

Design guidelines for timing closure

Thomas Zerrer

Slowest speed grades for AMD devices:

Product Version	FPGA Family	Link Speed / Link Width				
		Gen 1, 2.5 Gbps	Gen 2, 5 Gbps	Gen 3, 8 Gbps		Gen 4, 16 Gbps
		X1 / X2 / X4 / X8	X1 / X2 / X4 / X8	X1 / X2 / X4	X8	X1 / X2 / X4 (***)
64-Bit	Artix 7	-1	-2 for X1/X2*	--		--
64-Bit	Kintex 7	-1 / -2 for X8**	-1 / -2 for X4	--		--
256-Bit	Artix 7	-1	-2	--		--
256-Bit	Kintex 7	-1	-1 / -2 for X8	--		--
256-Bit	Virtex 7	-1	-1 / -2 for X8	-1 / -2 for X4	-3 / -2 ^x	--
256-Bit	Ultrascale	-1	-1	-1	-2 / -1 ^x	--
256-Bit	Ultrascale+ / MPSoC	-1	-1	-1	-1	-1

Table 1

(*) Gen 2 – X4 is not supported for the 64 Bit version for Artix, use the 256 Bit version instead. Artix does not support x8 links.

(**) Gen 2 – X8 is not supported for the 64-Bit version for Kintex, use 256-Bit version instead.

(***) Gen 4 is supported by AMD only for specific devices. Please check device datasheets, if Gen4 is supported.

(^x) Speedgrade is supported but with limitations (Maximum of 2 Read and 2 Write channels). See Chapter 5.5 for details.

This table has been validated with 8 independent read and 9 independent write channels except for speedgrades marked with (^x).

If more channels are used, it might be possible that a higher speedgrade has to be selected. Contact Smartlogic in this case for a recommendation.

- This table lists the minimum speedgrade required for the IP Core and for the AMD Hard IP
- Speedgrades with X only meet timing when using a maximum of 2 DMA Write and 2 DMA Read interfaces. Speedgrades without X were experimentally tested with 9 Write and 8 Read interfaces.
- The following Link Speed / Link width combinations need special attention in order to achieve timing closure:
 - 64-Bit core : Gen2-X4
 - 256-Bit Core : Gen3-X8

Slowest speed grades for Altera devices :

Product Version	FPGA Family	Link Speed / Link Width				
		Gen 1, 2.5 Gbps X1 / X2 / X4 / X8	Gen 2, 5 Gbps X1 / X2 / X4 / X8	Gen 3, 8 Gbps X1 / X2 / X4 X8		Gen 4, 16 Gbps X1 / X2 / X4 (***)
256-Bit	Arria 10	-3	-3	-3	-1 / -2 ^X	--
256-Bit	Cyclone 5	-8	-7*	--		--
256-Bit	Stratix 10	-3	-3	-3	-2	-2
256-Bit	Cyclone 10	-6	-6*	--	--	--

Table 2

(*) Cyclone 5 and Cyclone 10 do not support x8 links.

(***) Gen 4 is supported by Altera only for specific devices. Please check device datasheets, if Gen4 is supported.

(X) Speedgrade is supported but with limitations (Maximum of 2 Read and 2 Write channels). See Chapter 5.5 (UG) for details.

This table has been validated with 8 independent read and 9 independent write channels except for speedgrades marked with (X) .

If more channels are used, it might be possible that a faster speedgrade has to be selected. Contact Smartlogic in this case for a recommendation.

- This table lists the minimum speedgrade required for the IP Core and the Altera Hard IP
- Speedgrades with X only meet timing when using a maximum of 2 DMA Write and 2 DMA Read interfaces. Speedgrades without X were experimentally tested with 9 Write and 8 Read interfaces.
- The following Link Speed / Link width combinations need special attention in order to achieve timing closure:
256-Bit Core : Gen3-X8

Slowest speed grades for Lattice devices :

Product Version	FPGA Family	Link Speed / Link Width		
		Gen 1, 2.5 Gbps	Gen 2, 5 Gbps	Gen 3, 8 Gbps
		Speedgrade / Core Voltage		
256-Bit	Crosslink NX	-7 High Performance	-7 High Performance	N/A
256-Bit	Certus NX	-7 High Performance	-7 High Performance	N/A
256-Bit	Certus Pro NX	-7 1.0V	-7 1.0V	-9 1.0V

Table 3

This table has been validated with 4 independent read and 4 independent write channels.
 If more channels are used, it might be possible that a faster speedgrade has to be selected. Contact Smartlogic in this case for a recommendation.

