AN 359: POS-PHY Level 4 MegaCore Function Parameter Selection Calculator

© Sep 2008 AN-359-2.0

Introduction

This application note describes the Altera® POS-PHY Level 4 MegaCore® function parameter selection calculator, which is a Microsoft Excel-based tool. The calculator helps you determine the ideal parameter values to set in the MegaCore function, for your system to meet the specific performance requirements of the SPI-4.2 protocol.

The packet over SONET/SDH physical layer (POS-PHY) level 4 interface, first developed by the SATURN® Development Group, was later adopted by the Optical Internetworking Forum (OIF) as the system packet interface level 4-phase 2 (SPI-4.2). Thus, POS-PHY level 4 and SPI-4.2 are synonymous.

The POS-PHY Level 4 MegaCore function uses the SPI-4.2 interface for high-speed cell and packet transfers between physical (PHY) and link layer devices. The SPI-4.2 interface has the following features:

- Supports a data width of 16 bits
- PHY-link, link-link, link-PHY, or PHY-PHY connection in multi-gigabit applications, including:
  - Asynchronous transfer mode (ATM) and packet over SONET/SDH (STS-192/STM-64)
  - 10-Gbps Ethernet
  - Multi-channel Gigabit Ethernet

In compliance with the SPI-4.2 interface specification, the POS-PHY Level 4 MegaCore function allows you to implement transmitter and receiver functions.

This calculator may only be used with interfaces involving Altera devices and remains the property of Altera. It is being provided on an "as-is" basis and as an accommodation; therefore, all warranties, representations, or guarantees of any kind (whether express, implied, or statutory) including, without limitation, warranties of merchantability, non-infringement, or fitness for a particular purpose, are specifically disclaimed.

About the Calculator

The calculator file consists of three worksheets. The **Cover** worksheet provides warranty and copyright information, and usage recommendations. The **Parameters** worksheet contains the calculator.

The calculator fields of the **Parameters** worksheet are color-coded to indicate their use. Green-shaded fields with black text are input boxes for you to fill in based on your knowledge of the system, or as you make parameter decisions. Blue-shaded fields contain information returned by the calculator. Green-shaded fields with gray text are unused. The calculator changes the coloring of the fields according to the

information you provide. For example, if you select **No** for the **Transmitter is an Altera SPI-4.2 MegaCore function** field, the Lite Transmitter parameter field is grayed to indicate that it is not used. Nothing prevents you from modifying the blue fields (the results returned by the calculator). However, modifying the blue fields may lead to incorrect calculation results.

The calculator's default state is a SPI-4.2 link carrying 10 gigabits per second (Gbps) of Ethernet traffic from a transmitter MegaCore function to a receiver MegaCore function.

The calculator provides information for different options, such as minimum required frequencies. However, this information does not infer that the MegaCore function can meet these frequencies. For example, if you select an interface capable of carrying 16 Gbps of traffic, and select a 32-bit transmitter, the calculator indicates that the transmitter MegaCore function needs to run at an internal frequency of over 500 MHz. This calculation result does not infer that the transmitter MegaCore function can run at over 500 MHz.

For guidance on performance, refer to the *POS-PHY Level 4 MegaCore Function User Guide*. For detailed performance information, Altera recommends that you use the Quartus® II software to compile your design for the targeted FPGA.

Using the Calculator

This section describes the calculator, briefly describing each field or result, and explains how to use it. This section assumes that you are familiar with the POS-PHY Level 4 MegaCore function parameter options, and that you have the calculator open at the **Parameters** worksheet.

The calculator begins when you enter the expected traffic distribution and required system performance, and then guides you through the selection of various performance-affecting parameters, such as clock rates, buffer thresholds, and burst sizes.

Before you Begin

This application note and the accompanying calculator make the following assumptions:

- You are familiar with the SPI-4.2 protocol
- You have read the *POS-PHY Level 4 MegaCore Function User Guide*
- You have adequate knowledge about the device on the other side of the SPI-4.2 interface

Traffic Characteristics

Under **Traffic Characteristics**, define the type of traffic to be carried across the SPI-4.2 interface. Table 1 lists and describes the applicable fields.

