



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA




Fondazione
ICSC
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

FPGA Architecture

Andrea Triossi – University of Padova – INFN Padova

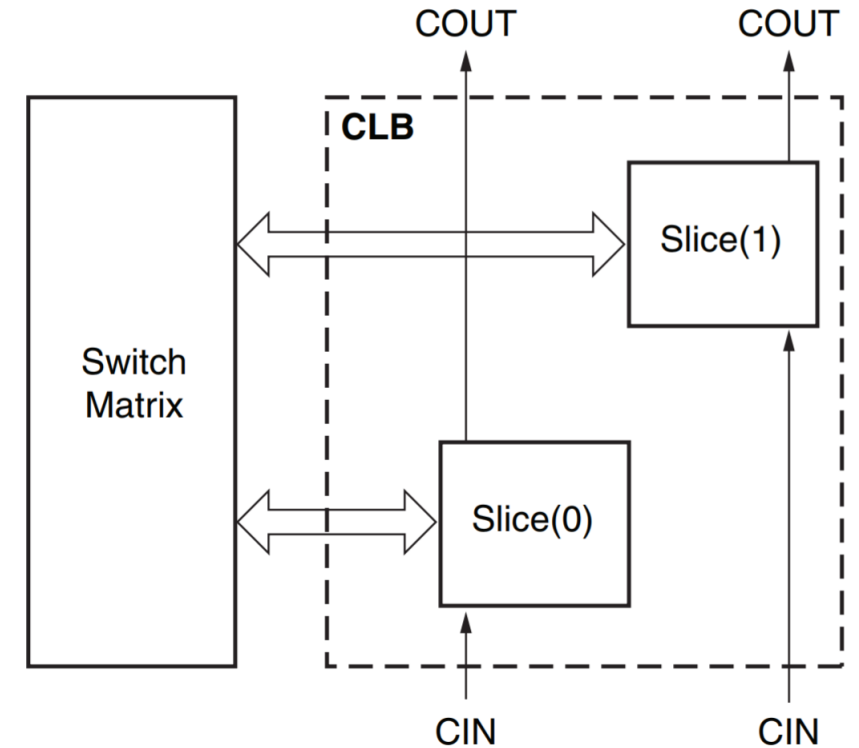
FPGA architecture

- Programmable logic blocks linked by programmable interconnections
- Programmable Input and Output
- Different names depending by the manufacturer
- Xilinx/AMD: Configurable Logic Block, Switch Matrix and I/O bank
- More than a million of CLB
- More than a thousand of I/O




FPGA architecture

- Each FPGA has a slightly different architecture
 - As running example, we use the 7 series by Xilinx/AMD
- Each CLB contains a pair of slices composed by
 - 6-input look-up tables (LUTs) and storage elements
 - Distributed memory and shift register logic capability
 - Wide multiplexers
 - Dedicated high-speed carry logic
- There are two kind of slice called SLICEL and SLICEM depending on the available features
- All the resources of a CLB are inferred by the compilation tool but they can also be instantiated directly



Look up table

- It is used to implement the truth table of a Boolean function
- It is equivalent to a multiplexer that use as inputs preconfigured values (configuration memory cells)
- 4 LUTs for each slice
- Each LUT has 6 independent inputs and 2 independent outputs -> it can accommodate up to 2 Boolean functions
- Functions with more inputs can be generated using multiplexers that combines LUT outputs



