




Next Generation FPGA / Gemini

John Bunton
CSIRO Astronomy and Space Science

14th February 2019 - C4SKA @ AUT



FPGA History

- **Start**
- **The Revolution of 2000**
- **Evolution – parameters for correlators and beamformers**
- **Second revolution - 2019**
 - **HBM**
 - **Multicore**
 - **RFSoc**

Gemini riding the revolution

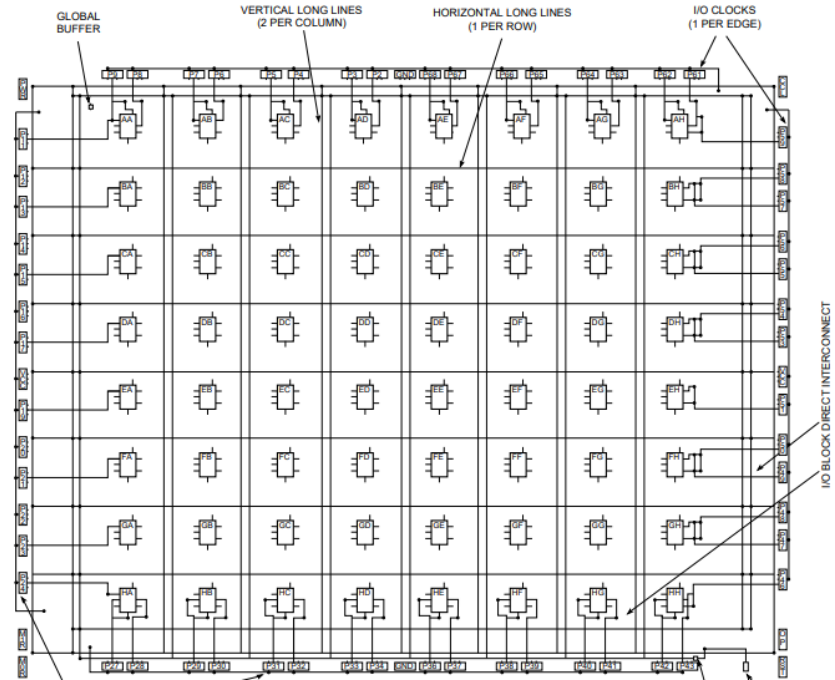
The GPU challenge

Early FPGAs (Last Century)



Model	Launch	Logic Block		
XC2018	1985	100	4 Variable	?
XC4000	1999	6272	4 variables	80MHz

Only Logic
No multipliers
No Serial I/O
No block memory



1999 to 2002 Revolution



Model	Year	Kilo LUTs	RAM 18kbit	DSP	TMACS	SERDES Gbps	CPU
Virtex-E	1999	36	4kit x 208				
Virtex-II	2000	93	168	164	0.034		
Virtex-II Pro	2002	88	444	444	0.11	120	PowerPC

In 3 years we went from a logic array to an FPGA capable of true signal process

High bandwidth on-chip RAM, True multipliers and high speed I/O, even on chip processor (PowerPC)

2005 to Now - Evolution



Model	Year	Kilo LUTs	RAM 18kbit	DSP	TMACS	SERDES Gbps
Virtex-4	2004	55	320	512	0.14	
Virtex-5	2006	150	1032	1056	0.3	90
Virtex-6	2009	298	2128	2016	0.7	238
Kintex-7	2011	359	1910	1920	0.7	400
Kintex UltraScale	2014	2533	4320	5520	2.4	1024
Virtex UltraScale+	2016	1728	5200 + Ultra	12288	6.0	4000

Initially backwards – no CPU, no SERDES (V4)

CABB correlator - VII-Pro for I/O and V4-55 for DSP

Sufficient I/O



Return of SERDES to Standard Part

- 2006 - Virtex 5

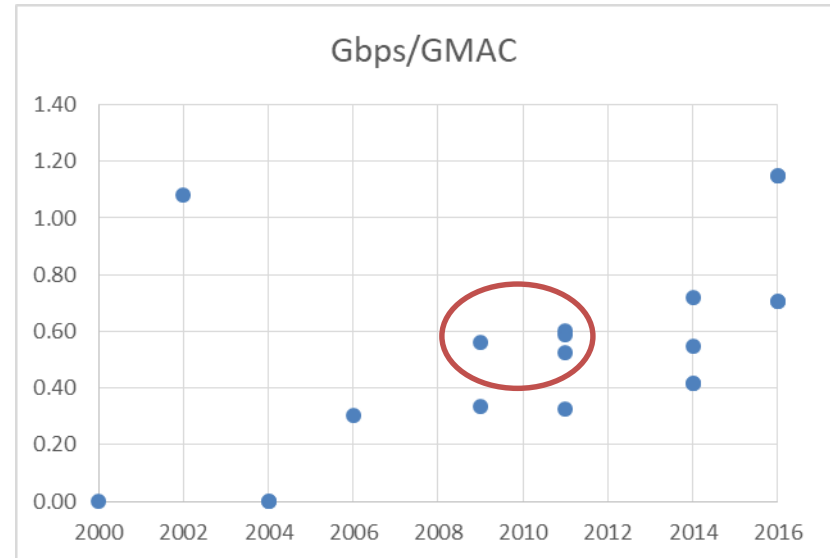
Sufficient SERDES that data connection by single ended links not needed

