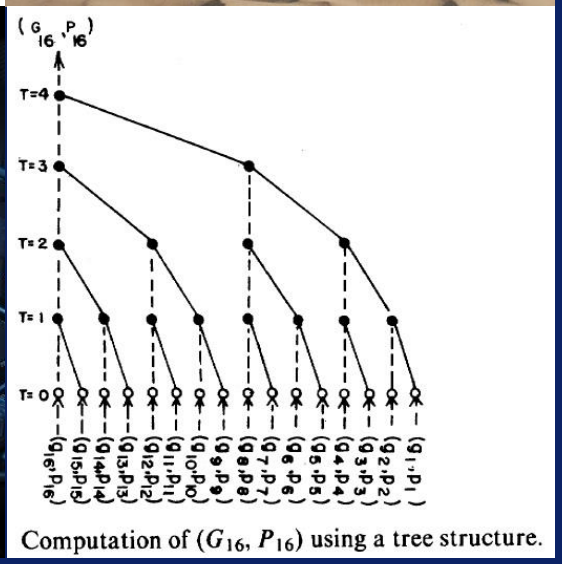

Computation of (G_{16}, P_{16}) using a tree structure.

Future Algorithmic Processing

Prof Simon Maskell
s.maskell@liverpool.ac.uk

UNIVERSITY OF LIVERPOOL

Moore's Law: According to NVIDIA's CEO

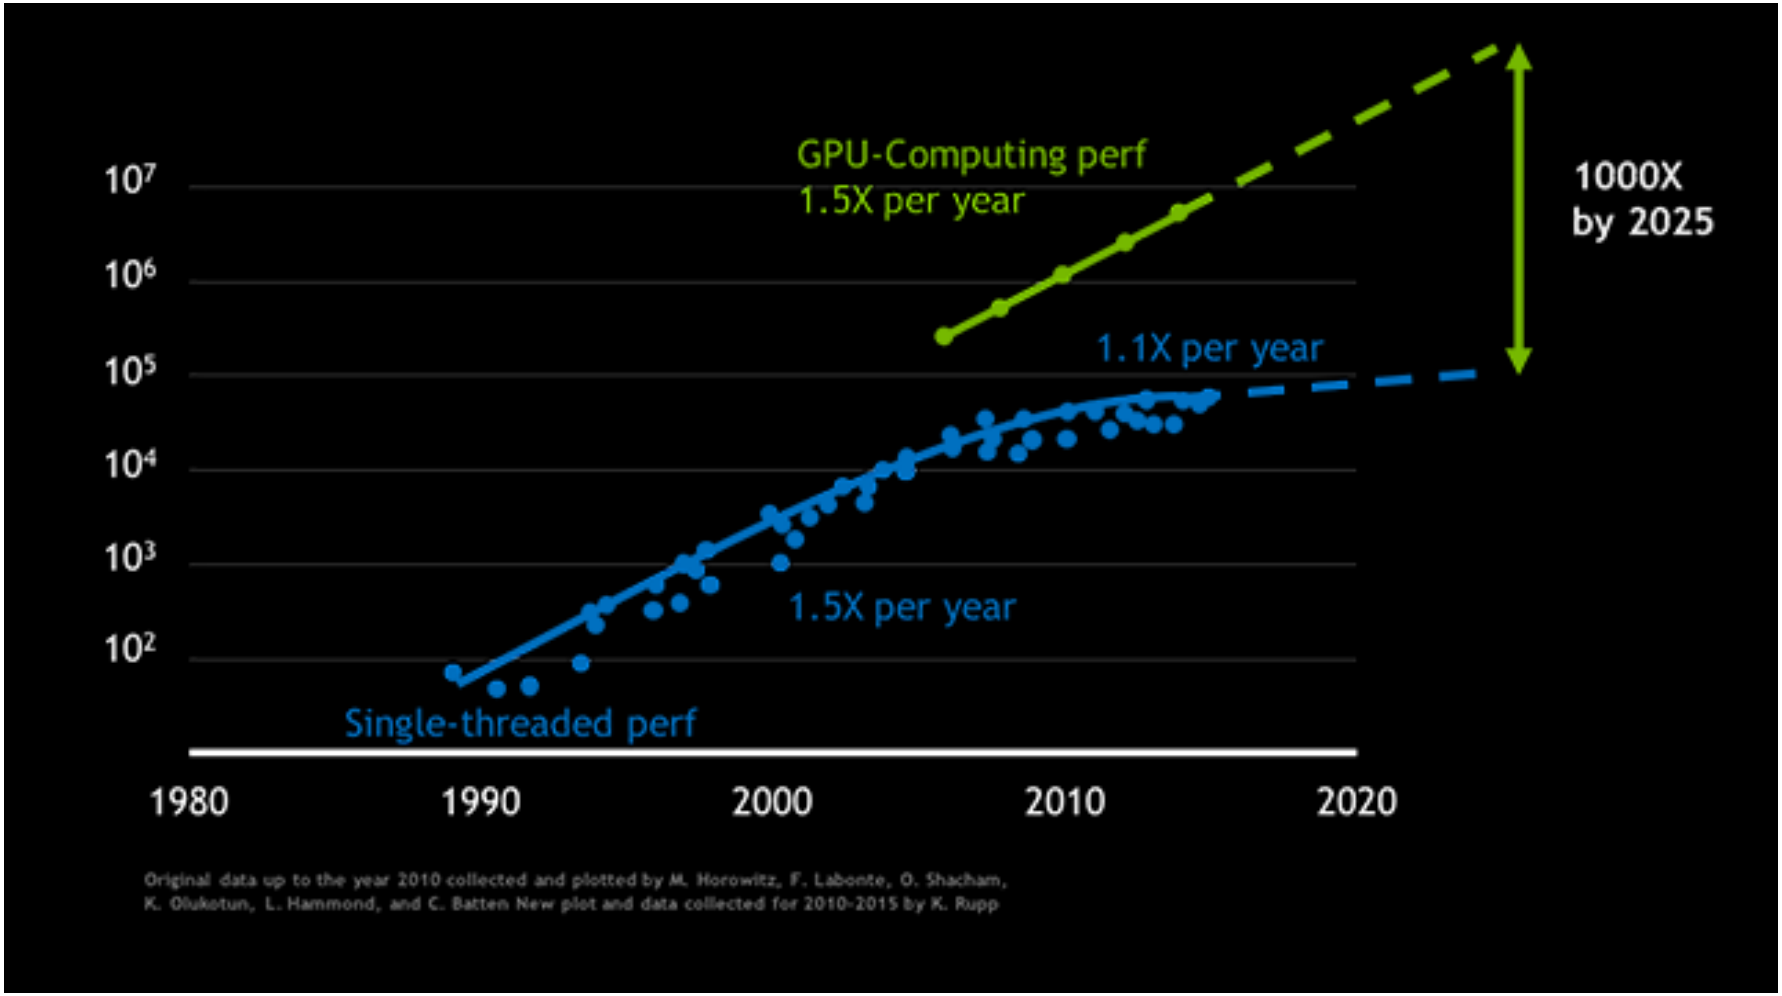

UNIVERSITY OF LIVERPOOL

A Recent Application of Parallelism

- 1,920 CPUs
- 280 GPUs

- Deep Learning
  - Convolutional Neural Net
  - Hill-Climbing

- Parallelism
  - Calculation of gradient as sum
  - Established implementation
    - Generically applicable
    - Quick to apply a scalable solution

- “Training took around 3 weeks”

PNY 12GB NVIDIA Tesla P100 PCI-E Passive GPU module

- 12GB PNY NVIDIA Tesla P100 module,
- PCIe 3.0 (x16),
- HBM2,
- GPU TBA,
- **3584 Cores**,
- 9.3 TFLOPS SP,
- **4.7 TFLOPS DP,**
- Passive

1.5x [1 core running at 3GHz]

- **£5,082.49 (www.scan.co.uk)**

UNIVERSITY OF LIVERPOOL

Solutions?


