A Simple View of SDP Compute Complexity

- fitting 2000 hrs in 20 min

C4SKA at AUT
2018-02-15
tn@compucon.co.nz

Outline

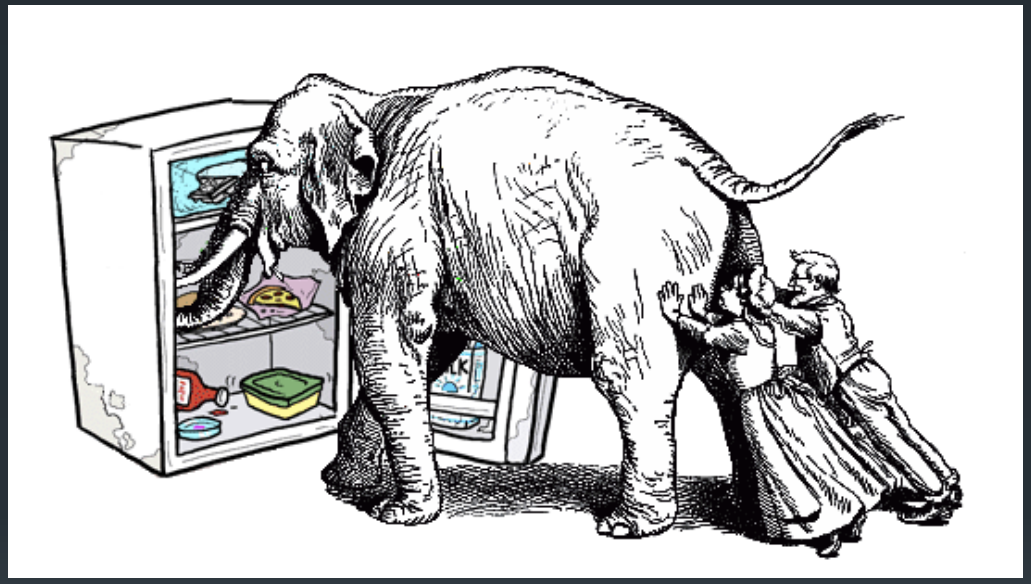

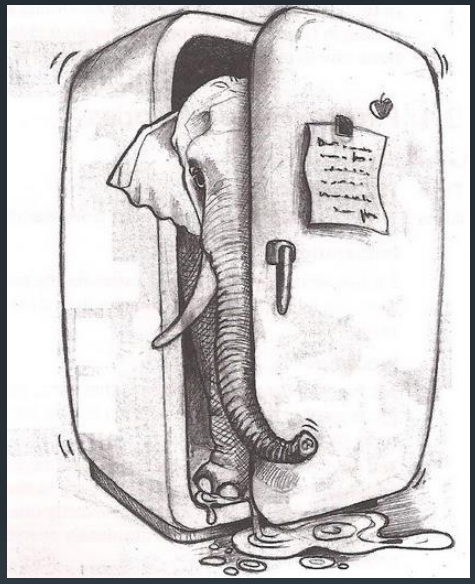

Image Credit http://uncyclopedia.wikia.com/wiki/File:Elephant_fridge.png

Technical Complexity

an order of magnitude better than the best

Image Credit https://thenounproject.com/term/telescope/137485/

- **Angular Resolution**
  - **Micro Arc Second scale**
- **Sensitivity**
  - **Under 200K**
- **SKA Physical Size**
  - **Phase 1 for 197 dishes in SA and 130,000 antennas in WA**
  - **Phase 2 is 10 times**
- **Challenges and Complexity**
  - **Epoch of Re-ionization for LOW**
  - **Calibration and Imaging for MID**

The Reality (8x discrepancy)