- 2009-11 Virtex-6, Kintex-7

All processing integer

Astronomy needs

0.6Gbps/GMAC




Internal Memory (High Speed)



Improved Internal memory depth

- 2016 UltraScale+
- UltraRAM 288kbit RAMs, Up to 45MB on chip
 - Plus 12MB of 18kbit RAMs



But CSIRO designs still needed external DRAM

- Mid Speed mainly big buffers for corner turn operation
- Increases with DSPs 1 per Virtex-5, 2 Kintex-7, 4 for UltraScale
- 8 DRAM per FPGA would be needed for Ultrascale+ ????

In 2016 a solution to the fast-large memory problem was Hybrid memory Cube

- DRAM with a SERDES interface.
- Intended for shared memory in multi-CPU system
- Must sacrifice external I/O to get DRAM bandwidth

Proof of concept built

- UltraScale+
- Four HMC arrowed
- Four 100G I/O
- Three MBO TX/RX



Unfortunately the HMC part used on the Proof of Concept board immediately went end-of-life

But both Xilinx and Altera announced FPGAs with direct attach DRAM: High Bandwidth Memory

- **Data from all columns of read available in one cycle**
- **I/O bandwidth up to 460GBps = 20 external DRAM**
 - **Sufficient for this generation of processing AND the next**
- **Available HBM provides at least 650kB per DSP (L2 cache)**
 - **High-speed memory 2-6kB per DSP (L1 cache)**

But HBM not available until late 2018

- To develop solution further a Non-HBM Ultrascale+ board built

One external DRAM for very deep slow memory

4x100G QSFP+, 3x300G Mid Board Optics



Gjjs Schoonderbeek ASTRON

Gemini HBM



Board under construction

- UltraScale+ HBM part procured
- Swap FPGA for HBM part

Due June 2019

3.7 Tbps to/from HBM

1.3 Tbps optical I/O

GPU challenge



GPUs said to easier to program

But available only in conjunction with server

- Interface PCI
- Can design FGPA to meet interface requirements
 - Great for high I/O systems, e.g. switches

GPUs floating point – FPGA fixed point

- Fixed point OK for radioastronomy correlators and beamformers

Multicore Processors



Popular around 2012

- Kalray 256 cores, Adapteva aimed of 1024 cores (delivered 64)
- Achilles heel – I/O, and programming

GPU multicore

- V100 adds 640 Tensor cores
 - over 100T FP16 op/sec but only 4x4 matrix operations (FP32add)
- Xilinx UltraScale+ is ~12T 18bit ops/sec,
 - for 8bit can double this, see talk by Norbert Abel
- GPU uses PCI4x16 = 252Gbps
 - Astronomy needs ~0.6Gbps/GMAC = 0.8T 8bit op/sec ???
 - I/O mismatch

FPGA Multicore - AI Engine



Single Instruction Multiple Data

- 512 bit wide data=64Bytes

Fixed and Float

- Plus Scalar RISC

Neighbour AI Engines share data (Systolic)

