

I/O timings for xCORE200

IN THIS DOCUMENT

- ▶ Basics
 - ▶ The xCORE200 IO circuitry
 - ▶ Source-synchronous clocks
 - ▶ Non source-synchronous clocks
 - ▶ A worked example: master I2S
 - ▶ Appendix - Timing components
-

In this document we show you how to estimate the IO timings for data input and output by the xCORE200. Most interfaces are slow enough that the timings work out in any case. Faster interfaces need their timings analysed to ensure operation under all circumstances.

On reading this document you will see that there are a couple of circumstances where it will “just work” because the clock timings are so forgiving:

- ▶ Application clocks of 15 MHz or slower typically just work
- ▶ For data and clock travelling in the same direction (ie, source-synchronous clocking where both output by the xCORE or both input by the xCORE), frequencies of up to 60 MHz can work (if they work electrically).

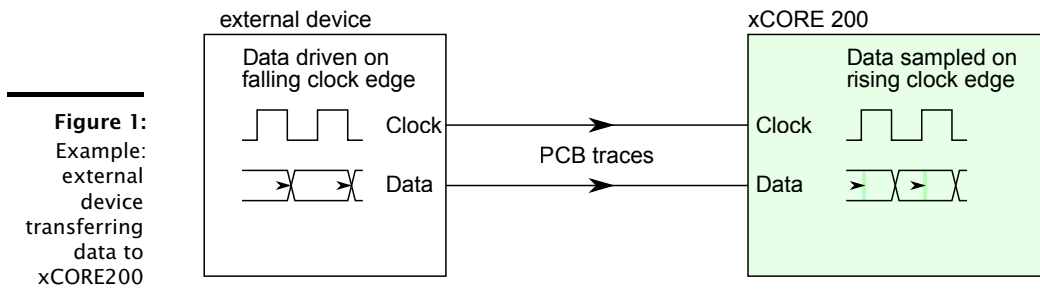
Faster clocks may also work, but an analysis is needed to show whether any programmatic delays are necessary to cover all corners. In particular, some timings depend on PVT variations (Process, Voltage, and Temperature). An analysis will ensure that an interface will work under all possible corners, rather than the corners that were tested.

1 Basics

IO refers to transferring digital data between devices. In this document, we are concerned by transferring a single bit of data between an xCORE200 device and some external device. We only consider *synchronous IO*, that is, we assume that there is a *clock*, and the data is transferred synchronous to the clock. The limitation to just use a single data-line is trivially generalised to multiple data lines, including data valid lines.

We assume that the data is transferred in a specific direction. That is, there will be two sides to the data transfer. One side will be *outputting* the data-bit by *driving* the data-net, the other side will be *inputting* the data-bit by *sampling* the data-net. The clock is generated by either side or by a third party. An example is shown in Figure 1 where the xCORE is using an externally provided clock to sample a data

signal, and the external device is generating the clock and the data. We assume that data is clocked in by the input side on the *rising edge* of the clock, and that the data is clocked out by the output side on the *falling edge* of the clock.



We use the term *application clock* to refer to the clock signal; this makes it clear that it is different from other clock signals in the system, such as the *Core clock* of the xCORE (typically 500 MHz), the *Reference clock* of the xCORE (typically 100 MHz), or the PLL input clock (typically 20-25 MHz).



There are three critical times in the analysis of the IO timings; one on the output side, and two on the input side. Every digital device will specify those times in one form or another, although terminology may differ:

Timing parameter	Brief description
T(CLKtoDATA)	Time at which output will change
T(setup)	Time at which input must be stable
T(hold)	Time until which input must be kept stable

On the output side, T(CLKtoDATA) refers to the time between the clock edge falling and the data appearing on the data line. T(CLKtoDATA) normally has a range of possible values that depend on environmental conditions, and we use the *minimum* and *maximum* T(CLKtoDATA) to define the first possible time where the data could become valid, and the last possible time where the data will be valid. In between these times, the data may be old, new, or something undefined. Normally, T(CLKtoDATA) has a positive value, indicating that the data will be valid after the falling clock edge. A negative value indicates that the data may or will be valid prior to the falling clock edge.

On the input side, T(setup) refers to the time at which the data must be valid for it to be clocked in correctly. T(setup) is the length time before the rising application clock edge. T(hold) refers to the amount of time that the data should stay valid for, in order for it to be clocked in correctly. T(hold) is the length of time after the rising edge of the application clock. Normally setup and hold times are positive numbers. By convention, the setup-time is the time that the data has to be stable *before* the clock edge, and the hold time is the time that the data has to remain stable *after* the clock edge. We use a *negative* setup-time to indicate that the data

only has to be stable *after* the clock edge; and a *negative* hold time to indicate that the data does not have to be held after the clock edge.



These times are summarised in Figure 2 and Figure 3. The important consideration is that the external device shall not change in the capture window of the xCORE (for input on the xCORE), or that the xCORE shall not change the data in the capture window of the external device (for output by the xCORE)

Note that this assumes an idealised setup - typically clock and data travel across the PCB, and signals will not abruptly change from a zero to a one but take time to rise and fall depending on the PCB characteristics. A discussion on this is outside the scope of this document. In this document we assume an ideal setup where there are no latencies incurred in PCB traces.