Storage element

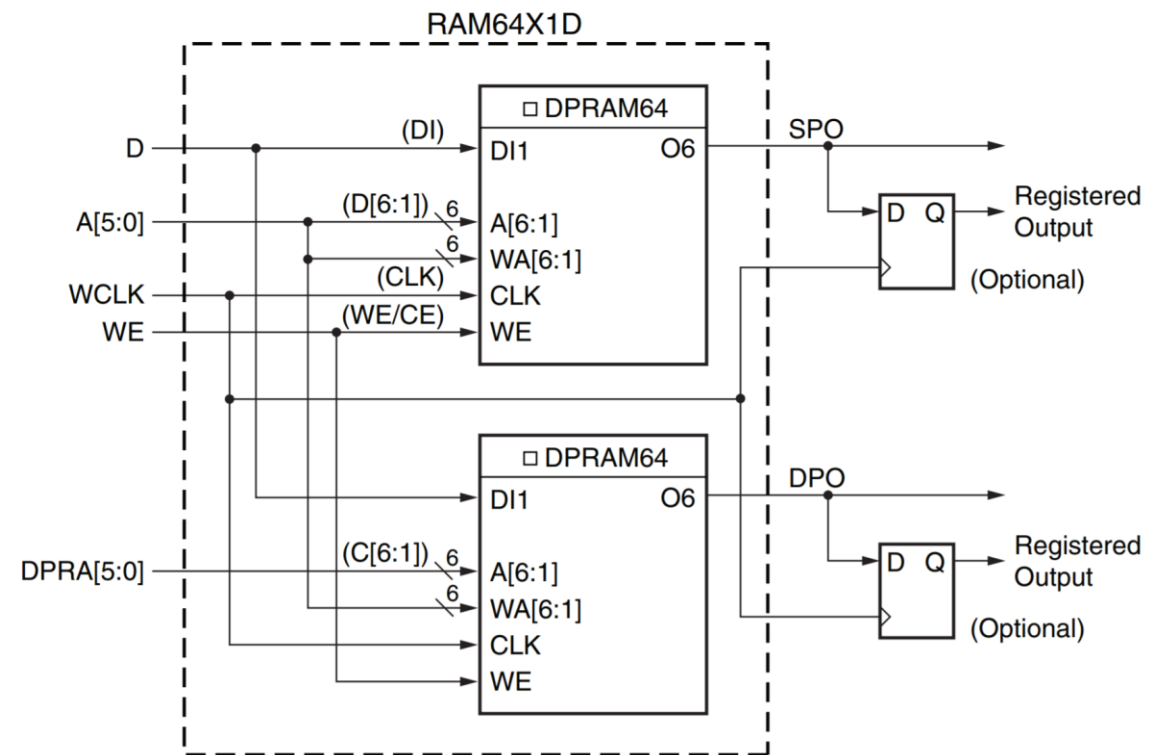
- 8 storage elements for each slice
 - 4 can be configured as D-FFs or latches (level sensitive)
 - 4 D-FFs (not available if the formers are configured as latches)
- Driven by a LUT or directly from input
- Control signals (CLK, CE, S/R) are common to all them
- Initial state independent of S/R

Distributed RAM

- LUTs can be implemented as synchronous RAM
- Single port
 - Common address port for synchronous write/asynchronous read
- Simple Dual port
 - One port for synchronous write
 - One port for asynchronous read
- Dual port
 - One port for synchronous write/asynchronous read
 - One port for asynchronous read
- Quad port
 - One port for synchronous write/asynchronous read
 - Three ports for asynchronous read

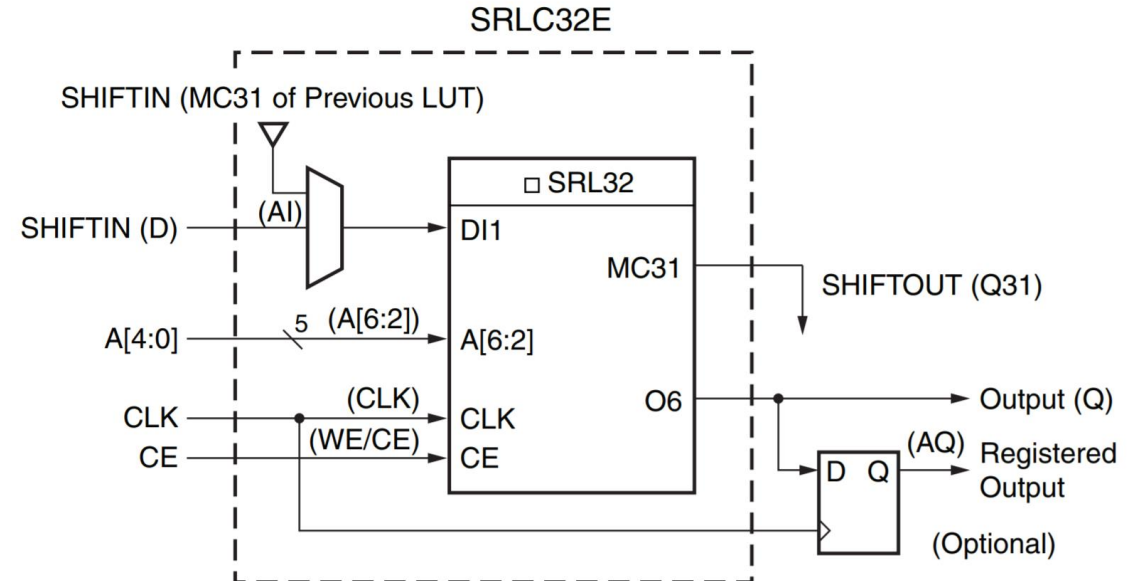
Distributed RAM

- Different sizes available: 32x1bit S, 32x1bit D, 32x2bit Q, 32x6bit SD, 64x1bit S, 64x1bit D, 64x1bit Q, 64x3bit SD, 128x1bit S, 128x1bit D, 256x1bit S
 - Multiplexers combine multiple LUTs when >64
- Only SLICEM can be configured as RAM
 - ROM can be built in SLICEL or SLICEM