- This table lists the minimum speedgrade required for Lattice devices

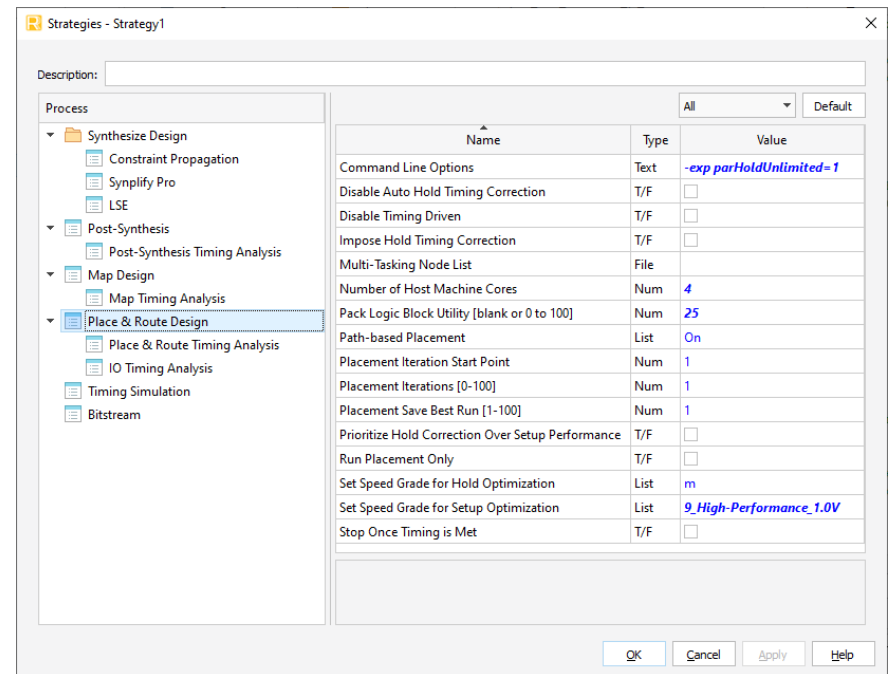
HardIP Situation:

- The PCIe HardIP has to be clocked with 250 MHz for Gen3
- The Lattice PCIe HardIP has very high input setup values (2 ns) and very high clock to out values (2 ns) for outputs. Since the TLP interface has to be clocked at 250 MHz only 2 ns timing margin remain for routing and logic.
- Smartlogic added a highly optimized logic block to meet timing closure on this critical path
- Starting with Radiant 2023.1 the user is required to specify the clock uncertainty of a PLL generated clock. For 250 MHz the uncertainty has to be 250 ps according to the Lattice data sheet for Certus NX. This uncertainty will further decrease the timing margin from 4.0 ns to 3.75 ns (effectively 2.0 ns to 1.75 ns).

Recommended Radiant settings to meet timing:

- The provided reference design has recommended P&R settings from Lattice. Add these settings to your project.

These settings work with Radiant 2023 and 2024.1
It might be possible that they change with newer Radiant Versions.



Specifying the clock uncertainty

- When a clock is generated with a PLL, the clock edges will jitter. Within the same clock domain jitter reduces the clock period and needs to be considered during timing analysis. Important is, that jitter has to be considered only for setup but not for hold. For hold jitter has only to be considered between two different clocks which jitter.
- Example to specify the uncertainty within the PDC file, only for setup:

```
set_clock_uncertainty -setup 0.25 [get_clocks {clk_125M}]
```

250 MHz external Oscillator

- If possible, work with a 250 MHz external oscillator and clock the PCIe HardIP with this free running clock
- In this case it is possible to set the uncertainty to a very low value (< 10 ps) which will greatly improve timing closure
- Other clocks can be derived out of this 250 MHz clock. The uncertainty has to be set correctly for the derived clocks, but this should be uncritical

125 MHz external Oscillator:

- It is possible to work with a 125 MHz oscillator, but Smartlogic does not recommend this, as it makes timing closure more difficult. But since the Versa Board NX has a 125 MHz oscillator, the following steps are recommended to close timing:
 1. Radiant seems to produce worse results, when the 250 ps uncertainty is specified correctly compared with timing slacks, when the uncertainty is set to a relaxed value of 25 ps.
 2. If you have entered a 25 ps uncertainty, make sure, that you will have a timing margin of 0.225 ps minimum.
 3. When running P&R, always keep in mind that the values reported in the Place & Route Report under „Cost Table Summary do not reflect the true timing but are only an estimate. Therefore always look in „Place & Route Timing Analysis“ – endpoint slacks , where the true timing slack is reported.
 4. For experimental / pre-liminary checks in the Lab, you can work with bitstreams that have only a timing slack of 0.100 ns to 0.250 ns. However, for production bitstreams, it is necessary that the timing slack includes the total uncertainty of the clock.