Table 1. Traffic Characteristics Fields

Field	Description
Min. Packet Size	Specifies the minimum packet size for each of the three weighted distributions: A, B, and C.
Max. Packet Size	Specifies the maximum packet size for each distribution.
Average Packet Size	Specifies the average packet size for each distribution.
Standard Deviation	Specifies the standard deviation for each distribution.
Weight	Specifies how the spreadsheet weights each of the distributions. Only the ratio of the weights is relevant (a weighting of 3, 2, and 1 is the same as a weighting of 6, 4, and 2).
Weight Packets by Size	Select **Yes** to weight each packet by its own size, because larger packets consume more bandwidth per packet.
Required Traffic Bandwidth	Specifies the bandwidth that the interface has to carry.
Per-Packet non-SPI-4.2 Overhead	Any packet overhead that the source or destination carries, but that the SPI-4.2 interface does not carry (for example, an Ethernet preamble).
Required Net Traffic Bandwidth	The calculator provides the net traffic bandwidth that the interface has to carry after accounting for the per-packet non-SPI-4.2 overhead at the specified traffic mix, with the required traffic bandwidth.

Distribution

You begin by defining a mix of traffic sizes, specified as the weighted sum of three normal (Gaussian) distributions. For each distribution, specify the minimum, maximum, and average packet size, as well as the standard deviation. You finish by specifying the relative weighting of each distribution. The graph labelled **Normalized Packet Size Distribution** in the calculator shows the resulting distribution to be used for subsequent calculations.

As an example, Figure 1 shows the packet size distribution for a single normal distribution with packet sizes ranging from 32 to 2,048 bytes—with an average of 128 bytes, and a standard distribution of 100.

Figure 1. Simple Normalized Packet Size Distribution

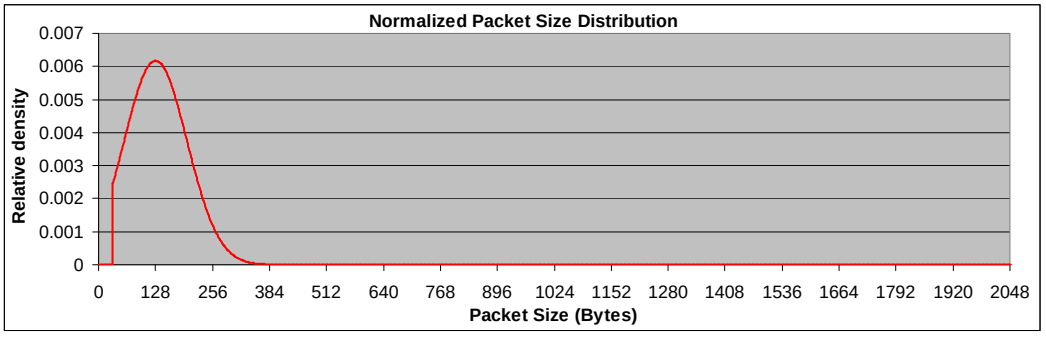

This application note uses the default values shown in Figure 2, and the distribution shown in Figure 3, as examples. As you can see in Figure 3, the distribution reveals that the bulk of the traffic is made up of 128-byte packets, and that larger packet sizes are less common. The spike of 64-byte packets represents a frequently occurring control packet in the system.

Figure 2. Example Settings Values

	A	B	C
Min Packet Size	32	32	32
Max Packet Size	2048	2048	1024
Average Packet Size	128	256	64
Standard Deviation	100	1000	1
Weight	2	2	1

Figure 3. Example Normalized Packet Size Distribution

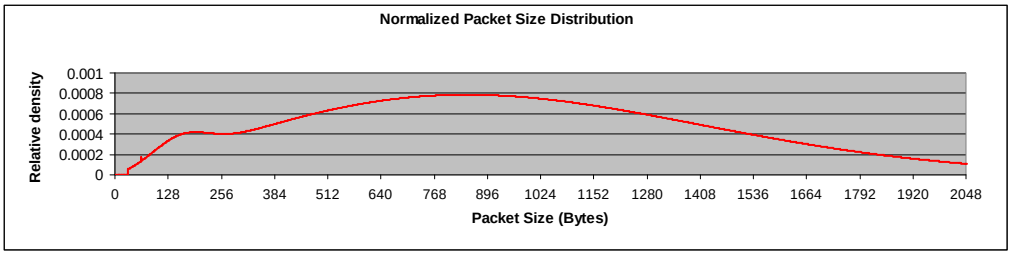

Weight Packets by Size

Now that you have entered the expected packet size distribution, you have to decide whether or not packets should be weighted by size. The default value is **Yes**. See Table 1 on page 3 for the applicable field. When you change this field to **No**, you can see the immediate effect in the **Normalized Packet Size Distribution** graph.