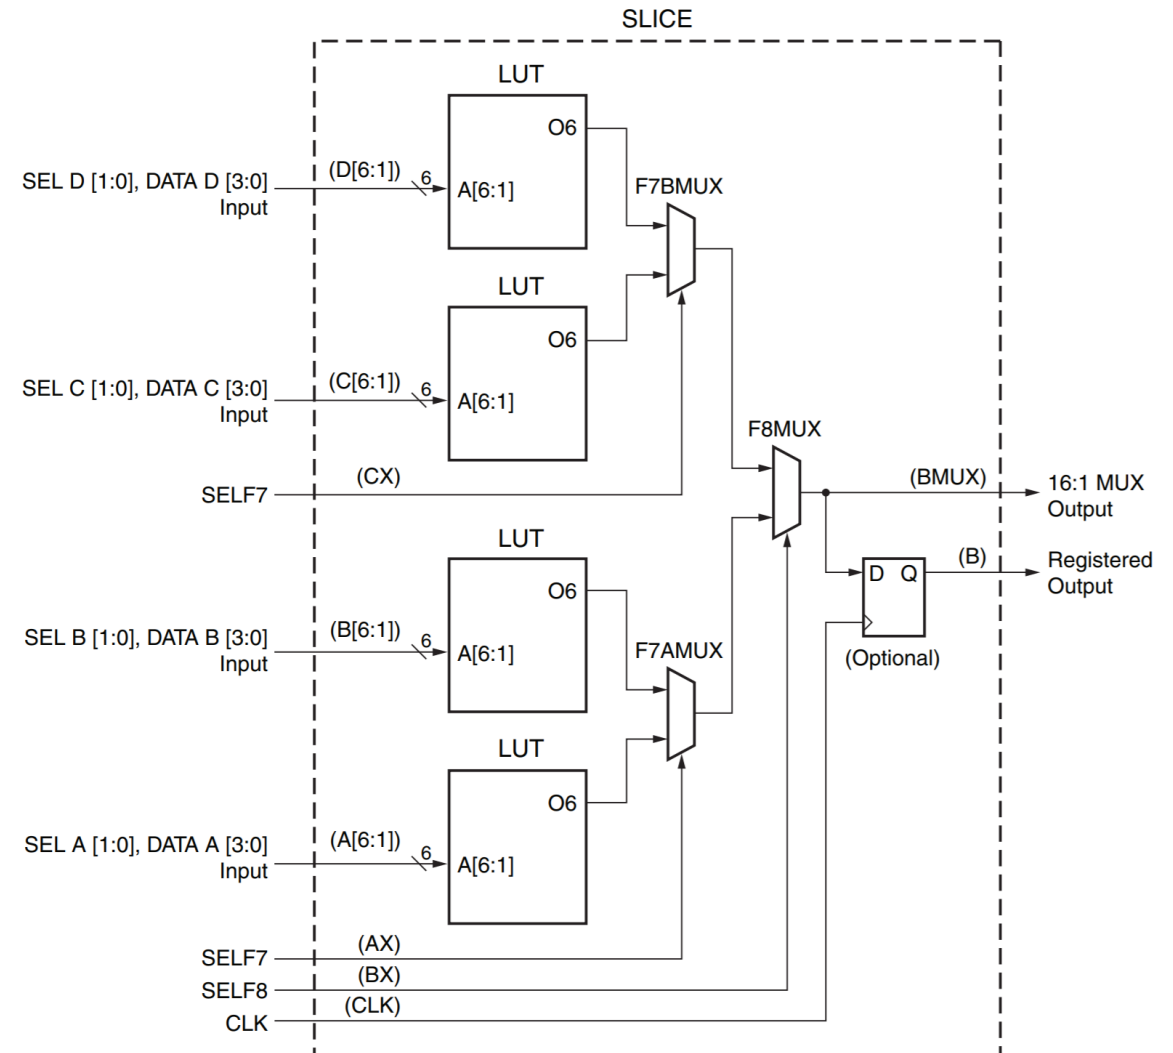
Shift register

- SLICEM can implement shift registers without using the available FFs
- Each LUT can be configured as a shift register long up to 32 bit (32 clock cycle delay)
 - Cascading is possible
- Synchronous write/asynchronous read
 - Any of the 32 bits can be read out (shorter shift registers)
- S/R not supported