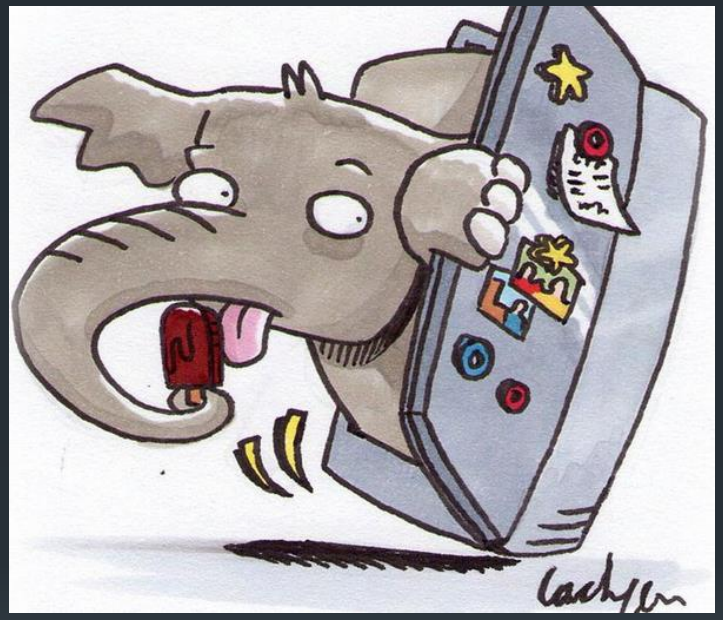

Image Credit: https://viccareers.com/2014/11/02/how-do-you-fit-an-elephant-in-a-fridge-coping-with-curveball-interview-questions/

- **Elephant is SDP**
  - **259 PFLOPS on 10% HU (parametric modelling)**
  - **60 GFLOPS/watt**
- **Fridge is Budget**
  - **60M Euro (US$72M)**
- **Global PC Industry 2017-11**
  - **#1 Green= 17 GFLOPS Linpack /watt, SDP NOT = LNP**
  - **Moore's Law hitting the wall**
  - **US$390M got 33 PFLOPS in 2013 (Tienhe2 on Xeon and Phi)**

Towards SDP CDR

this is a simple view of WBS

Elephant is bigger than Fridge by 16 times
8 from time based offering
2 from tech complexity

