2011 Signal Processing CoDR: Technology Roadmap

W. Turner SPDO

14th April 2011

Technology Roadmap

Objectives:

- Identify known potential technologies applicable to the SKA
- Provide traceable attributes of technology in terms of:
  - Performance
  - Cost
  - Thermal Dissipation
- Where possible project attributes of technology to the time frame of the SKA
- Provide an overview of potential future technologies that may be applicable to the SKA within the time frame of the SKA1 or SKA2.
- List ‘also ran’ technologies that have been considered but have been considered unsuitable in their current format

Technology Readiness

- Concept descriptions have focused on established processing technologies:
  - x86 multi-core
  - GPU
  - FPGA
  - ASIC
- Networks
  - 40 G bps and 100 G bps Ethernet
  - 10 G bps or DWDM
- Roadmap for these technologies to 2015 reasonably well understood and stable.
- More risky technologies (TRL 3 and below):
  - Storage Class Memory
  - Printed circuit board optical interconnect

These technologies are at a technology readiness level of 7 or above for existing Radio Telescope implementations SKA implementation using these technologies in 2015 will decrease this to technology level of the order of 4.

General Purpose Processor

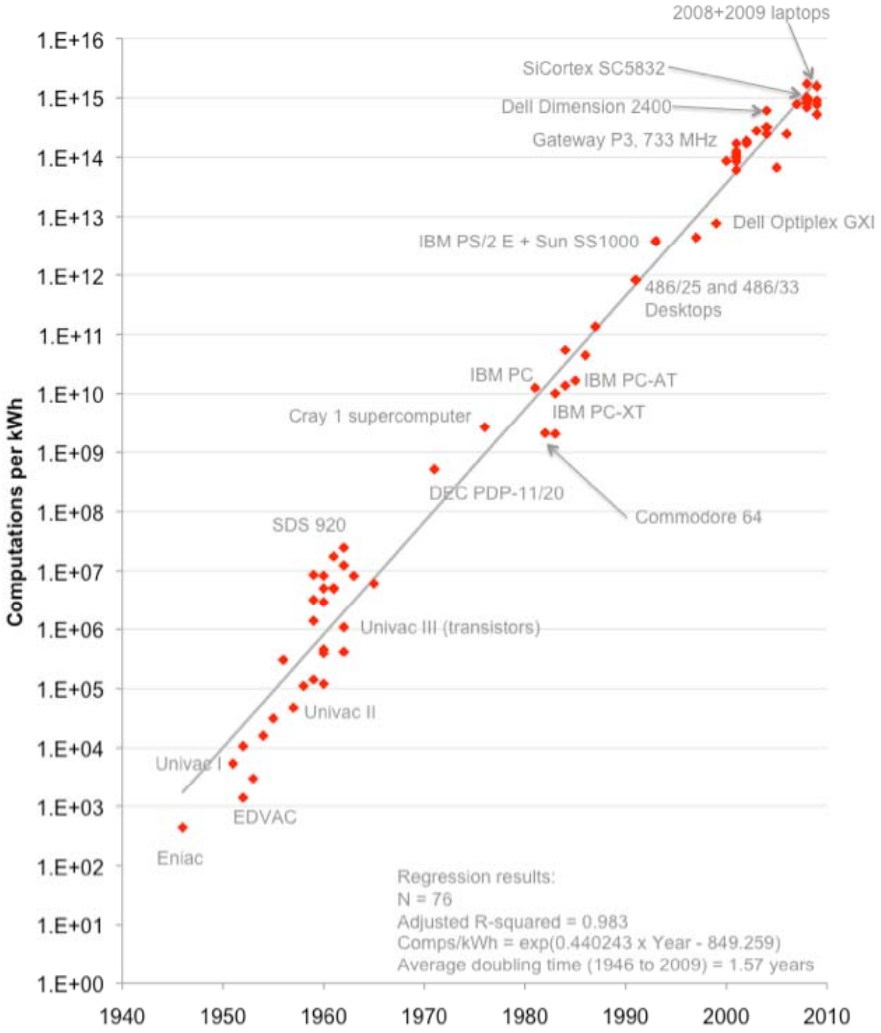

(J G. Koomey Stanford)

Projecting from this graph suggests
2.7 x 10^{16} computations per kW hour by 2015
or alternatively 7.5 x 10^{15} computations per second per Mega watt dissipation.