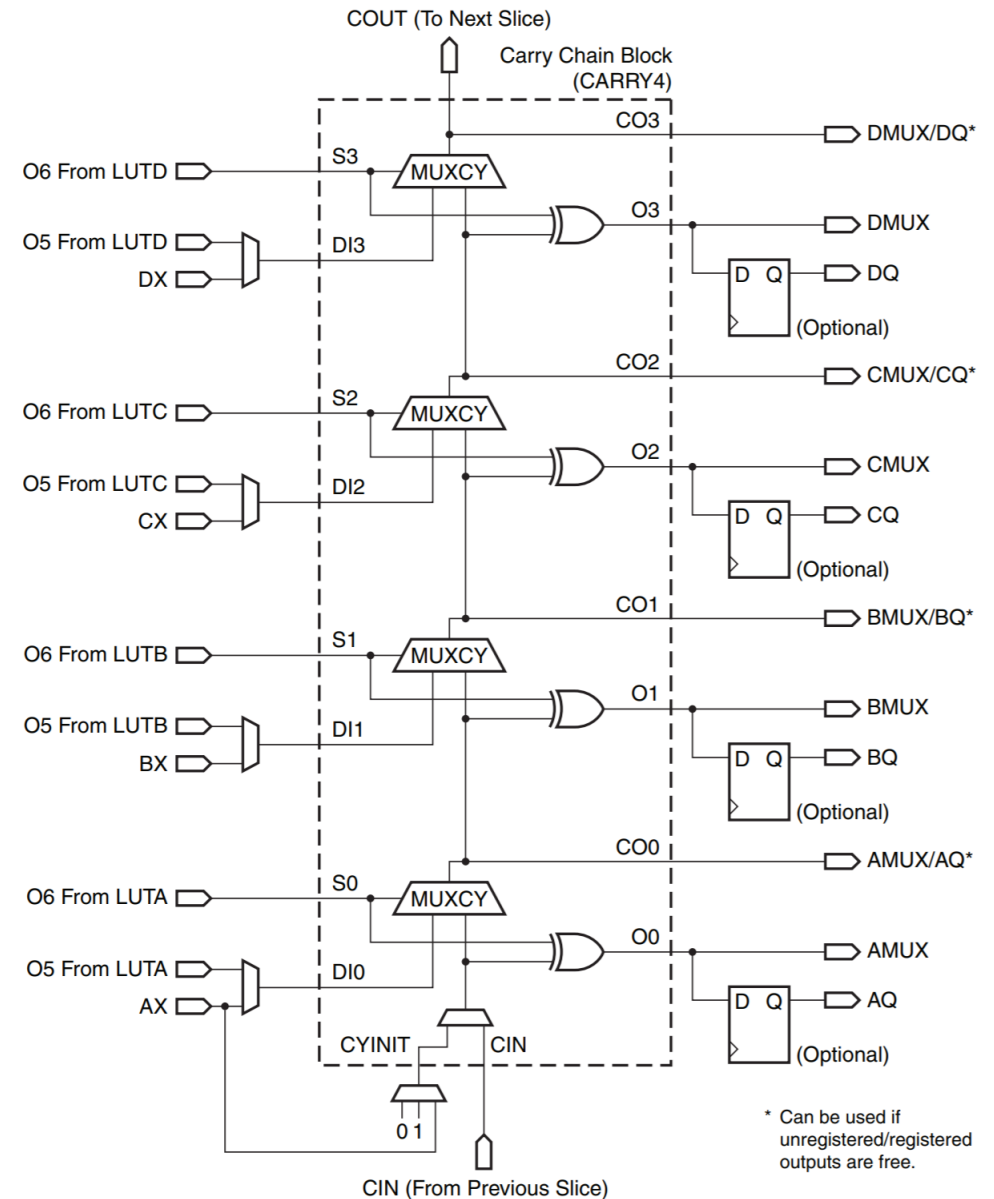
Multiplexer

- A LUT can host a 4:1 multiplexer (4 inputs and 2 select lines)
- One slice can host then up to four 4:1 multiplexers
- They can be combined in a bigger multiplexer (up to a 16:1 in a single slice) by means of additional mux out of the LUT



Carry logic

- Used for addition and subtraction
- 4 bits per slice (slices can be cascaded)
- Sx is the propagate
- Dlx is the generate
- CYINIT control add/subtract
- Ox is the sum
- COx is the carry
- Equivalent to a lookahead carry



Ultrascale CLB

- The two independent slice are combined in a single one
 - Carry logic expanded to 8 bits
- More flexibility on the controls
 - Four clock enables per CLB
 - Optional inverter on S/R line
- More flexibility on distributed RAM
 - Dedicated clock
 - Write enable independent from storage element clock enable and can be combined with other inputs to create up to 8 independent write enables
- All storage elements output always available with 4 output per LUT:
 - Direct LUT output
 - Multiplexed combinatorial/sequential output
 - Two storage elements output
- All storage elements can be FF or latches

Clock distribution


- Clock must arrive to each synchronous FF almost at the same time (very low skew)
 - Dedicated I/O and routing resources
 - Clock Management Tiles (CMT) provide clock frequency synthesis, deskew and jitter filtering
- Global clock tree
- The FPGA is divided in clock regions
 - It includes all synchronous elements in an area of 50 CLB and one I/O bank
- Buffers are used to propagate the clocks across the device giving access to the clock trees

Clock buffers

- Global clock buffers (BUFG)
 - Drive global clock lines
 - Clock enable functionality
 - Multiplexer between two clocks
 - Horizontal clock buffers (BUFH)
 - Allow access from global lines to a single clock region
 - Clock enable for clocks that span a single region
 - I/O clock buffers (BUFIO)
 - Provide access to all the I/O resources in the same I/O bank
 - Regional clock buffers (BUFR)
 - Drive regional clock trees
 - Clock divider functionality
 - Multi-clock region buffers (BUFMR)
 - Allow regional and I/O clocks to span up to three vertically adjacent regions
- In conjunction are used to clock programmable serializer/deserializer