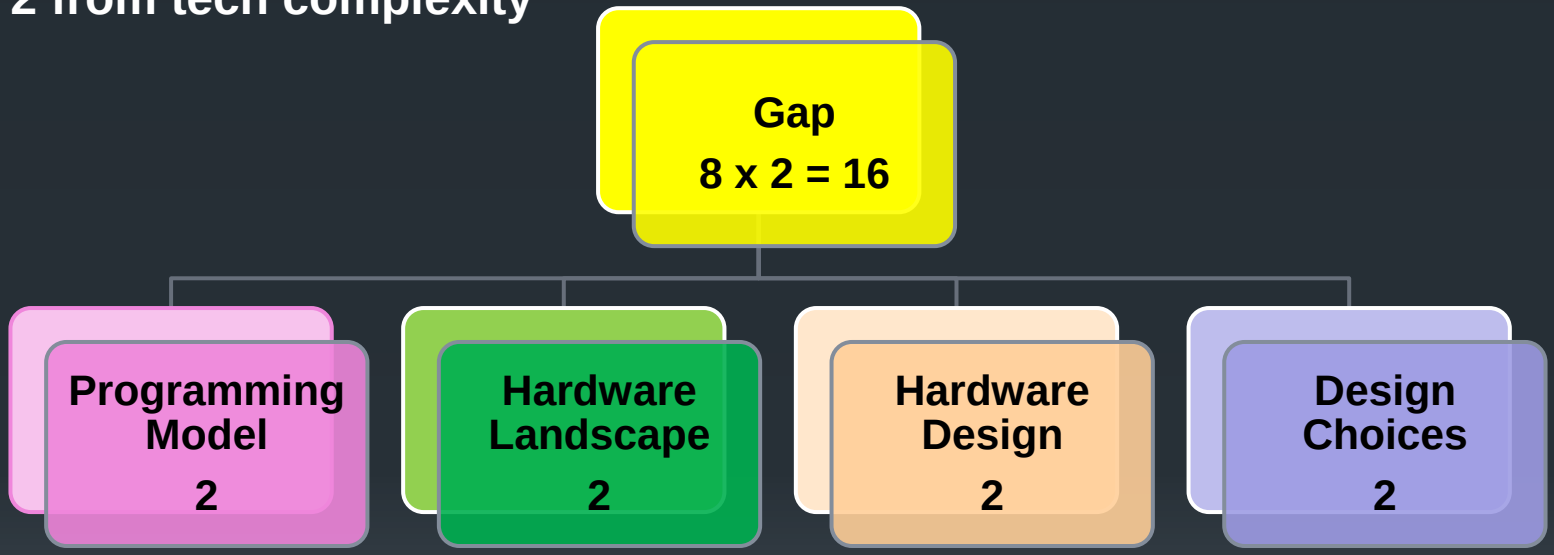

Remember the colour scheme for solutions

Measuring the Elephant

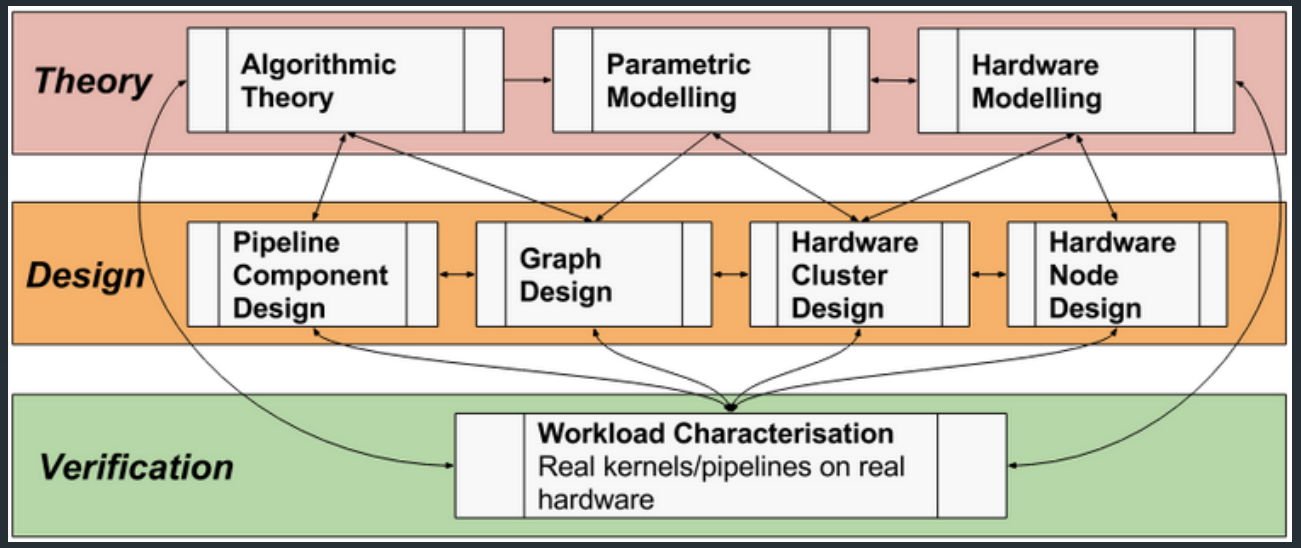

Diagram credit: SDP

Investigation TSK-1808 by Compucon

Continuum Imaging Pipeline (simplified)