AXI spans rows and columns

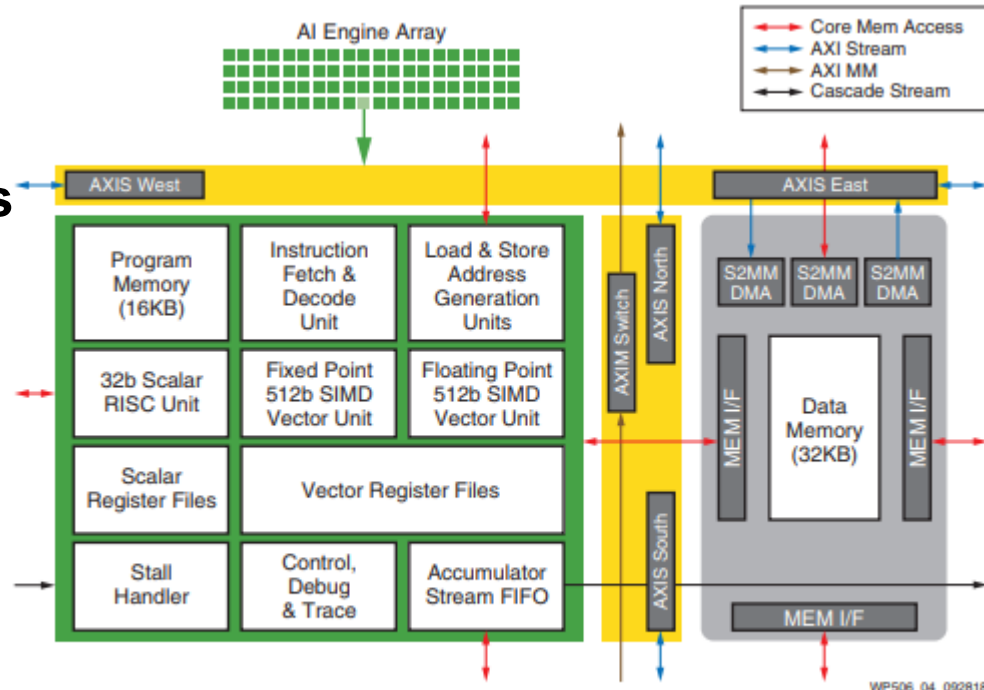


Figure 4: Detail of AI Engine Tile

WP506_04_092818

As well as standard FPGA fabric

Up to 400 AI engines

- SIMD, VLIW up to 256 8bit integer ops/clock per AI engine
- Up to 100 Top/sec 8-bit with large accumulators
- 25 Top/s 16-bit
- Native floating point added
 - Up to 6.4 Tflops FP32 - SDP applications???

SERDES I/O 1.6Tbps

- Still I/O limited
- Astronomy will use smaller cheaper part to reduce mismatch

VERSAL Memory limits



NO HBM part announced

- Presume this will happen once all the major changes have their teething problems ironed out.
- Need Versal HBM to fully utilise in astronomy

In the mean time only large memory is DRR

- In built DDR controllers – 256 bit bus width (128 smaller part).
 - Really need almost an order of magnitude more.

Some functions eg Correlator Low are more compute intensive

- Mix Versal AI with Ultrascale+ HBM in next gen design

RFSoc - Real World Interface



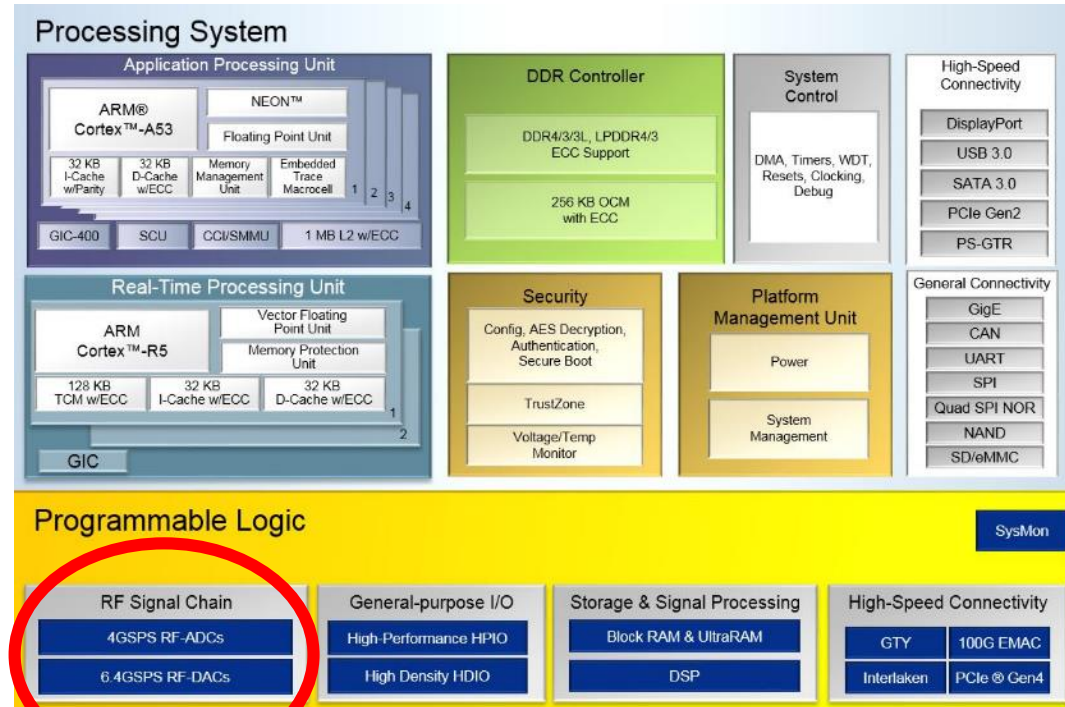
Effectively cut down
Versal

- No AI

Added ADC/DACs

- Interface DONE
- Low power
- 2 or 4 GSps ADC

Compute for FB
SERDES for Outputs



Questions / Discussion?

Thank-you!

Additional features



For a while FPGAs included hard coded PCI and Ethernet

VERSAL adds hard coded DDR controllers

Dual core ARM Cortex A72 and R5 added

- **256kB memory,**
- **Ethernet (x2); UART (x2); CAN-FD (x2); USB 2.0 (x1); SPI (x2); I2C (x2)**
- **Move Command and Control from separate servers to ARMs?**
- **Phase computation?**
- **??**