



FPGA memory performance

Sensor to Image GmbH
Lechtorstrasse 20
D 86956 Schongau

Website: www.sensor-to-image.de
Email: email@sensor-to-image.de

Sensor to Image GmbH

- Company

- Founded 1989 and privately owned
- Company goal: development, production and service of OEM image processing hardware tool
- Production of 1000-2000 boards/year
- Team of 9 developers, 2 people in production and 3 in administration
- Full service from training → development → layout → software → sample production → redesign → series production → support → repair is possible

- Development

- High speed and impedance controlled layout, pitch > 0.8 mm
- FPGA design tools experience with Altera, Xilinx, Lattice and MicroSemi tools
- FPGA simulation and verification with ModelSim 10.x
- Embedded C-Compilers for FPGA based 8051, Coldfire, Microblaze/NIOS, ARM in Zynq/Cyclone V SoC and BLACKFIN running Linux
- Capability for writing PC drivers VC6 – VC2015

- Production

- Semiautomatic Pick & Place machine with manual BGA placement and part handling down to size 0402
- Vapor phase oven for RoHS conform soldering PCB size <= 30x30 cm
- Thermal box for testing systems up to 50x50x50 cm under operation from -40° C up to 90° C
- Mixed signal oscilloscopes with bandwidth of 5 GHz

Certificated & Applications

- EN ISO9001:2000 and EN ISO 14001 certified procedures since 1996
- CE Production to EN 61000-6-3:2007 and EN 61000-6-4:2007 standard
- Production to MIL-STD-810E 501.3, ...502.3, ... 514.4-All, ... 516.4 and VG95328, VG95373 and VG96903
- Production to Marine standard EN 60945:2002
- Asia OEM application: inspection of copper in PCB production, ≥ 700 PCIe line scanning systems
- German OEM application: A0 printer & scanner electronics, ≥ 2000 CIS & USB3 IF boards
- American OEM application: 3D OEAM camera, standard CMOS Sensor FPGA pre-processing and embedded Linux, ≥ 1000 system
- ...
- And you are welcome to verify this by yourself by visiting us in Schongau



Agenda

- Overview
- System setup
 - OPB, PLB bus
 - Avalon bus
 - AXI bus
 - Systems based on hard- and soft placed memory controller
- Results
 - Typical system setup
 - Some samples on bandwidth vs. memory configuration
 - Tips&Tricks to extend this presentation for your application

