5.4 Sustainable energy consumption: Data storage

Tom Thomson

University of Manchester

ESM 2022 – lecture 19 Sept. 2022

Outline – first part 5.4

- Magnetic data storage – energy considerations

- What are the data storage requirements?
- What are the available magnetic data storage options?
- Team exercise – looking at energy efficiency

MANCHESTER 1824

The University of Manchester

Hierarchy in data storage

- The world of data storage is increasingly diverse as different optimisations point to different solutions
- Driving factors:
  - Cost
  - Cost
  - Cost
- Other driving factors
  - Malware (ransomware)
  - Business continuity
  - Regulatory compliance
  - Energy consumption

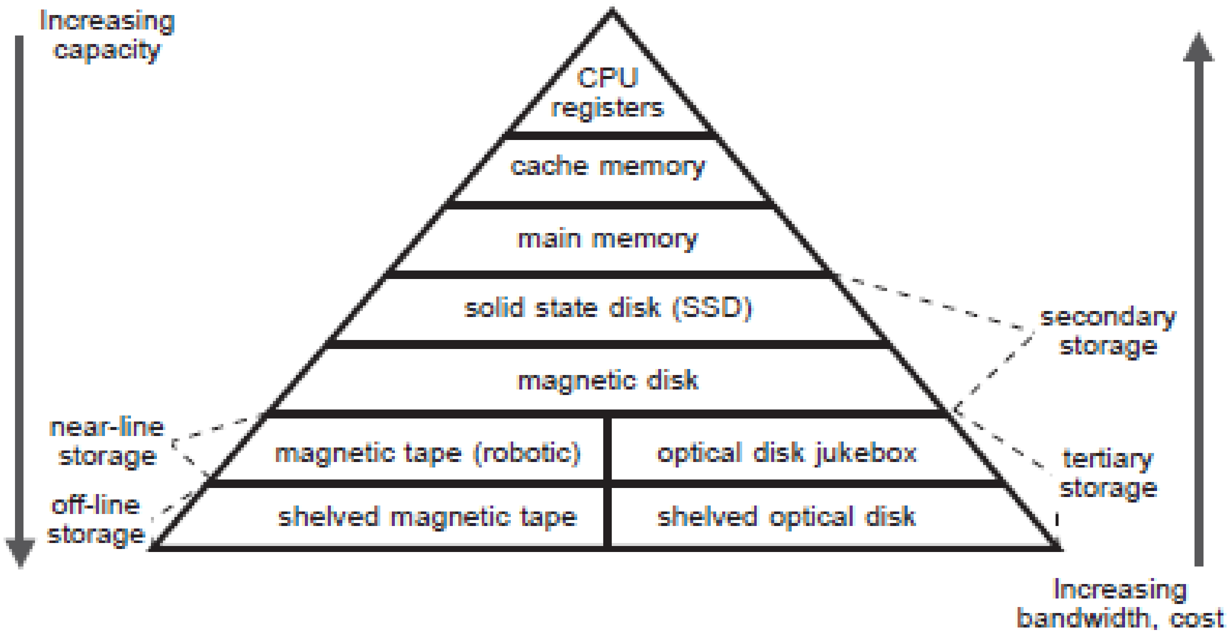

Millar PhD thesis (UC Berkeley) 1995

Data storage use cases

- Increased complexity
- Energy needs to be considered at system level
- Our research is focused on best performing devices

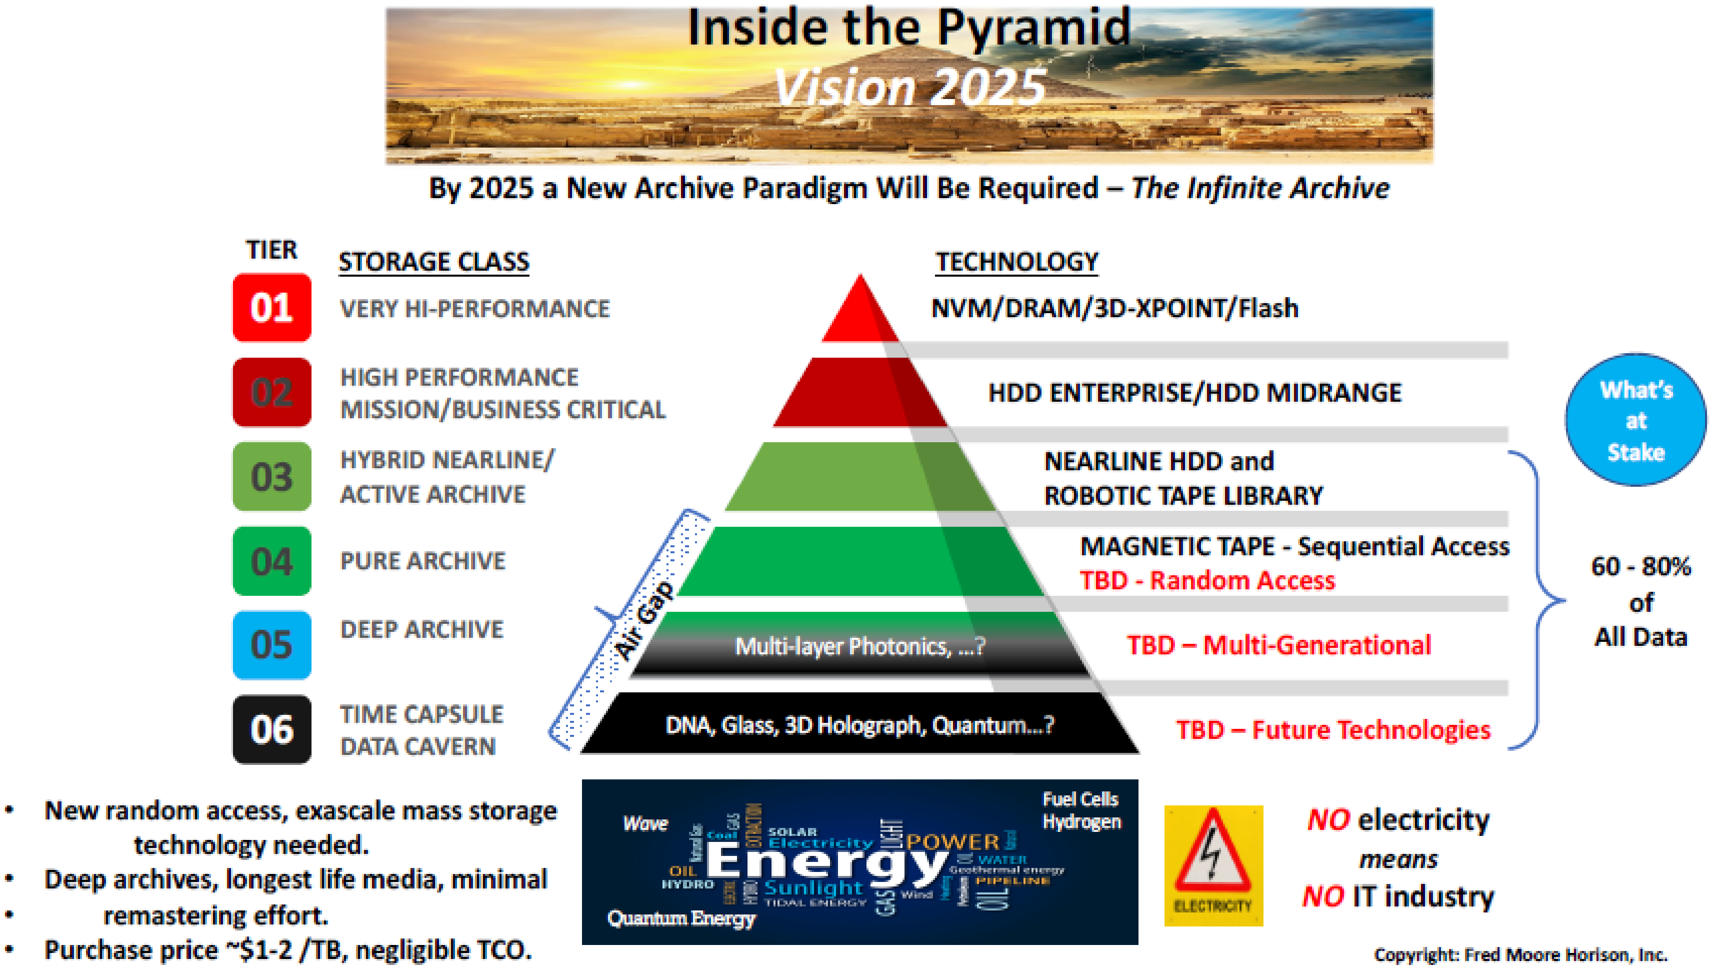

MANCHESTER 1824
The University of Manchester

A modern data center – San Jose, CA

- What is the noticeable thing?

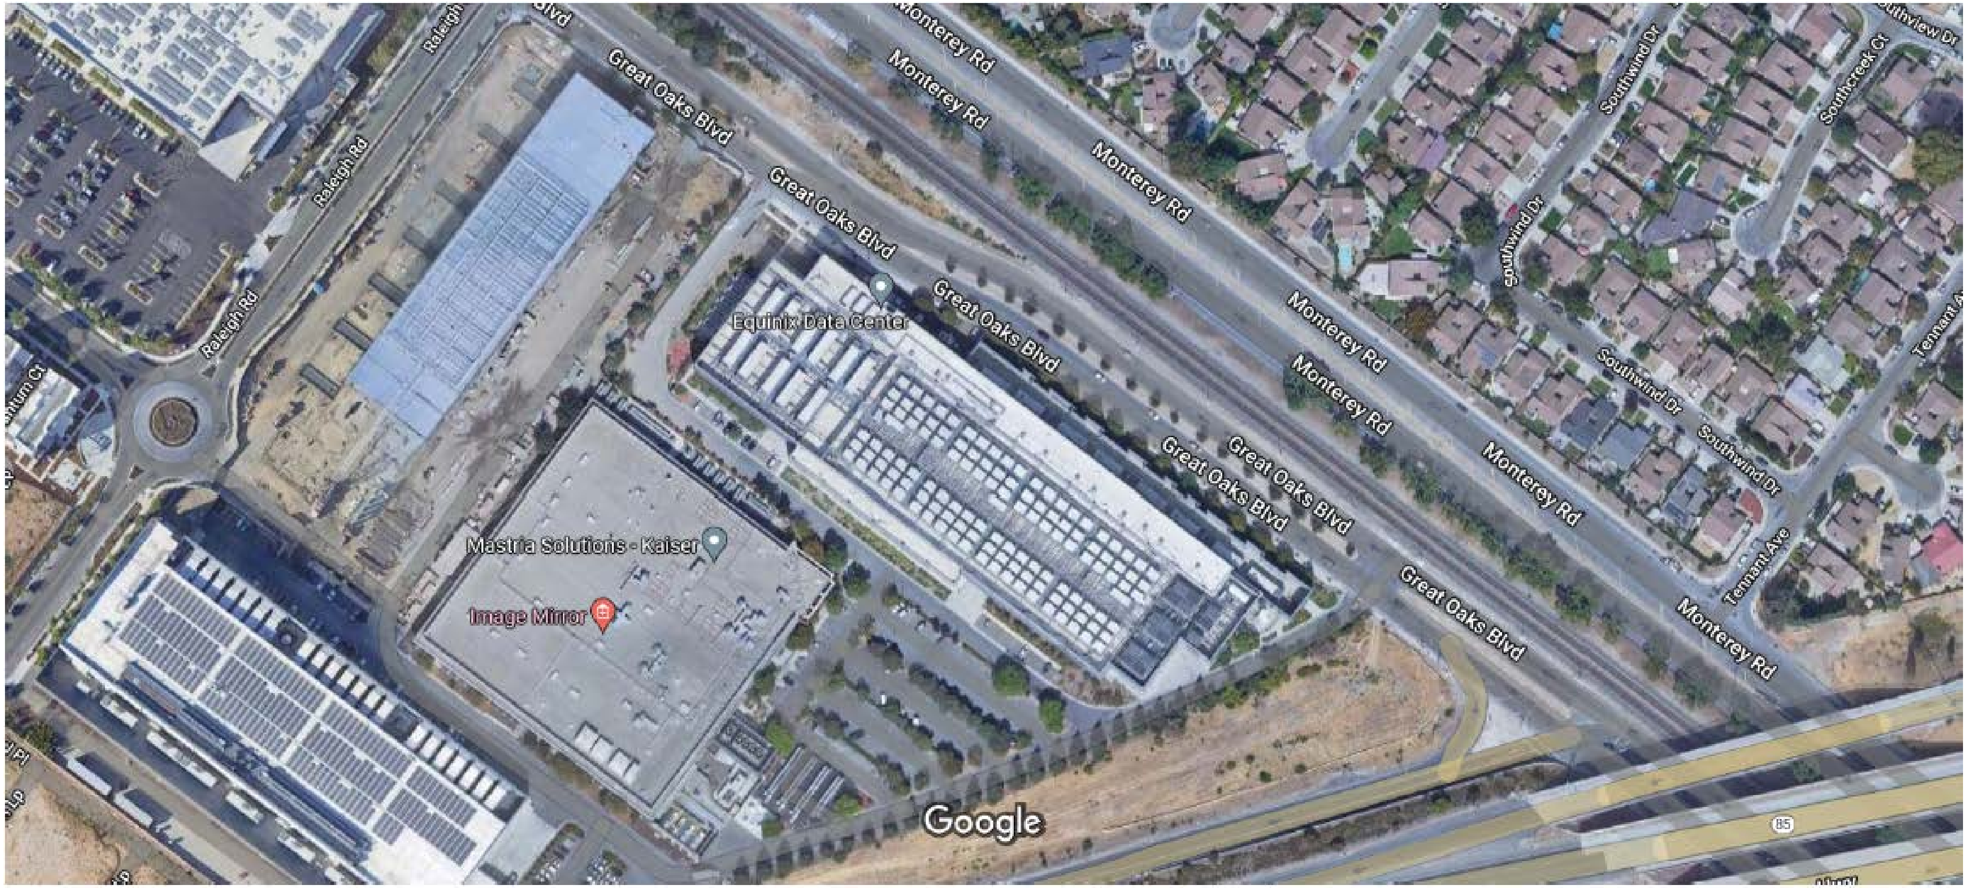

Imagery ©2022 AMBAG, CNES / Airbus, Maxar Technologies, U.S. Geological Survey, USDA/FPAC/GEO, Map data ©2022 50 m

MANCHESTER
1824
The University of Manchester

A modern data center – San Jose, CA

- Air conditioning units

- Global data centre electricity use in 2020 was 200-250 TWh, or around 1% of global final electricity demand

Imagery ©2022 AMBAG, CNES / Airbus, Maxar Technologies, U.S. Geological Survey, USDA/FPAC/GEO, Map data ©2022 50 m

The University of Manchester

Magnetic data storage

- Think of data storage as a single bit state machine – what are the energy considerations:
  - Energy to maintain a state
  - Energy to change a state (i.e. 0 -> 1)
  - Energy consumption of the system (Local & Global)
  - Total energy cost of building the device in the first place
- What has been happening to data storage energy use?

