores**, 2-3GHz each, superscalar out-of-order
- **64 GBytes** of HBM3, 4 stacks, **2 TB/s peak bandwidth**
- 8 channels of DDR5 DRAM, each at 6.4 Gbps → 409 GB/s per socket, **512 – 1024 GB per node**
- Fast interconnect for parallel meshing
- Fast in-network storage (e.g. burst buffer)

Key takeaways

- Meshing is hard and getting harder
- Hardware is becoming much more parallel, with much deeper, complex memory hierarchies
- New memories are going to be an order of magnitude **faster** and **smaller** than we're used to with DDR DRAMs in DIMMS
- → Will potentially need radical new approaches to meshing

No more "business as usual"!

For more information

Bristol HPC group: https://uob-hpc.github.io/

Email: S.McIntosh-Smith@bristol.ac.uk

Twitter: @simonmcs