Typical S2I FPGA system

- Machine Vision FPGA IP vendor

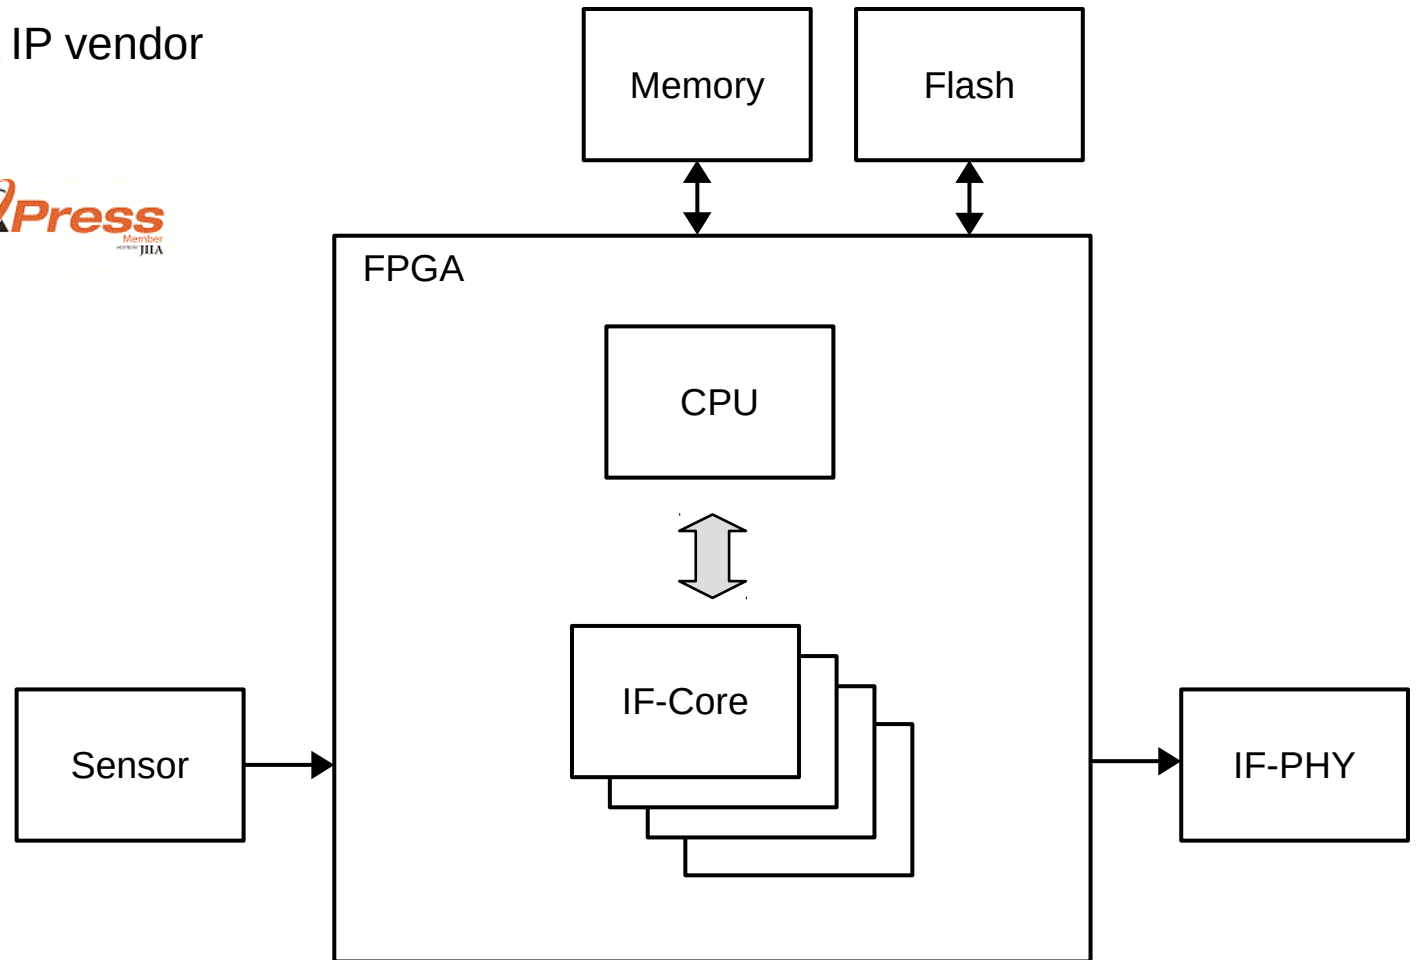


- FPGA's:

- Altera
- Xilinx
- Lattice
- MicroSemi

- Memory buses:

- proprietary
- OPB, PLB
- Avalon
- Wishbone
- AXI on Altera, Microsemi and Xilinx



Basic DDR3 functions

Configuration

- 256 Meg x 4
- 128 Meg x 8
- 64 Meg x 16

Timing cycle time

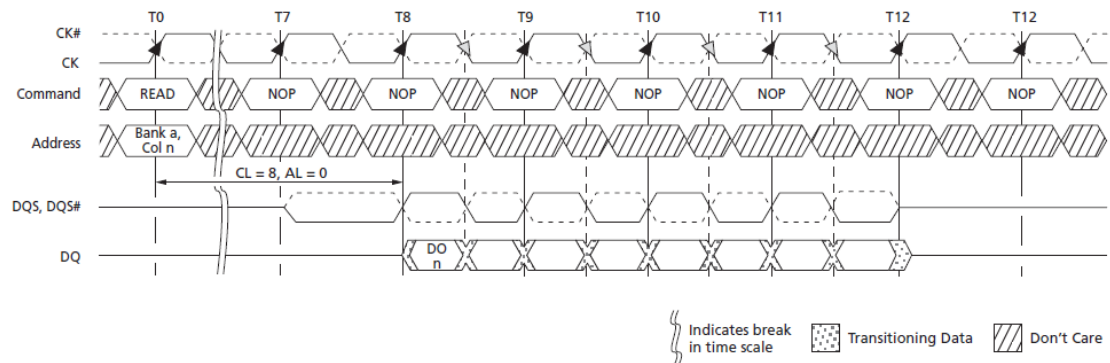
- 1.07ns @ CL = 13 (DDR3-1866) -107 Layout 933 MHz
- 1.07ns @ CL = 12 (DDR3-1866) -107E Layout 933 MHz
- 1.25ns @ CL = 11 (DDR3-1600) -125 Layout 800 MHz
- 1.5ns @ CL = 9 (DDR3-1333) -15E Layout 666 MHz
- 1.87ns @ CL = 7 (DDR3-1066) -187E Layout 533 MHz

Speed Grade	Data Rate (MT/s)	Target tRCD-tRP-CL	tRCD (ns)	tRP (ns)	CL (ns)
107	1866	13-13-13	13.91	13.91	13.91

Figure 68: READ Latency

Page 159 data sheet
MT41J64M16

16bit DDR3-1600
Read 64bit: RL 11, CL 11, Data
24*1.25ns = 30ns, 8Byte
net data rate: 2 GBit/s
max data rate: 30 GBit/s



- Notes:
1. DO n = data-out from column n.
 2. Subsequent elements of data-out appear in the programmed order following DO n.

Memory needs an interface on a FPGA bus interface !?