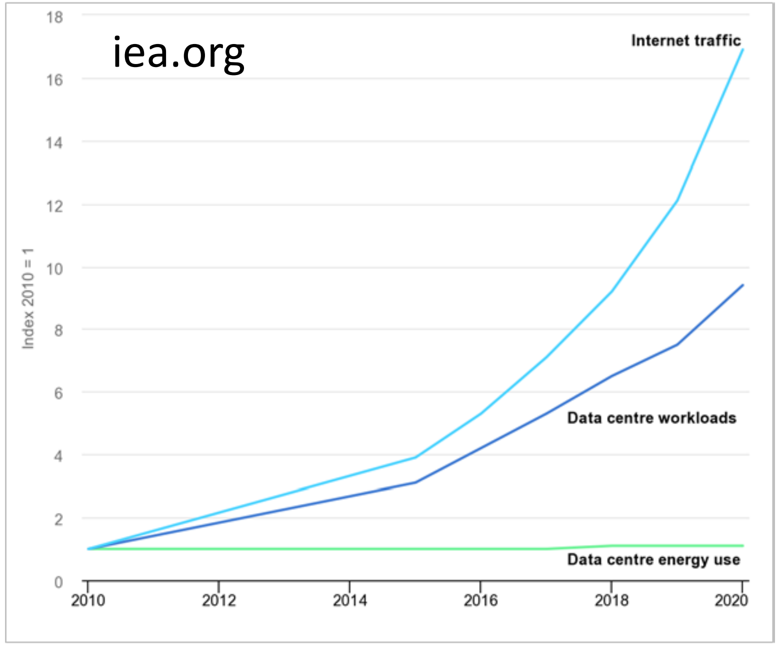

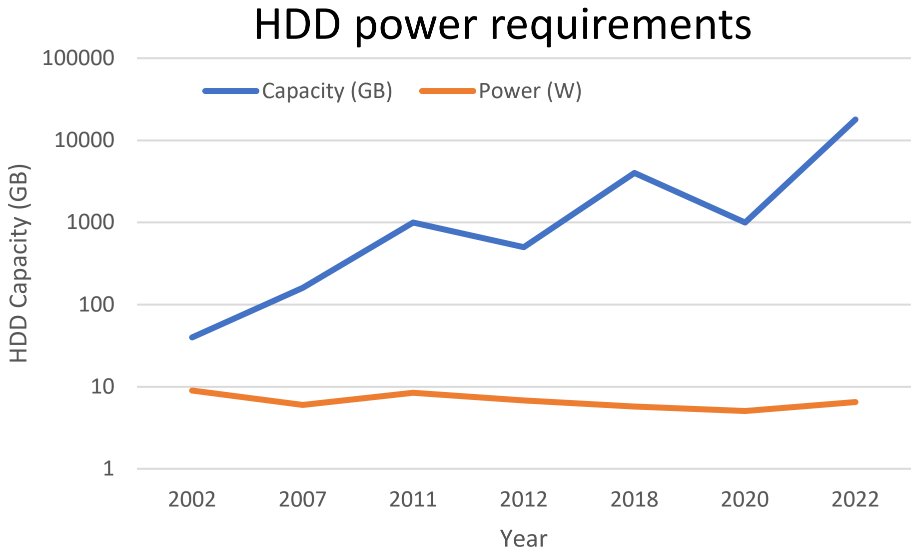

Magnetic – HDD: the current workhorse

- HDD - Simple scaling
  - Power consumption per device remains constant
  - But increase in areal density slowed to a trickle
  - Now adding more disks from 2/3 -> 9/10
  - A $20b+ industry

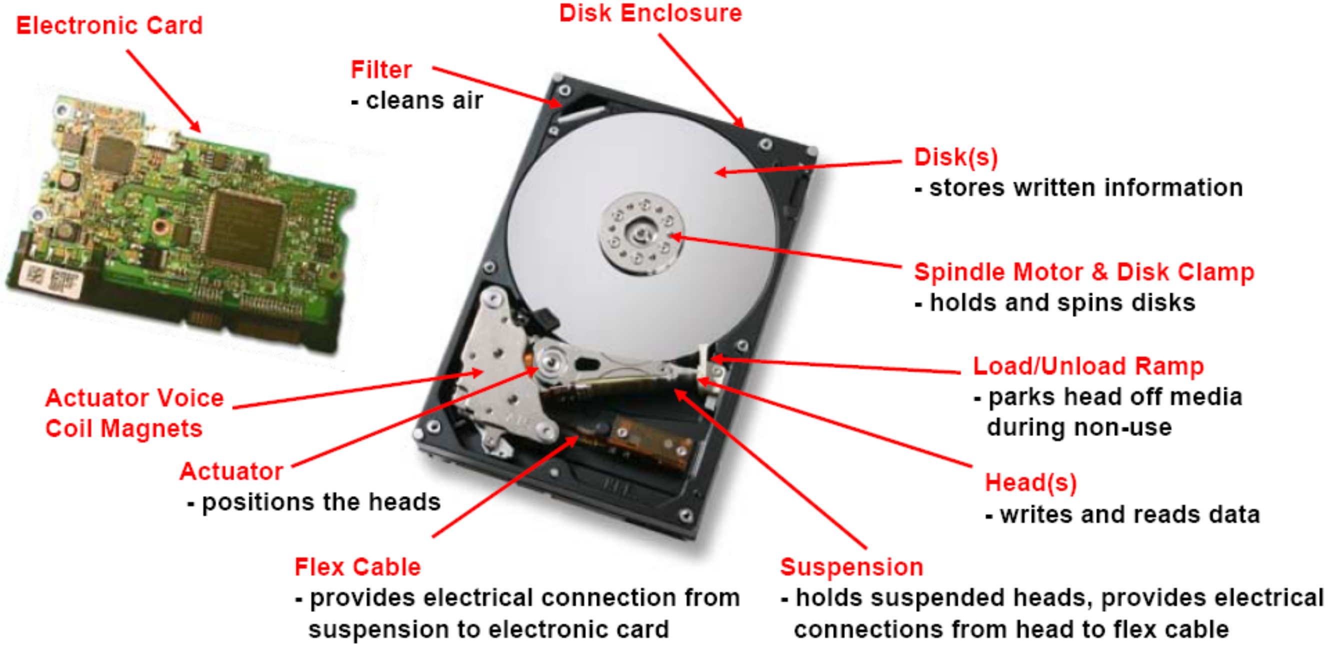

Electronic card

- Provides data interface to disk controller
- Control operation of disk drive (spindle, actuator, position servo)
- Encodes written data and decodes read back data
- Provide read/write signals to heads via flex cable

The University of Manchester

HDD – power consumption

- Electronics accounts for a surprising fraction of the power consumption
- Fine detail of control systems matter for energy consumption