The ideal way to implement data transfer uses a *source-synchronous clock*. That is, the device that drives the data also drives the clock. The alternative is that the clock is driven by the device that inputs the data, or by some third party. Section 4 describes the timings of source-synchronous clocked systems; Section 5 describes the timings of systems that are not source-synchronous.

2 The xCORE200 IO circuitry

IO circuitry comprises two sorts of components: logic that is running synchronous to the core clock, and IO drivers that are asynchronous. These are shown in Figure 4.

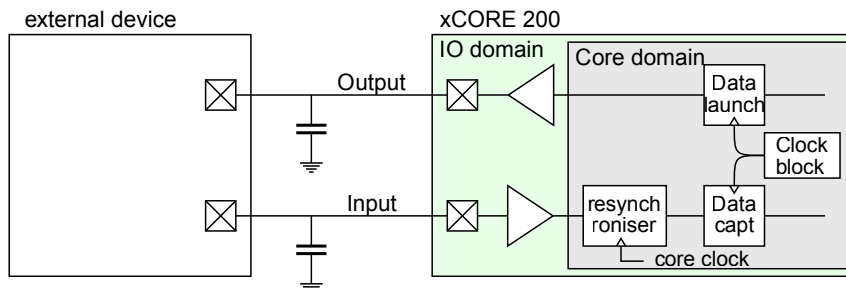


Figure 4:
xCORE200 IO
elements

On the output path, data is launched on a clock-edge, it then has to traverse the IO drivers, which in turn have to drive the net, until the signal is stable. On the input side, the external device has to drive the net, it then has to traverse the input buffer, a resynchroniser which makes the signal synchronous to the core clock, until it is captured on a clock edge. There are multiple clock blocks available to deal with different application clocks.

The IO driver logic introduces a delay and uncertainty that is independent of the core clock, but that does depend on environmental factors such as the core voltage, IO voltage, temperature, and process variation (“PVT variation”). The timings depend on the pins that you use; some pins are better matched and/or faster than others. In the main body of this document we show how to calculate worst case $T(\text{setup})$, $T(\text{hold})$, and $T(\text{CLKtoDATA})$ for xCORE200 devices. These calculations are based on three device parameters, the input skew, output skew, and round trip time. These three device parameters are characterised for xCORE200 devices as follows:

Parameter	Description	Min	Typ	Max	Units
$T(\text{iskew})$	Input skew			2.0	ns
$T(\text{oskew})$	Output skew			2.7	ns
$T(\text{RTT})$	Round trip time	3.0		11.3	ns

$T(\text{oskew})$ is defined as measuring the time difference between two signals that are launched by the xCORE at the same time (Figure 5). $T(\text{iskew})$ is defined as the time difference between two signals entering the xCORE at the same time (Figure 6). $T(\text{RTT})$ finally is defined as the time difference between a signal being launched by the xCORE, travelling through a loopback, and then back in (Figure 7). Note that all three only relate to the IO domain, not to the core domain.

We will explain later how to use these parameters. The values above are worst-case values assuming a very small external load of 2 pF. Appendix A lists tighter timings

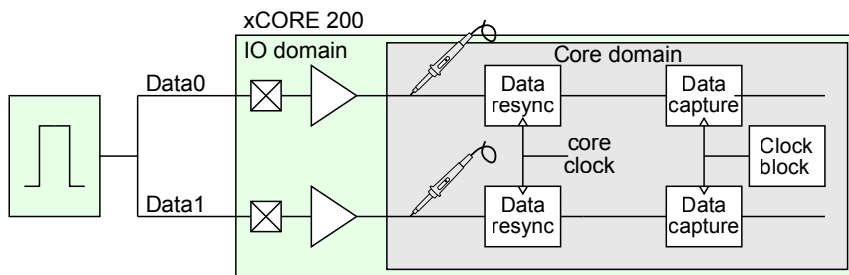


Figure 5:
Definition of $T(iskew)$

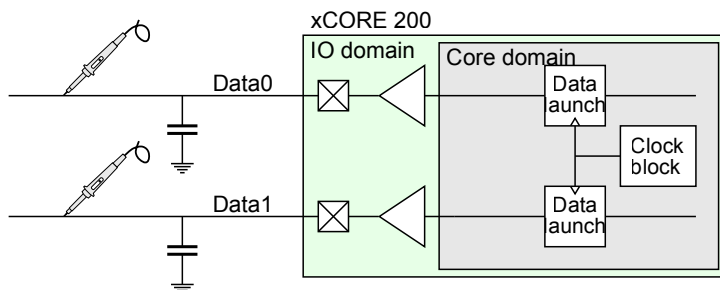


Figure 6:
Definition of $T(oskew)$




Figure 7:
Definition of $T(RTT)$

for these parameters for specific groups of IO pins on specific tiles, and explains in which situations the minimum and maximum values may be encountered. This will enable you to make tighter designs, and guide you in using the pins that have tighter IO timings. Appendix A also gives timings for higher loads.