- what is on the market off the shelf
- built the memory interface yourself

Advanced eXtensible Interface Bus (AXI)

AXI, the third generation of AMBA interface defined in the AMBA 3 specification, is targeted at high performance, high clock frequency system designs and includes features that make it suitable for high speed sub-micrometer interconnect:

- separate address/control and data phases
- support for unaligned data transfers using byte strobes
- burst based transactions with only start address issued
- issuing of multiple outstanding addresses with out of order responses

"Low power is also important to ARM, and the CoreLink Network Interconnect is no exception. The RTL has been optimized to make extensive use of automated clock gate insertion by synthesis tools. Implementation trials have shown that as many as 95% of the flops are clock gated when idle"

<http://www.arm.com/products/system-ip/interconnect/axi/index.php>

- AXI4 - for high-performance memory-mapped requirements.
- AXI4-Lite - for simple, low-throughput fixed 32 bits 32-bit address width memory-mapped communication
- AXI4-Stream - for high-speed streaming data.
- AXI4 - Each master and slave connection can independently use data widths of 32, 64,128, or 256 bits wide
- Burst lengths up to 16/256 (AXI3/4)
- Built-in clock-rate and AXI4-Lite conversion

Altera Avalon

Avalon specification defines the following seven interfaces:

- Avalon Streaming Interface: Avalon-ST, an interface that supports the unidirectional flow of data, including multiplexed streams, packets, and DSP data.
- Avalon Memory Mapped Interface: Avalon-MM, an address-based read/write interface typical of master–slave connections.
- Avalon Conduit Interface: Avalon-Conduit, an interface type that accommodates individual signals or groups of signals that do not fit into any of the other Avalon types. You can connect conduit interfaces inside a Qsys system or to other modules in the design or to FPGA pins.
- Avalon Tri-State Conduit Interface: Avalon-TC, an interface to support connections to off-chip peripherals. Multiple peripherals can share pins through signal multiplexing, reducing the pin count of the FPGA and the number of traces on the PCB.
- Avalon Interrupt Interface, an interface that allows components to signal events to other components.
- Avalon Clock Interface, an interface that drives or receives clocks.
- Avalon Reset Interface, an interface that provides reset connectivity.

A single component can include any number of these interfaces and can also include multiple instances of the same interface type, e.g.:

- Avalon-MM
- Avalon-ST
- Avalon-Conduit
- Avalon-TC
- Avalon-Interrupt
- Avalon-Avalon Clock

Lattice/OpenCores Wishbone

The WISHBONE interconnection makes System-on-Chip and design reuse easy by creating a standard data exchange protocol. Features of this technology include:

- ✓ Simple, compact, logical IP core hardware interfaces that require very few logic gates.
- ✓ Full set of popular data transfer bus protocols including: R/W cycle, BLOCK transfer cycle, RMW cycle
- ✓ Modular data bus widths and operand sizes.
- ✓ Supports both BIG ENDIAN and LITTLE ENDIAN data ordering.
- ✓ Variable core interconnection methods support point-to-point, shared bus, crossbar switch, and switched fabric interconnections.
- ✓ Handshaking protocol allows each IP core to throttle its data transfer speed
- ✓ Supports single clock data transfers
- ✓ Supports normal cycle termination, retry termination and termination due to error
- ✓ Modular address widths
- ✓ MASTER / SLAVE architecture for very flexible system designs
- ✓ Multiprocessing (multi-MASTER) capabilities. This allows for a wide variety of System-on-Chip configurations
- ✓ Arbitration methodology is defined by the end user (priority arbiter, round-robin arbiter, etc.)

Xilinx OPB, PLB

The OPB Bus Structure is used as the OPB interconnect for Xilinx FPGA based embedded processor systems. The bus interconnect in the OPB v2.0 specification is essentially a distributed multiplexer implemented as an "AND" function in the master or slave driving the bus and an "OR" to combine the drivers into a single bus → no real multi master, not useful for UMA design

The Xilinx 128-bit Processor Local Bus (PLB) v4.6 provides bus infrastructure for connecting an optional number of PLB masters and slaves into an overall PLB system. It consists of a bus control unit, a watchdog timer, and separate address, write, and read data path units. It contains an optional DCR slave interface to provide access to its bus error status registers.

- PLB arbitration support for up to 16 masters
- PLB address and data steering support for up to 16 masters
- 128-bit, 64-bit, and 32-bit support for masters and slaves
- PLB address pipelining
- Three-cycle arbitration
- Four levels of dynamic master request priority
- DDRx memory controller available
- Only supported in ISE development environment

Which bus?

Sensor to Image chooses AXI, as:

- (free) choice of FPGA vendor as supported by Altera, Microsemi and Xilinx
- Another approach to build a unified IP cores
- Another approach to reuse IP cores
- Support to many memory types from 1=one bus interface
- High bandwidth can be reached on AXI4
- ...

Altera Nios II

- 32-bit RISC embedded soft processor
- Little Endianness
- Interconnect: Avalon (AXI3, AXI4)
- 2 Variants:
 - Nios II/f
 - (Nios II/s)
 - Nios II/e