5 Major Algorithms → 61% of runtime

Data Coalescing / De → 39%

Unable to get a single figure on HU vs 10% of 259 PFLOPS

Feed Peanuts to Elephant

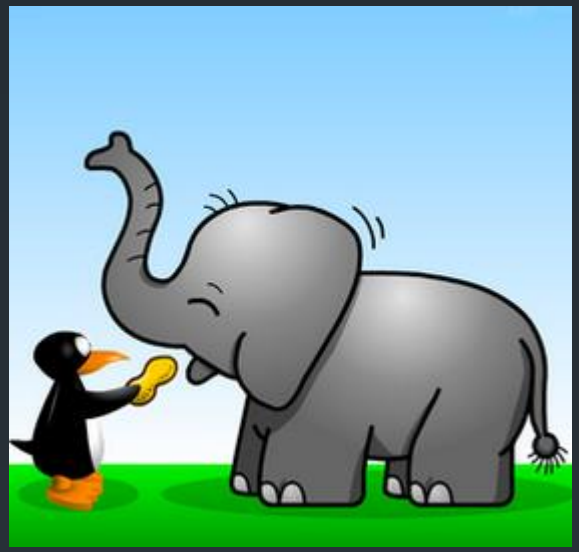

Image Credit: https://www.districtadministration.com/article/how-many-peanuts-should-penguin-feed-elephant

Investigation TSK-1916 (Compucon)

- **Continuum Imaging Pipeline**
- **Deconvolution MSMFS Clean (99%)**
- **5% of pipeline runtime**
- **Investigated SP vs DP, GPU vs CPU**

Elephant response was encouraging

- **Memory & Runtime improved for SP kernels (good)**
- **GPU HU < 1% SP vs 10% SDP sizing assumption (further investigation)**

Solution 1- Pink

- Continue peanut research path with multicore & GPU

TSK- 2090 (Compucon)- in progress

Image Credit
http://www.balkanstartups.me/

Three programming models

1. **Python Numpy- array single core**
2. **Numba- multicore CPU, 40% speedup with 2 cores (i3-7100)**
3. **PyCUDA for GPU, 6 times speedup with 768 cores (GTX1050Ti)**

Cal & Imaging Pipeline Scaling

- **Good Scaling for frequency, baseline, time slice, polarisation, facet**
- **Poor Scaling due to global synchronization blocks w/in pipeline (scattering and gathering)**

Solution 2- Green

Wise shopping for larger fridge

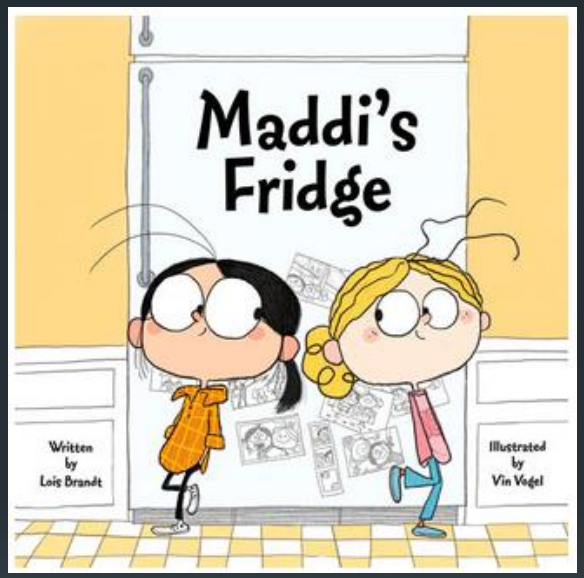

https://www.goodreads.com/book/show/20949046-maddi-s-fridge

- **Moore's Law for CPU Headroom Performance**
  - 2014 14nm
  - 2018 10nm
  - 2022 7nm
  - 2026 5nm
- **Reduce Bottlenecks by Micro-Architecture by 2020**
  - Memory Footprint x2
  - Memory Bandwidth x2
  - PCIe Bottleneck x2
- **Beyond 2024 for 50 years**
  - Algorithmic Accelerators
  - FPGA
  - In-Memory Computing (PCM)
  - Quantum Computing