In the rest of this document we use the above worst-case values for $T(iskew)$, $T(oskew)$ and $T(RTT)$ to calculate example IO timings.

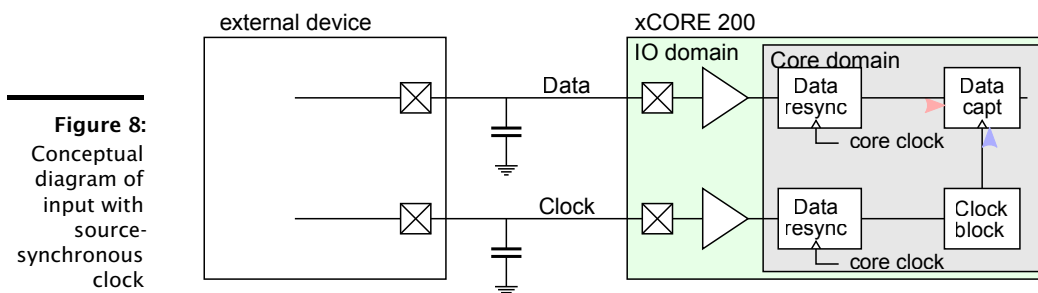
3 Source-synchronous clocks

The timings of source-synchronous clocked systems depend on the *skew* between pins only. The *skew* measures the difference in timings for signals that travel in the same direction, and is up to $T(oskew)$ (2.7 ns) depending on the conditions and the direction of the signal.

When designing a source-synchronous clocked system, you can use the skew times listed in this document as a start, but you should add the skew introduced by external components. All nets (clock and data) should be matched for maximum performance.

3.1 Input: Calculating setup and hold times for an external application clock

This section describes how you can compute the setup and hold time for capturing data relative to an external application clock. This is the situation where both clock and data originate outside the xCORE200, as shown in Figure 8. The blue arrow is the timing path of the clock, and the red arrow is the timing path of the data. The difference between the two is the skew that governs the uncertainty of where the data will be captured.



The table below shows how to compute the setup and hold timings for the xCORE200 processor. The setup- and hold-time depends on the *input skew* of the xCORE200 and *core clock frequency* of the xCORE. and the table shows setup- and hold-times for a 500 MHz device, and a formula for the general case. The general formulas show that $T(iskew)$ is required on both sides of the eye. An extra $T(coreclock)$ is required at the end of the eye because of resynchronisation. The whole window is shifted one $T(coreclock)$ to the right because the clock is delayed by one $T(coreclock)$. The signals are shown in Figure 9.

Parameter	500 MHz	For any core-clock
T(setup)	0 ns	$T(iskew) - T(coreclock)$
T(hold)	6 ns	$T(iskew) + 2 \times T(coreclock)$

$T(\text{is skew})$ is at most 2 ns. For specific sets of pins tighter $T(\text{is skew})$ can be found in Appendix A, reducing the eye size.

For a part that runs on a 500 MHz core clock we show the setup and hold times in the diagram below, in this example with a 50 MHz application clock. The capture point is notionally 3 ns after the application clock edge, with 3 ns uncertainty either way. This is the *minimum* required eye, and the data may be presented earlier, for example before the clock edge. This will show a more familiar eye that straddles the clock edge.

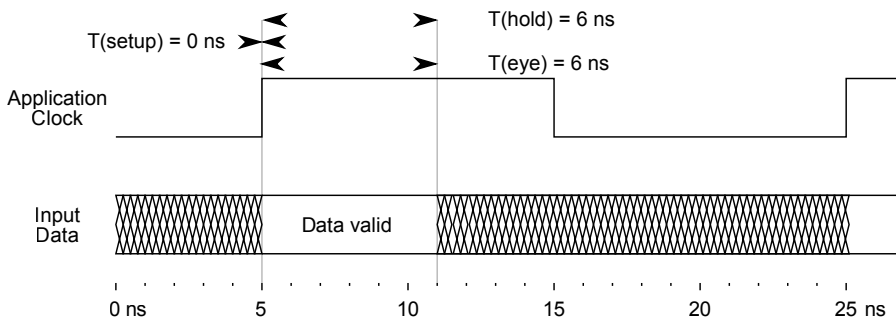


Figure 9:
Setup and hold time for a 50 MHz external clock.

We can calculate the eye by adding the setup-time and the hold-time. Given that $T(\text{setup})$ is $T(\text{is skew}) - T(\text{coreclock})$ and that $T(\text{hold})$ is $T(\text{is skew}) + 2 \times T(\text{coreclock})$, the eye is $T(\text{is skew}) - T(\text{coreclock}) + T(\text{is skew}) + 2 \times T(\text{coreclock}) = 2 \times T(\text{is skew}) + T(\text{coreclock})$ or at most 6 ns for an arbitrary pair of pins on a 500 MHz xCORE200 part.