250 MHz 90 degree phase shifted clock

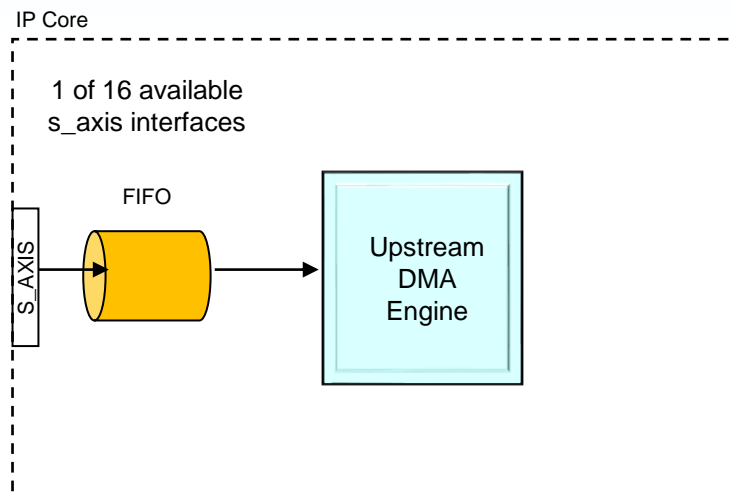
- Due to the HardIP situation, Lattice introduced a 250 MHz phase shifted clock, in order to increase the timing margin of the output signals and to decrease the timing margins slightly for the input signals of the HardIP. See X4 user guide of the PCIe HardIP for details.
- However in this case since two PLL generated clocks are introduced, the uncertainty has to be specified for setup **and** hold, which imposes new timing challenges:

```
set_clock_uncertainty 0.25 [get_clocks {clk_250M}]
```

The Smartlogic IP Core offers the option to work with the Lattice recommended 90 degree phase shifted clock and without this clock. The recommended approach is to work with only the 250 MHz clock which has several advantages (fewer clocks, uncertainty needs only to be specified for setup).

In order to work with 250MHz only, connect the toplevel ports of the IP (entity pcie_top_hcc) as follows

```
PORT MAP (  
    clk_usr_i      => clk_250M,  
    clk_usr_ps90_i => clk_250M,  
    ...  
)
```

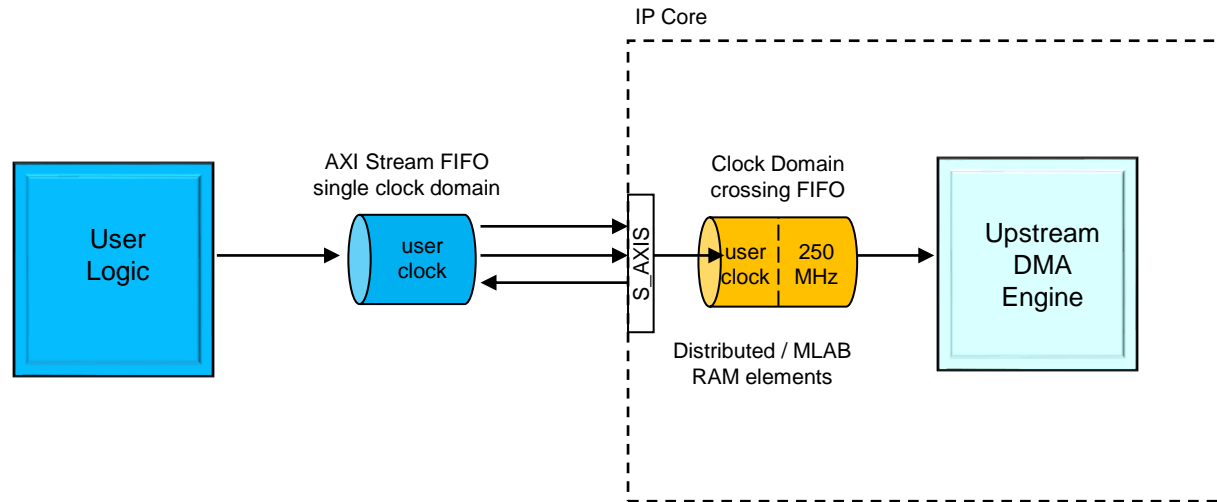



The upstream side of the core has up to 16 axi stream interfaces. The number is user configurable and each axi stream interface has its own data fifo with adjustable depth at compile time.