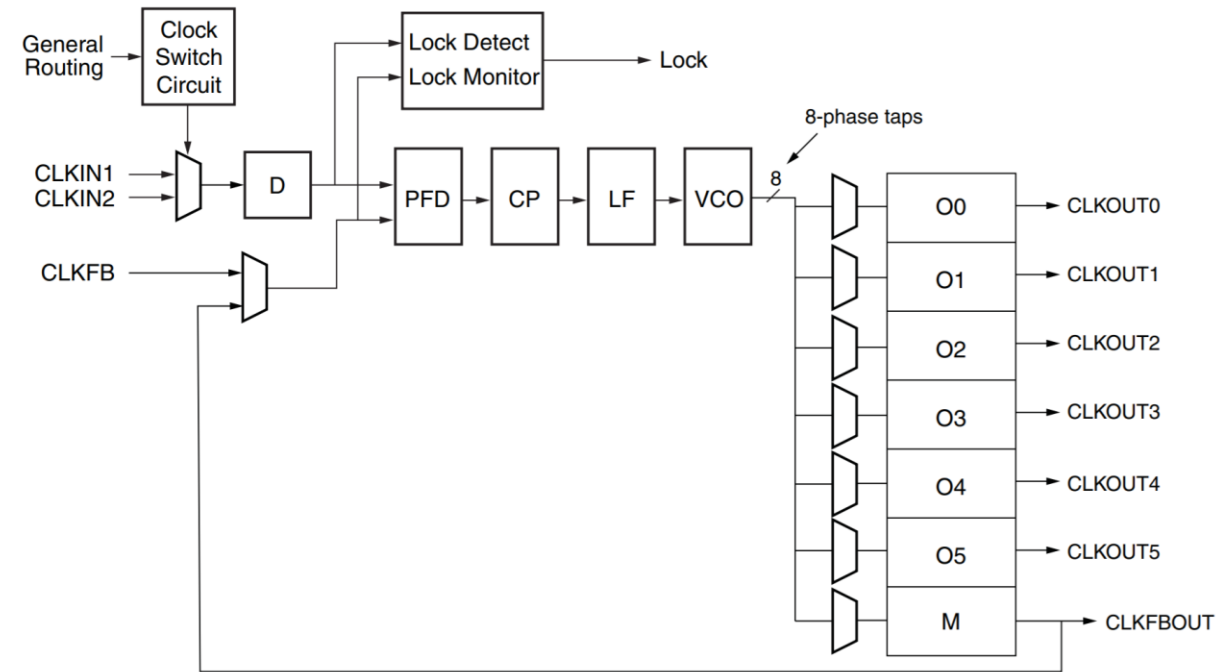
Clock management tile

- Contains a Phase-Locked Loop (PLL) and a Mixed-Mode Clock Manager (MMCM)
- MMCM is a PLL with the addition of some functionalities:
 - High performance connection to BUFIO and BUFR
 - More clock outputs
 - Cascading of the output dividers
 - Fractional divide and multiply
 - Fine and dynamic phase shifting



Clock management tile

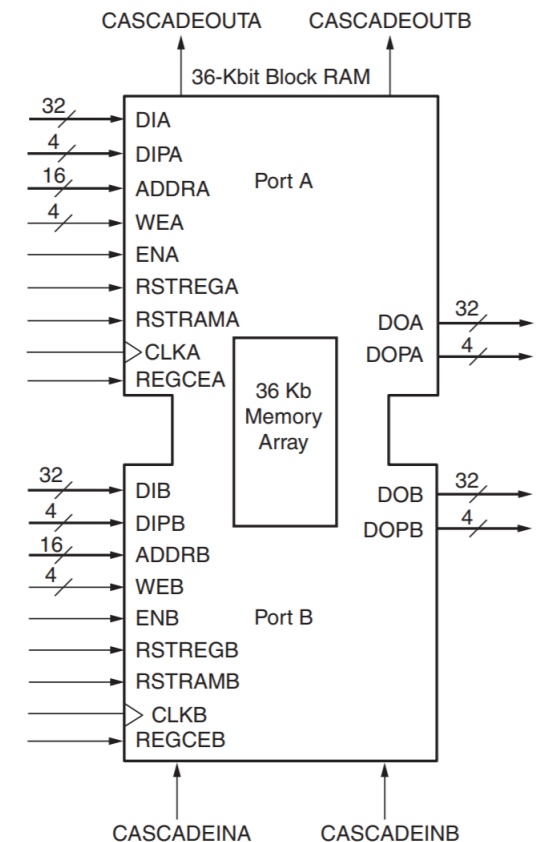
- PLL is a feedback control system that automatically adjusts the phase and frequency of a locally generated signal to match those of an input signal
- Programmable counter divider reduces the input frequency
- A phase-frequency detector compare the input clock with the feedback one
- Its output drive a charge pump and a loop filter to generate a reference voltage for the voltage-controlled oscillator
- VCO produces 8 phases that can be sent to any output counter
- A dedicated counter is used to drive the feedback path