Architecture Change		Fabrication Process	Release Date	Energy scaling	Delay Scaling
Tick	Shrink/derivative (Penryn)	45nm	2008	0.5	> 0.7
Tock	Microarchitecture (Nehalem)		2009		
Tick	Shrink/derivative (Westmere)	32nm	2010	0.5	> 0.7
Tock	Microarchitecture (Sandy Bridge)		2011	0.5	> 0.7
Tick	Shrink/derivative Ivy Bridge	22nm	2012	0.5	> 0.7
Tock	Microarchitecture Haswell		2013	0.5	> 0.7
Tick	Shrink/derivative Rockwell	16nm	2014	0.5	~1
Tock	Microarchitecture TBD		2015	0.5	~1

Intel's Tick Tock Roadmap

$C_{\_(x86\_2015)} \approx$ 490 G FLOPS
Power ≈ 130 Watts per processor chip
Cost ≈ 230 Euros

Graphics Processing Units/ Streamed Processing

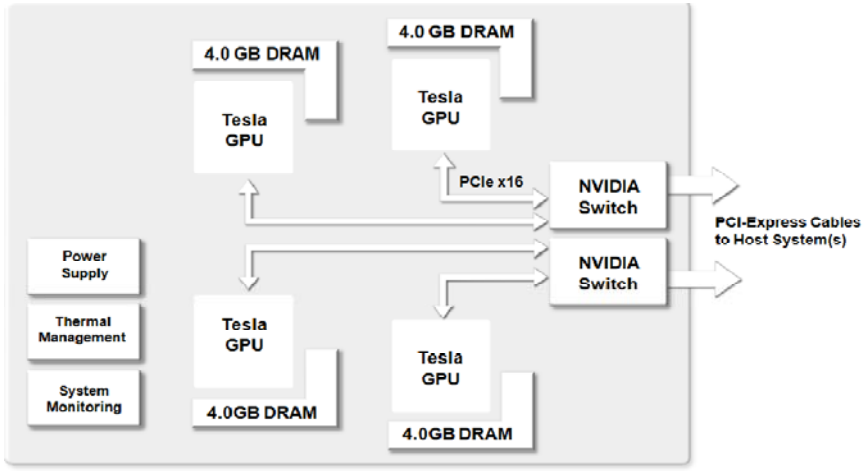

Nvidia Tesla Architecture

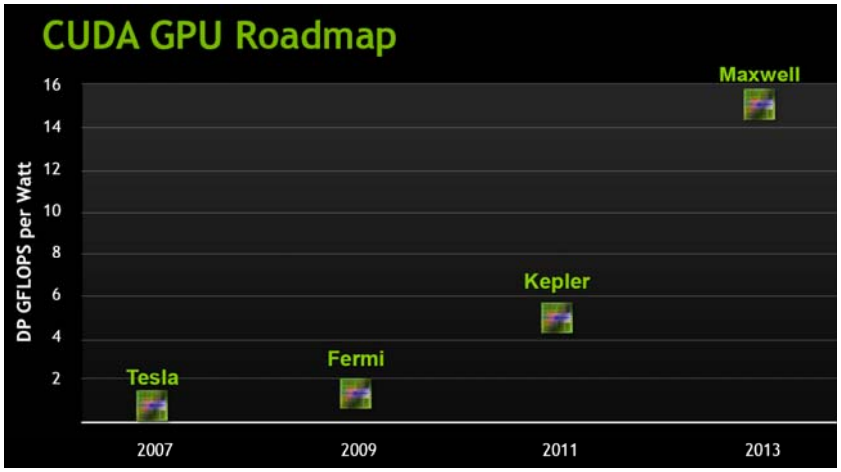

Nvidia processing Power per Watt Roadmap

- Nvidia and ATI are currently the key providers of GPU processing engines
- GPU acceleration used on the current world's fastest computer in the Top 500 (2.6 P Flops)
- Intel are entering the market with many core technology code named"Knights Bridge" and "Knights corner:

http://download.intel.com/pressroom/archive/reference/ISC_2010_Skaugen_keynote.pdf