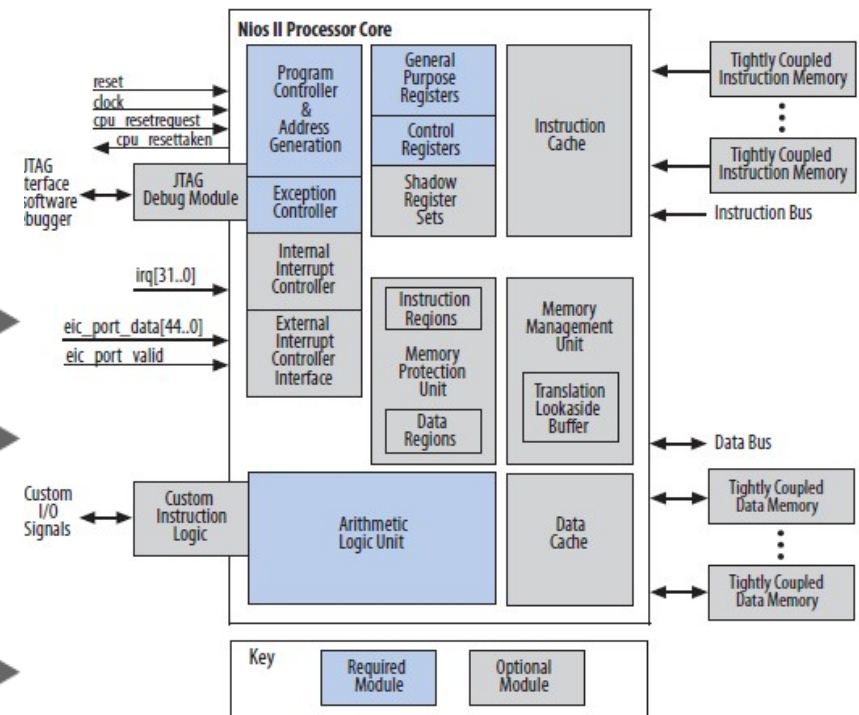



Figure 2-1: Nios II Processor Core Block Diagram



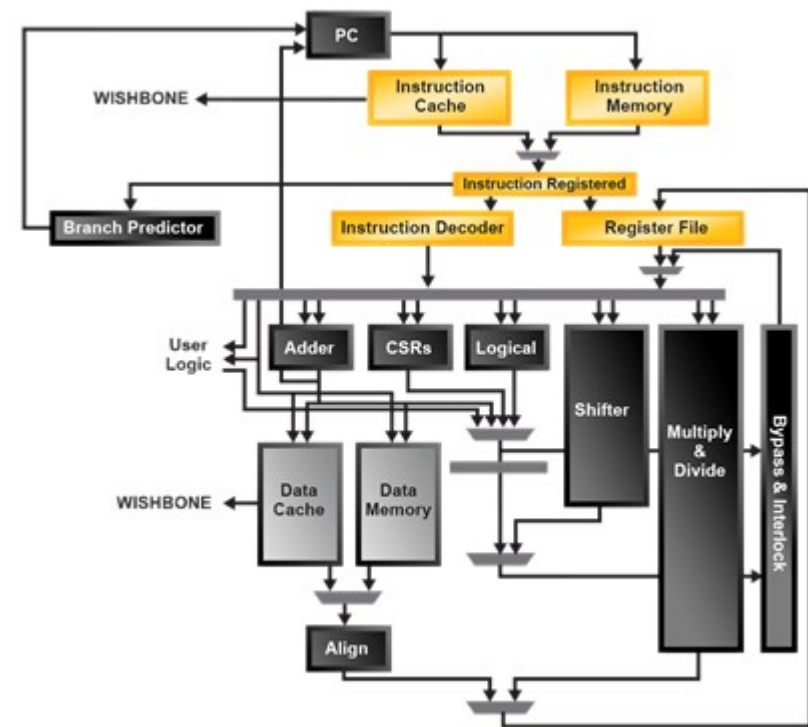
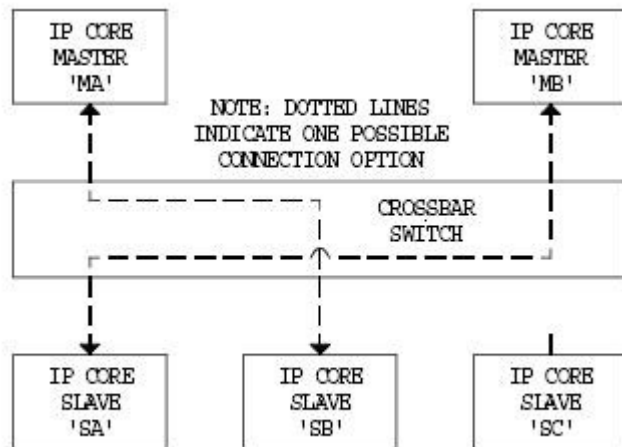
ALTERA Cyclone V SoC

- 28 nm Low Power (28LP) process
- Hard Processor System (HPS)
 - 32 bit dual-core ARM® Cortex™-A9 MPCore
 - HPS-FPGA Bridge (ARM AMBA AXI3)
 - I/O peripherals
- FPGA Fabric
 - Dynamical & partial reconfiguration
 - Hard-IPs (DSP, PCIe, Memory Controller)
 - Complex logic block
 - 10 Kb Block RAM with ECC



Lattice Mico32

- 32-bit Harvard, RISC architecture "soft" microprocessor
- free with an open IP core licensing agreement
- Memory controllers: DDR, DDR2, DDR3, async. SRAM, BRAM
- Interconnect: **Wishbone**



Xilinx Microblaze

- 32-bit RISC embedded soft processor
- Big/Little Endianness
- Interconnect: PLB / LMB / AXI3, AXI4
- High configurable
- 6 Variants