3.1.1 Adjusting setup and hold timings


If the setup and hold times are too early or too late to reliably capture the data, then they can be adjusted in increments of a single *core-clock cycle*, ie, in steps of 2 ns on a 500 MHz part. This is implemented by one of two mechanisms:

- ▶ The data signal can be delayed by up to *five* core-clock cycles by setting a delay on the data port. Setting the delay to X on the data port, will subtract X core-clock cycles from the hold-time, and add X core-clock cycles to the setup-time. Use the `set_pad_delay` function.
- ▶ An external application clock signal can be delayed by up to *4096* core-clock cycles by setting the rise-delay and fall-delay on the clock block. Setting the rise- and fall-delays to X , will add X core-clock cycles to the hold-time, and subtract X core-clock cycles from the setup-time. The application clock should never be delayed by more than one application clock cycle. Use the `set_clock_fall_delay` and `set_clock_rise_delay` functions. The delay is limited to half a clock cycle.

Setting the data-delay or the application-clock rise- and fall-delay brings the eye forwards and backwards; it does not affect the size of the eye.

An example is shown in Figure 10 for an external clock, where the data is delayed by one core-clock. As it is delayed in the chip, it needs to be presented on the pads one core-clock-cycles, or 2 ns, earlier to meet the setup-time. The delay also reduces the hold-time by 2 ns, effectively bringing the eye forward by 2 ns.

Figure 10:
Setup and hold times for data given an external clock, with a delay of one core-clock cycles in the data.



It is tempting to think that any device that meets this eye will meet the required setup- and hold-times. However, given that the eye-window can only be moved in steps of a single core-clock cycle, it is not possible to meet all setup- and hold-time combinations that sum to the minimum eye opening. However, any device that provides an eye-opening that is at least one core clock cycle longer than the minimum eye, can always have its setup and hold times met.

3.1.2 Maximum clock rates

The application clock should never be faster than half the core-clock-frequency. For example, if the xCORE is running at 100 MHz (rather than the normal 500 MHz), then the application clock is limited to 50 MHz.

The application clock is also limited by the speed at which signals can be supplied to the xCORE. This is limited by the drive strength, the analogue characteristics of the pad, and inductances, capacitances, and resistance of the PCB traces. In addition, the data has to be stable during the eye.

For example, for a specific driver and PCB design, you may have calculated that 10 ns is sufficient for the data-signal to stabilise after an external clock. In that case the minimum application clock period is $6 \text{ ns} + 10 \text{ ns}$, or 66 MHz (assuming a 500 MHz part, and assuming that the eye can be lined up with the application clock). This is visualised in Figure 11.

3.1.3 Input window summary

Setup and hold times for an external clock for various data delays are shown below:

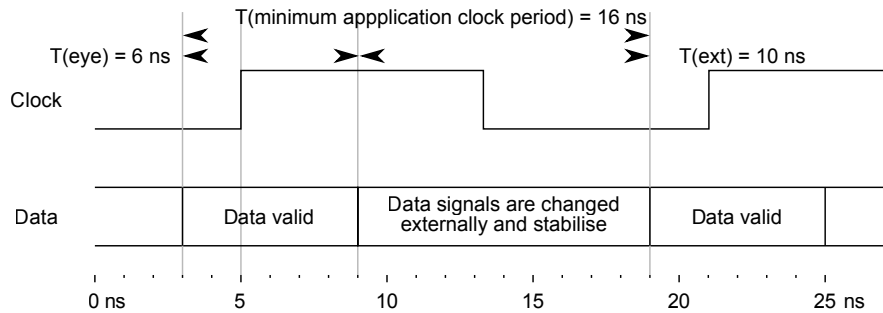


Figure 11: Establishing the minimum externally application clock period.

Core freq	Data delay	T(setup)	T(hold)
500 MHz	2	4 ns	2 ns
	1	2 ns	4 ns
	0	0 ns	6 ns

Assuming that the data is input on a port *p*, the second line can be achieved by calling `set_pad_delay(p, 1)`.

3.2 Output: Calculating clock-to-data times for an internal application clock

The output timings of an xCORE are straightforward if an internal clock is used. This is source-synchronous clocking where both data and clock will leave the synchronous core at the same time, as shown in Figure 12. The blue arrow is the timing path of the clock, and the red arrow is the timing path of the data. The difference between the two is the skew that governs the uncertainty of where the data will be launched.

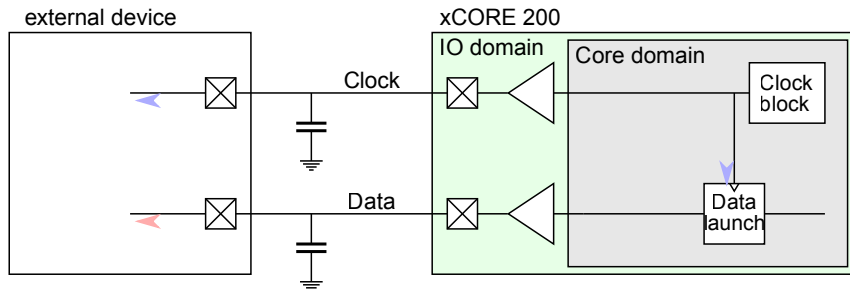



Figure 12: Conceptual diagram of output with source-synchronous clock

The only uncertainty is the output skew. The maximum output skew on an arbitrary pair of pins on an xCORE200 is 2.7 ns, hence the Clock-to-data time is between -2.7 ns and +2.7 ns. That is, the data may precede the clock by as much as 2.7

ns; and the data will be stable 2.7 ns after the falling clock edge, as shown in Figure 13. Better timings can be achieved by picking a tile and bank with a lower output skew, as listed in Appendix A.