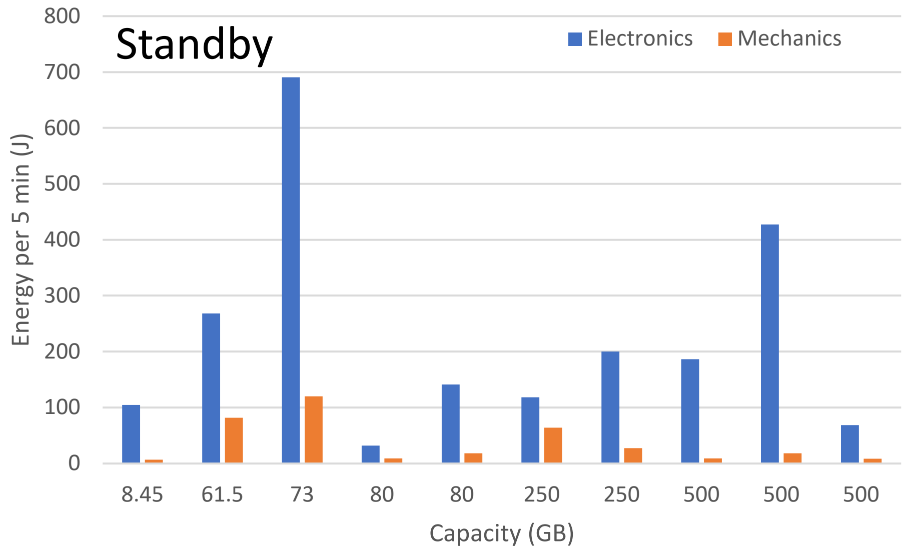

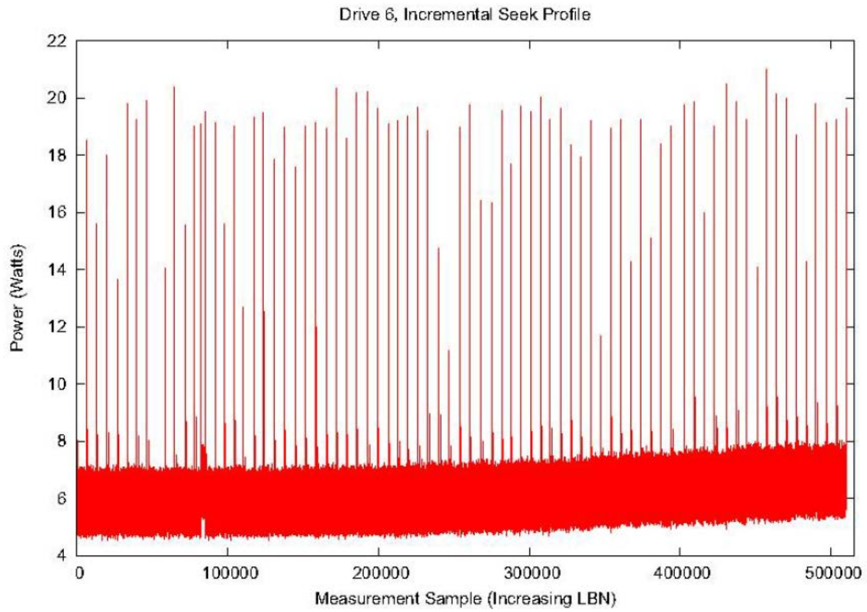

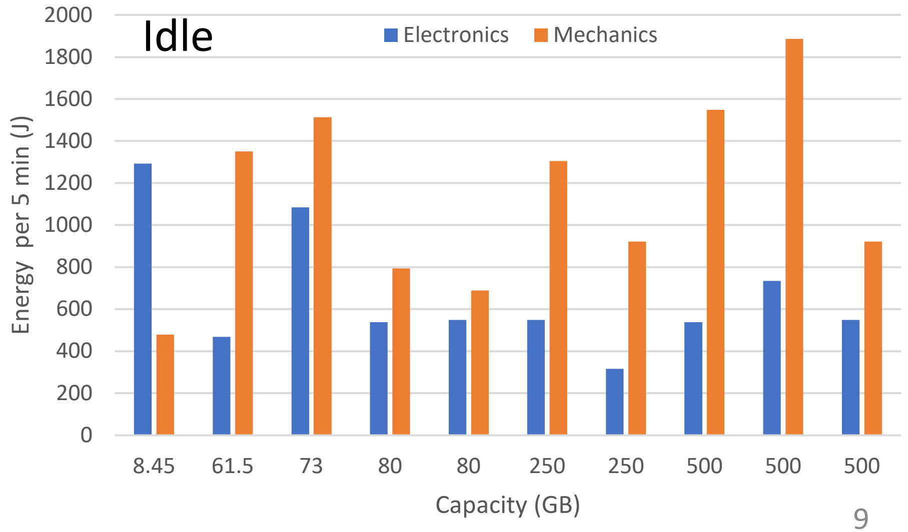

Hylick et al. ISBN 978-1-4244-281 2008 IEEE Xplore

Magnetic data storage basics

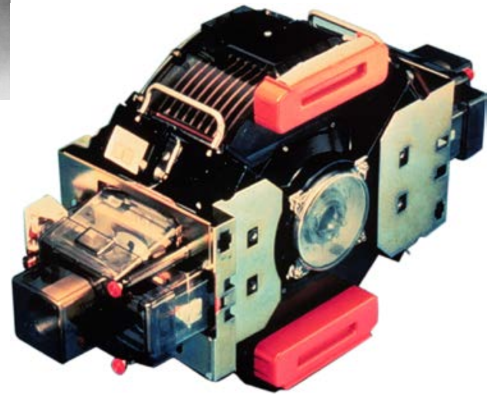

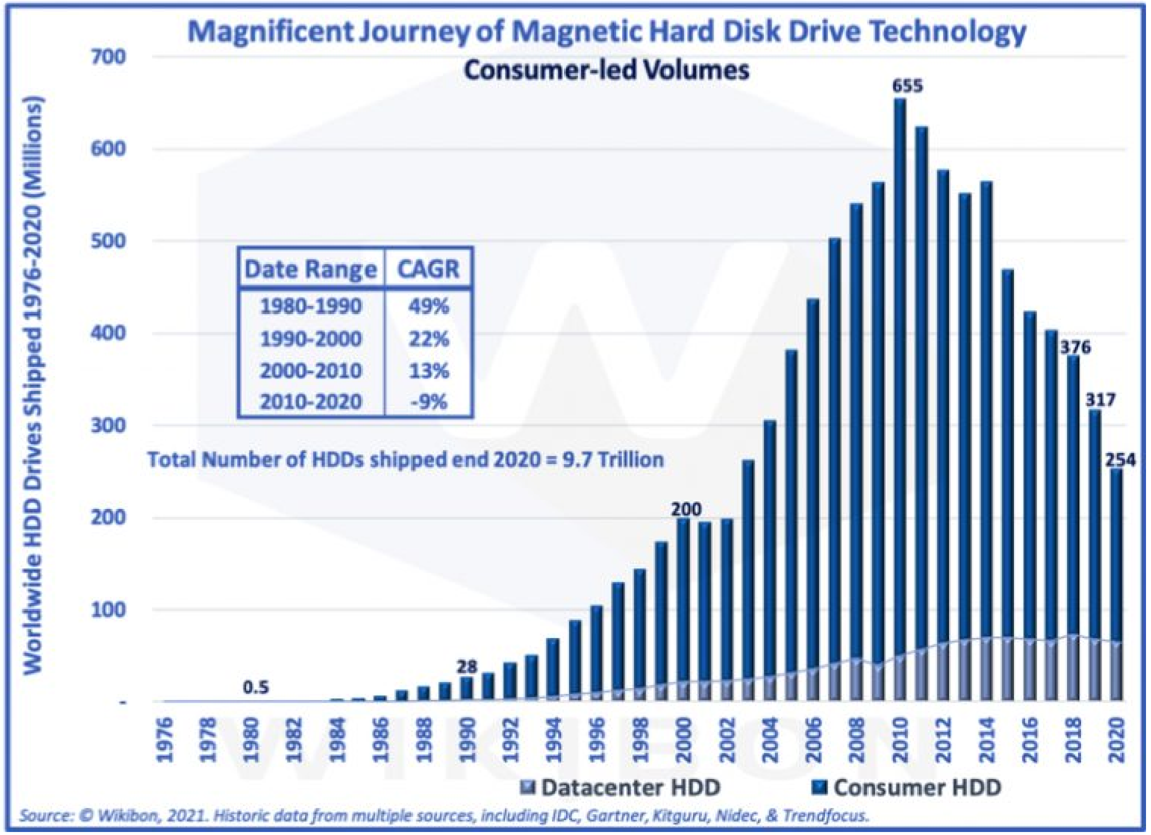

Data storage - Trilemma

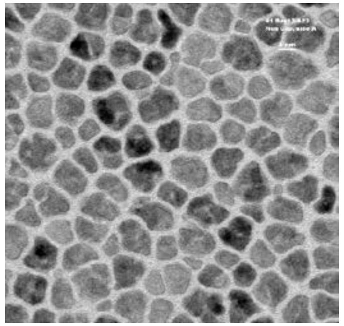

Grain = 8nm

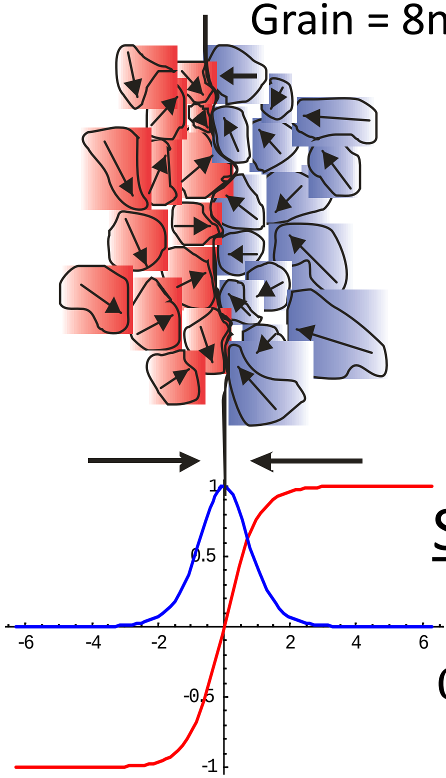

Writability

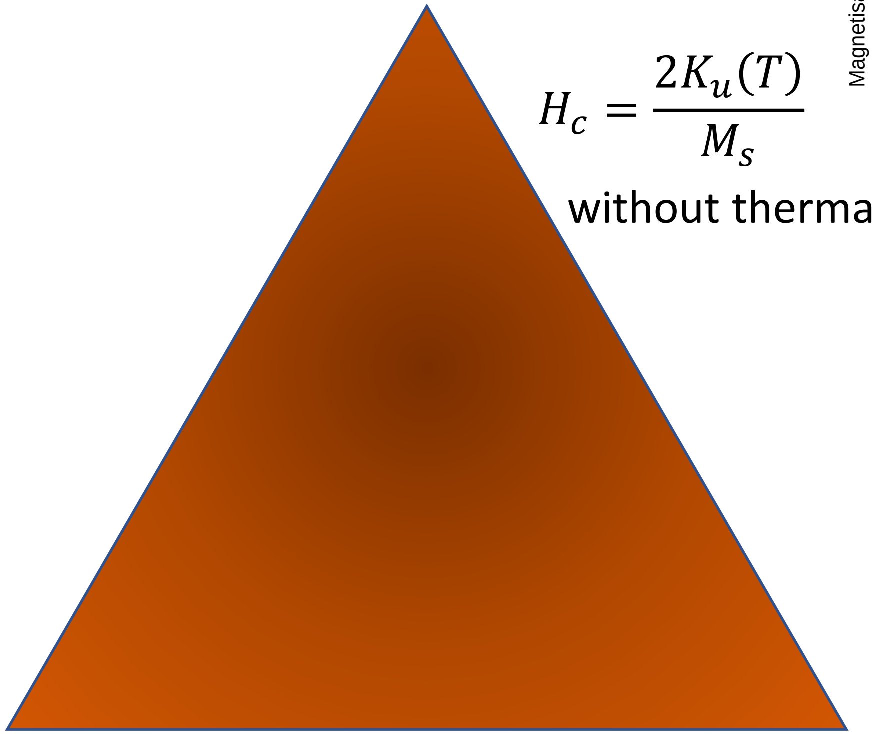

$$H_c = \frac{2K_u(T)}{M_s}$$

without thermal activation

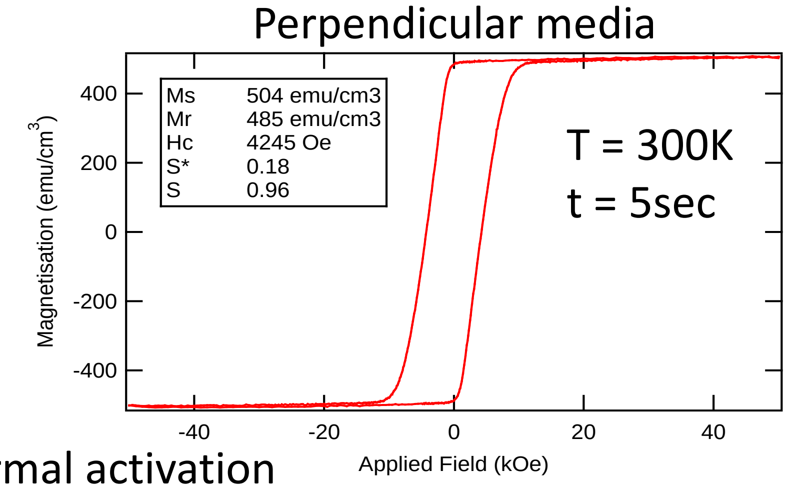

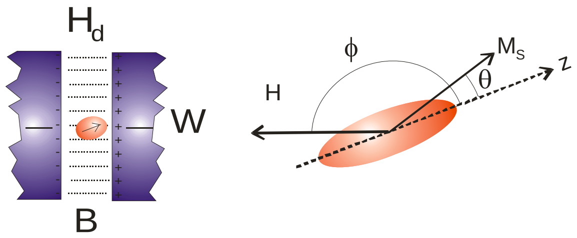

Signal to noise

$$SNR \propto log_{10}(N)$$

Grain volume (V)

Thermal stability

$$\Delta E = \frac{K_u(T)V}{K_B T} \approx 60$$

at T= 300K

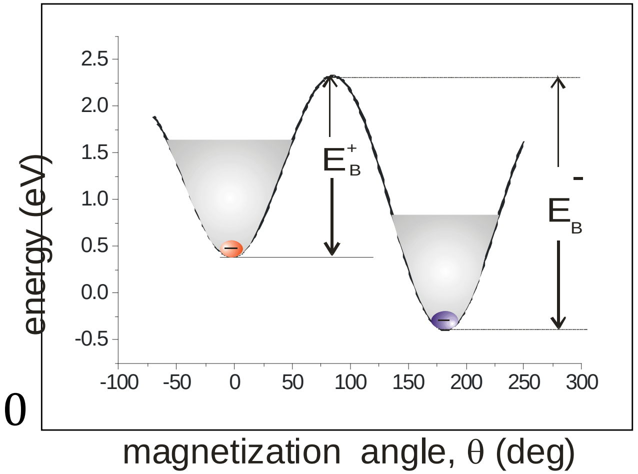

The University of Manchester

Longitudinal Magnetic Recording – Past (1956-2006)

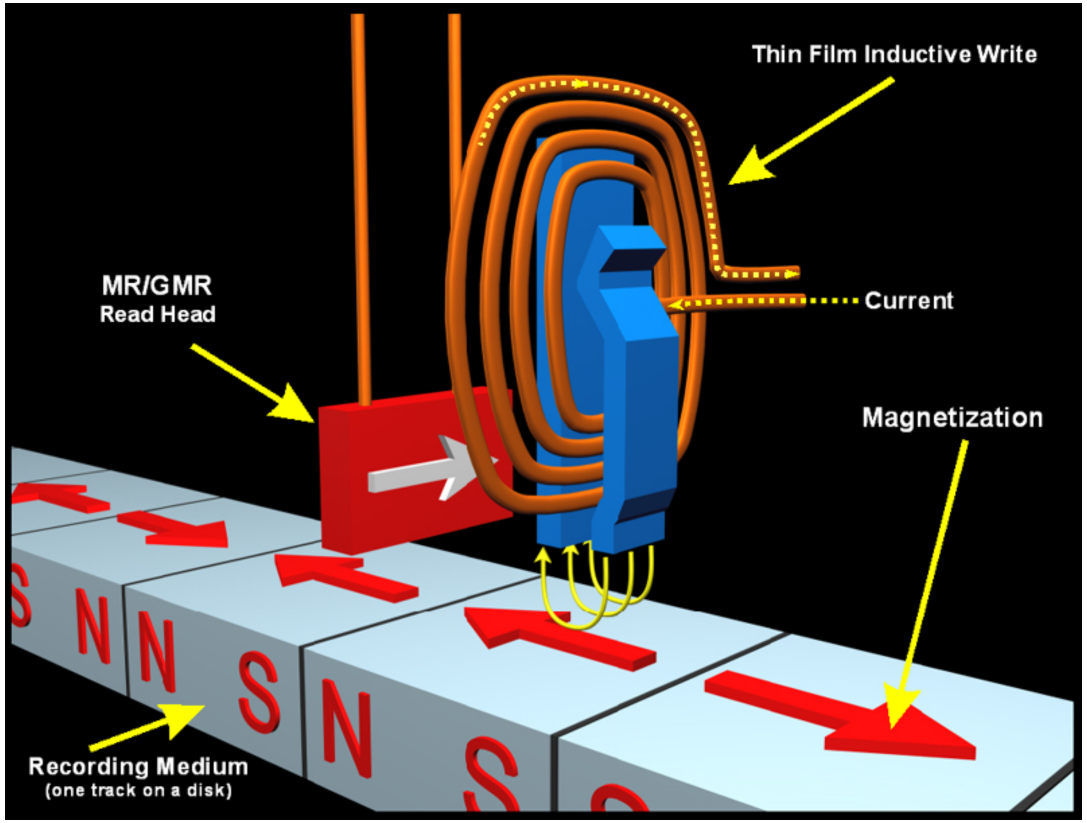

- Disk
  - Ultra smooth surface
  - Thin magnetic coating
  - Protective overcoat
  - Surface lubricant
- Inductive Write Element
  - Soft magnetic poles
  - Copper write coil
  - Alternate coil current to write magnetic transitions
- Resistive Read Element
  - GMR sensor to detect magnetic transitions

- Disk rotates under a slider that has an integrated read/write head at its trailing end
- Very close slider-to-disk surface proximity critical for high resolution recording
- Information is stored in magnetic transitions written onto the disk's thin magnetic coating
- The magnetization is in the plane of the disk surface

Perpendicular Magnetic Recording - Current

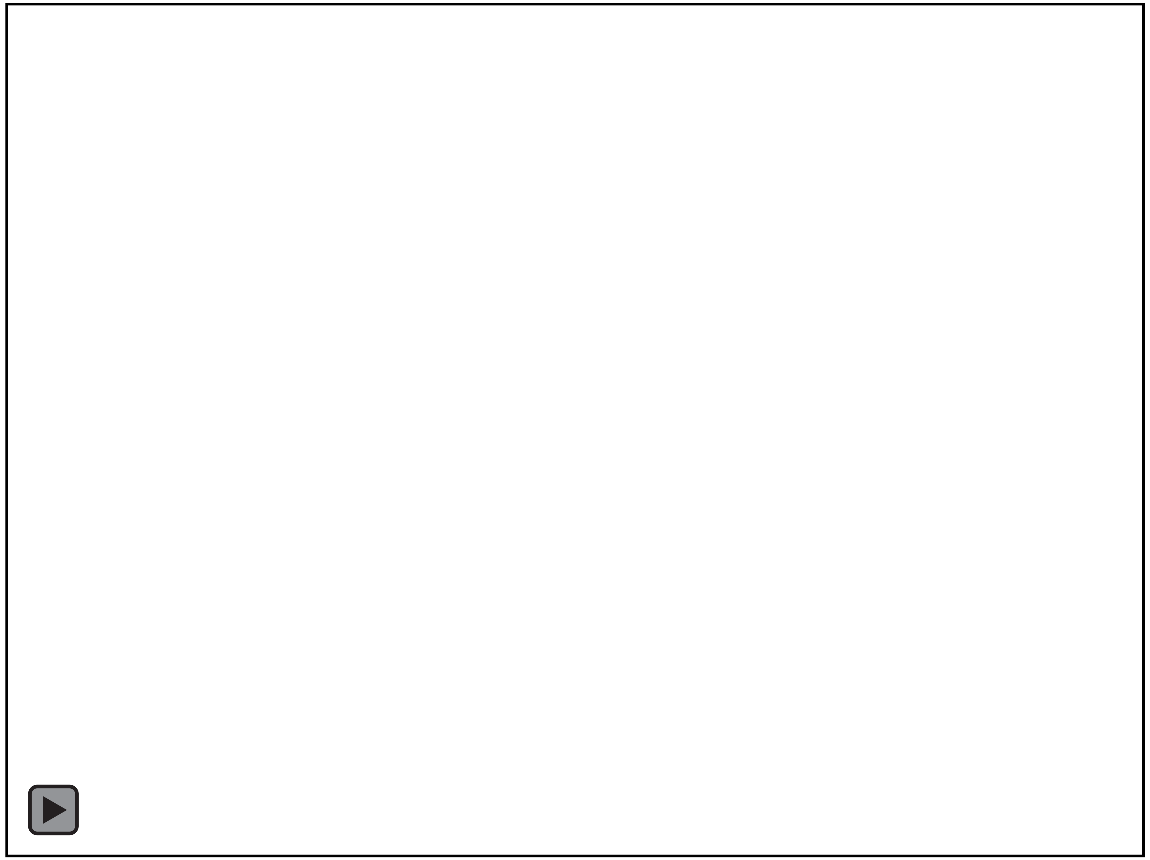

- Disk
  - Ultra smooth surface
  - Thicker magnetic coating
  - Protective overcoat
  - Surface lubricant
- Inductive Write Element
  - Modified design
  - Soft magnetic poles
  - Copper coil
  - Alternate coil current to write magnetic transitions
- Resistive Read Element
  - TMR to detect magnetic transitions

- Disk rotates under a slider that has an integrated read/write head at its trailing end
- Very close slider-to-disk surface proximity critical for high resolution recording
- Information is stored in magnetic transitions written onto the disk’s thin magnetic coating
- The magnetization is perpendicular to the disk surface

HDD enhancements – extending areal density

- Areal density is around 1 TB/in2 (Wood)
- Adaptive fly height
- Helium filled
- Shingled writing
- Signal processing

Wood et al. “Recording Technologies for Terabit per Square Inch Systems” IEEE Trans Magn. 38 1711 (2002)

Energy assisted recording

HAMR: Heat assisted magnetic recording
MAMR: Microwave assisted magnetic recording

HAMR

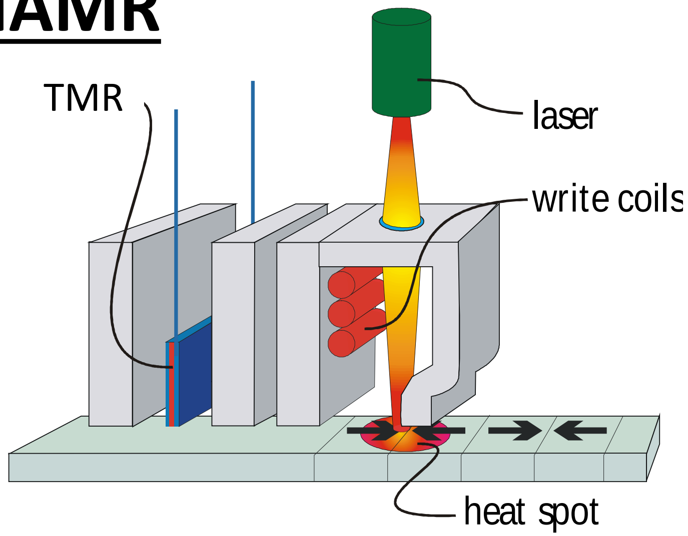

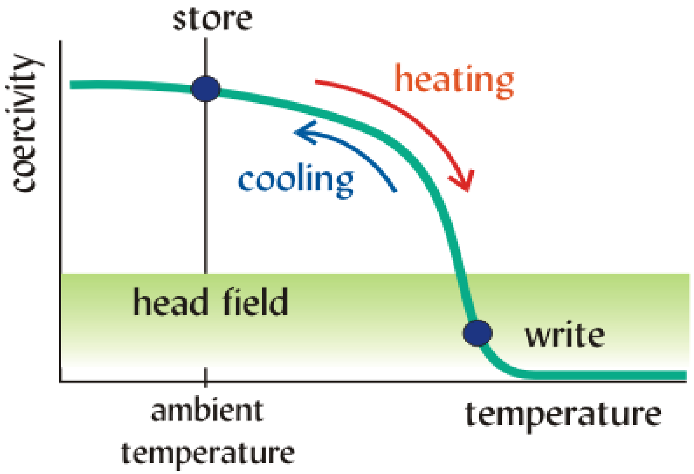

MAMR

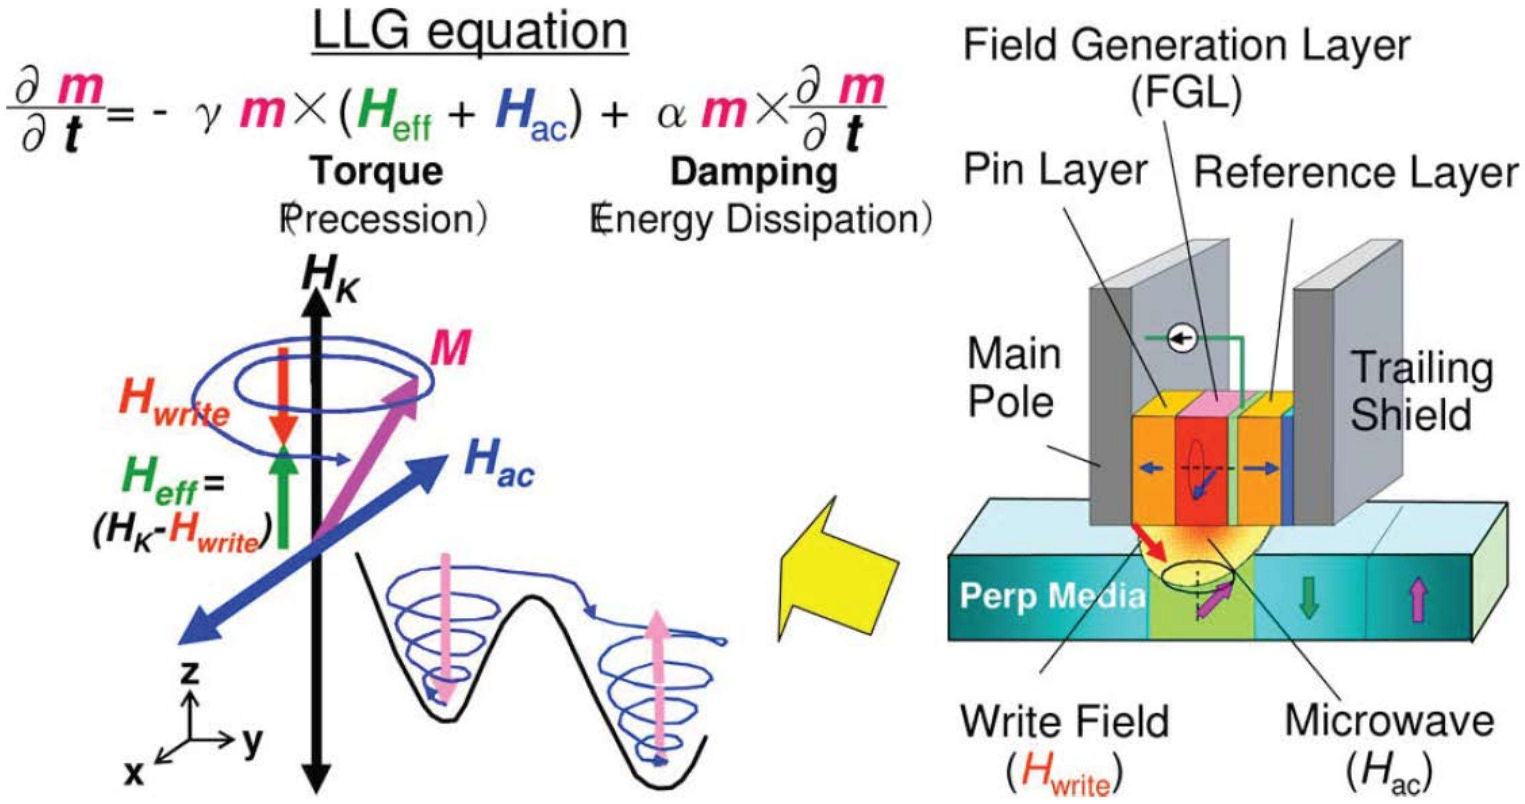

Video

Shiroishi et al. IEEE Trans Magn. 48 3816 (2009)