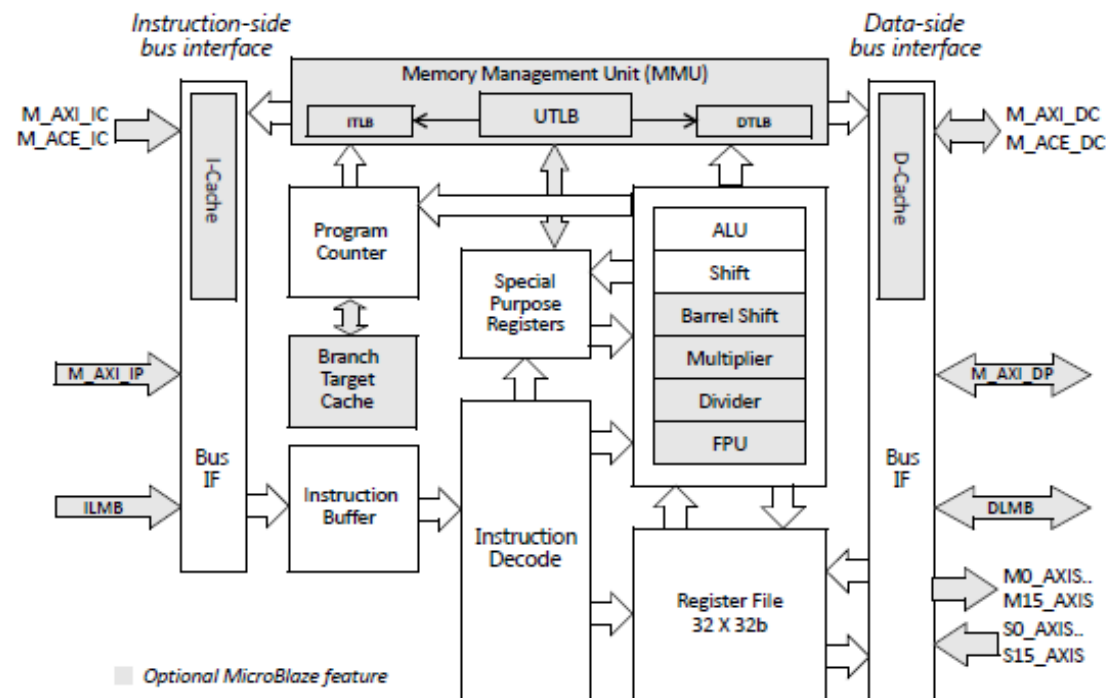



Figure 2-1: MicroBlaze Core Block Diagram

Table A-4: Device Utilization - Artix-7 FPGAs (XC7A200T fbg676-3)

Configuration	Device Resources		
	LUTs	FFs	F _{max} (MHz)
Minimum Area	648	214	234
Maximum Performance	3879	3064	154
Maximum Frequency	906	545	234
Linux with MMU	3653	3218	142
Low-end Linux with MMU	3108	2582	153
Typical	2070	1777	190

Xilinx Zynq

- 28 nm technology
- Processor System (PS)
 - 32 bit dual-core ARM® Cortex™-A9 MPCore
 - ARM AMBA AXI3 Interconnect
 - I/O peripherals
 - Memory interfaces
- Programmable Logic (PL)
 - Dynamical & partial reconfiguration
 - Complex logic block
 - 36 Kb Block RAM
 - 48 bit DSP



Design sample on Xilinx AXI in S2I standard architecture

To keep clock crossing low and FiFo running well, it is important to obey some basic rules on DDRx design:

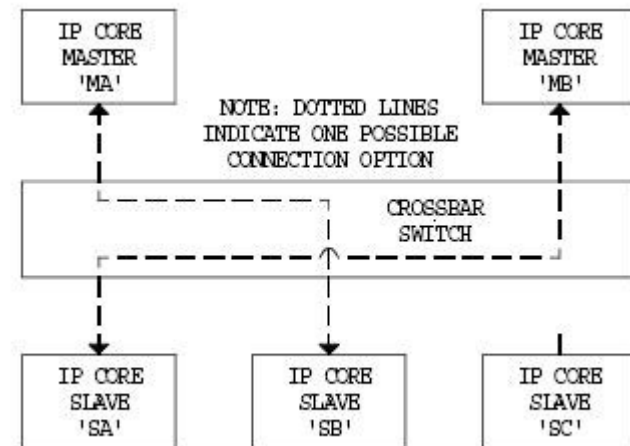
- ✓ run internal and external buses synchronous speed
- ✓ run internal and external buses at synchronous width
- ✓ adapt to the FPGA by 2/4/.. times data widths at 1/2/1/4/.. speed

This enables some basic setup like:

- ✓ external DDRx setup at N bits@M MHz → FPGA internal bus width at 2N bits@M MHz
- ✓ external DDRx setup at N bits@M MHz → FPGA internal bus width at 2*2N bits@M/2 MHz
- ✓ external DDRx setup at N bits@M MHz → FPGA internal bus width at 4*2N bits@M/4 MHz
- ✓ ...