If open-drain outputs are used, then an extra 5 ns skew should be accounted for, in addition to the extra times required by the external resistor to pull the signal high.

4 Non source-synchronous clocks

For non source-synchronous clocks, the timings are governed by the *round trip time* through an output pin and an input pin. We define the *round trip time* as the sum of time taken by a signal to traverse the output path of the xCORE200, and the time taken by a signal to traverse the input path of the xCORE200. Electrically, this path will never be taken, but logically it will be taken as one signal (eg, the clock) will travel out whereas the causally dependent signal (eg, the data) will travel in. The *round trip time* is difficult to control tightly. Systems that need to work under a set of environmental conditions need to be closed so that they work with both the best case round trip time $T(RTT_{min})$, and worst case round trip time, $T(RTT_{max})$.

When designing the system, you can use the round-trip times listed in this document as a start, but you should add the time required to drive the input and output net; taking into account the load on the net.

4.1 Input: Calculating setup and hold times for an internal application clock

When data is input on an internal application clock, we have to contend with a round-trip time out of the xCORE and back in, as shown in Figure 14. The blue arrow shows the timing path of the clock causing the data to be driven by the device; the red arrow shows the timing path of the clock sampling the data on the xCORE. The capture window has to be aligned so that it captures the data late enough to allow for the round trip delay; and the window has to be large enough to deal with uncertainty of the minimum and the maximum round trip times.

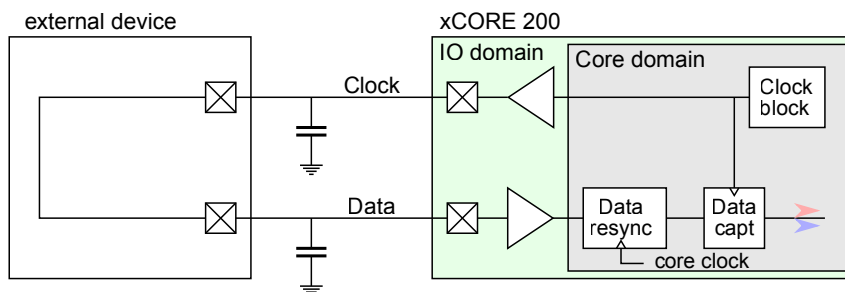


Figure 14:
Round trip of
clock and
data

The minimum and maximum round trip times for XCORE200 are 3.0 and 11.3 ns. (See Appendix A for specific banks that have tighter round trip times.) In addition we have a delay of between 4 and 5 core clocks for resynchronising the data and capturing the data. For a 500 MHz part this brings us to a minimum round trip time of $3.0 + 4 \times 2 = 11.0$ ns, and a maximum round trip time of $11.3 + 5 \times 2 = 21.3$ ns. That is an eye of $21.3 - 11.0 = 10.3$ ns, that closes 11.0 ns before the rising clock edge.

The general formulas for setup and hold timings are given in the table below. The setup time is governed by the maximum round trip time, $T(RTT_{max})$, plus the time

for travelling through the resynchronisers assuming the signal arrives just after the core-clock edge. The hold time is governed by the minimum round trip time, $T(RTT_{min})$, plus the time for travelling through the synchronisers assuming the signal arrives just before the core-clock edge. The hold time is negative because the data can be released before the clock edge:

Parameter	500 MHz	For any core-clock
T(setup)	21.3 ns	$T(RTT_{max}) + 5 \times T(\text{coreclock})$
T(hold)	-11.0 ns	$-T(RTT_{min}) - 4 \times T(\text{coreclock})$

The resulting eye limits the clock to 97 MHz; at which point the data valid window occupies the whole clock cycle, and there is no longer any time for the external device to modify the data. In practice, the external device will have a clock-to-data window, that will reduce the eye and the maximum frequency.




Figure 15:
Setup and hold times of an xCORE on an internal application clock

It is important to note that for clocks in excess of $1/(2 \times 21.3 \text{ ns}) = 23.4 \text{ MHz}$ the eye opening may overlap with the falling edge of the application clock. An external device with a short clock-to-data timing will modify data in the eye. In these cases, the data should either be delayed or sampled on the *falling* edge in order to move the eye away from the window where the data changes. Figure 15 above has a clock period of 50 ns (20 MHz), if one imagines the falling edge of the clock shifting closer to the rising edge, then it will end up in the data valid window, which may cause the device that drives the data to update during the data valid window.

4.2 Output: Calculating clock-to-data times for an external application clock

When data is output on an external application clock, we have to contend with a round-trip time in to the xCORE and back out, as shown in Figure 16. The blue arrow shows the timing path of the clock causing the data to be driven by the xCORE; the red arrow shows the timing path of the clock sampling the data inside the device. The capture window has to be aligned so that it captures the data late enough to allow for the round trip delay; and the window has to be large enough to deal with uncertainty of the minimum and the maximum round trip times.