Magnetic - tape

- Poor relation – but still a useful part of the data storage infrastructure
- Physical separation of media and device
- Almost exclusively Linear Tape Open (LTO) technology
- "Tape has a significantly lower environmental impact as there is no need to have it constantly powered-on during data storage, thereby reducing CO2 emissions generated during its lifecycle by 95% compared to hard disk drives (HDDs)" - FujiFilm Recording Media USA.
- A $5b industry

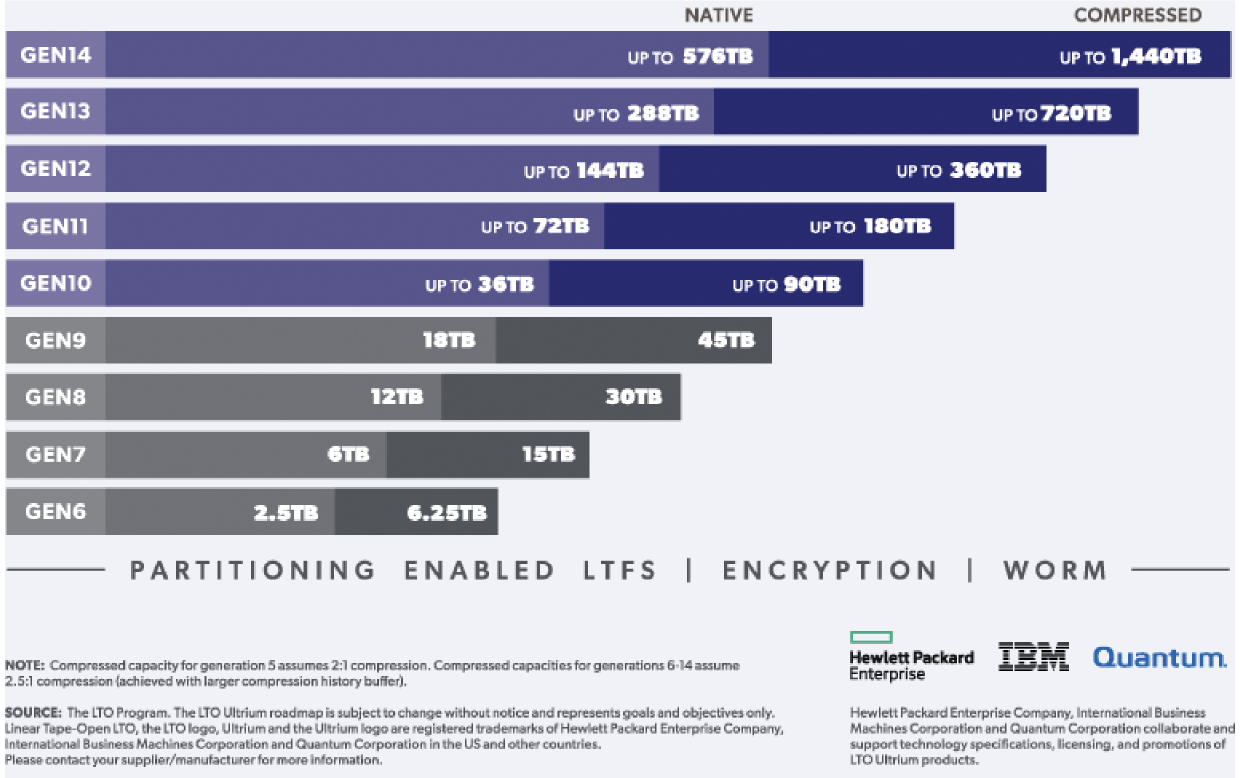

MANCHESTER 1824
The University of Manchester

Magnetic - Roadmaps

TT leaves the industry…

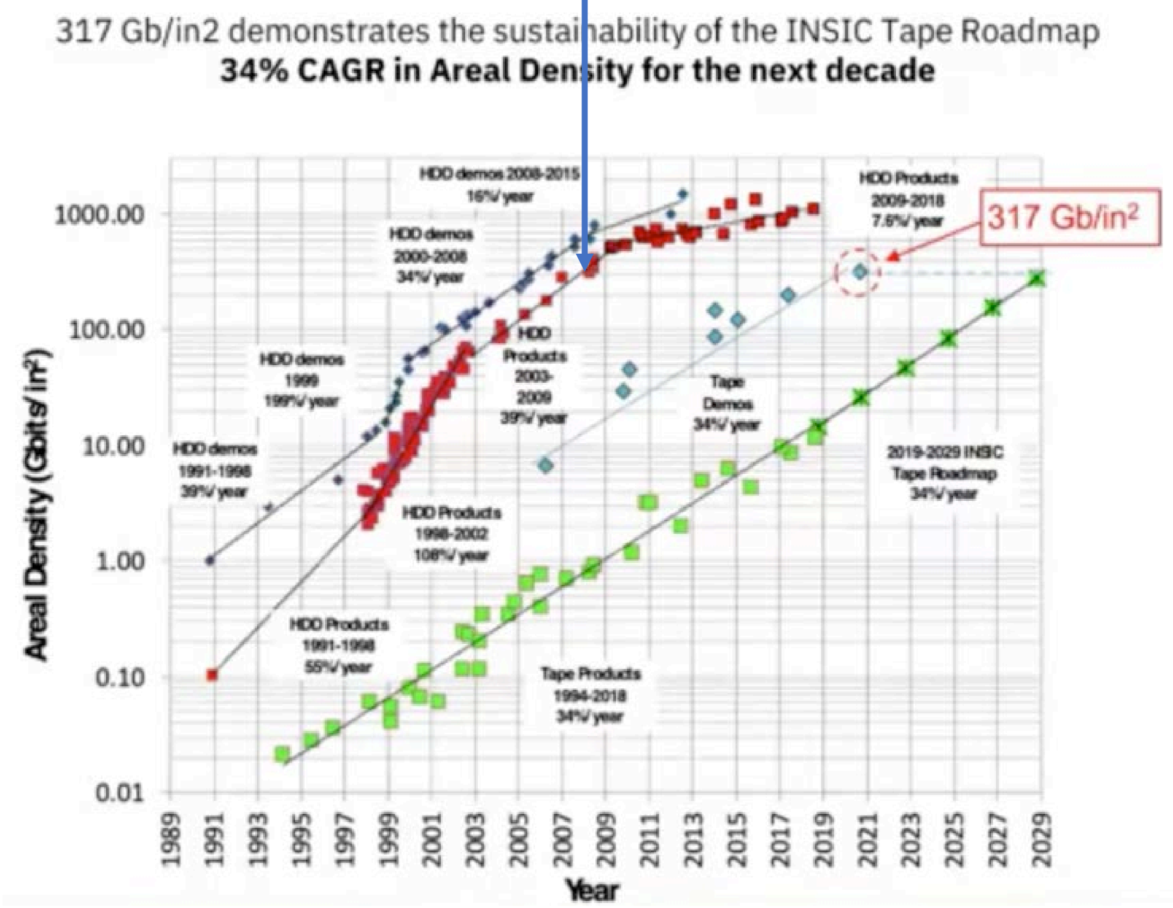

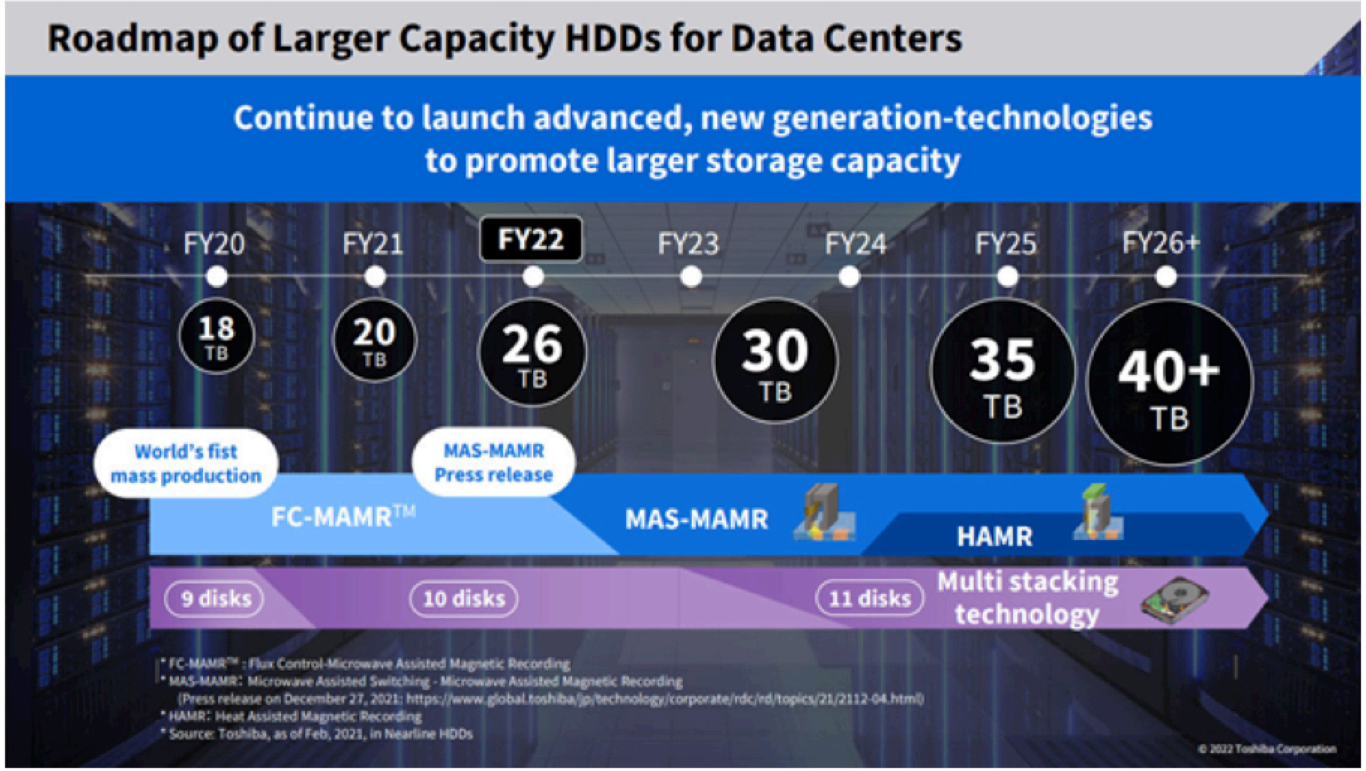

Magnetic - MRAM

- Three types of MRAM
  - Toggle switch MRAM
    - Some specialist applications
  - Spin Transfer Torque (STT-)MRAM
    - In production, mostly for embedded systems
    - 2 terminal device
  - Spin Orbit Torque (SOT-)MRAM
    - Advanced research / pilot line
    - 3 terminal device
- Small $$$ (so far…)

Magnetic Random Access Memory

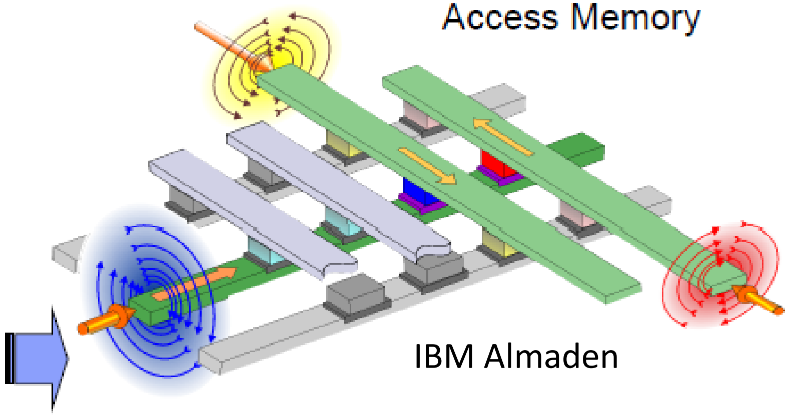

IBM Almaden

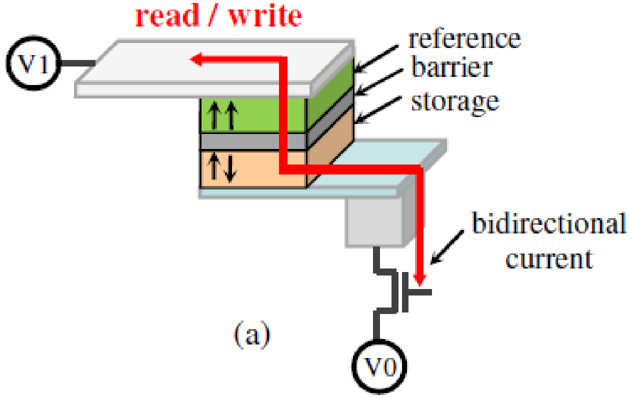

STT-MRAM

SOT-MRAM

Jabeur et al. (Spintec) ELELIJ 6, 1 (2017).