hadoop


MPI


OpenMP


NVIDIA
CUDA


Spark
Lightning-Fast Cluster Computing

Solution

Re-usable Generic Implementation

Adding Up a List of Numbers

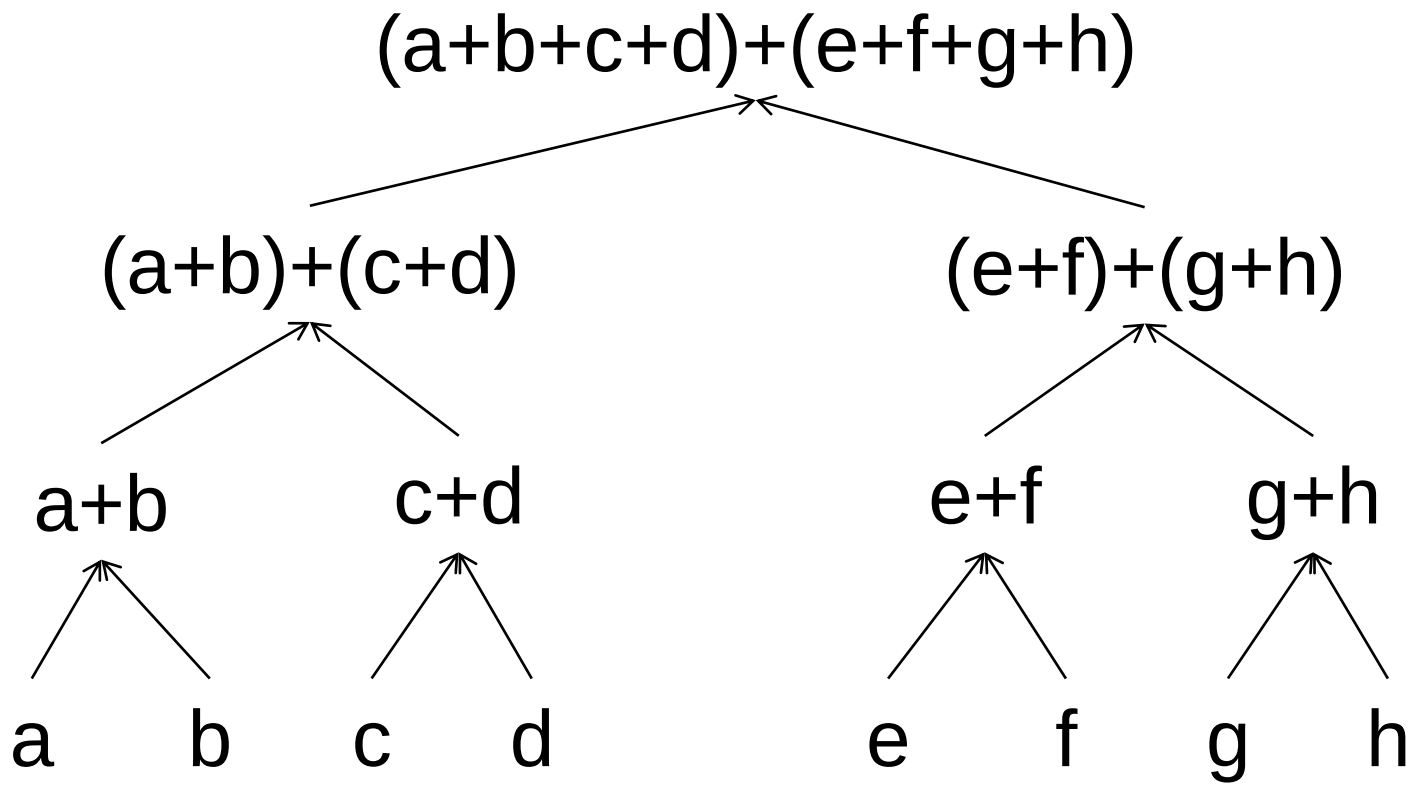

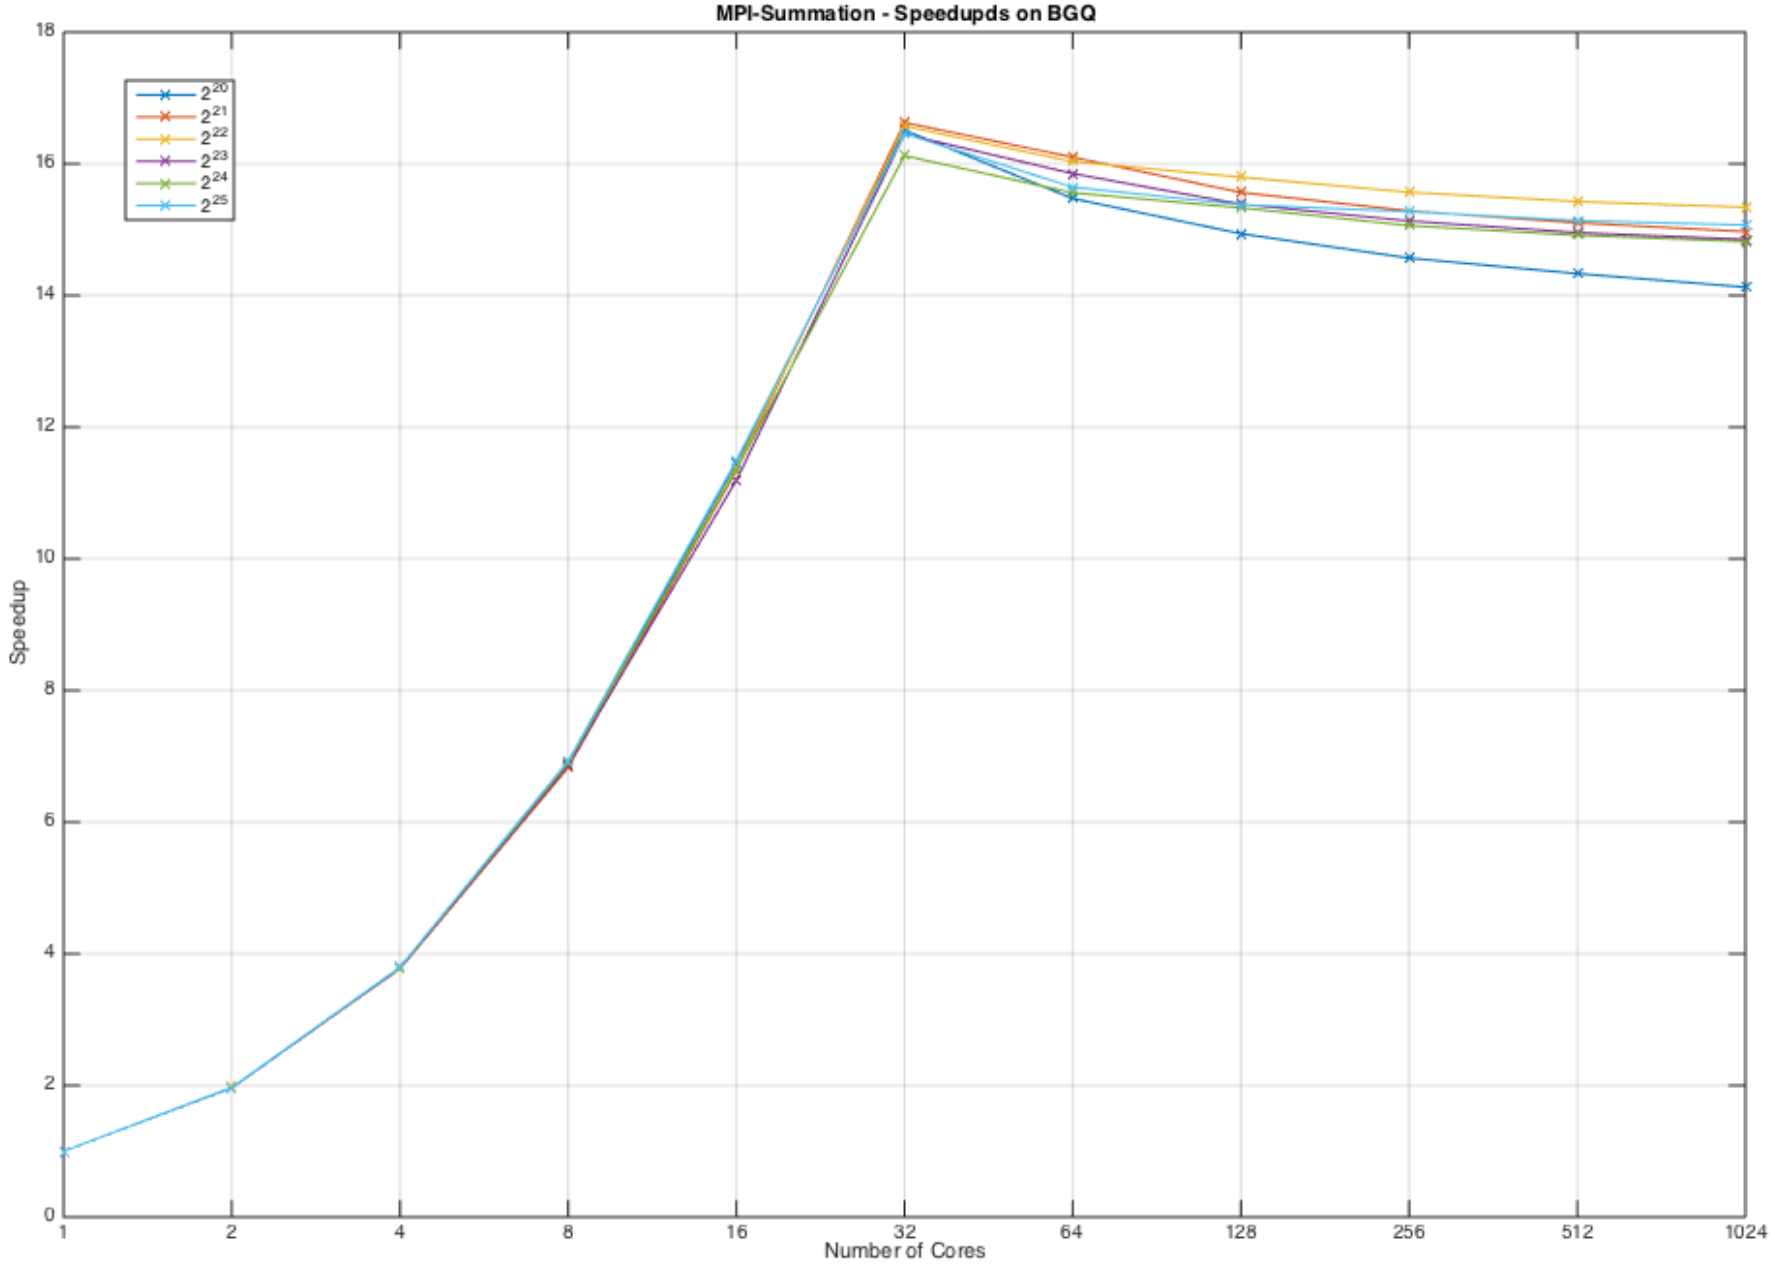

UNIVERSITY OF LIVERPOOL

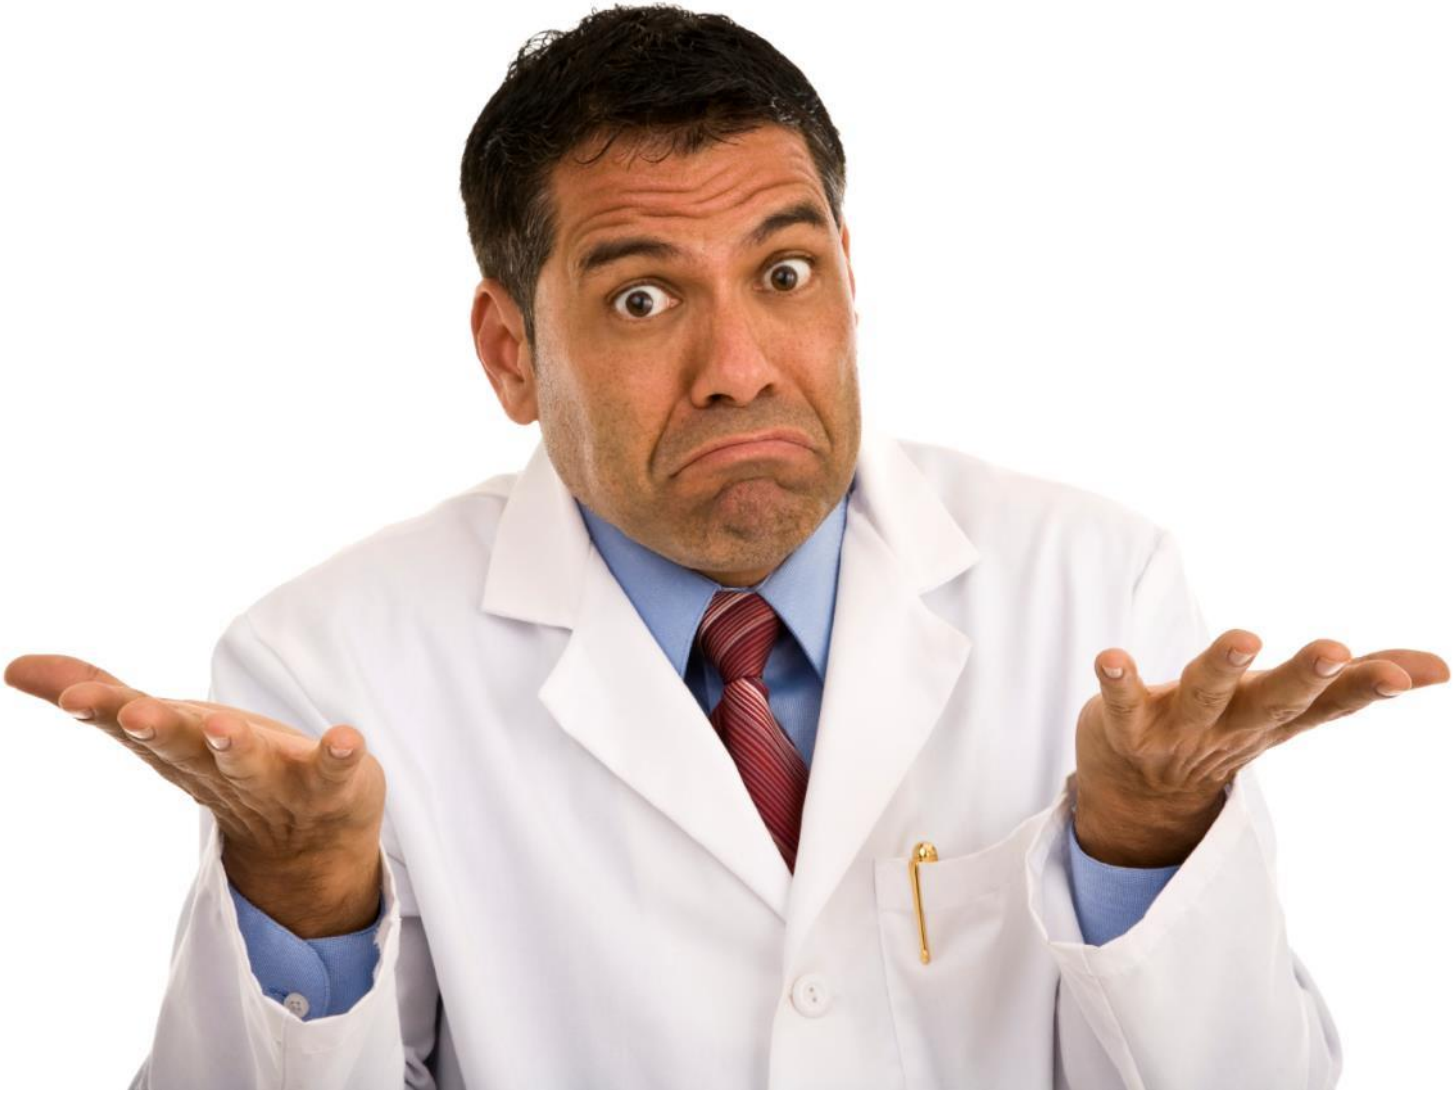

UNIVERSITY OF LIVERPOOL