$$F_{VCO} = F_{IN} \frac{M}{D} \quad F_{OUT} = F_{IN} \frac{M}{D \cdot O}$$

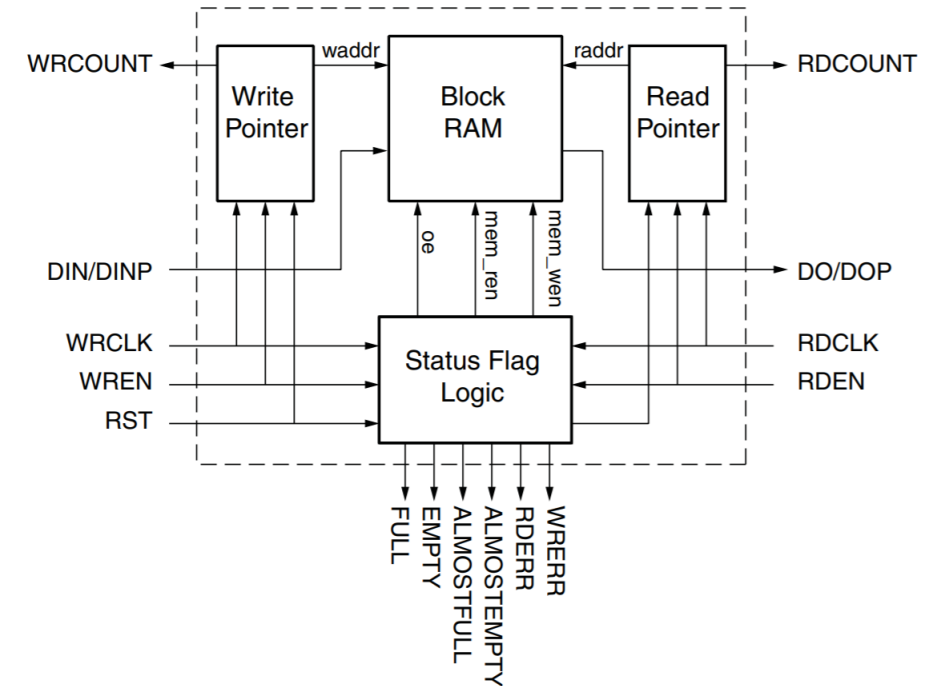
Block RAM

- In addition to distributed RAM there are a large number of block RAMs implemented in specific hardware
- Cascadable 36Kb block RAMs that contain two independently controlled 18Kb RAMs
- They can be configured as single or dual port RAMs, ROMs, synchronous FIFOs and dual clock FIFOs
- Simple or true dual port
- Data can be read/written from/to either or both ports
- Write operations are synchronous
- Optional output register
- Byte-wide write enable
- Single-bit errors correction (double-bit detection) in data read out




Built-in FIFO

- Block RAM can be used as FIFO
- Standard or first-word fall-through
- Synchronous or dual-clock
 - Very useful to cross clock domains
- Cascading or parallel connection allowed



Digital Signal Processor

- Add/subtract logic unit, multiplier and accumulator
- Optional internal registers depending on used functionalities
- Internal features:
 - 25 bit pre-adder (can be used as just an input multiplexer)
 - 25 x 18 bit multiplier
 - Three 48 bit datapath multiplexers followed by a three inputs adder/subtractor or a two inputs logic unit
 - Pattern detector (it could be used for checking overflow/underflow conditions)
- Cascade of multipliers to achieve larger products
- Adder/subtractor/accumulator can be split in up to four smaller units




Microprocessor

- 7 series FPGA + Single/Dual core ARM Cortex-A9 processor = Zynq 7000
 - System on Chip: the processor is the main device and the programmable logic part is considered a peripheral device
 - Different programming environments for the two devices
- Soft processors like Xilinx/AMD MicroBlaze, Lattice Mico32, RISC-V...
 - Completely implemented in logic
 - Less performance but more customizable
 - Massively parallel computing


I/O resources

- Input/Output Block (IOB):
 - Wide variety of standard interfaces (single-ended or differential)
 - Programmable control of output strength and slew rate
 - On chip termination with digitally-controlled impedance
 - Integrated pull-up/pull-down resistor