MANCHESTER 1824
The University of Manchester

NAND Flash – SSD/PCIe/NVMe

- Electronics overhead as least as much as for HDD
- Read operation approximately constant with bits/cell
- Write operation – depends on bits/cell (longer program times required)
- A number of energy modeling tools are now available [1]
- A $70b industry

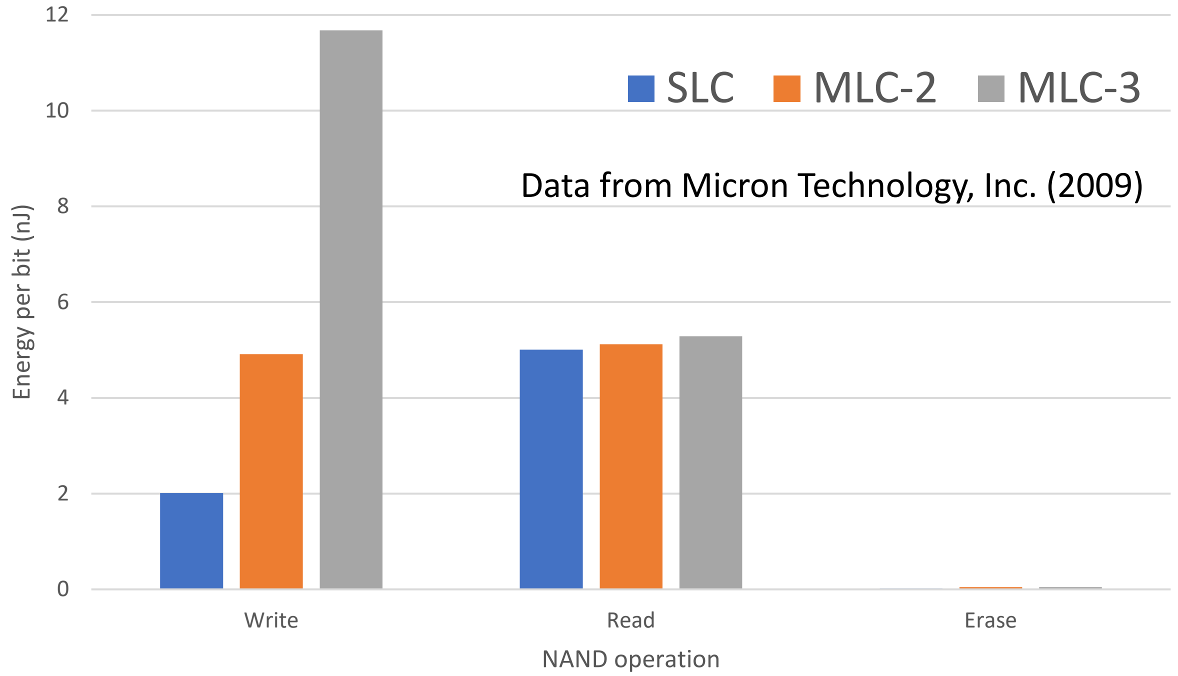

[1] IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, VOL. 32, NO. 7, JULY 2013

Developing new low-energy data storage

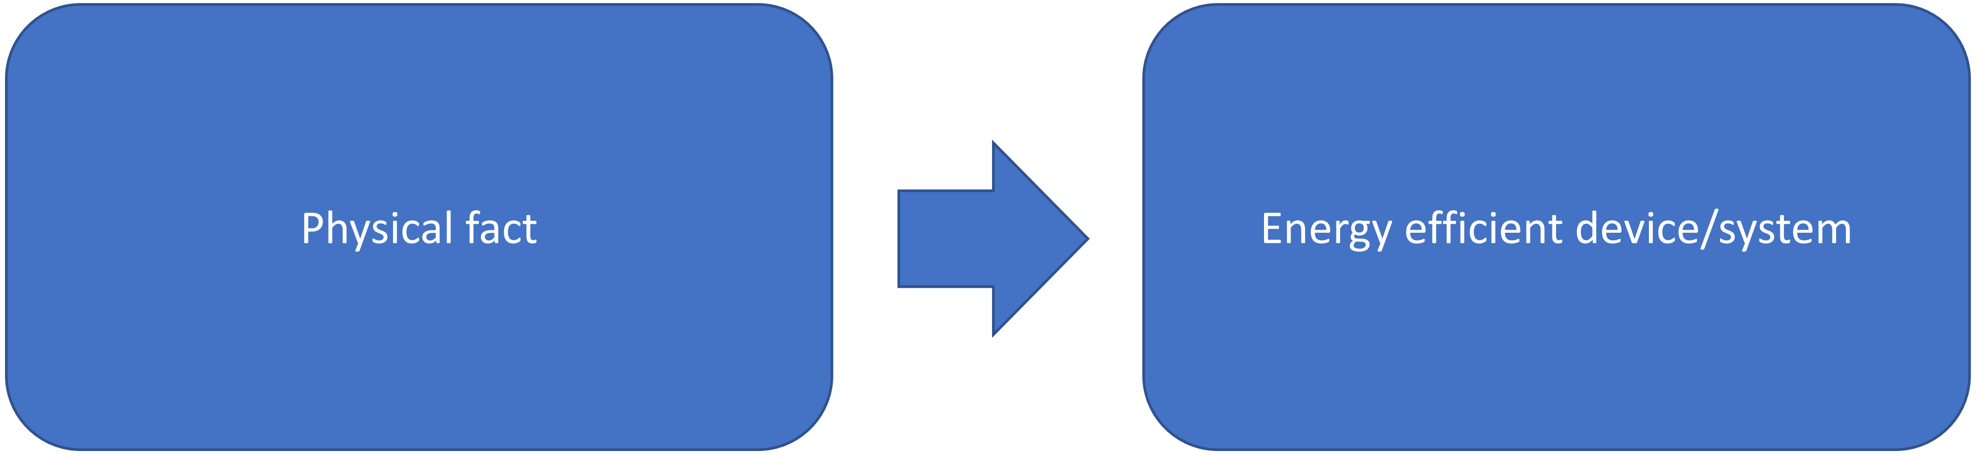

Developing new low-energy data storage

Movement of spin does not induce Joule heating

→

Magnonics/Skyrmions/MRAM/Optical switching

Developing new low-energy data storage

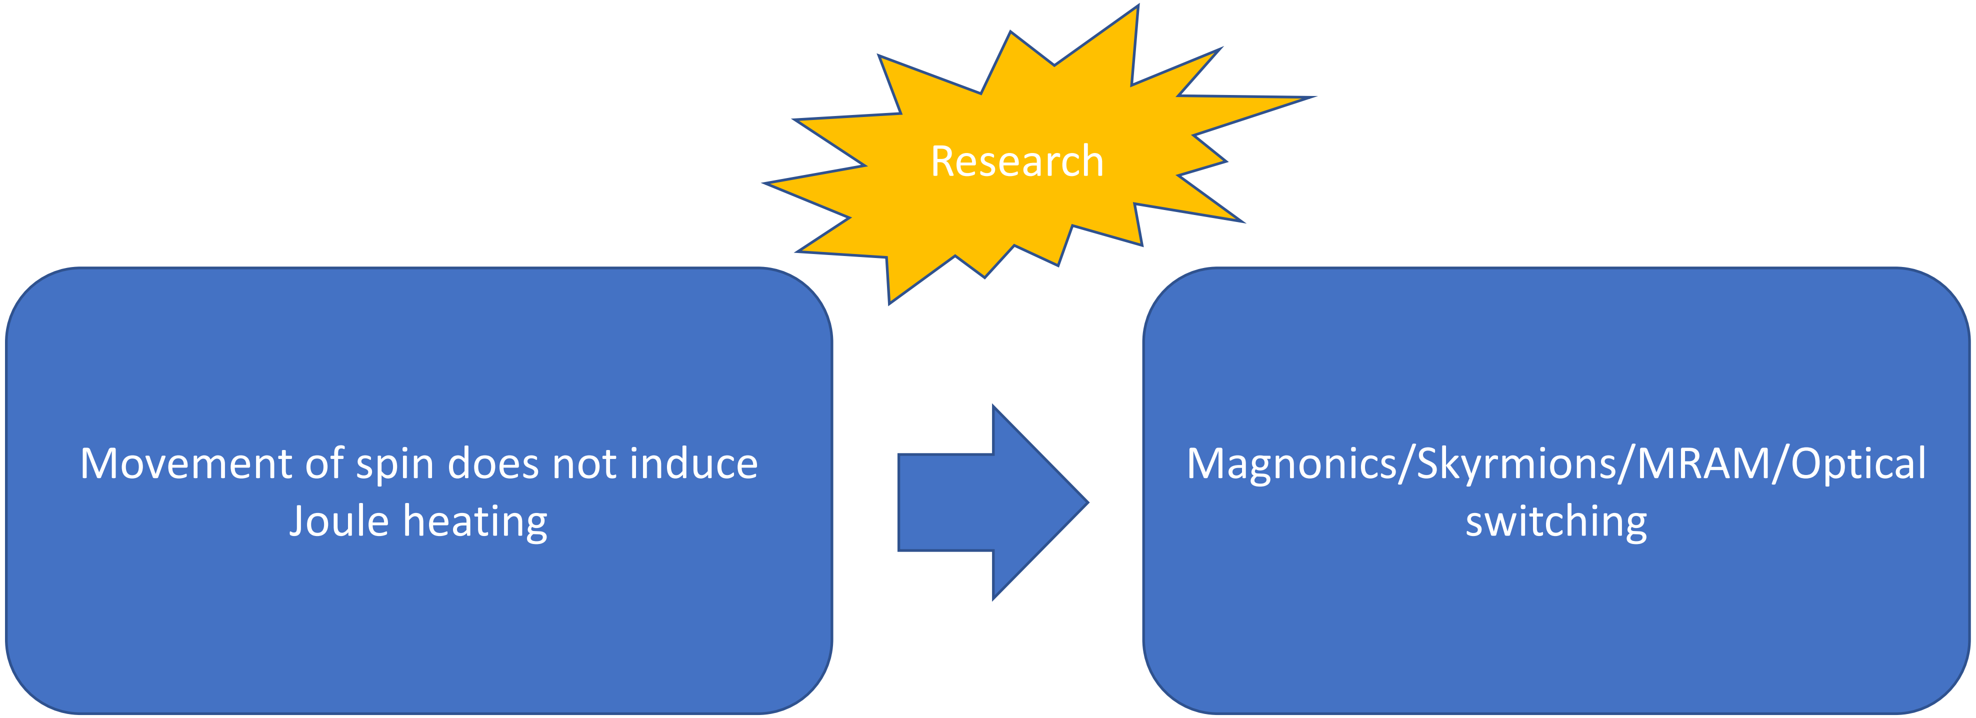

Where are we today? – Writing energy

- Optically switchable magnetic tunnel junction (MTJ) memory device (68)
- Electrically switchable spin valves using mechanisms of spin- transfer torque (50–55 triangles)
- Spin–orbit torque (56–61 squares) and electric- field-induced switching (62–64 circles) is shown

- The red line show the eqn with the characteristic timescale of switching dynamics t_{c0} = 1 ns and the static switching energy U_{c0} = 10 fJ
- The shaded area indicates the target specifications by future technologies

$$U(\tau) = I^2 R\tau = U_{c0}\left(2 + \frac{\tau}{t_{c0}} + \frac{t}{\tau_{c0}}\right)$$

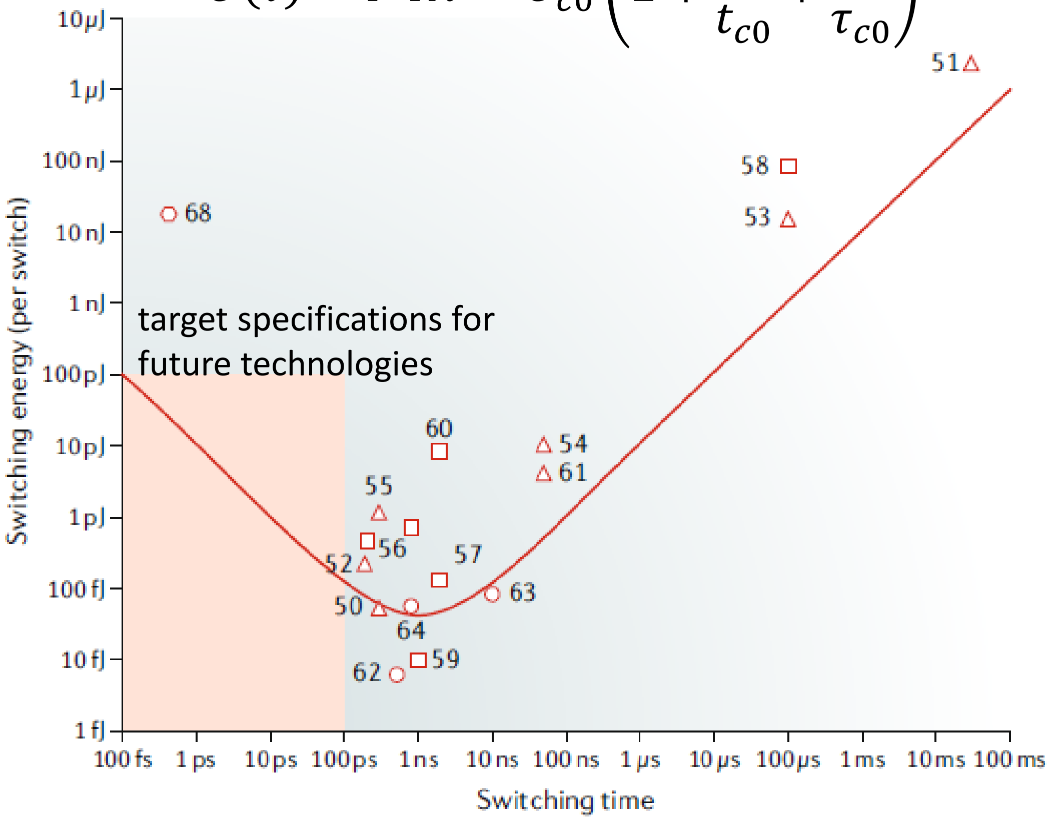

Kimel & Li Nature Review Materials 4 189 (2019)

Team exercise

- Time 30 mins
  - Team 1/5/9 – Race-tracks/domain walls -> Skyrmions
  - Team 2/6 – SOT MRAM
  - Team 3/7 – Magnonics
  - Team 4/8 – All optical switching

- Research questions
  - Is there an implementation scheme for your data storage/computation?
  - How are estimates for energy consumption obtained? - what counts is energy at the wall plug
  - Could it be manufacturable?
  - Is the target area the correct size/shape?

Teams

1) Project-Computation: Reducing energy of computation (mixed)
2) Project-Internet: Reducing the power consumption of the internet
3) Project-cars: “Sustainable Magnetic Materials for Future Electric Application”
4) Project-fridge: magneto-calorics
5) Project-Skyrmion
6) Project-Altermagnets
7) Project-energy: Multiferroics
8) Project-water
9) Project-solarwind

The University of Manchester

Team time

MANCHESTER
1824
The University of Manchester

5.5 Sustainable energy consumption: Computing

A big thank you to Prof. Steve Furber for the background information on silicon systems

Tom Thomson

University of Manchester

ESM 2022 – lecture 19 Sept. 2022

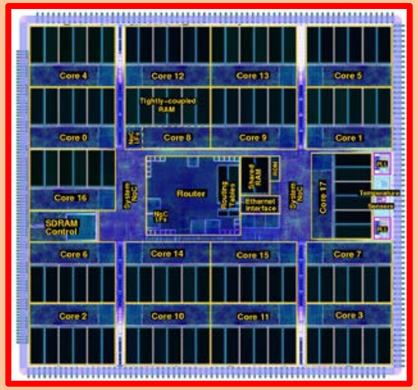

GRAPHENE NOWNANO CDT

HENRY ROYCE INSTITUTE

UKRI Engineering and Physical Sciences Research Council

NEST Nano Engineering Spintronic Technologies

Manchester Baby

- First stored program computer (1948)
- Recognised as an IEEE Milestone (2022)

