SDP Compute for 50 Years- an awareness

C4SKA at AUT
2017-02-09
tn@compucon.co.nz (personal opinion)

Outline

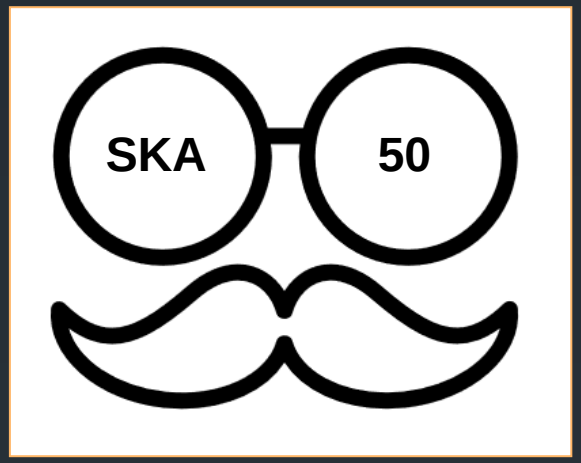

Image Credit
https://thenounproject.com/term/old-man/191197/

- Costs of SKA1 in Euro
  - Design and Construction ~ 1B
  - Operating = 115M * 50 years
  - Electricity ~ 8%
  - SKA2 budget ~ 10 times?
- Challenges on SDP
  - Wave after wave of revamping
  - 9 Overviews
  - 1 Philosophical View- not that!

SKA Process Overview

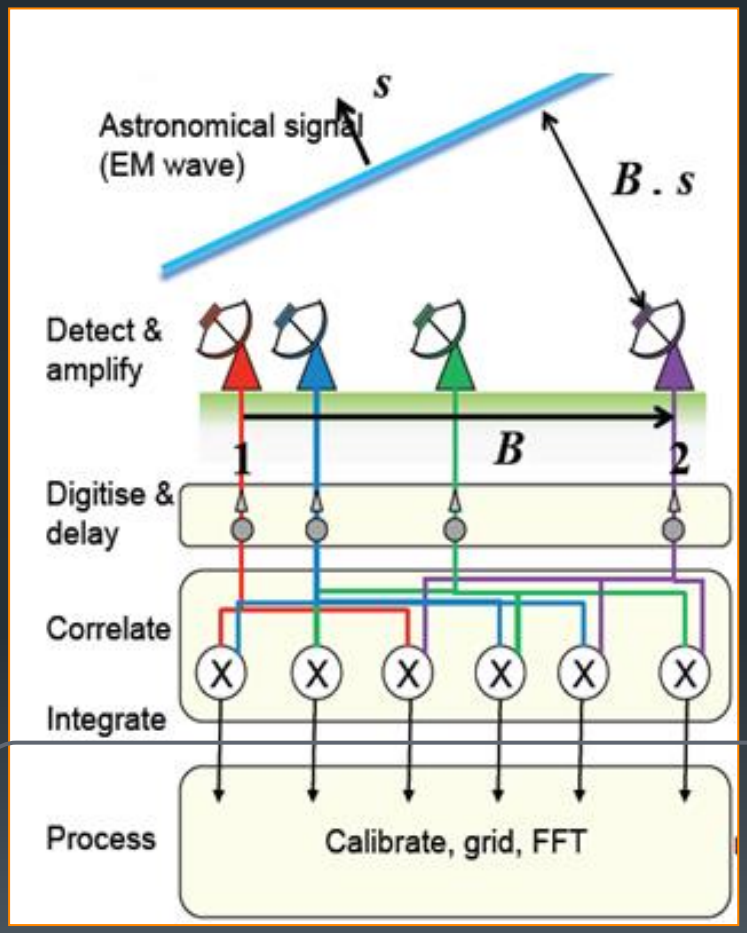

CSP

SDP

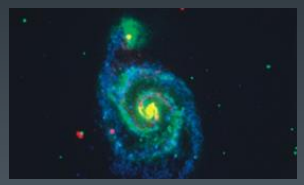

Image credit: SDP Consortium Update 2016-06-21 Jeremy Cole page 12

SDP Science Objectives Overview

Table 1: Available sustained compute rates and CAPEX costs associated with target power caps.