UNIVERSITY OF LIVERPOOL

Adding Up a List of Numbers

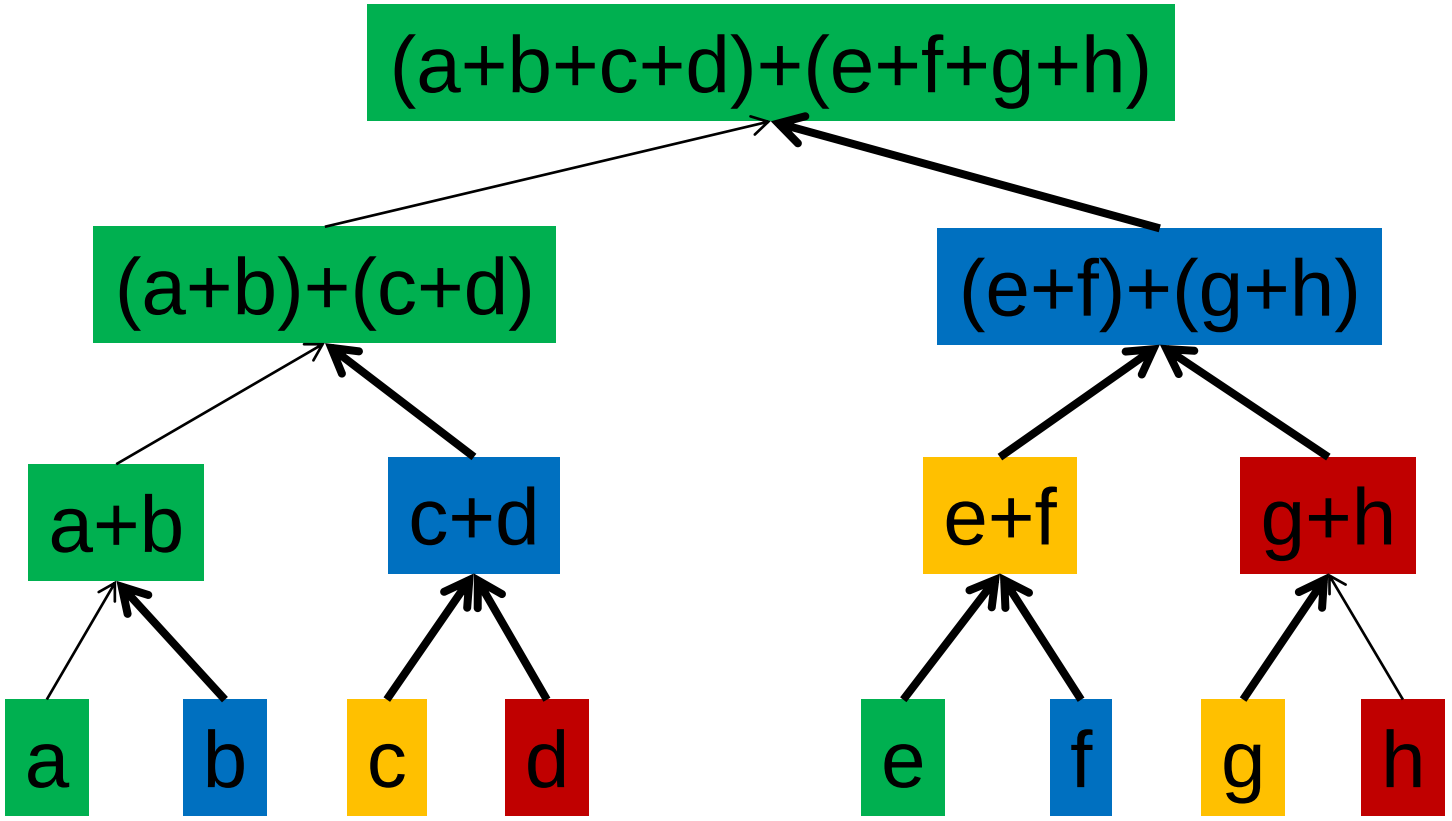

Adding Up a List of Numbers

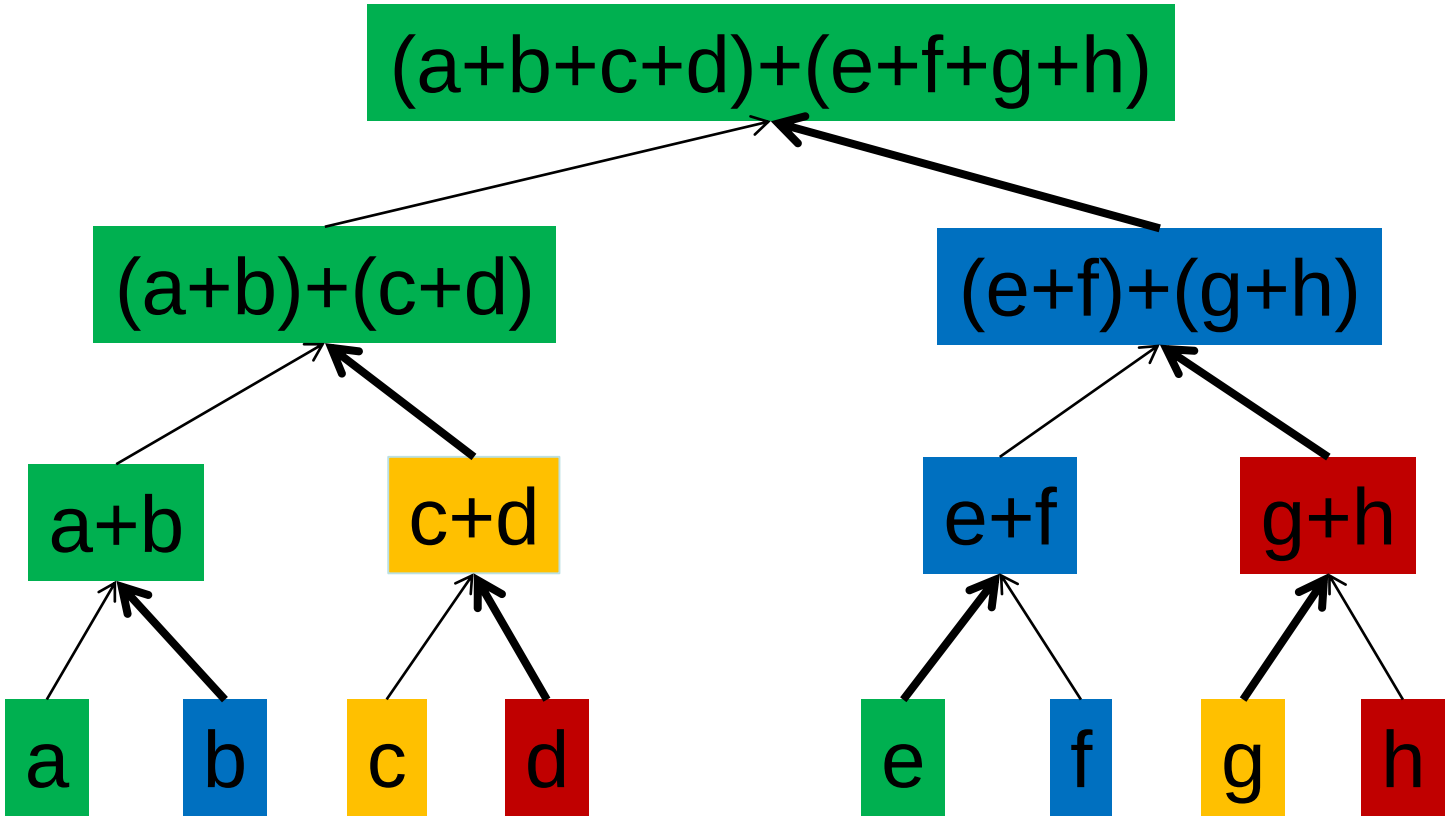

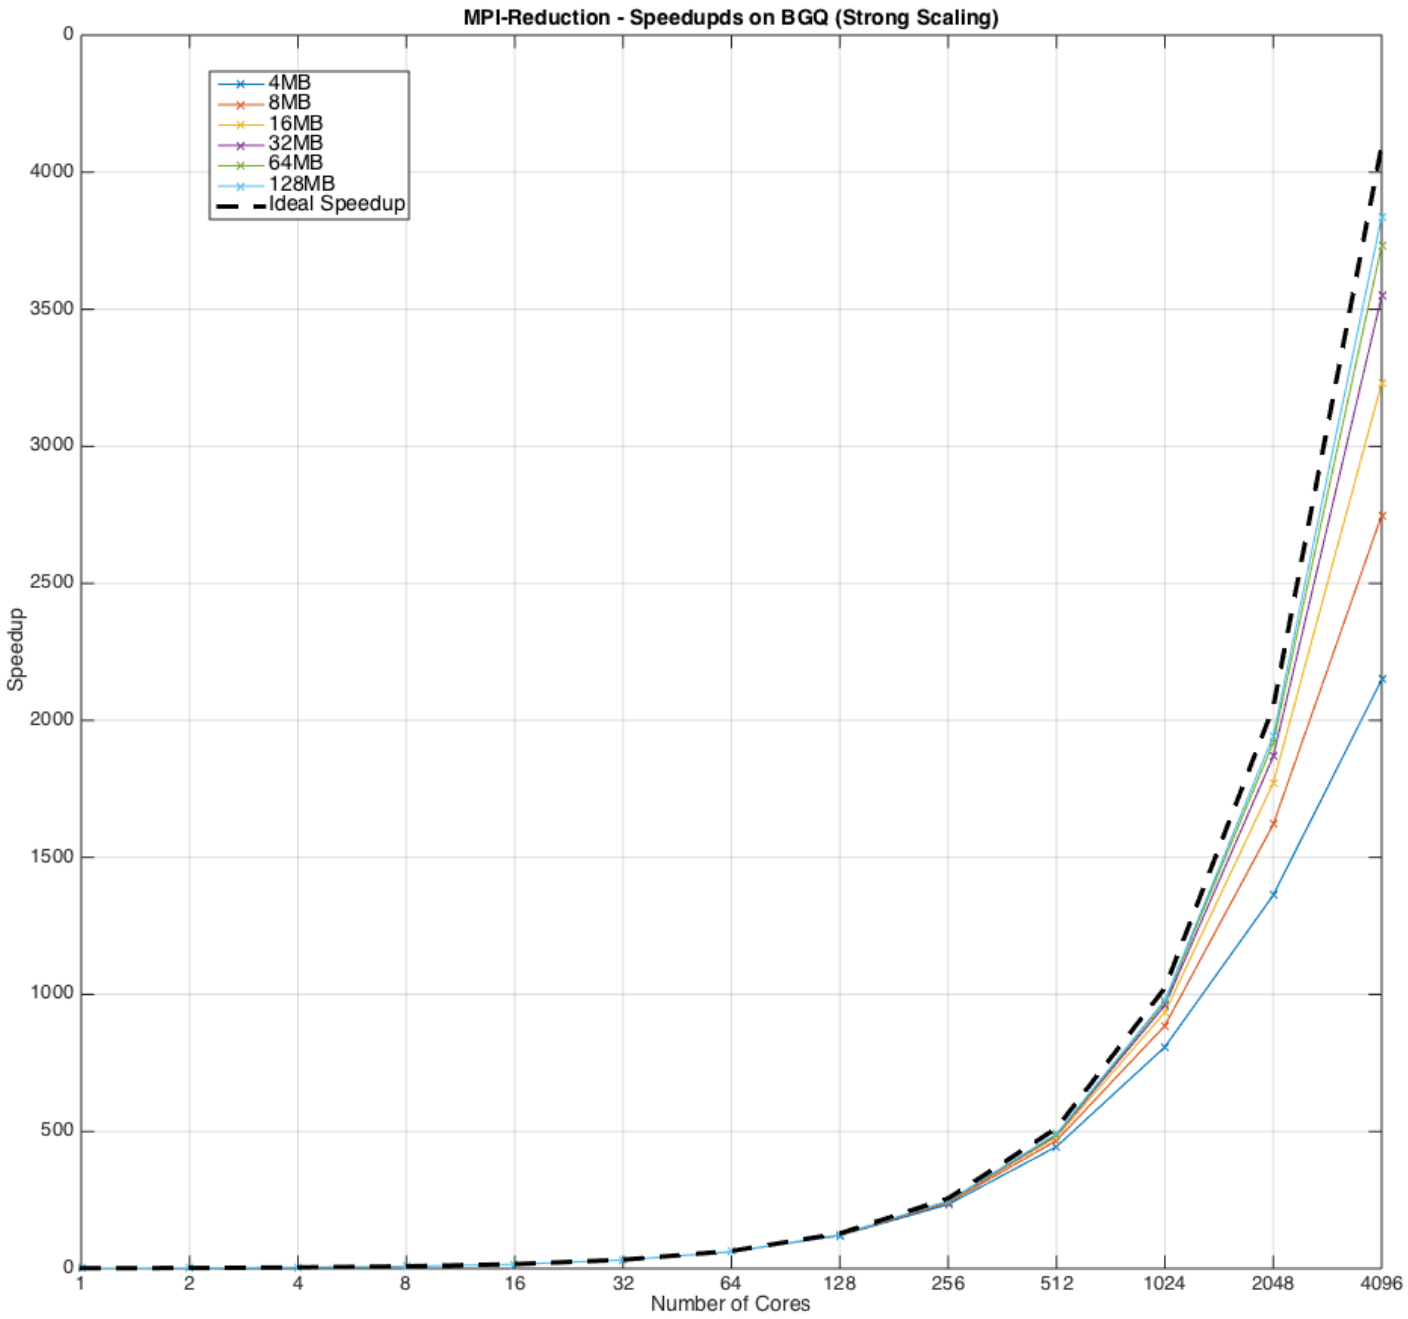

UNIVERSITY OF LIVERPOOL

UNIVERSITY OF LIVERPOOL

Implications and/or Conclusions

- Parallel algorithmic implementations are of increasing importance