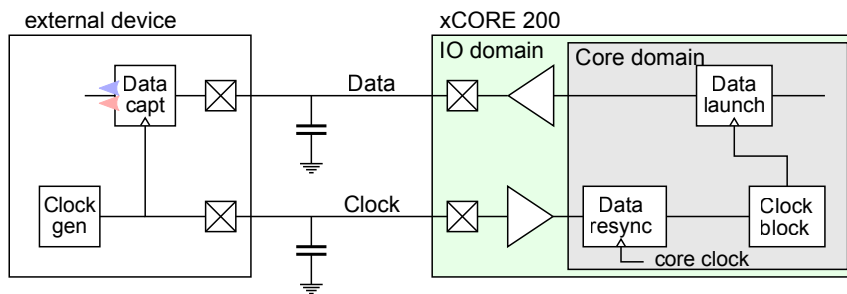


Figure 16:
Round trip of clock and data

The general formulas for setup and hold timings are as follow, and are illustrated in Figure 17:

Parameter	500 MHz	For any core-clock
$T(\text{CLKtoDATA}_{\text{max}})$	21.3 ns	$T(\text{RTT}_{\text{max}}) + 5 \times T(\text{coreclock})$
$T(\text{CLKtoDATA}_{\text{min}})$	11.0 ns	$T(\text{RTT}_{\text{min}}) + 4 \times T(\text{coreclock})$

Banks with tighter timings can be found in Appendix A. One can also use a source-synchronous clock to achieve far superior timings.

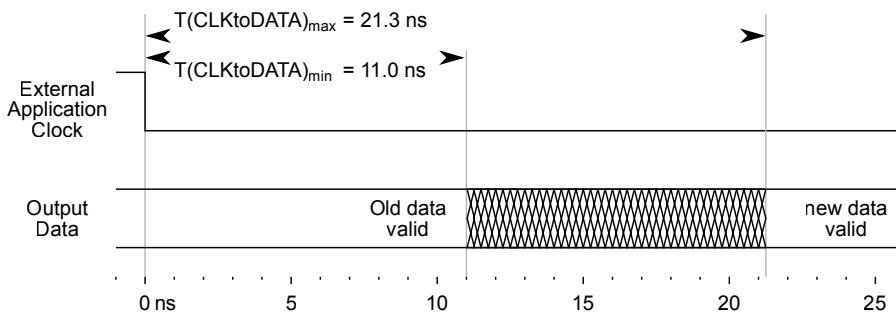


Figure 17:
 $T(\text{CLKtoDATA})$ on externally generated clocks.

If open-drain outputs are used, then an extra 2 ns margin on the max time should be accounted for, and 3 ns should be subtracted from the min time; in addition to any time required by the external resistor to pull the signal high

5 A worked example: master I2S

Below we show the timings of a worked example: I2S. I2S uses a bit-clock (signal I2S_BCLK) to clock DAC-data out to the CODEC (signal I2S_DAC), to clock ADC-data in from the CODEC (signal I2S_ADC), and to clock out a Left-Right clock (I2S_LRCLK) that indicates whether the data transferred is a left-sample or a right sample. There are a master and a slave, and in this example the xCORE will be the master that produces I2S_BCLK and I2S_LRCLK (and I2S_DAC), and the external device will be the slave. The bit-clock is produced by the xCORE from a master-clock (signal I2S_MCLK). I2S transmits 2 x 32 bits of data for each audio-frame; a 32-bit left value and a 32-bit right value.

We assume:

- ▶ On the xCORE there are five 1-bit ports connected to the I2S_MCLK, I2S_BCLK, I2S_LRCLK, I2S_DAC and I2S_ADC signals.
- ▶ The LR-clock will operate at 192 KHz.
- ▶ The bit-clock will operate at 12.288 MHz (LR-clock x 64)
- ▶ The master-clock will operate at 24.576 MHz (LR-clock x 128)