Power Cap*	LOW 4 MVA	MID 4.5 MVA	MID 10 MVA
Sustained Compute Load Total (PFLOPS)	27.8	29.3	70.5
Hardware CAPEX Estimate (M€)**	38.9	47.2	106.7

*Note: This is "apparent load", Power Cap includes PUE=1.2 and Power Factor=0.85

**Note: Power Estimate and Cost Estimate excludes EoR uv archive for LOW.

- **System Sizing paper (SDP-038, Cambridge, 2016-0721)**
  - **Assumed 25% hardware utilisation efficiency (remember this)**
  - **Not based on peak computational capacity requirements**

SDP Design Overview

(not to scale)

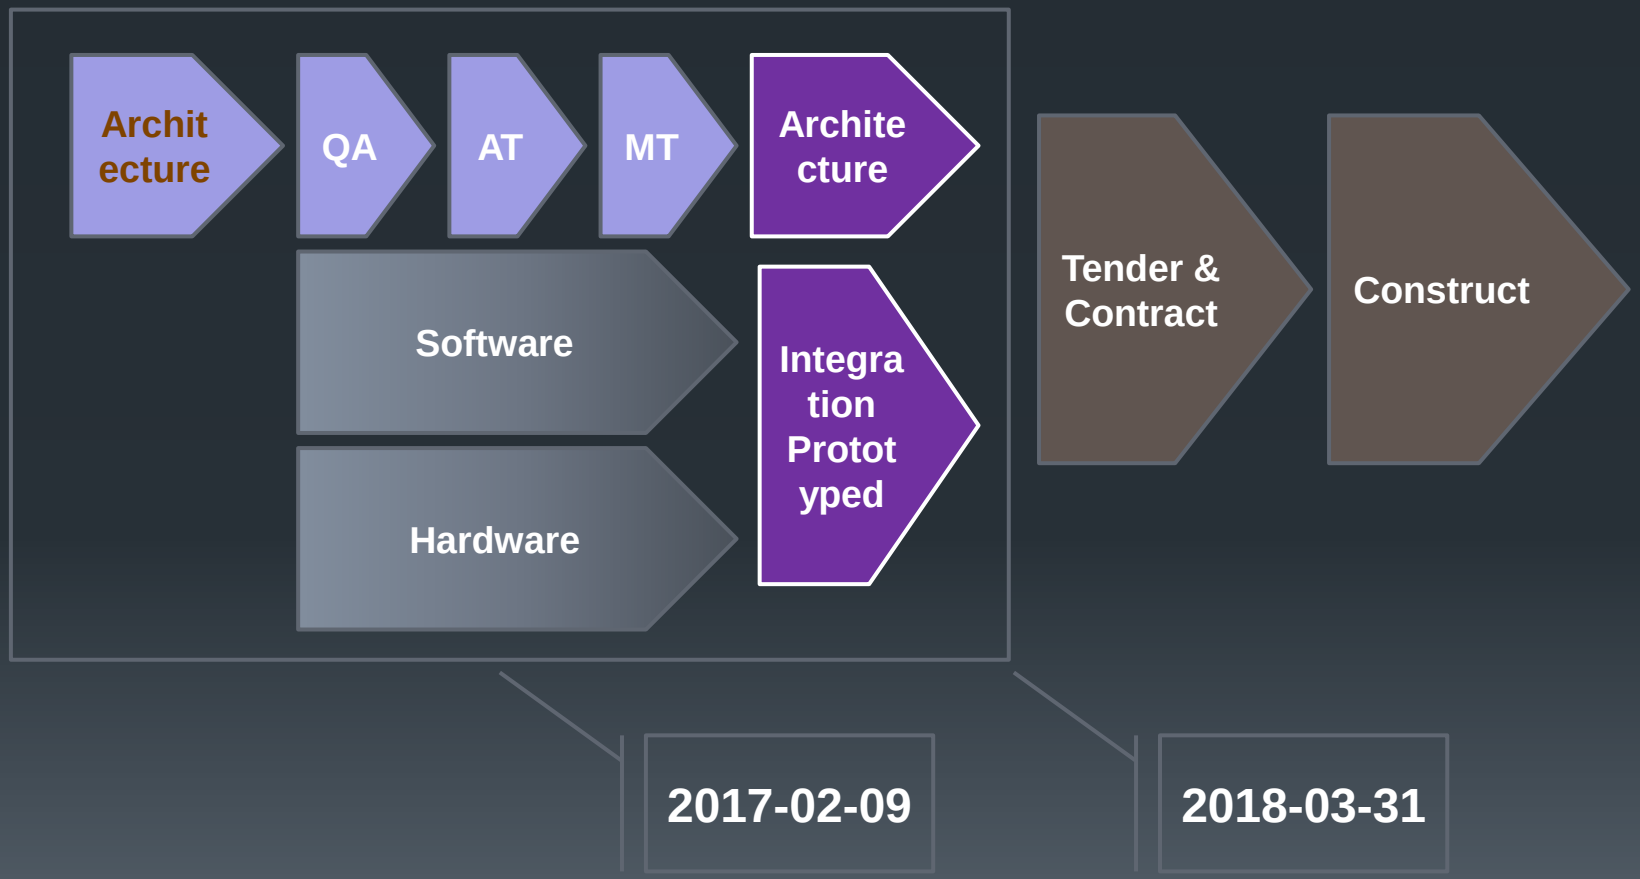

SDP Algorithm Overview

Continuum Imaging most demanding & reduction process is iceberg of the tip below

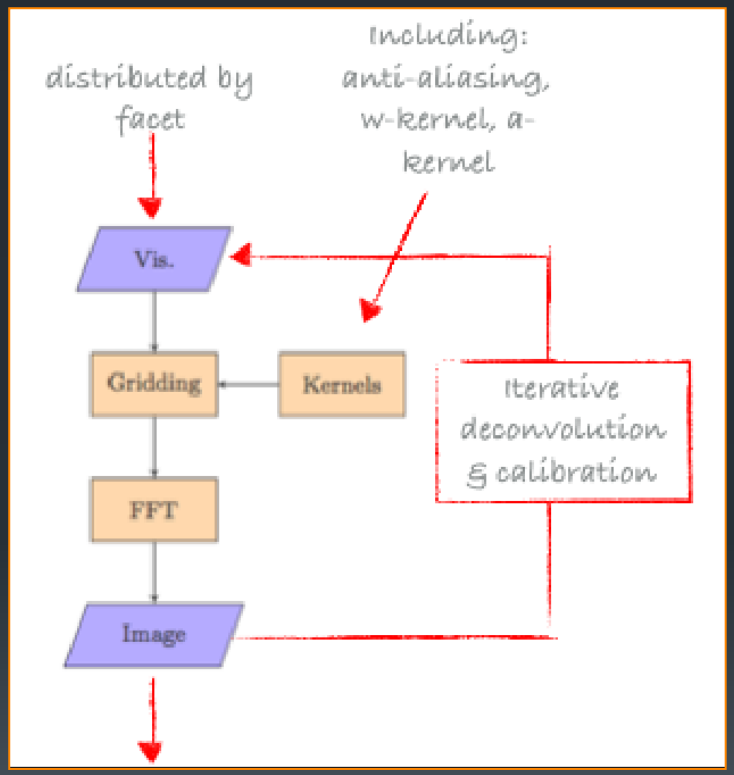

SKA Imaging Challenge Anna Scaife page 37
SDP-030 figure 12 & 16

- Algorithm AI (dataset size dependent)
  - Gridding: $Nk^2 > 100$
  - FFT: $5/32 \log_2 N < 5$
  - Cleaning ?
- x86 Hardware AI (fixed)
  - CPU < 5
  - Xeon Phi ~ 5
  - GPU > 25
- Challenge
  - Assumption of 25% HWU?