With UMA architecture and these bus members

- ✓ CPU AXI master
- ✓ Interface IP core AXI master
- ✓ AXI based DDRx memory interface
- ✓ Some AXI-lite control connections



FPGA design (and cost) limitations

A: FPGA type and speed grade

B: max. bitrate on the FPGA due to the data sheets “ds181_Artix_7_Data_Sheet.pdf”, XC7A50T-1CPG236I, around 100US\$ single piece, and “ds182_Kintex_7_Data_Sheet.pdf”, XC7K70T-1FB484I, around 200US\$ single piece

C: memory width placed on the PCB

D: clock rate the memory is operated. This has to be less or equal to the B column due to the note in the data sheets “When using the internal VREF, the maximum data rate is 800 Mb/s (400 MHz).”, where the RED marked path is the only indication from my point of view between bit- and clock rate. You can play free with this GREEN marked values

E: this is the practical FIFO bandwidth out of this memory configuration, which is 66% memory efficiency due to accumulated AXI components and another 50% off due to write/read operation when doing FIFO operation. This value has to fit into your needed system bandwidth!

F: this is the multiplier between external and internal bus width, which determines the value in column H for the external → internal clock divider as well. GREEN so play with it as well.

G: resulting internal AXI bus width due to column C and F

H: resulting internal AXI clock divider due to column F

I: resulting internal AXI clock speed due to column F

K: S2I recommended max. AXI speed for G based AXI bit width

FPGA bandwidth values

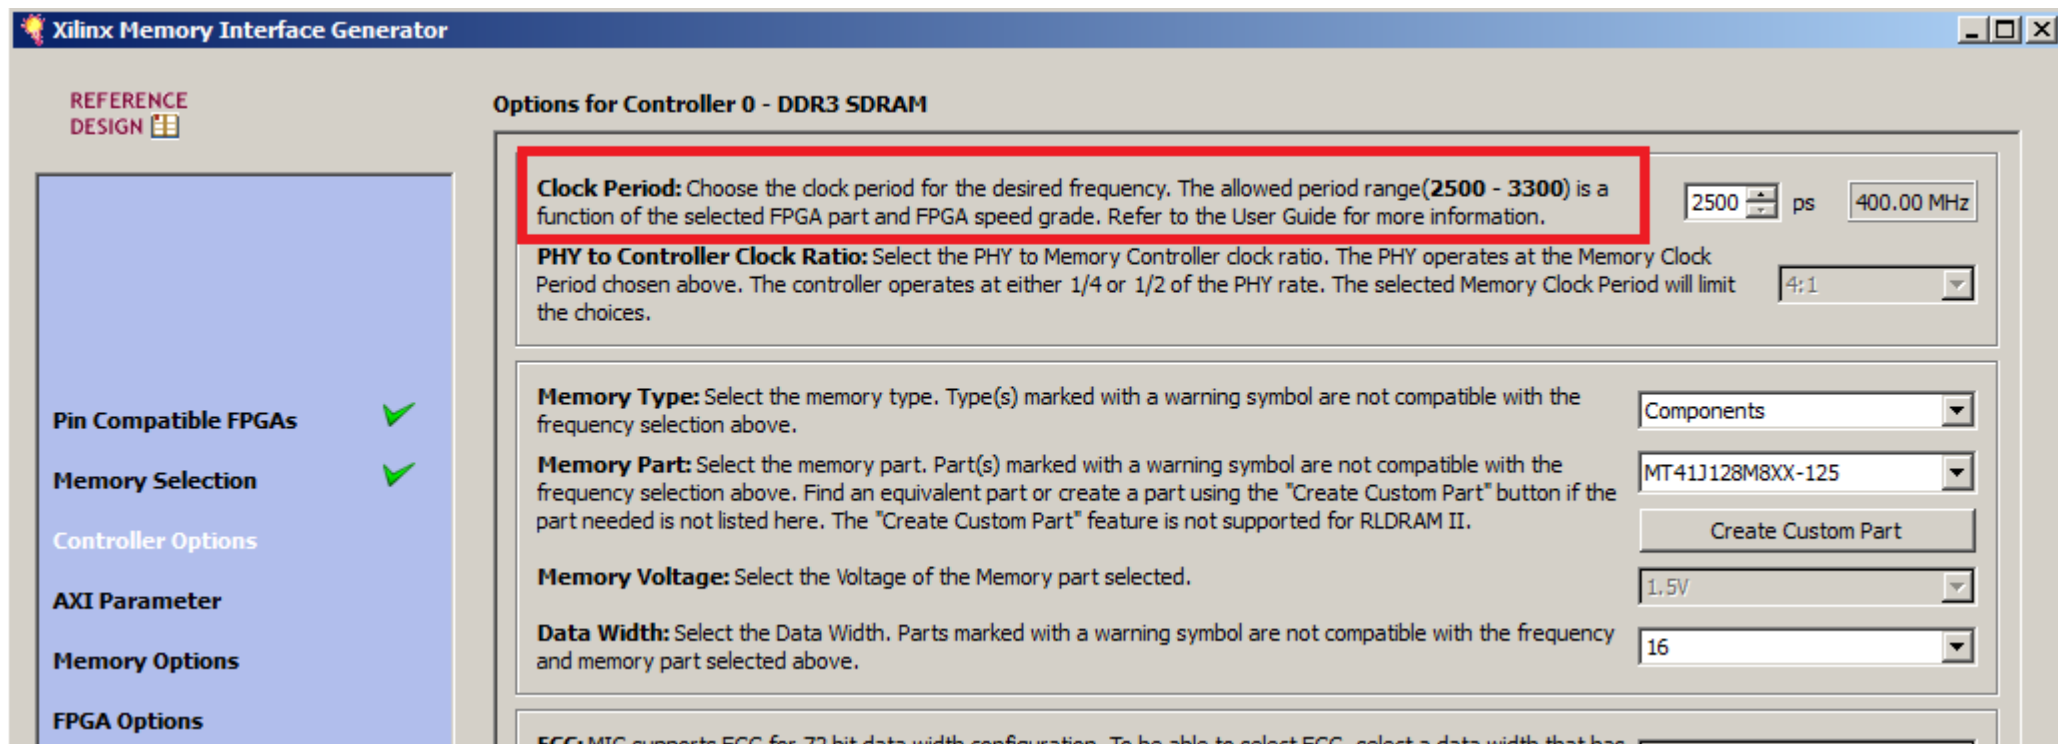
FPGA type	memory clock [MHz]	memory width	practical FIFO bandwidth [GBit]	internal bus multiply 2/4/8	internal AXI bus [bit]	internal F DIV 1/2/4	internal AXI speed [MHz]	FPGA AXI max [bit]	FPGA AXI max [MHz]
Artix7-2	400	16	4,2624	2	32	1	400	128	150
Artix7-2	400	16	4,2624	8	128	4	100	128	150
Artix7-2	333	16	3,548448	4	64	2	166,5	128	150
Artix7-2	333	16	3,548448	8	128	4	83,25	128	150
Artix7-2	200	16	2,1312	4	64	2	100	128	150
Artix7-3	533	16	5,679648	8	128	4	133,25	128	150
Kintex7-2 HP	666	32	14,193792	8	256	4	166,5	256	250
Kintex7-2 HR	533	32	11,359296	8	256	4	133,25	256	250
Kintex7-3 HP	800	32	17,0496	8	256	4	200	256	250
Kintex7-3 HR	533	32	11,359296	8	256	4	133,25	256	250