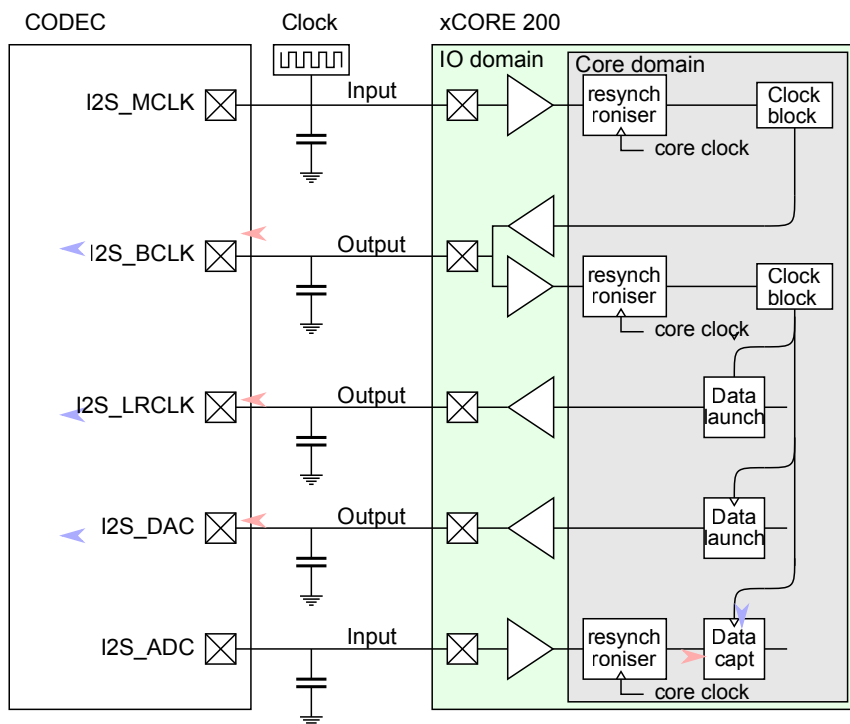


Figure 18:
I2S clock
block
organisation

We assume that the master-clock is sampled on a one bit port, this master clock is divided down and drives the bit-clock, which in turn drives a clock-block that is used to clock out the LRCLK and DAC data, and used to clock in the ADC data. This is shown in Figure 18. From top to bottom note that:

- ▶ I2S_BCLK is not source-synchronous to I2S_MCLK (there is a long blue path from the Clock generator into the xCORE and back out through I2S_BCLK, and a short red path outside the xCORE).
- ▶ I2S_LRCLK is not source-synchronous with I2S_BCLK (the blue arrow is a long path into and back out of the xCORE, the red arrow is a short path near the device).
- ▶ I2S_DAC is not source-synchronous with I2S_BCLK (the blue arrow is a long path into and back out of the xCORE, the red arrow is a short path near the device).
- ▶ I2S_ADC is source-synchronous to I2S_BCLK (the blue path is the clock path into the xCORE, the red path is the data signal into the xCORE). Note that as the clock does not emanate from the DAC, there is some delay caused by the CODEC in the red path. Note that the reason it is source synchronous is because the BCLK is in this setup resampled at the pin.

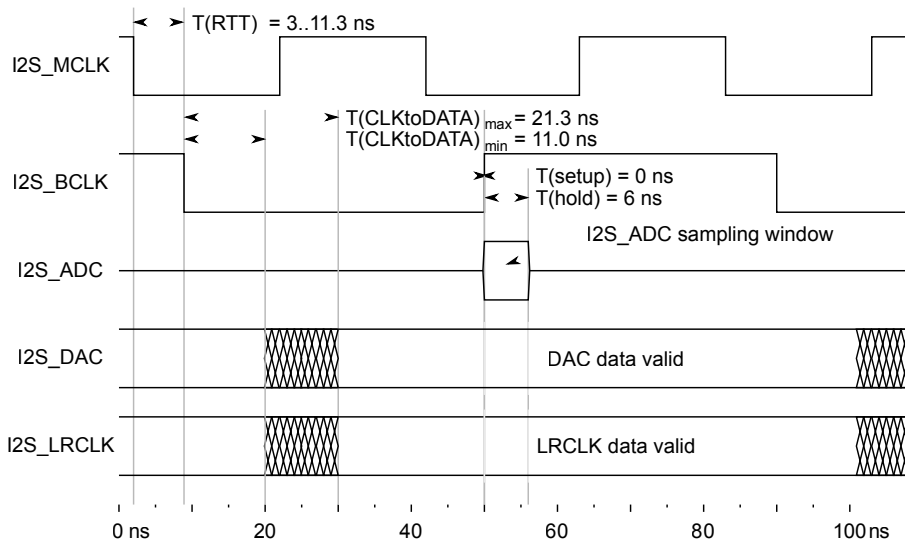


Figure 19:
I2S signal
timings

Hence, I2S_BCLK is a round-trip-time behind I2S_MCLK. In addition, I2S_LRCLK and I2S_DAC will be delayed relative to I2S_BCLK (Section 4.2), and I2S_ADC will incur a small skew (Section 3.1). The time between the MCLK edge and the BCLK edge is variable amongst voltage, process, and temperature, but constant in a given environment.

The timings are shown in Figure 19. Since the external CODEC will typically sample the I2S_DAC on the rising edge this will have a good margin (close to 20 ns); the

CODEC will typically present the ADC on the falling edge which has a larger margin (around 40 ns).

Note also that as I2S_MCLK is resynchronised to the xCORE, I2S_BCLK will be running at 12.288 MHz, but *synchronous to the core-clock*. That is, I2S_BCLK will have a period of 81.380... ns, but this will be a sequence of clocks that are 80, 82, 82, 80, 82, 82, 80, ... ns apart to, *on average*, create a clock of exactly 12.288 MHz. This introduces jitter on the BCLK; this normally does not matter, but if the CODEC is sensitive to jitter on I2S_BCLK then it can be resynchronised to the master clock using a D-type flip flop clocked by I2S_MCLK. Note that doing so will alter the setup, hold, and clock-to-data timings, because the clock edges of BCLK are shifted.

6 Appendix - Timing components

The tables below show the expected worst case minimum and maximum Round-Trip-Times for specific groups of pins and ports. In the case of 4-tile devices, tile 2 behaves the same as tile 0 and tile 3 the same as tile 1. All times are given in ns (where 1 ns = 0.000,000,001 s).