SDP Hierarchy Overview

Prescribed by Execution Framework SDP-015

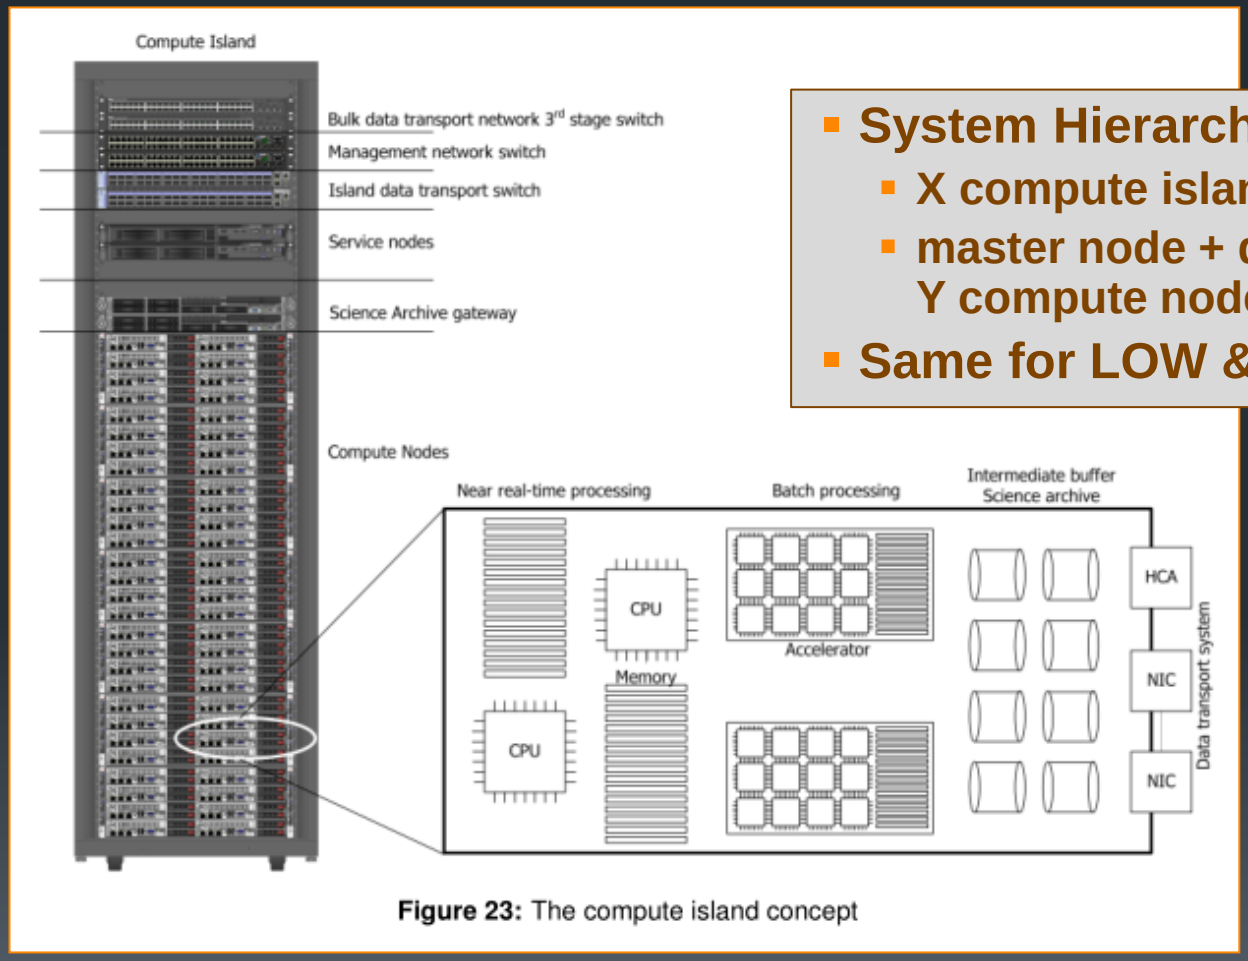

Figure 23: The compute island concept

- **System Hierarchy**
  - **X compute islands**
  - **master node + data buffer + Y compute nodes**
- **Same for LOW & MID**

Source: SDP-018 Data Processor Platform (ASTRON 2016-0715)

Hardware Parameters Overview

Meet $, wattage, HWU, AI

Table 1

Target	LOW	MID
MVA	4	4.5
PF	0.85	0.85
PUE	1.2	1.2
MW	2.83	3.19
PFLOPS	27.8	29.3
PF/MW	9.81	9.19
GF/w	9.81	9.19

Table 2

Target	LOW	MID
PFLOPS	27.8	29.3
TFLOPS	20	20
Eff-base	0.25	0.25
Node	5560	5860
per rack	33	33
Racks	168.5	177.6
Eff-discr	0.26	0.26
Racks	162.0	170.7

Table 3

Target	LOW	MID
MW	2.83	3.19
Node	5560	5860
W/node	510	544

Table 4 on AI attainable on # GPU

Assume GPU capable of 1TB/s implying 5.5PB/s LOW 5.8PB/s MID

PFLOPS	27.8	29.3
AI	No. of GPU	
0.5	55600	58600
1	27800	29300
1.25	22240	23440
1.66	16747	17651
2.5	11120	11720
5	5560	5860

Racks	168.5	177.6
Nodes /R	33	33
Nodes	5560	5860
No. of GPU/Node		
AI = 1.25	4.0	4.0
AI = 1.66	3.0	3.0
AI = 2.5	2.0	2.0
AI = 5	1.0	1.0

- Hardware Sizing Guidance
  - 1TB/s
  - 5TFLOPS SP
  - AI = 5
  - 25% HWU
  - ~500W
  - >5000 nodes @