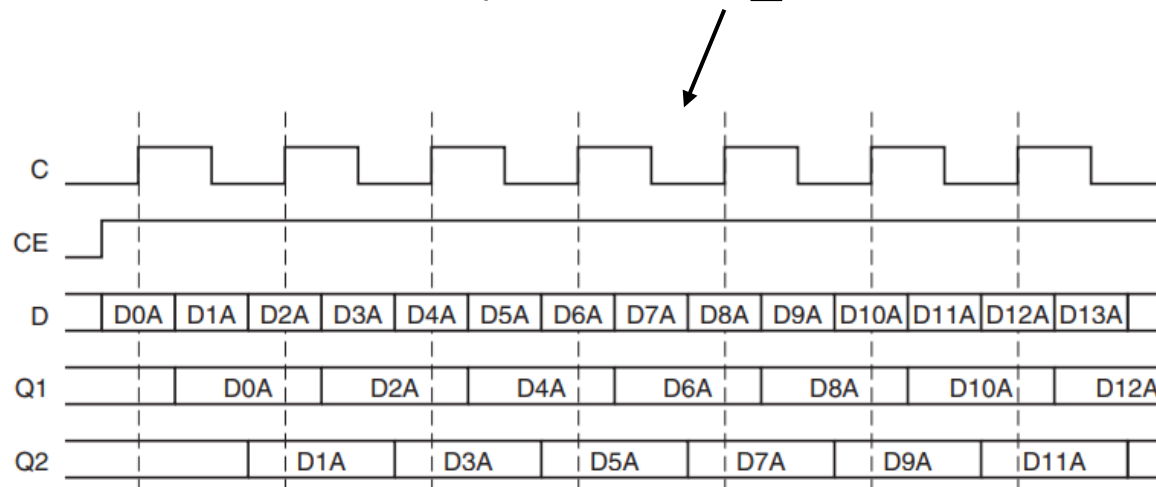
I/O resources

- IDELAY/ODELAY:
 - Programmable delay (31-tap)
 - Delay resolution continuously calibrated
 - Both can be accessed also from the internal logic



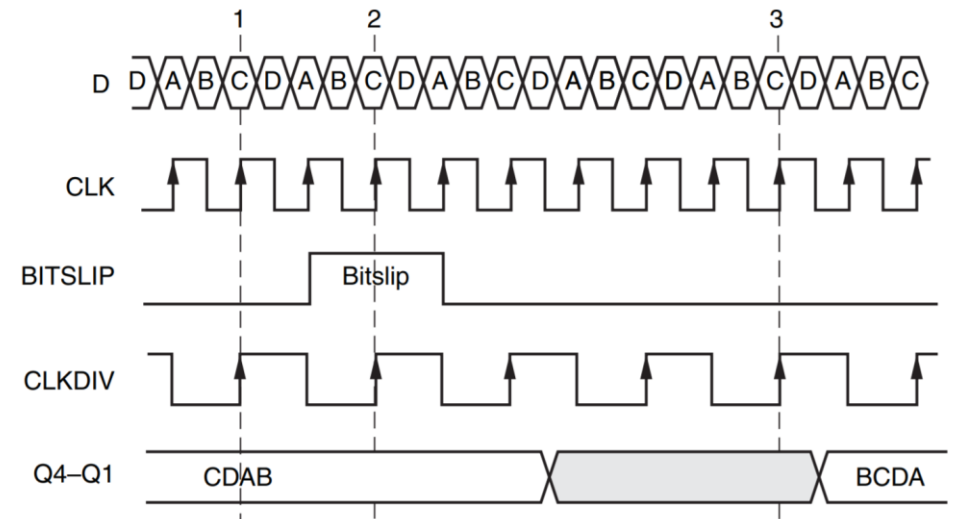
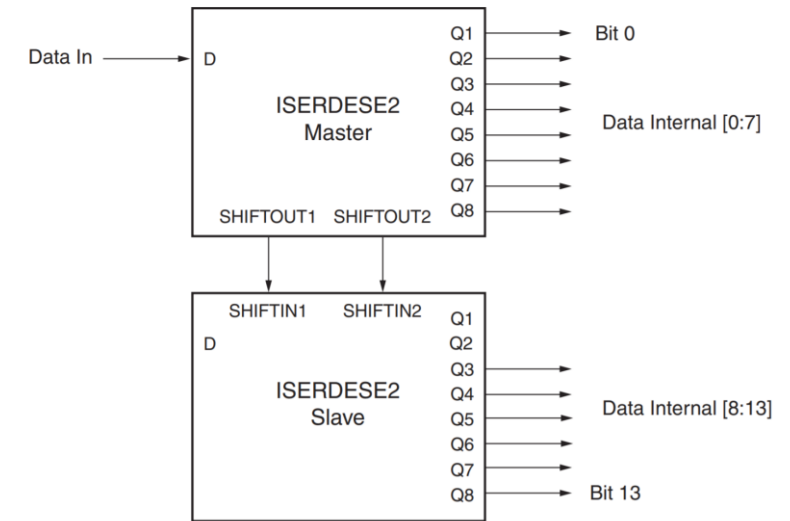
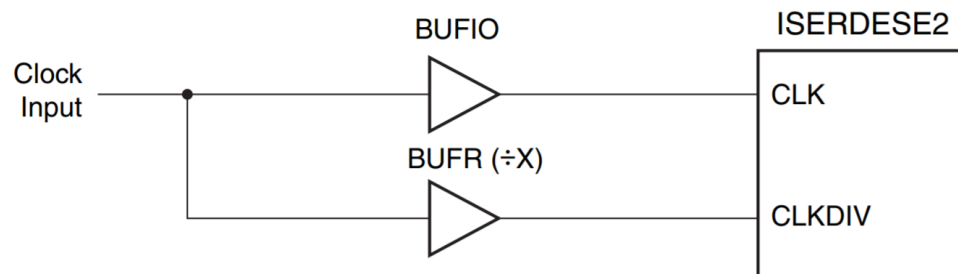
I/O resources