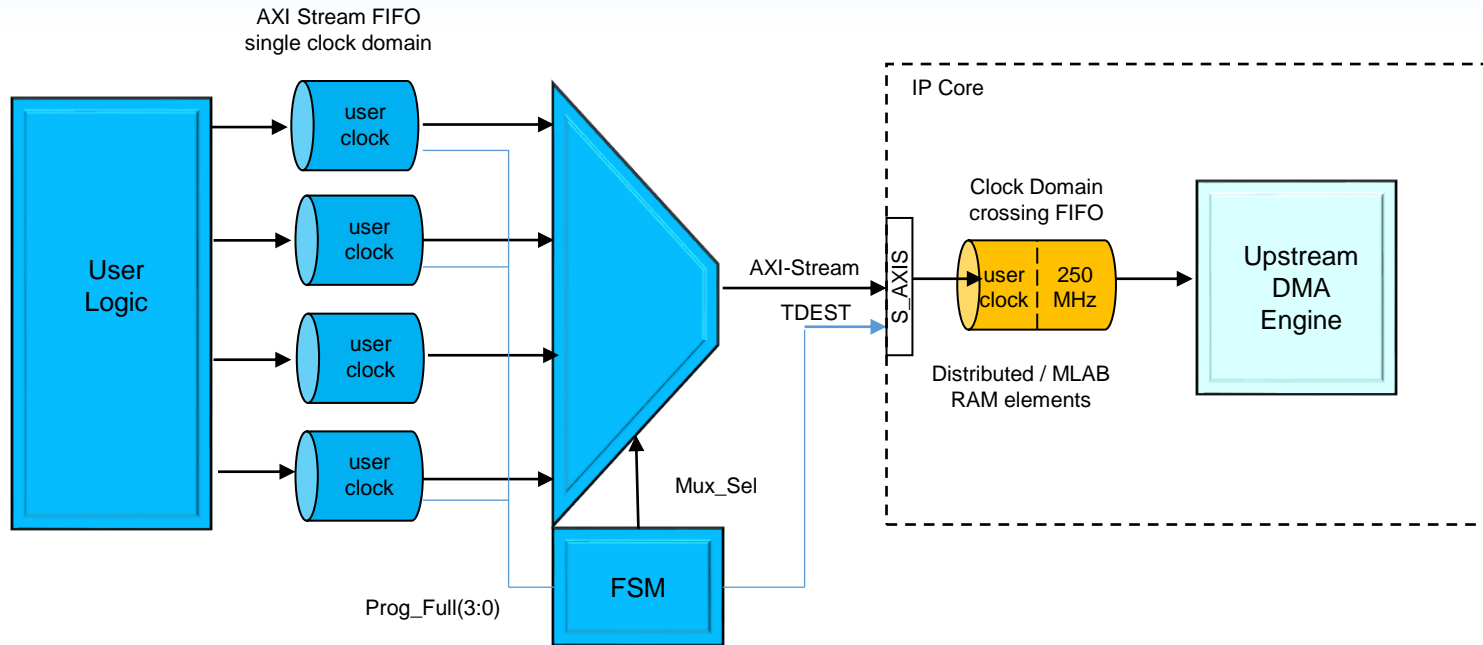
The FIFO can be built up with either BlockRAMs or with distributed RAMs.