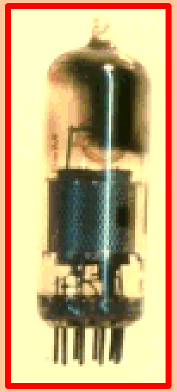

Feature size 10 x 10^6 nm

MANCHESTER 1824

The University of Manchester

SpiNNaker CPU (2011) – Neuromorphoic computing

- Billions of transistors on a single chip

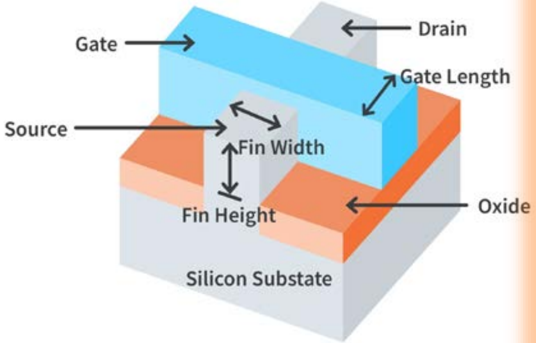

- Research wafers with 2 nm min feature size – IBM 2021

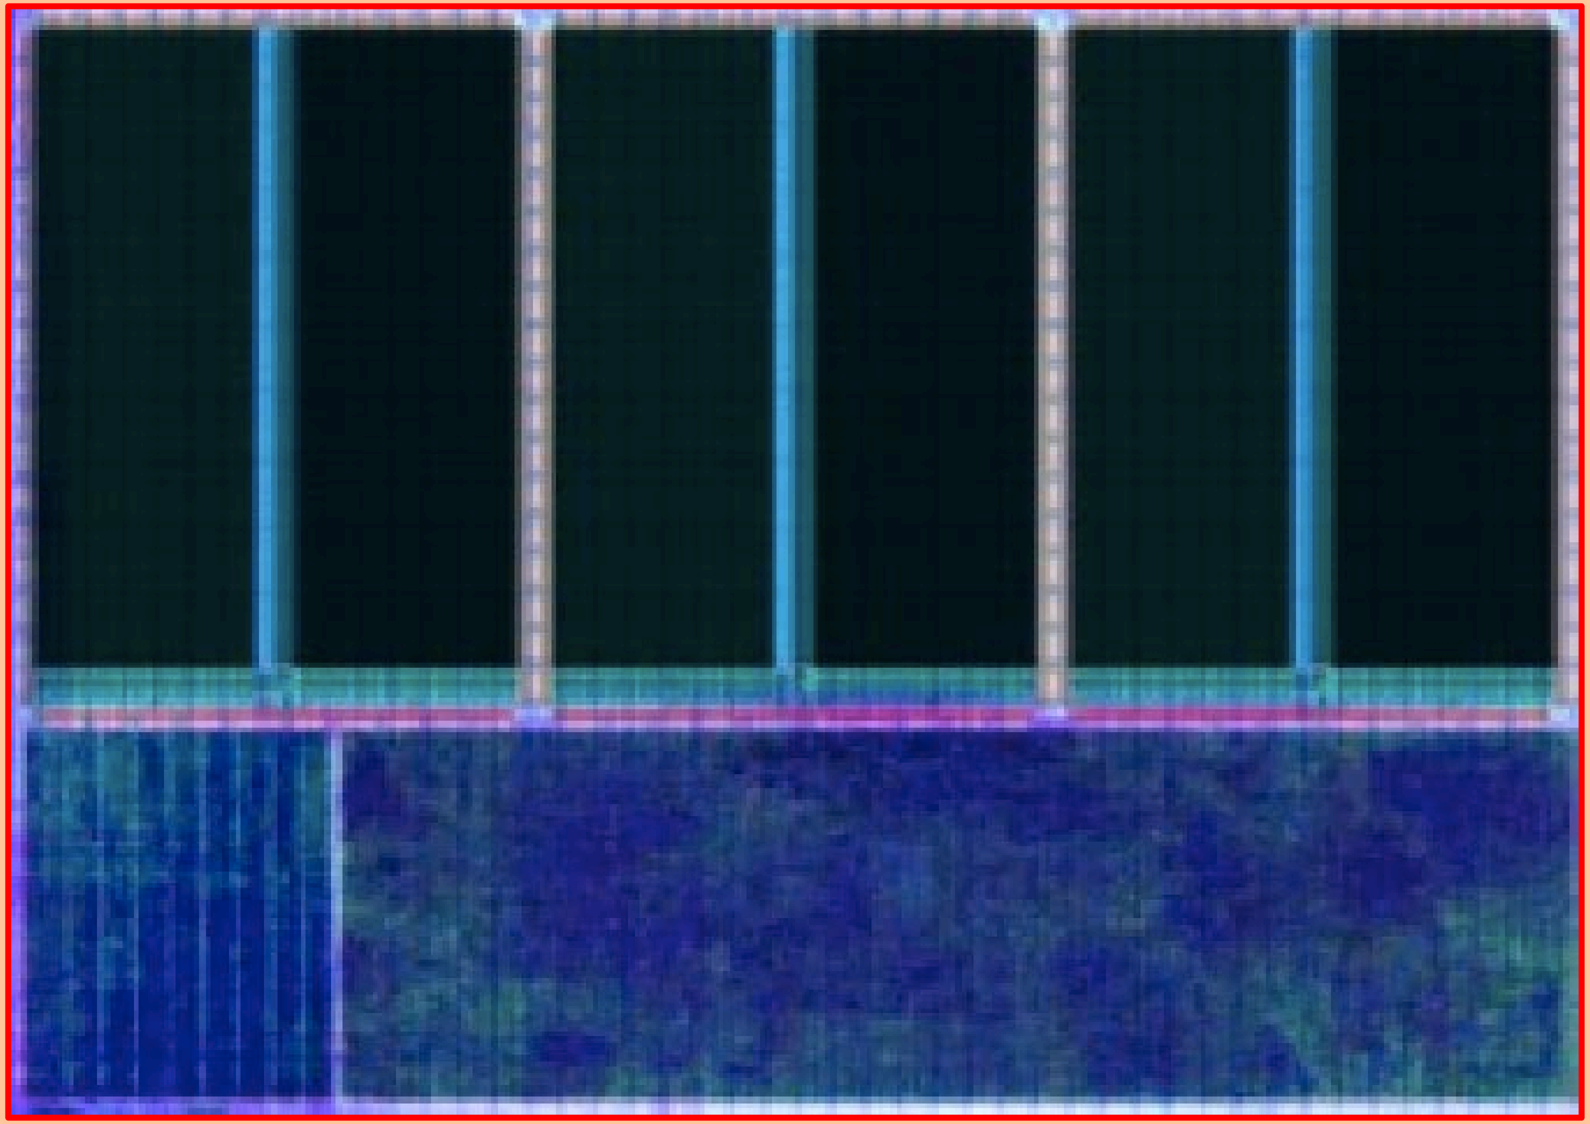

MANCHESTER 1824
The University of Manchester

Seven decades of progress

- ***Baby:***
  - filled a medium-sized room
  - used 3.5 kW of electrical power
  - executed 700 instructions per second
- ***SpiNNaker ARM968 CPU node:***
  - fills ~3.5mm^2 of silicon (130nm)
  - uses 40 mW of electrical power
  - executes 200,000,000 instructions per second

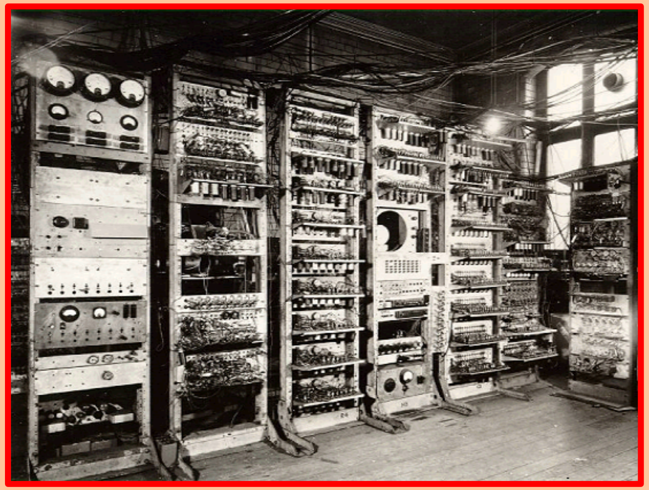

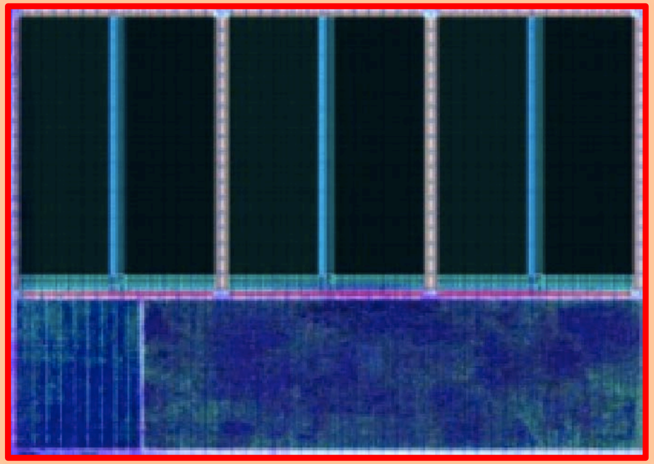

Energy efficiency

- Baby:
  - 5 Joules per instruction

- SpiNNaker ARM968:
  - 0.000 000 000 2 Joules per instruction

25,000,000,000 times better than Baby!

(James Prescott Joule born Salford, 1818)

The University of Manchester

Multi-core CPUs

- High-end uniprocessors
  - diminishing returns from complexity
  - wire vs transistor delays
- Multi-core processors
  - cut-and-paste
  - *simple* way to deliver more MIPS
- Moore's Law
  - more transistors
  - more cores

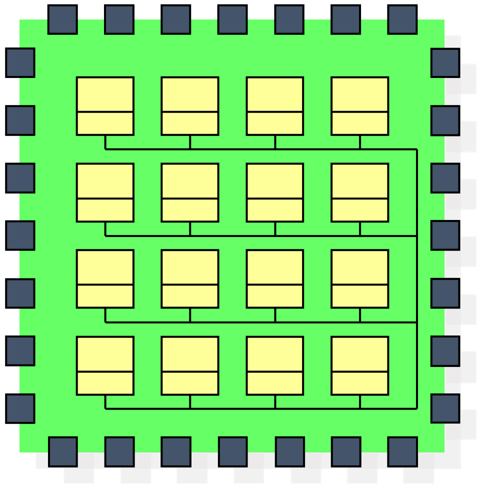

- General-purpose parallelization
  - an unsolved problem
  - the 'Holy Grail' of computer science for half a century?
  - but imperative in the many-core world
- Once solved…
  - few complex cores, or many simple cores?
  - simple cores win hands-down on power-efficiency!

Back to the future

- Imagine…
  - a limitless supply of (free) processors
  - load-balancing is irrelevant
  - all that matters is:
    - the energy used to perform a computation
    - formulating the problem to avoid synchronisation
    - abandoning determinism

- How might such systems work?

Neuromorphic computing - *SpiNNaker* project

- Multi-core CPU node
  - 18 ARM9 processors
  - to model large-scale systems of spiking neurons

- Scalable up to systems with 10,000s of nodes
  - over a million processor
  - 50-100kW

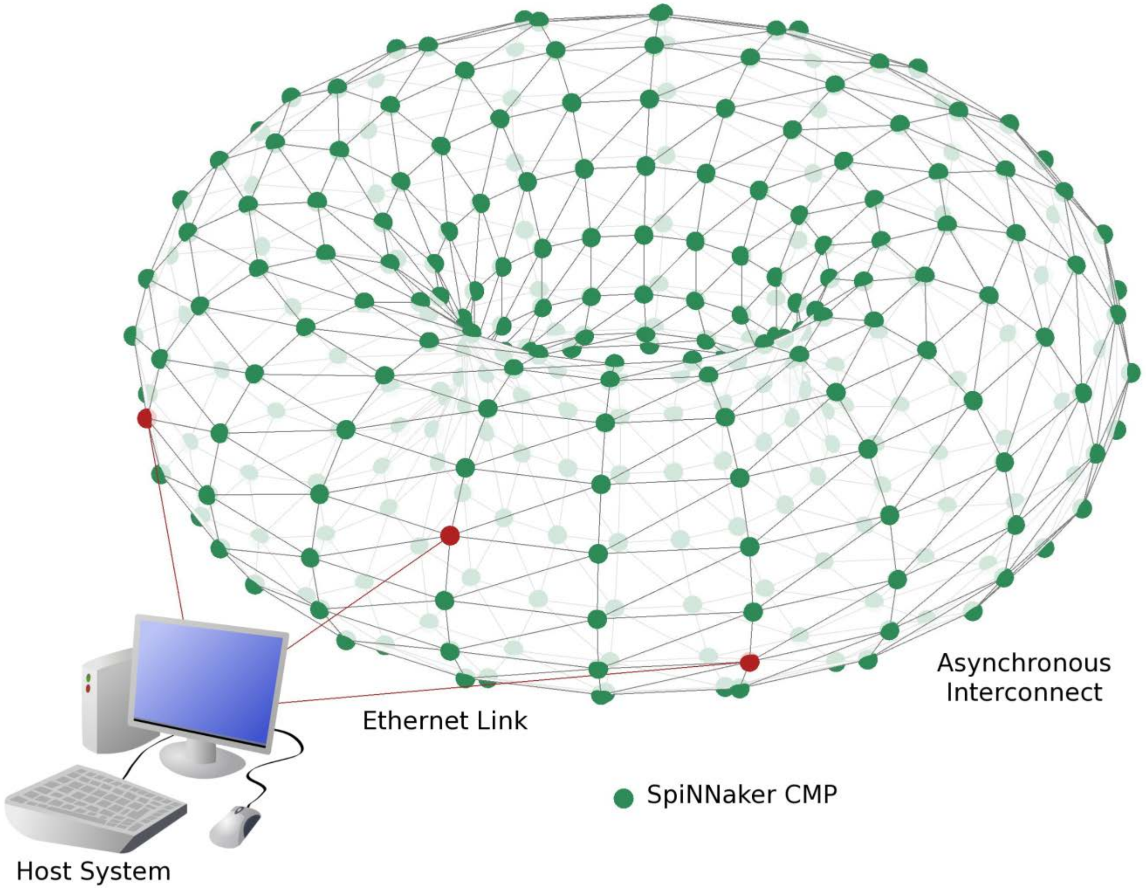

Technology Scaling

- 90nm SpiNNaker CPU node

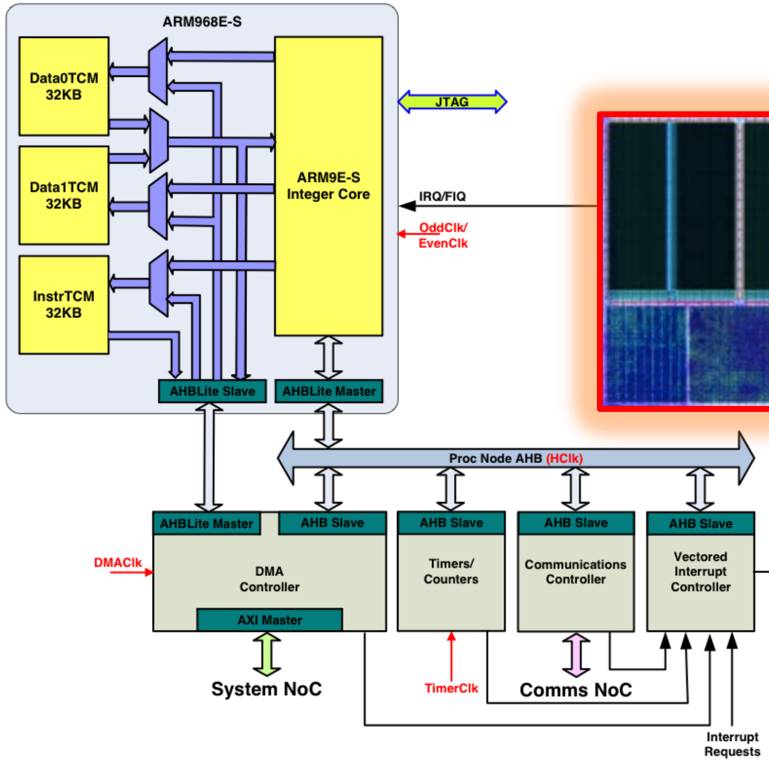

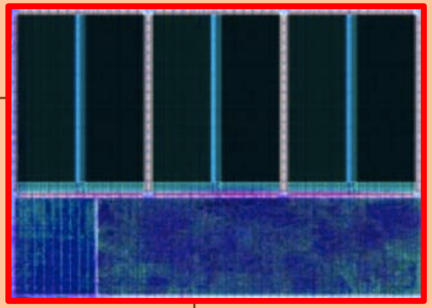

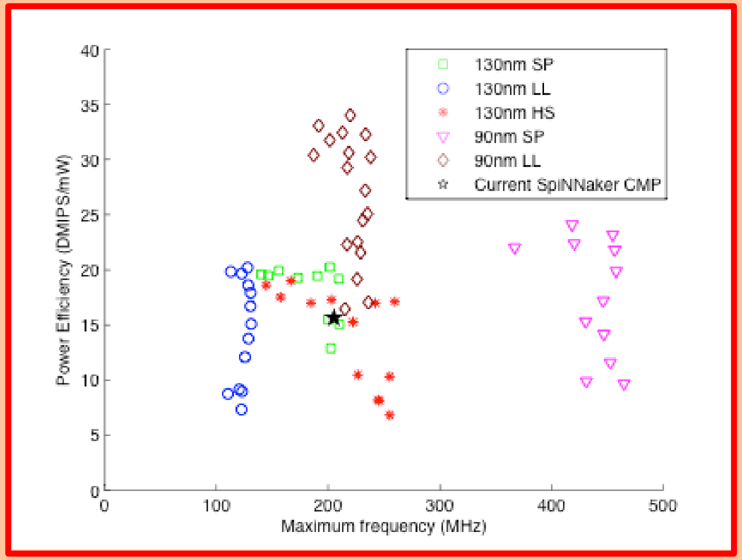

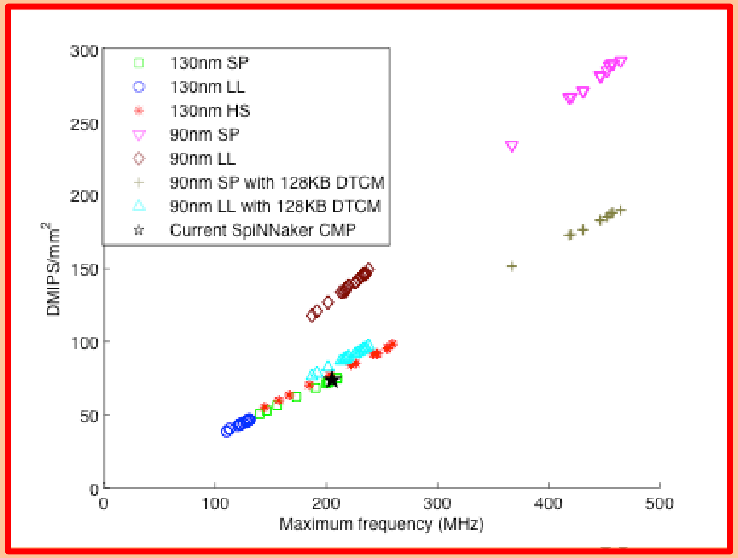

The Exascale objective

- 10^{18} FLOPS at 10MW
  - 100,000 MFLOPS/W
  - 30x current state-of-the-art
- Key ideas:
  - use process advances for efficiency, not speed
  - simplify processors, localize memory
  - 3D integration
    - single package many-core node
- Energy is the real cost of computing!