The minimum round trip time is defined as the fastest possible input path plus the fastest possible output path. The maximum round trip time is defined as the slowest possible input path plus the slowest possible output path. Paths can be slow or fast depending on the following parameters:

- ▶ Variation in the process will cause some devices to be slower and some devices to be faster. The values shown in the tables below assume the fastest possible silicon, and the slowest possible silicon. “Typical silicon” will have RTTmin and max times of 4 and 8 ns.
- ▶ Variation in temperature will cause a device to run slower or faster. For the IO measurements, hot devices will run slowest, and cold devices will run fastest. Designs that run in a controlled temperature (for example, always hot, or always cold) will have RTTmin and RTTmax times that are closer together.
- ▶ Variation in core voltage and (to a much lesser degree) IO voltage will cause a device to run slower or faster. The minimum time is for 0.9V core and 3.0V IO, and the maximum time is for 1.1V core and 3.6V IO. Designs with tightly controlled voltages (eg, 1.0V +/- 2%) will have RTTmin and RTTmax times that are closer together.
- ▶ By design, some IO pins are systematically faster than others. For example, using pins X0D12..X0D23 to input data using an internal clock will result in an eye that is 2.7 ns smaller than picking random pins from X1D00..X1D71.
- ▶ Driving rising edges and falling edges have different timings.

Extra time should be allowed for external loads, and this will add to both minimum and maximum Round-Trip-Times. Skew may also increase with an increased external load. Note that six IO pins (X1D26/27/40/41/42/43) have notionally 8 mA drivers and will respond faster to higher loads than the other IO pins (that are notionally 4 mA drivers).

The following tables assume a 2 pF external load.

Overall timings:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
Any IO pin		3.0	11.3	2.0	2.7

Timings per tile:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D71	All	3.0	10.8	2.0	2.1
X1D00..X1D71	All	3.0	11.3	1.8	2.7

Timings of pairs of banks, each pair comprising eight 1-bit ports and a 16-bit port:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D23	1A..H 16A	3.0	10.3	1.8	1.6
X0D24..X0D43	1I..P 16B	3.2	10.8	1.8	1.8
X1D00..X1D23	1A..H 16A	3.1	11.3	1.8	2.4
X1D24..X1D43	1I..P 16B	3.1	9.8	1.3	1.5

Timings of individual banks, each bank comprising four 1-bit ports and an 8-bit port (two 4-bit ports):

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D11	1A..D 4A..B 8A	3.5	10.3	1.0	1.3
X0D12..X0D23	1E..H 4C..D 8B	3.0	8.7	0.6	1.2
X0D24..X0D35	1I..L 4E..F 8C	3.2	9.7	0.8	1.7
X0D36..X0D43	1M..P 8D	3.7	10.8	1.1	1.3
X1D00..X1D11	1A..D 4A..B 8A	3.8	11.3	1.2	1.3
X1D12..X1D23	1E..H 4C..D 8B	3.1	9.1	0.8	1.3
X1D24..X1D35	1I..L 4E..F 8C	3.1	9.8	1.3	1.4
X1D36..X1D43	1M..P 8D	3.1	9.6	1.1	1.5

The following tables assume a 30 pF load

Overall timings:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
Any IO pin	All	3.8	13.8	2.0	3.5

Timings per tile:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D71	All	4.4	13.3	2.0	1.9
X1D00..X1D71	All	3.8	13.8	1.8	3.5

Timings of pairs of banks, each pair comprising eight 1-bit ports and a 16-bit port:

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D23	1A..H 16A	4.5	12.9	1.8	1.5
X0D24..X0D43	1I..P 16B	4.7	13.3	1.8	1.6
X1D00..X1D23	1A..H 16A	4.6	13.8	1.8	2.2
X1D24..X1D43	1I..P 16B	3.8	12.4	1.3	2.6

Timings of individual banks, each bank comprising four 1-bit ports and an 8-bit port (two 4-bit ports):

Pins	Ports	RTTmin	RTTmax	Tiskew	Toskew
X0D00..X0D11	1A..D 4A..B 8A	5.0	12.9	1.0	1.2
X0D12..X0D23	1E..H 4C..D 8B	4.5	11.1	0.7	0.9
X0D24..X0D35	1I..L 4E..F 8C	4.7	12.3	0.8	1.6
X0D36..X0D43	1M..P 8D	5.1	13.3	1.1	1.3
X1D00..X1D11	1A..D 4A..B 8A	5.3	13.8	1.2	1.2
X1D12..X1D23	1E..H 4C..D 8B	4.6	11.7	0.8	1.1
X1D24..X1D35	1I..L 4E..F 8C	3.8	12.4	1.3	2.5
X1D36..X1D43	1M..P 8D	3.8	12.1	1.1	2.6



Copyright © 2017, All Rights Reserved.

Xmos Ltd. is the owner or licensee of this design, code, or Information (collectively, the "Information") and is providing it to you "AS IS" with no warranty of any kind, express or implied and shall have no liability in relation to its use. Xmos Ltd. makes no representation that the Information, or any particular implementation thereof, is or will be free from any claims of infringement and again, shall have no liability in relation to any such claims.