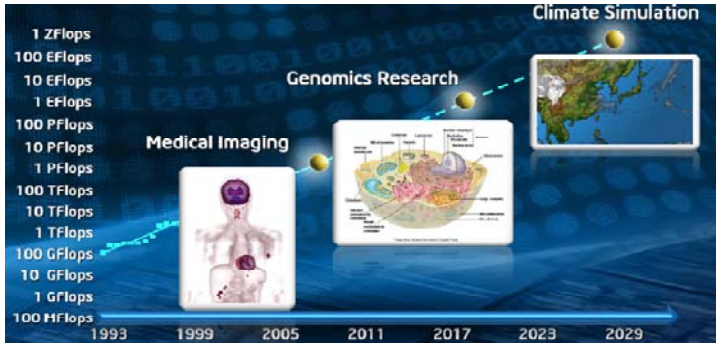

$C_{\_(GPU\_2015)} \approx$ 20 T FLOPS (Needs Host server)
Power ≈ 130 Watts per processor chip
Cost ≈ 12k Euros for 4 off GPUs in rack mount unit

Field Programmable Gate Array

UNIBOARD

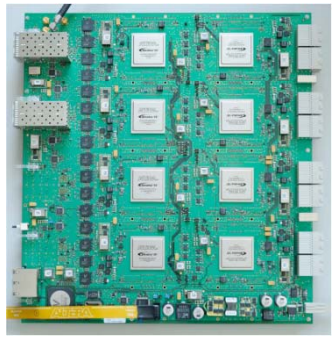

Red Back 2

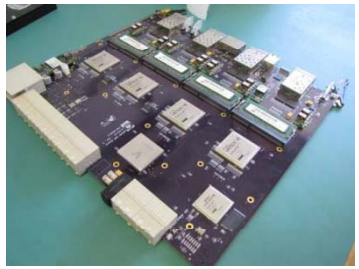

ROACH 2

Two main manufacturers of FPGAs:

- •Xilinx used on Casper Roach and ASKAP Red back boards
- •Altera used on UNIBOARD.

•New generation of FPGA every 3 years: expect two further generations by 2015.

•Taking Xilinx as an example:

- •Virtex 7 expected to be available Q2 2012
- •The top of the range Virtex 7 devices support 3960 DSP slices and are likely to clock at up to speeds of 600MHz.

600 MHz ×3960 DSP slices ≈2.4 T MACS*

•Assume Virtex 8 available in 2015 that doubles this performance for the same thermal dissipation

•Note a CMAC can be performed in 1 clock cycle for 4 bit data

$C_{\_(FPGA\_2015)}$≈ 4.8 T MACS per device (18 bit)
Power ≈ 120 Watts per device
Cost ≈ 8k Euros unit price (actual price is significantly lower)

* MACS Multiply Accumulate per second

ASIC

Used on ALMA and WIDAR Correlators Low power consumption & cost but high NRE

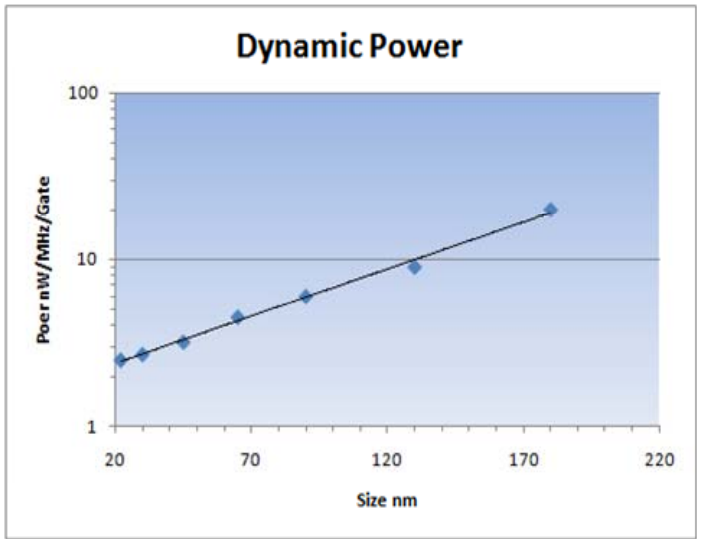

Projected dynamic thermal dissipation in nW/MHz/Gate

Scenario

- An optimised 4 bit multiplier accumulator can be constructed from ~ 500 gates
- 65nm technology the order of 400,000 gates can be implemented per millimetre square
- 22nm technology this increases to 6400 multiplier accumulators per square millimetre
- The processing power of these multipliers will depend on how fast they can be clocked and the thermal dissipation for the device
- Taking a reasonable but arbitrary clock rate for the ASIC of 400 MHz and assuming 16 mm^2 area for the multipliers