FPGA design tool restrictions 1v2

1) limitation in FPGA AXI Bus version: if you use a SOC device the PS based AXI bus is implemented in version AXI3, which has a limitation to a BURST memory transaction of 16 clock cycles. This will reduce the pure memory bandwidth of the formula in column E from 0.666 to <0.4 . This can be fixed by a PL based AXI4 → AXI3 bridge, but additional space/setup/risk is needed

The efficiency number of 0.666 assume a BURST memory transaction of ≥ 64 clock cycles

2) limitation due to input mask of VIVADO 2015.4: in this mask of MIG setup, you are NOT free to set the selected DDR3 to work at any clock rate as shown in the RED marked area which means you can run the selected DDR3 “only” between 300 to 400MHz



Options for Controller 0 - DDR3 SDRAM

Clock Period: Choose the clock period for the desired frequency. The allowed period range(2500 - 3300) is a function of the selected FPGA part and FPGA speed grade. Refer to the User Guide for more information. 2500 ps 400.00 MHz

PHY to Controller Clock Ratio: Select the PHY to Memory Controller clock ratio. The PHY operates at the Memory Clock Period chosen above. The controller operates at either 1/4 or 1/2 of the PHY rate. The selected Memory Clock Period will limit the choices. 4:1

Memory Type: Select the memory type. Type(s) marked with a warning symbol are not compatible with the frequency selection above. Components