Technology Trend Overview

1965 – 2015 Moore's Law

Cost per transistor has been on a linear decline during Moore's Period

http://wccftech.com/intel-10nm-cannonlake-ice-lake-tiger-lake-cpu/

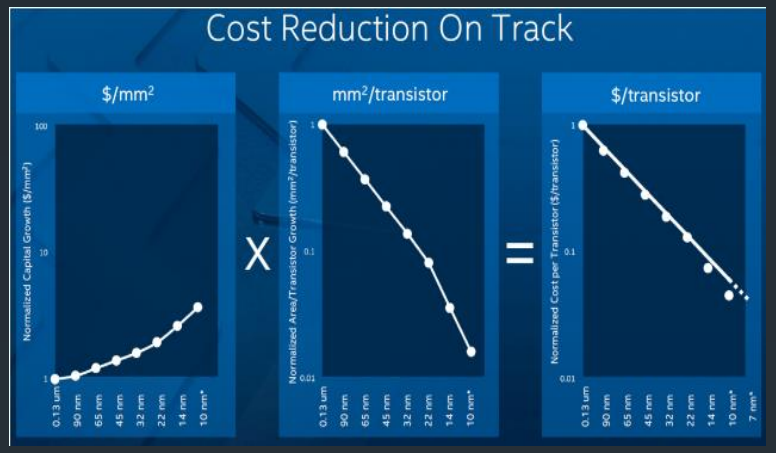

2015 – 2030 Moore's Wall

Reaching physical limits

Trends and Challenges in Big Data, Ion Stoica, UC Berkeley, PDSW-DISCS 16, slide 43

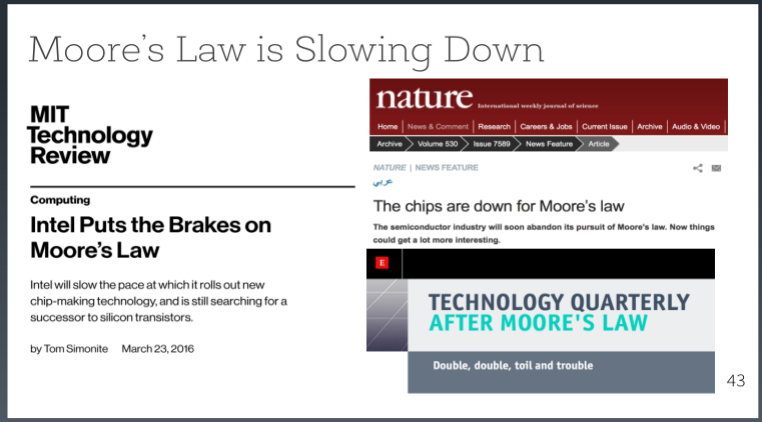

Near Term Technology

- CPU
  - 25% performance improvements
  - Mostly by increasing number of cores, SOC and SIP (such as EMIB)
- Memory
  - 35 per year
  - Stacked technologies
  - Attain 1TB/s in 2017 (AMD Vega HBM2)
- Network
  - 40% per year
  - 100/200/400GbE NIC on horizon

Long Term Technology

- 2017 – 2030 Evolutionary
  - Software catching up on hardware
  - Allows expansion of science objectives
  - SKA2 on scope, dynamic range & resolution
- 2030 – 2070 Revolutionary
  - New programming model
  - FPGA
  - Quantum
  - Memory Processing
- Immediate Thought
  - Sub-arraying is a good idea

New Zealand 50 Years

- Views for SDP
  - High or Low language abstraction?
  - Open or Closed SW standards?
  - More emphasis on HW knowledge
- NZ well positioned
  - Contribute to SKA
  - Spin off via academic-industry collaboration projects
    - Collate range of expertise
    - Promote open architecture

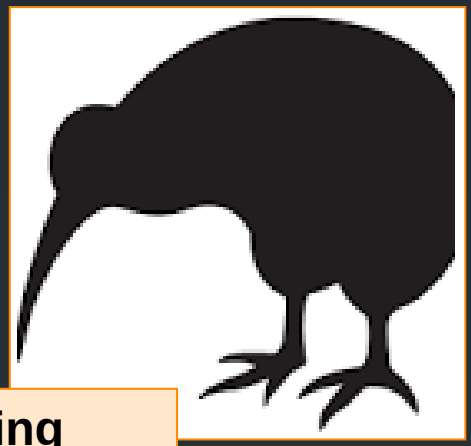

Being aware of situation