  - Non-trivial to come up with

    Future progress is likely to involve small number of “large leaps”

- A popular generic implementation ≠ an optimal specific solution
  - Deep Learning may be rightly popular, but it’s not a panacea!

- Need to focus development on meeting needs of key communities
  - Algorithm developers really need to understand those needs

Workshop Questions

- What algorithms do you have available and how effective are they?
  - What makes the status-quo challenging?
    - Data is sparse
    - Ambiguity is pronounced
    - Performance cannot afford to be compromised

- What scenarios would you like to tackle?
  - What does Nirvana look like?
    - Navigating though the impenetrable literature to an out-the-box good algorithm
    - Running existing algorithms 50,000 times faster (6 mons = 52,560 x 5 mins)
    - Making the computer output better (and in line with human intuition and experience)

- [What do we need in terms of data to achieve this?]
  - What do we not have?
    - Access to a shared repository of data from which we can all beneift
    - More data related to rare events
    - Enough human time to analyse (retrospectively) uninteresting data

UNIVERSITY OF LIVERPOOL

Starter Questions

- What is your first name?
  - Simon
- Who do you work for?
  - University of Liverpool
- What two words should form part of our discussion?
  - Utility, Bayes

Algorithmic Workshop

Verification and Validation
Trust, Confidence, Law
Robustness to Attack
Black swans
HMIs / Man+Machine

Experimental design
Algorithms 101
Data repositories

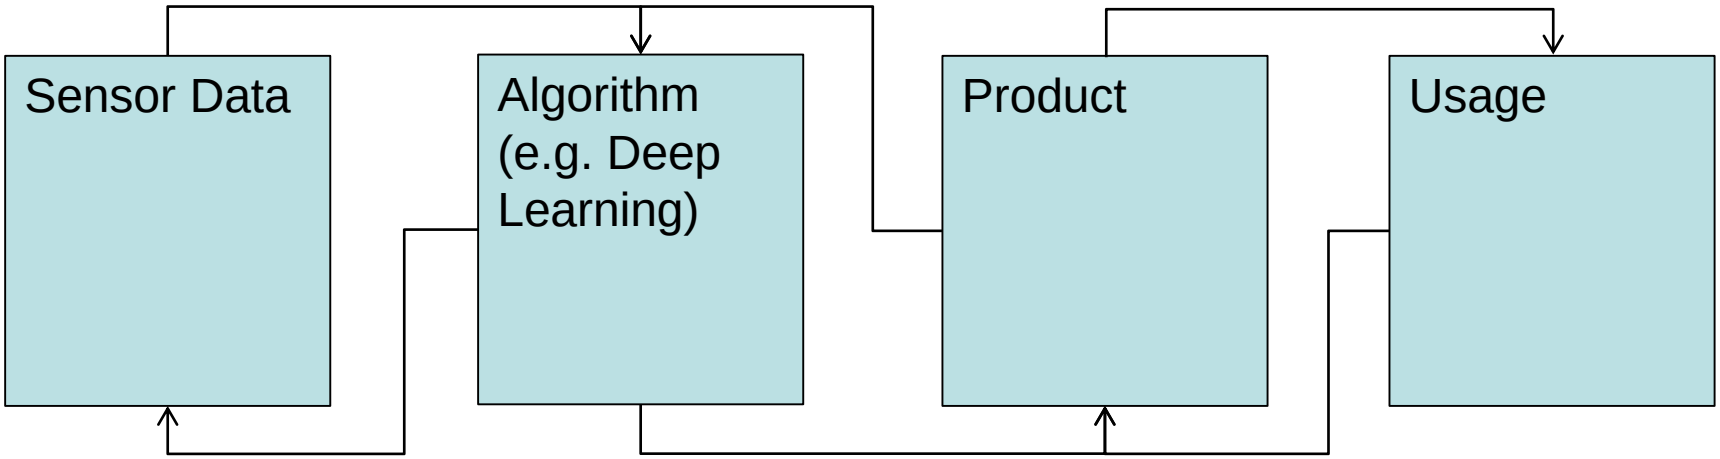

Measure ≠ Inform
IP Protection
Code Quality
Development Time