Multi-chip packaging by UNISEM Europe

unisem

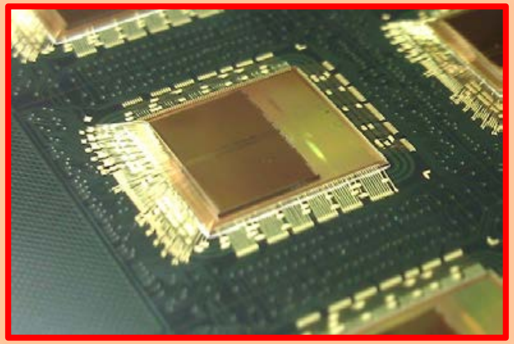

FLOPS = floating-point operations per second
Exa = 10^{18}

Power consumption – Si circuits

- CMOS power consumption
  - voltage change on a gate capacitance requires *charge transfer*, and therefore power consumption
  - once a gate is charged it can maintain its level without any additional charge movement
- CMOS circuitry **only** consumes power when switching states
  - well, until leakage starts to bite!

Power consumption

$$\Delta P = \frac{1}{2} \times C_{total} \times f_{clock} \times V_{DD}^2 \times \alpha$$

where:

P	= dynamic power consumption
C_{total}	= total node capacitance
f_{clock}	= switching frequency of device clock
V_{DD}	= supply voltage
α	= activity: mean no. transitions/clock cycle

e.g. for clock tree $\alpha = 2$

Power consumption

$$\Delta P = \frac{1}{2} \times C_{total} \times f_{clock} \times V_{DD}^2 \times \alpha$$

- Reduce V_{DD}?

$$t_d \propto \frac{V_{DD}}{(V_{DD} - V_t)^2}$$

T_d = circuit delay
V_t = threshold voltage

- Use parallelism to offset increases in circuit delay

MANCHESTER
1824
The University of Manchester

Power consumption

$$\Delta P = \frac{1}{2} \times C_{total} \times f_{clock} \times V_{DD}^2 \times \alpha$$

- Reduce f_{clock}?
  - time to complete computation ~ $1/f$
  - power ~ f
  - so energy per operation independent of f
  - reduced f only helps if it allows lower V_{DD}

The University of Manchester

Power consumption

$$\Delta P = \frac{1}{2} \times C_{total} \times f_{clock} \times V_{DD}^2 \times \alpha$$

- Reduce C_{total}?
  - use smaller, simpler circuits
    - e.g. ARM core rather than Pentium
  - don't over-size gates and buffers
    - in particular, reduce drive off critical path
  - use on-chip rather than off-chip memories
    - off-chip capacitances >> on-chip

Power consumption

$$\Delta P = \frac{1}{2} \times C_{total} \times f_{clock} \times V_{DD}^2 \times \alpha$$

- Reduce α?
  - don’t switch more than is necessary
  - gate clocks
    - turn off processor when job-list is empty
      - don’t sit in an idle loop!
    - ‘event-driven’ style of design
      - in the limit, use asynchronous design

Power consumption - leakage

- Transistor off current isn't zero!

$$I_{off} \propto 10^{\left(-{V_t}/_{100mV}\right)}$$

  - V_t is the transistor threshold
- When V_{DD} = 5 V, V_t = 0.7 V, I_{off} ~ pA
  - x 1,000,000 transistors = 1 μA
- In deep submicron CMOS V_{DD} is lower
  - e.g. 130 nm, V_{dd} = 1.2 V, V_t = 0.3 V, I_{off} ~ 10 nA
  - x 100,000,000 transistors = 1 A
- **This is a big problem! - is there an alternative approach (Maybe…)**

Magnetics in computing

- Move up the hierarchy – think function rather than individual transistors or logic gates.
- Some analogies with quantum computing where you do an experiment and get an answer.

- Skyrmions – neuromorphic computing – some work in Manchester
- Spin torque oscillators (STO)
- Magnonic devices

MANCHESTER 1824
The University of Manchester

Brain-inspired computing

Tailored hardware

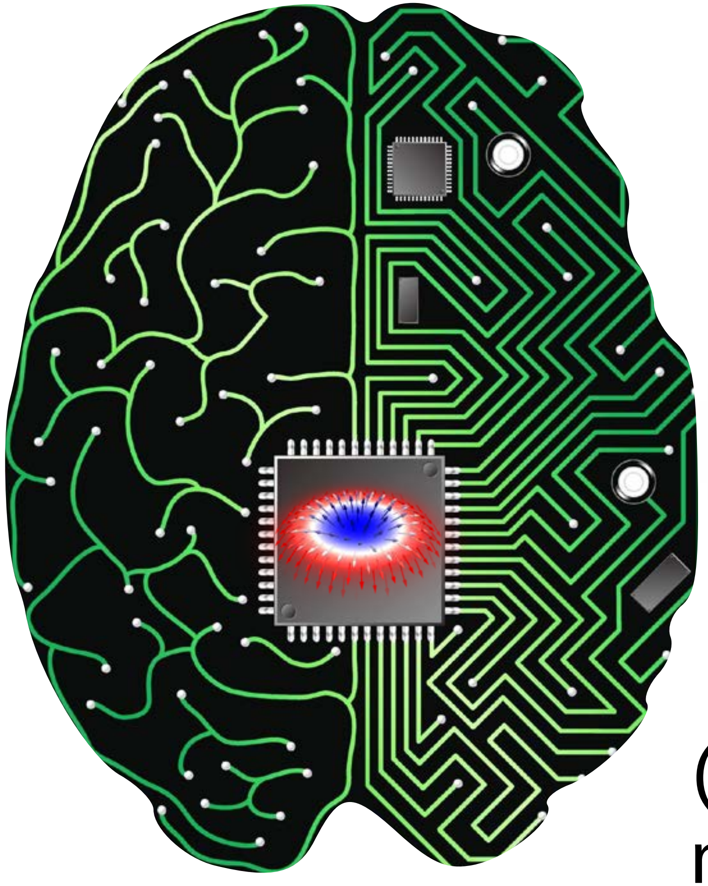

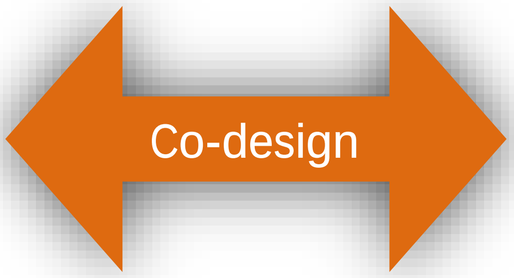

Spiking Neural Networks

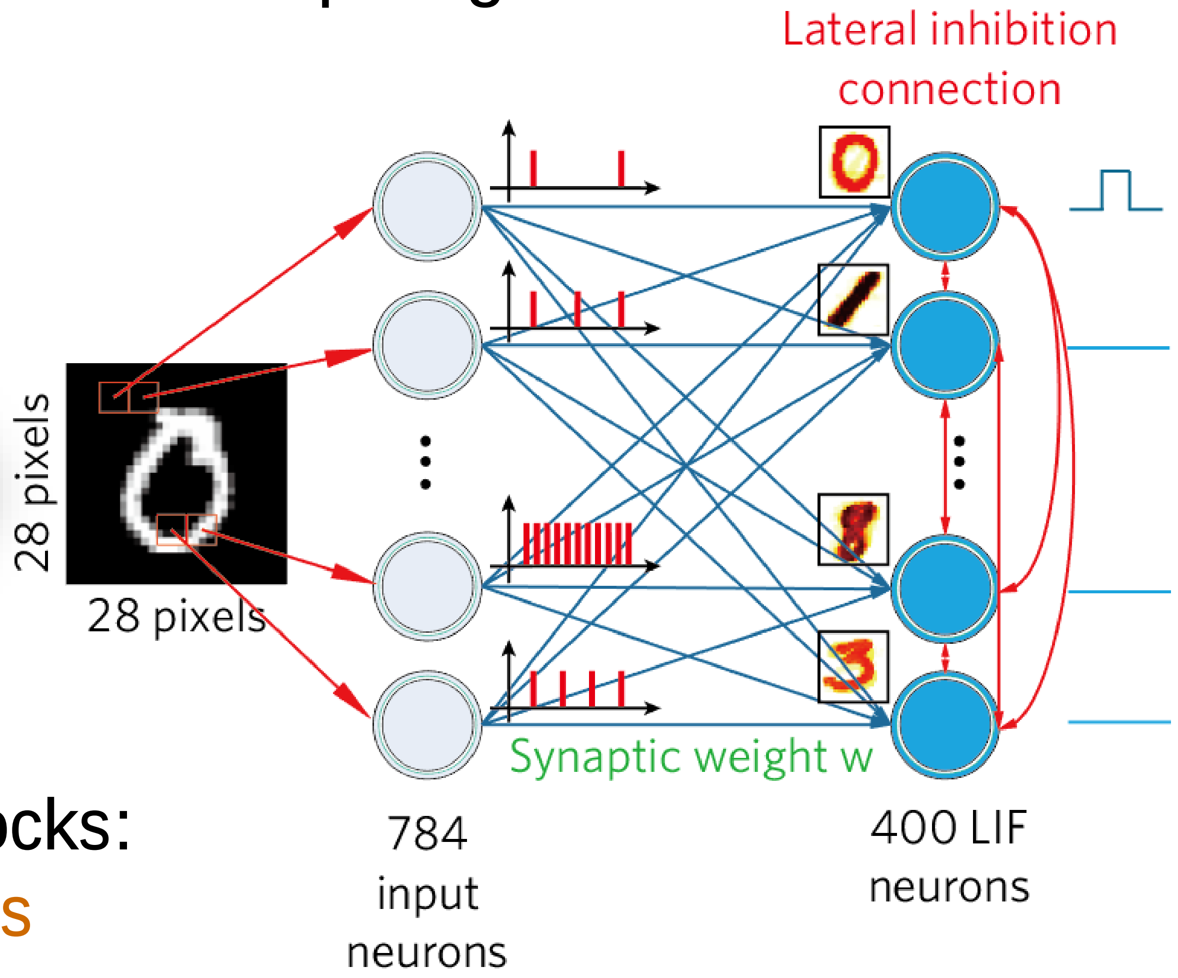

(some) building blocks:
neurons & synapses

Christoforos Moutafis

The University of Manchester

Highlight: Skyrmionic MML Nanosynapse for Deep Spiking Neural Networks

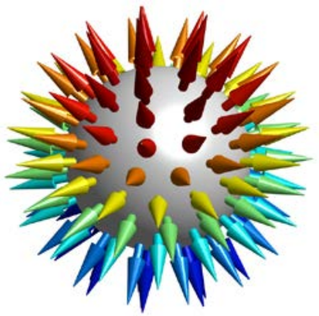

- Can we emulate Synaptic behaviour with topological quasi-particles in nanomagnets in realistic conditions?
  - Plasticity
  - Non-volatility
- Embedded in an SNN and Deep SNN framework to achieve superior classification accuracy in the MNIST handwritten data
- Skyrmionic synapse can be a potential candidate for future energy-efficient neuromorphic edge computing

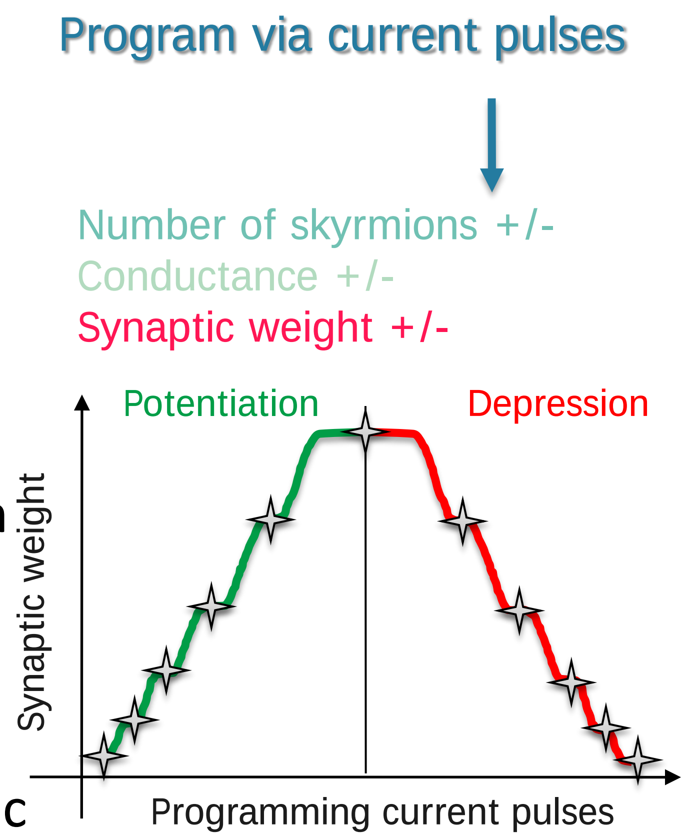

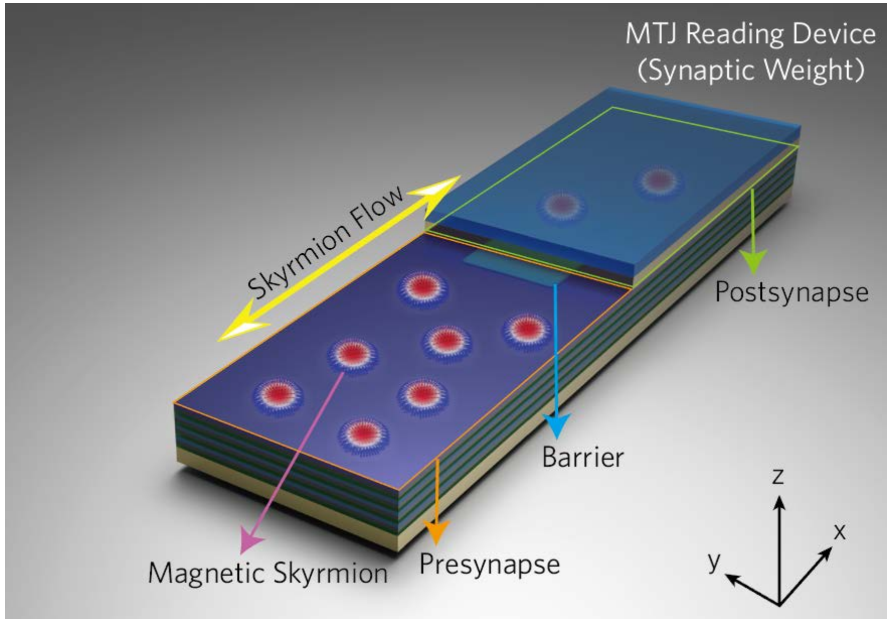

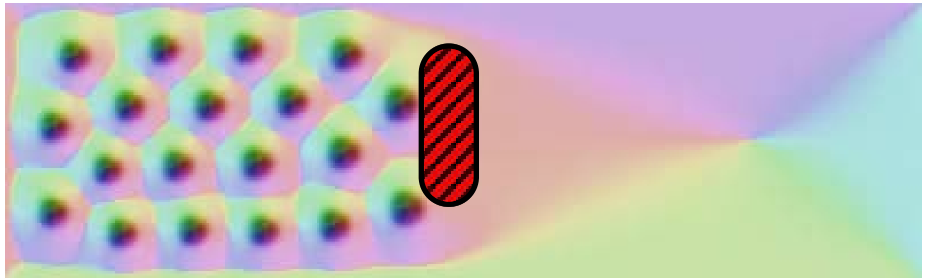

[1] R. Chen, C. Li, Yu Li, J. J. Miles, G. Indiveri, S. Furber, V. F. Pavlidis, C. Moutafis, **Phys. Rev. Applied 14, 014096 (2020)**