$C_{\_(ASIC\_2015)} \approx 400\ MHz \times 16\ mm^2 \times 6400\ multipliers/mm^2 \approx 40\ T\ MACS$

$Power \approx 400MHz \times (6400 multipliers/mm^2 \times 500 gates \times 16 mm^2) \times 2.4 \times 10^{-9} \div 2 \approx 25\ Watts^*$

Cost ≈ estimated die cost 1 Euro per die and ~ 15 Euro packaging

*1.6 T MACs per Watt assuming 50% gates switch per clock cycle

Scalability: Moore's Law

- Moore's Law
  - The number of transistors that can be placed inexpensively on an integrated circuit doubles approximately every two years.
- Breaking Moore's Law
  - Lithography - reduced dimensions makes mask production very difficult
  - Process technology complexity and maintaining yield
  - Length of interconnects on chip, leading to increasing propagation delays and parasitic capacitance
  - Reduced gate oxide thickness below 1 nm, leading to fluctuations in doping profiles (100 atoms long gate length, less than 100 dopant atoms)
- Rock's Law
  - the tooling cost for semiconductor die manufacture doubles every two years

Scalability: Amdahl's Law

- Amdahl's Law
  - The speedup of a program using multiple processors in parallel computing is limited by the sequential fraction of the program
  - Can't assume doubling the number of processors or processing cores will also double the available processing power
  - May need to bench mark scalability

$$Perf = \frac{1}{(1-f) + (f/n)}$$

Where:
n is the number of processors
f is the fraction in parallel ($0 \leq f \leq 1$).

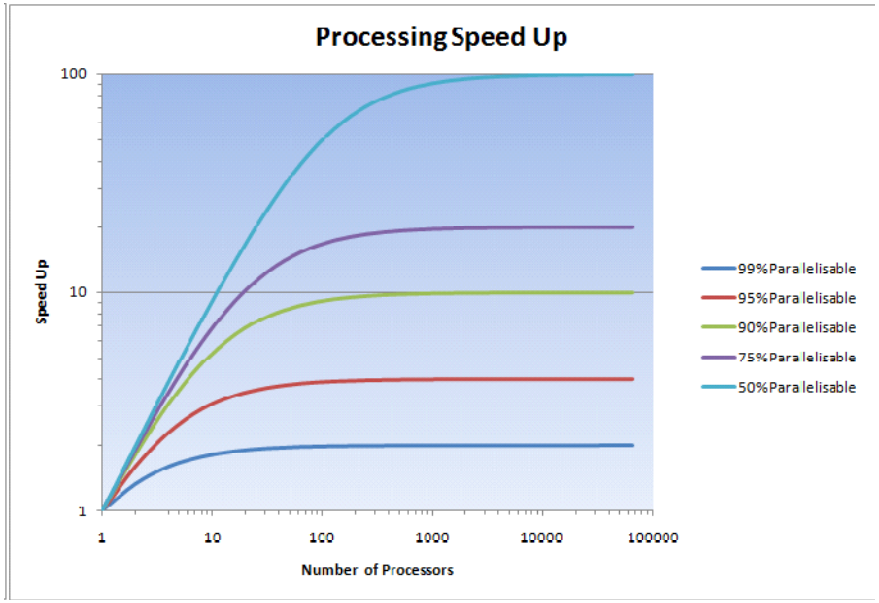

Conclusions

- Reasonably well defined roadmap to 2015 for:
  - Processors
    - X86
    - GPU
    - FPGA
    - ASICS
  - **Networks**
    - Infiniband
    - Ethernet

Multi-sourced
Used on existing projects
Large User base
Support Tools available

- Interesting Technologies being developed may be applicable for SKA2
  - Storage Class Memory
    - Potential to store & export “voltage” signals for Non-Imaging Processing
    - Low cost, low power, high reliability
  - Optical Interconnect
    - Reduces power consumption
- Need to regularly update the roadmap