- ILOGIC/OLOGIC:
 - Combinatorial (asynchronous)
 - Edge-triggered D-type flip-flop
 - Level sensitive latch
 - IDDR mode (OPPOSITE_EDGE or SAME_EDGE or SAME_EDGE_PIPELINED)



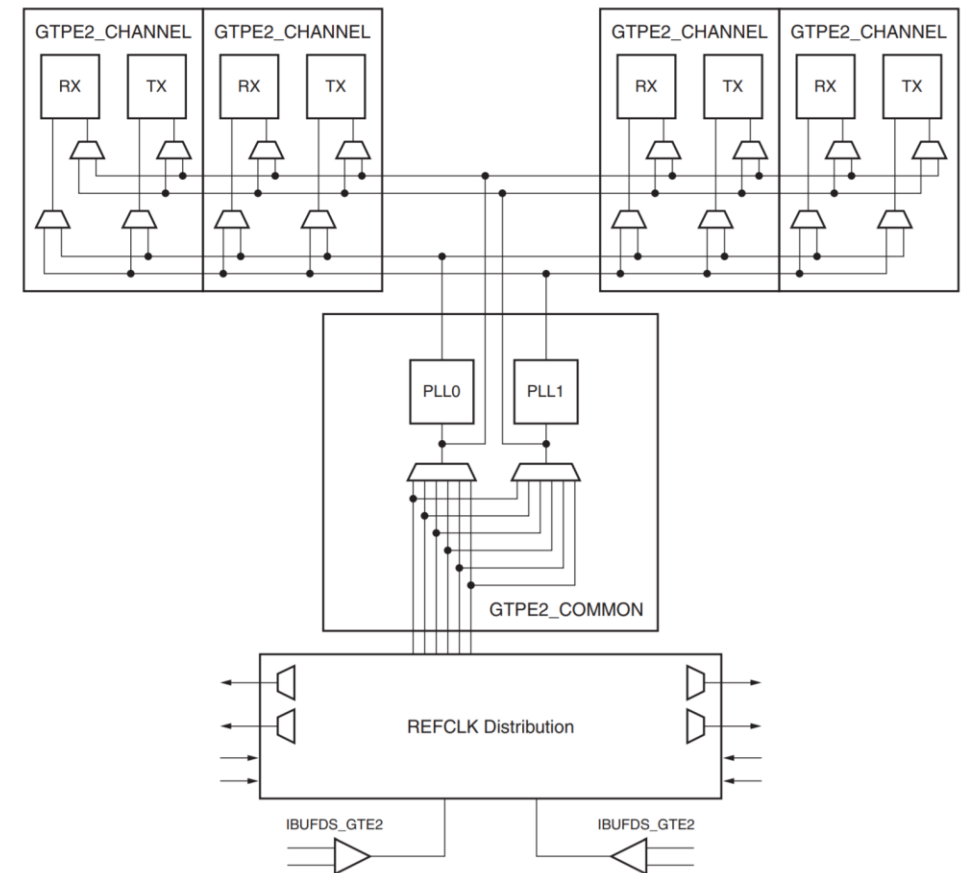
I/O resources

- ISERDES/OSERDES:
 - Serial to parallel converter and vice versa
 - SDR or DDR up to 1:8 ratio
 - Up to 1:14 in cascade mode
 - Bitflip to reorder data words in the parallel domain
 - Dedicated support for memory interfaces



Fast transceiver

- Gigabit line rates
- Physical layer compatibility with many protocols (PCI Express, SGMII, Serial Rapid IO, XAUI, SATA, etc.)
- Four channels and a common block are clustered together and are called a Quad
 - Each channel consists of a transmitter and a receiver
 - The common part contains two PLLs



Analog to digital converter

- XADC is a dual 12-bit 1 MSPS ADC plus several sensors
- Simultaneous sampling of the two ADCs
- External inputs (unipolar/bipolar)
- Single channel/automatic sequencer mode
- Optional external trigger
- Optional external multiplexer
- ADC conversion stored in registers accessible by a synchronous port called Dynamic Reconfiguration Port (DRP)
- Optional averaging and threshold alarm
- If not instantiated is anyway used in its default configuration