FIFO Type	FIFO depth	Timing	Comment
Distributed / MLAB RAM	AMD : 6* Altera : 5* Lattice : 4*	Fast clock to out	Use this RAM type for timing critical designs and to save BRAMs.
BRAM	AMD : 9 or 10 Altera : 9 Lattice : 9	Slow clock to out	For Timing critical designs this RAM type is not recommended. You may try depths of 9 but this is not guaranteed to meet timing

(*) This number is the depth of the distributed RAM primitive. It might be tolerable to use a slightly higher number (e.g. one more) but this will use logic resources very inefficiently



- In case the distributed / MLAB fifo depth is not sufficient, the user may add an additional AXI Stream FIFO in the datapath in front of the S_AXIS interface of the core
- In case the User clock is below 250 MHz, the timing for this FIFO is relaxed and it should be possible to build this FIFO with BRAMs. A further advantage of this FIFO is, that it is a single clock domain FIFO
- Suitable FIFOs can be found in the IP catalog of the FPGA vendor
- In case that no AXI Stream FIFO is available, it is possible to instantiate a FIFO, where the inverse of the empty is connected as tready. The read input is ready, when the FIFO is not empty. Make sure, that the FIFO is configured as „Fallthrough“ FIFO.



- The TDEST Inputs of each AXI Stream interface can be used to reduce the number of physical interfaces of the IP Core, while maintaining the number of destination databuffers in host memory.
- Sometimes it is overlooked, that each AXI Stream slave interface can reach ALL destination data buffers (up to 64).
- Therefore it is possible to reduce the number of interfaces by designing a FIFO mux structure within user logic. This will greatly improve timing closure in the critical 250 MHz paths.
- Note : The TDEST inputs are only available in the HCC and ABD IP Core. The Flex IP Core does not have the TDEST inputs for s_axis interfaces. For exact Timing see chapter 2.1 of the User guide.

Feature	Parameter name	Recommended value	comment
Number of Upstream interfaces (s_axis)	Write_Data_Interfaces_in_use_g	1-9	Higher values may be possible but are not guaranteed
Number of downstream interfaces (m_axis)	Read_Data_Interfaces_in_use_g	1-8	
RAM elements for s_axis data FIFOs	DMA_Write_Fifo_params_c.dFIFO_bram in dma_pkg.vhd	false	
RAM elements for m_axis data FIFOs	DMA_Read_Fifo_params_c.dFIFO_bram in dma_pkg.vhd	false	
Disable address fifo almost empty interrupts (upstream)	DMA_Read_Implement_irq_sg_ae_regs_c in dma_pkg.vhd	false	In this case the user has no ringbuffer support. True might be possible but is not guaranteed.
Disable address fifo almost empty interrupts (downstream)	DMA_Read_Implement_irq_sg_ae_regs_c in dma_pkg.vhd	false	

It is also recommended to enable physical optimizations and higher P&R efforts within Vivado / Quartus / Radiant

Table 1 marks some speedgrades with „X“. In this case, the following settings are valid

Feature	Parameter name	Recommended value	comment
Number of upstream interfaces (s_axis)	Write_Data_Interfaces_in_use_g	1-2	Higher values might be possible but are not guaranteed
Number of downstream interfaces (m_axis)	Read_Data_Interfaces_in_use_g	1-2	
RAM elements for s_axis data FIFOs	DMA_Write_Fifo_params_c.dFIFO_bram in dma_pkg.vhd	false	
RAM elements for m_axis data FIFOs	DMA_Read_Fifo_params_c.dFIFO_bram in dma_pkg.vhd	false	
Disable address fifo almost empty interrupts (upstream)	DMA_Read_Implement_irq_sg_ae_regs_c in dma_pkg.vhd	false	In this case the user has no ringbuffer support. True might be possible but is not guaranteed.
Disable address fifo almost empty interrupts (downstream)	DMA_Read_Implement_irq_sg_ae_regs_c in dma_pkg.vhd	false	

It is also recommended to enable physical optimizations and higher P&R efforts within Vivado / Quartus / Radiant

Software Settings that ensure low FIFO depths :

In case of several s-axis stream interfaces the channels are transmitted in round robin fashion, where each channel is allowed to transmit the amount of data contained in its associated IncrementLineOffset register. If the incrementLineOffset registers are set with high values, the FIFO buffers need also higher capacity, since they have to survive without reaching a „Full“ until the time they are selected and can be emptied.

In case of more than one active interface, we advise the following settings:

- The IncrementLineOffsets for DMA Write should be set to 0x200.
Note for Video Applications: If 0x200 does not match a complete line, use „stream“ mode
- Channels that are only transmitting data from to time to time can be set to lower values but it must be a power of 2.
- Higher priority channels or channels that have the double data amount as others might be set to 0x400

Example : Video Data Transmission Y (16-Bit), Cr (8-Bit), Cb (8-Bit) and Audio data:

Y Channel 0x400

Cr Channel 0x200

Cb Channel 0x200

Audio channel 0x100 or 0x200