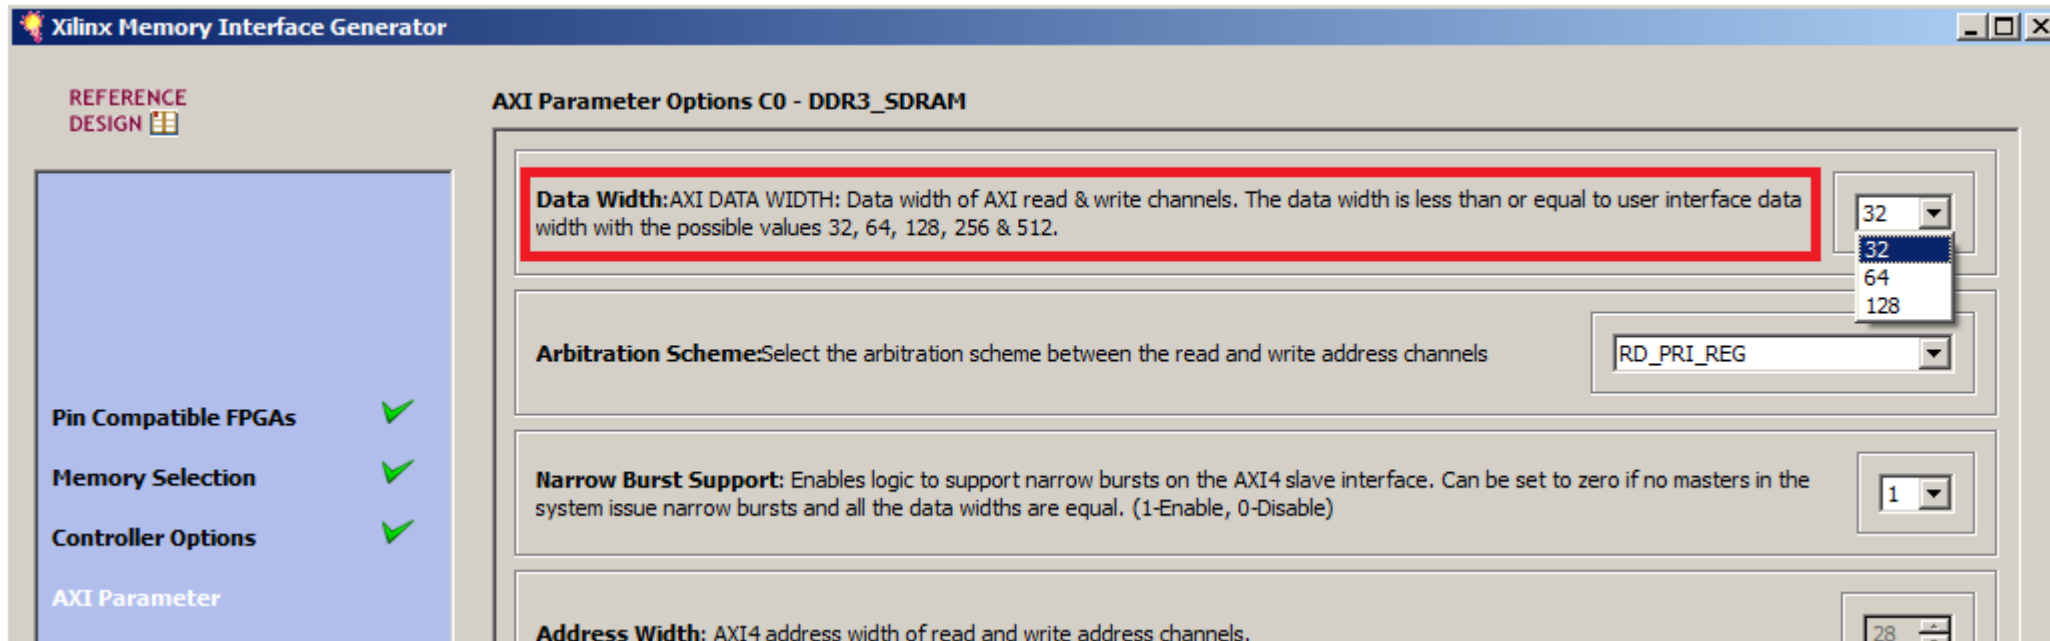
Memory Part: Select the memory part. Part(s) marked with a warning symbol are not compatible with the frequency selection above. Find an equivalent part or create a part using the "Create Custom Part" button if the part needed is not listed here. The "Create Custom Part" feature is not supported for RLD RAM II. MT41J128M8XX-125

Memory Voltage: Select the Voltage of the Memory part selected. 1.5V

Data Width: Select the Data Width. Parts marked with a warning symbol are not compatible with the frequency and memory part selected above. 16

FPGA design tool restrictions 2v2

3) another limitation due to input mask of VIVADO 2015.4: in this mask of MIG setup, you are NOT free to set the AXI bus width based on the selected 16bit DDR3 to work at any width as shown in the RED marked area which is part of the EXCEL table to set the bus width multiplier to 2/4/8, which corresponds to the values in the drop down table of 32/64/128.



The screenshot shows the Xilinx Memory Interface Generator (MIG) software interface. The title bar reads "Xilinx Memory Interface Generator". On the left, there is a "REFERENCE DESIGN" section with a list of categories: "Pin Compatible FPGAs", "Memory Selection", "Controller Options", and "AXI Parameter". Each category has a green checkmark. The main area is titled "AXI Parameter Options C0 - DDR3_SDRAM". It contains several parameter settings:

- Data Width:** AXI DATA WIDTH: Data width of AXI read & write channels. The data width is less than or equal to user interface data width with the possible values 32, 64, 128, 256 & 512. This section is highlighted with a red box, and the dropdown menu is open, showing options 32, 64, and 128.
- Arbitration Scheme:** Select the arbitration scheme between the read and write address channels. The dropdown menu shows "RD_PRI_REG".
- Narrow Burst Support:** Enables logic to support narrow bursts on the AXI4 slave interface. Can be set to zero if no masters in the system issue narrow bursts and all the data widths are equal. (1-Enable, 0-Disable). The dropdown menu shows "1".
- Address Width:** AXI4 address width of read and write address channels. The dropdown menu shows "28".

The clock value of the AXI bus configured here is NOT part of MIG configuration and you have to take care by yourself to adjust the rate external- and internal clock rates due to the selected memory speed, the AXI bus width, the basic rules mentioned mentioned before. **Here other restriction apply as well, as eg. μ Blaze on AXI can run only at full or half AXI speed.**

Which mechanisms ensure high bandwidth?

- ✓ Read and understand DDRx datasheet
- ✓ Long burst on DDRx
- ✓ Right scheduler on DDRx
- ✓ Right memory organisation on UMA architecture
- ✓ Match DDRx and FPGA memory controller settings
- x UMA, unified memory architecture = CPU + processing on 1 memory bank
- x Short burst
- x DDRx



sensor to image

Thanks for your interest and time!

Sensor to Image GmbH
Lechtorstrasse 20
D 86956 Schongau

Website: www.sensor-to-image.de
Email: email@sensor-to-image.de

Tel.: +49 8861 2369 0
Fax : +49 8861 2369 69
EU-VAT-ID : DE-812693714
Register Court ID Munich: HRB125437