Highlight: Topological Filtering for Next-Generation Non-Volatile Interconnects

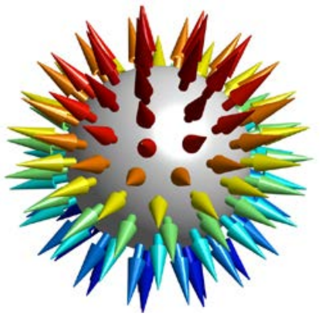

- Missing component? A skyrmionic interconnect device that exploits topological selectivity to achieve signal multiplexing
- Nucleation electrically within 500 ps (following [1])
- **Paradigm shift**, multiple skyrmionic textures / quasi-particles for multiple information carriers
- Exploring stability / metastability of topological and non-topological quasi-particles important for future work
- The topological properties of skyrmionic quasiparticles such as magnetic skyrmions and skyrmioniums enable their applications in future low-power, ultradense nanocomputing and neuromorphic systems

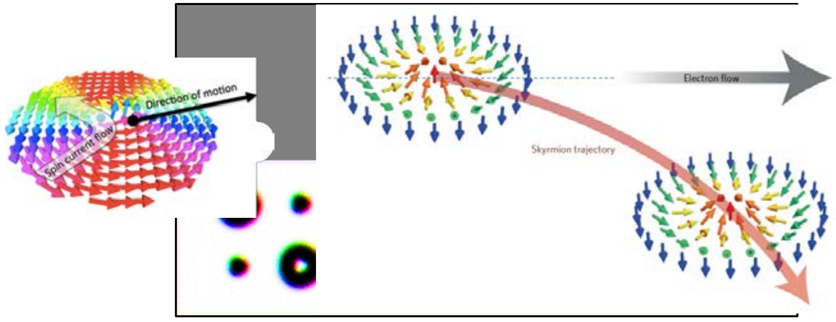

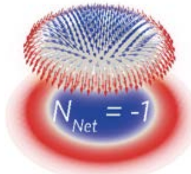

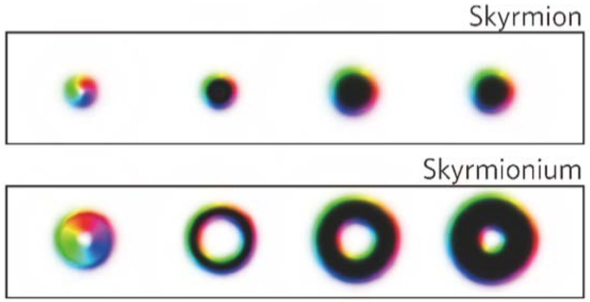

Signal 1
The presence of a skyrmion encodes logic “1”
The absence of a skyrmion encodes logic “0”

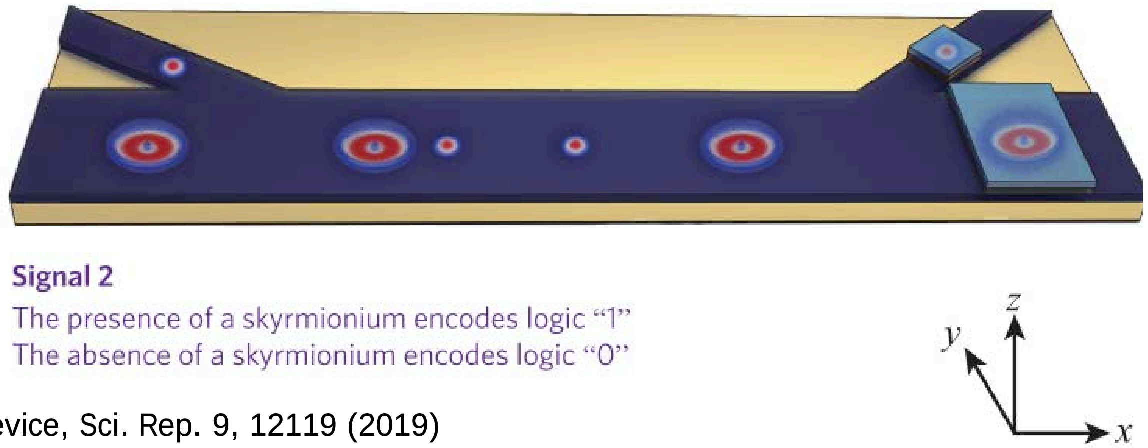

Signal 2
The presence of a skyrmionium encodes logic “1”
The absence of a skyrmionium encodes logic “0”

[1] B. Göbel, A. F. Schäffer, J. Berakdar, I. Mertig & S. S. P. Parkin, Electrical writing, deleting, reading, and moving of magnetic skyrmioniums in a racetrack device, Sci. Rep. 9, 12119 (2019)
[2] R. Chen, Y. Li, V. F. Pavlidis, C. Moutafis, Skyrmionic interconnect device, Physical Review Research 2, 043312 (2020)

MANCHESTER 1824
The University of Manchester

Skyrmions for Nanocomputing

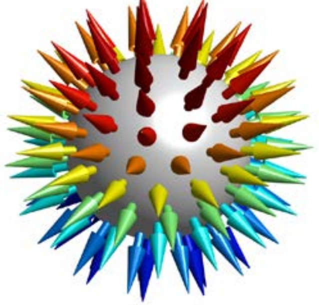

- Neuromorphic Computing: Explore concepts for pattern extraction / classification tasks, e.g. nanoscale multilayer skyrmion-based synapses for deep spiking neural networks [1]
- Interconnects: Encoding sequences of information with distinct skyrmionic textures for multiplexing/demultiplexing signals [2].
    - Multiple topological (& non-topological) spin textures as information carriers
    - Many challenges both at the device (e.g. which device design?) and system level (e.g. scalability)

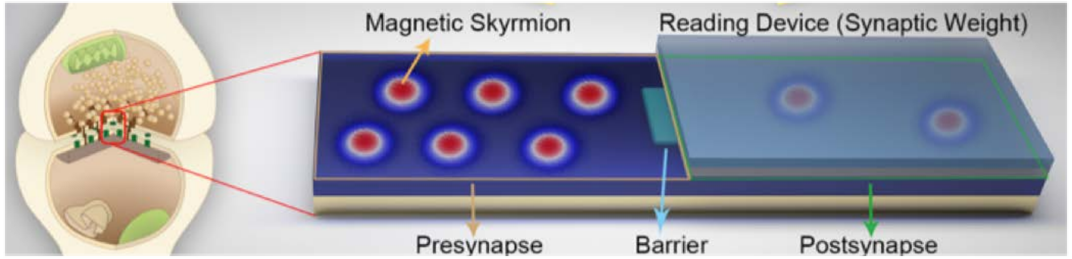

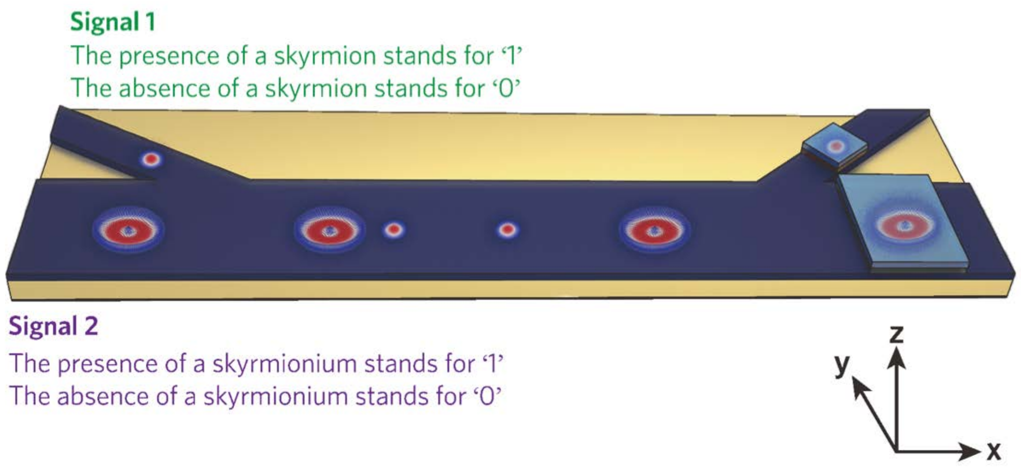

[1] R. Chen, C. Li, Y. Li, J. J. Miles, G. Indiveri, S. Furber, V. F. Pavlidis, C. Moutafis, Nanoscale RT Multilayer Skyrmionic Synapse for Deep Spiking Neural Networks, **Physical Review Applied 14, 014096 (2020)**

[2] R. Chen, Y. Li, V. F. Pavlidis, C. Moutafis, Skyrmionic interconnect device, **Physical Review Research 2, 043312 (2020)**

Skyrmionics for Neuromorphic Technologies
EPSRC (EP/V028189/1)

MANCHESTER
1824
The University of Manchester

Neuromorphoric computing with STOs

- Neuromorphic computing with spin torque nano-oscillators (STOs).
  - A fixed input current gives an oscillating voltage across the junction.
- Reservoir computing with STO using time multiplexing in pre- and post-processing, here recognizing the particular spoken digit as ‘1’.
- Top: schematic of the use of coupled nano-oscillators for vowel recognition.
- Bottom: the input is represented by the frequencies of two microwaves applied through a stripline to the oscillators. The natural frequencies of the oscillators are tuned by d.c. bias currents.
- These can be tuned so that the synchronization pattern between the oscillators corresponds to the desired output.

Grollier et al. Nature Electronics 3 360 (2020)

MANCHESTER 1824
The University of Manchester

Magnonics - Spin wave computing

- Magnonics addresses the physical properties of spin waves and utilizes them for data processing
  - Scalability down to atomic dimensions, operation in the GHz-to-THz frequency range
- Magnonics is definitely in the research phase but some proof-of-concept prototypes have been developed
- Computation operations with the Boolean digital logic and unconventional approaches, such as neuromorphic computing.

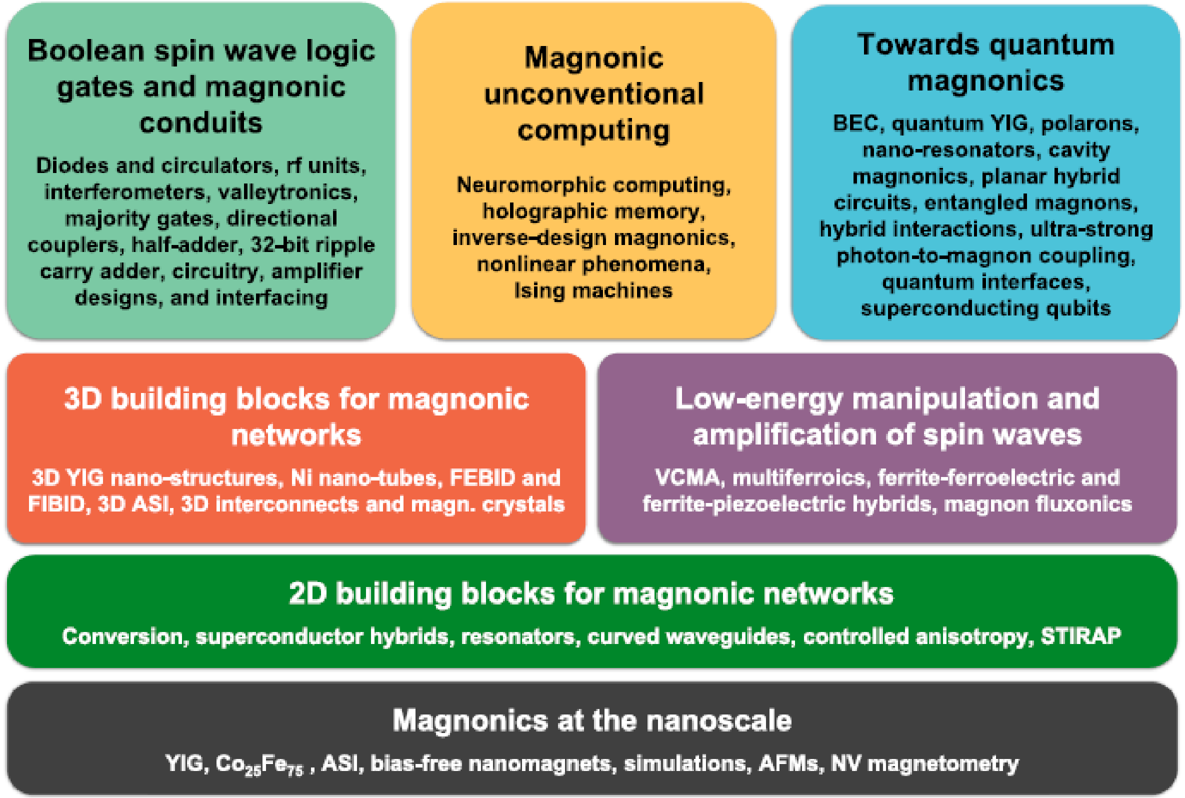

Chumak et al. IEEE Trans Magn. 58 0800172 (2022)

Magnonics – computational functionality

- The operational principle of the magnonic half-adder
- Schematic view of the magnonic half-adder
  - Typical parameters:
    - YIG waveguide width, w = 100 nm;
    - thickness, h = 30 nm;
    - edge-to-edge distances between waveguides, d1 = 450 nm, d2 = 210 nm;
    - angle between waveguides, ϕ = 20°;
    - gaps between coupled waveguides, δ1 = 50 nm, δ2 = 10 nm;
    - lengths of coupled waveguides, L1 = 370 nm and L2 = 3 µm.
- Red and black arrows show the flow path of magnons from the inputs to the logic gates.

Wang et al. Nature Electronics 3 765 (2020)

Summary

- Magnetic data storage and computation from an energy perspective
- Need to think about total energy budget of system as well as that of devices
- Magnetic devices offer new paradigms in computation but a long way to go
    - Several schemes for neuromorphic devices
    - Spin waves have promise
    - Synergies with quantum computing?

The University of Manchester

Thanks & Acknowledgements

- The whole NEST team
- The Funding agencies
- Our collaborators everywhere:

HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

CENIDE CENTER FOR NANOINTEGRATION DUISBURG-ESSEN

SINTEF

SPring-8

PAUL SCHERRER INSTITUT PSI

MESO SYS

Science & Technology Facilities Council ISIS Neutron and Muon Source

GRAPHENE NOWNANO CDT

HENRY ROYCE INSTITUTE

We are pleased to acknowledge our funders: EPSRC (EP/V007211/1, EP/S033688/1, EP/V028189/1, EP/L01548X/1, EP/S019367/1, EP/P025021/1)

UKRI Engineering and Physical Sciences Research Council

NEST Nano Engineering Spintronic Technologies

Questions

MANCHESTER 1824
The University of Manchester

Team presentations