3- Peach. 4- Purple

Hardware Design (Compucon 2105)

Each SDP consists of m nodes with each having n CPU and n . p GPU
Find X subject to argmax{X.C1} ^ argmin{X.C2} ^ argmin{X.C3} where X = (n + n . pi) per node and Ci refers to 1 = FLOPS, 2 = $, 3 = wattage

SP would improve the chance over DP as indicated by Amdahl

Design Choices (Peter Wortmann Memo 38 of 2017-12)

- Taylors Term = 5 (astronomer recommendation)
- Snapshot Length = 600s (for w-stacking & w-snapshots)
- Simultaneous Pipelines for MID (imaging and calibration)
- Sub-bands = 7 for LOW and 4 for MID (total independence)

Towards SDP CDR 2018-10

Solving complexity with simple views

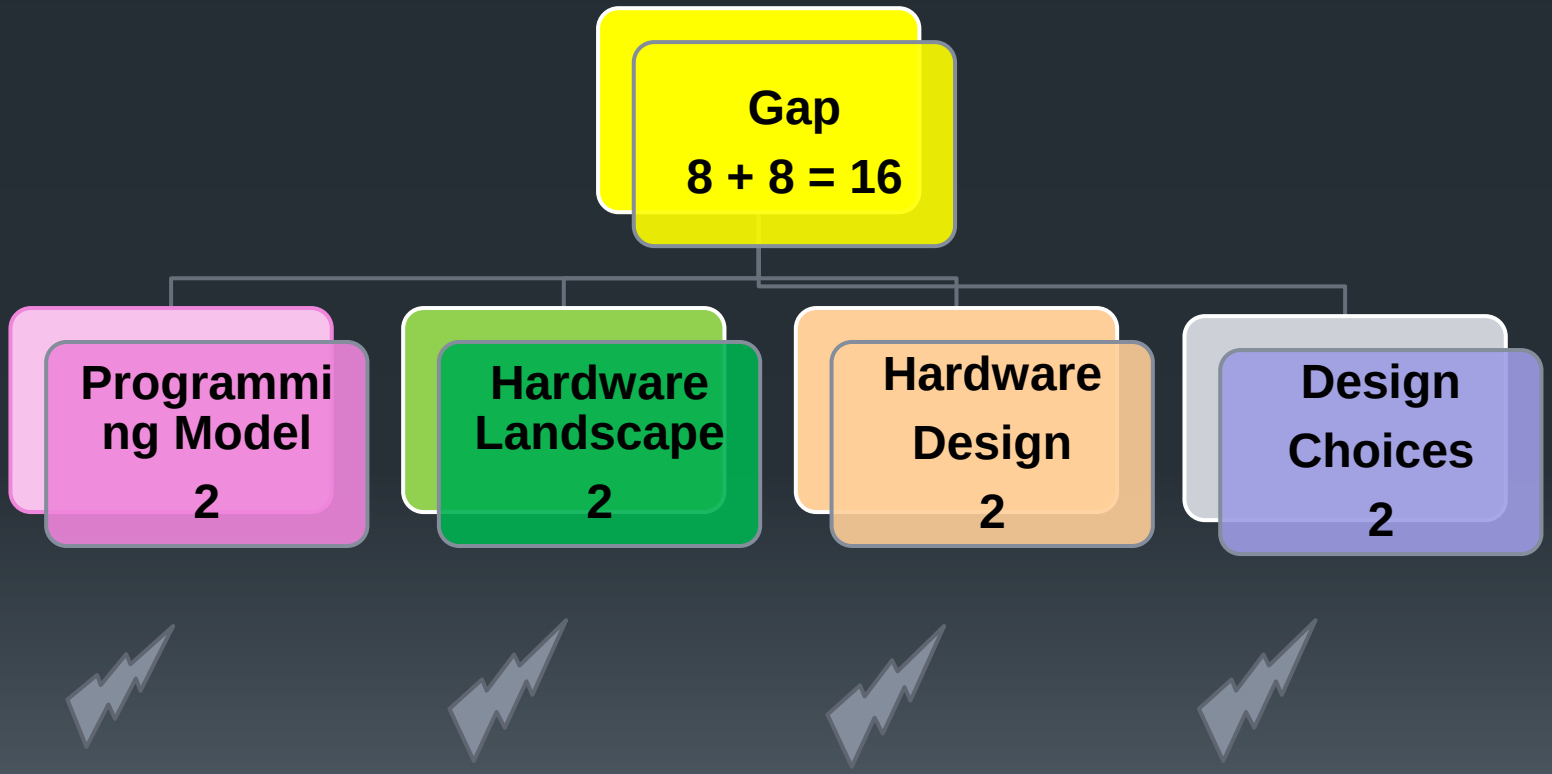

See the RHS Diagram!

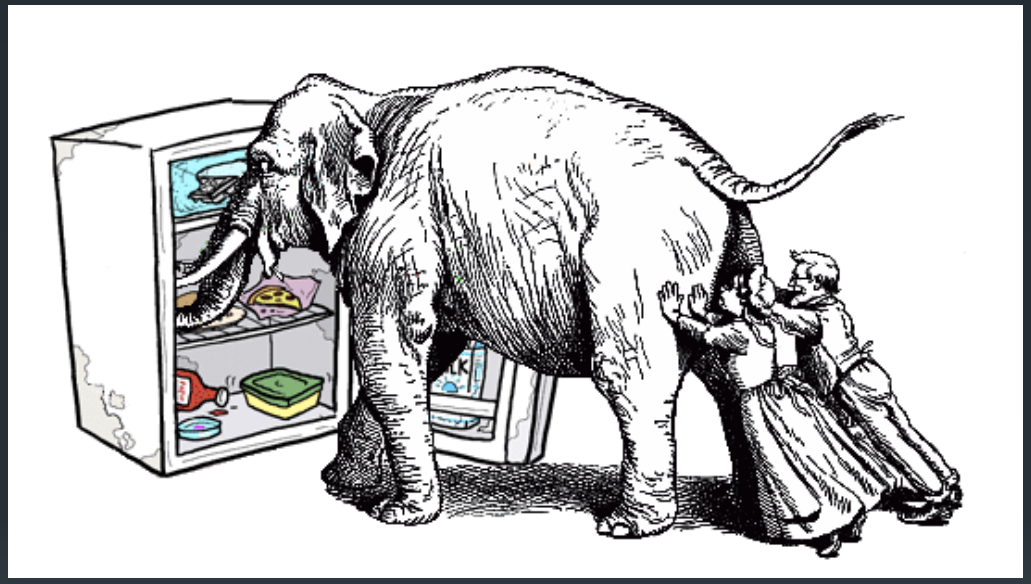

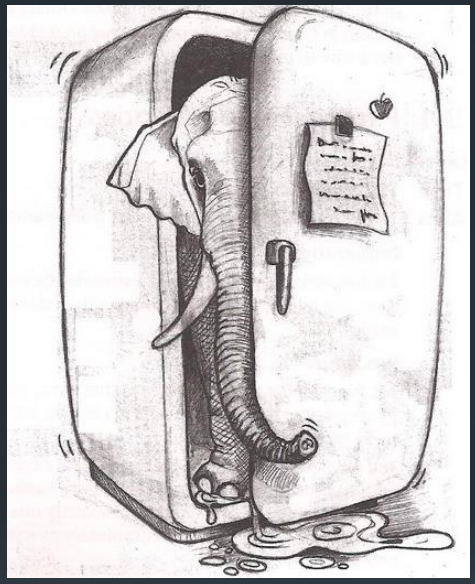

Image Credit http://uncyclopedia.wikia.com/wiki/File:Elephant_fridge.png