Consider the case where your distribution has 50% of the packets as 64-byte packets, and 50% of the packets as 640-byte packets. By selecting **Yes** for the **Weight Packets By Size** field, you declare to the calculator that 50% of the packets are 640 bytes, and 50% are 64 bytes. So about 90% of the bandwidth is used by 640-byte packets. The calculator does this conversion for you, and the graph shows the amount of bandwidth used by each packet size. By selecting **No** for this field, you declare to the calculator that 50% of the packets are used by each packet size, which in turn means that only about one tenth of the packets are 640-byte packets.

Required Traffic Bandwidth

The **Required Traffic Bandwidth** field specifies the bandwidth that the interface has to carry. The default value for this field assumes that the interface is carrying traffic from a 10-Gbps Ethernet link, thus the field is set to 10,000 megabits per second (Mbps).

Per-Packet Non-SPI-4.2 Overhead

The per-packet non-SPI-4.2 overhead field allows you to account for some packet overhead that the source or destination carries, but that the SPI-4.2 interface does not carry.

As an example for this application note, the interface is said to be carrying 10 Gbps of Ethernet traffic that has a 12-byte inter-packet gap, and an 8-byte preamble. For every packet, there are 20 bytes (12+8) in the Gigabit Ethernet domain that are not carried across the SPI-4.2 interface. Some of the 10 Gbps Ethernet link is wasted by overhead that the SPI-4.2 interface does not need to carry, so the SPI-4.2 interface can afford to operate at a lower frequency.

The default value of 20 set in the **Per-packet non-SPI-4.2 overhead** field indicates that the resulting **Required Net Traffic Bandwidth** (the actual traffic carried by the SPI-4.2 interface) is 9,668 Mbps, not 10,000 Mbps. Table 1 on page 3 shows the applicable fields.

When a non-zero amount of per-packet non-SPI-4.2 overhead is carried, as in the previous 10-Gbps Ethernet example, small packets waste a higher fraction of the link bandwidth on overhead than larger packets do. A large number of small packets means that the source or destination link is less efficient and, in turn, the SPI-4.2 interface sees less real traffic bandwidth.

When the **Weight Packets By Size** option is set to **No**, there is a skew towards these inefficient small packets that causes the calculator to incorrectly assume that a significant portion of the source link is being wasted, and that the SPI-4.2 interface sees only the non-wasted portion.

If for the 10-Gbps Ethernet example, the **Weight Packets by Size** field is set to **No**, the skew towards small, inefficient packets causes the calculator to incorrectly conclude that the net required bandwidth is 9.22 Gbps, instead of 9.67 Gbps.

Transmitter and Receiver Characteristics

Under **Transmitter and Receiver Characteristics**, provide information about the SPI-4.2 transmitter and receiver. Table 2 lists and describes the applicable fields.

Table 2. Transmitter and Receiver Characteristics Fields (Part 1 of 2)

Field	Description
Transmitter / Receiver is Altera SPI-4.2 MegaCore function	Specifies whether the receiver, transmitter, or both sides of the interface are occupied by POS-PHY Level 4 MegaCore functions. When MegaCore functions are used, the calculator provides most of the required data. The POS-PHY Level 4 MegaCore function operates either as a receiver where data flows from the SPI-4.2 interface to the Atlantic interface, or as a transmitter where data flows from the Atlantic interface to the SPI-4.2 interface.
Data Path Width	The data path width for each MegaCore function. Used to determine the required clock frequencies, and core latencies.
Atlantic Data Path Width	The Atlantic interface data path width. Determines the required clock frequencies. The Atlantic interface depends on the internal data path width. Refer to the *POS-PHY Level 4 MegaCore Function User Guide*, for information on the Atlantic data widths supported for each internal data path width.
Buffer Mode	The buffer mode for each MegaCore function. For Shared buffer with embedded addressing, all ports share a single Atlantic buffer with an 8-bit address field that supports up to 256 ports. For individual buffers, the POS-PHY Level 4 MegaCore function provides a FIFO buffer for each port. The individual buffers supports up to 16 ports.

Table 2. Transmitter and Receiver Characteristics Fields (Part 2 of 2)

Field	Description
Lite Transmitter	For transmitter MegaCore functions, selects the **Lite Transmitter** parameter, which lowers the effective bandwidth rate on the SPI-4.2 data bus, but it greatly reduces the logic consumption. If you turn off **Lite Transmitter**, the transmitter packs the packets more tightly together, and thus increases the effective bandwidth rate on the SPI-4.2 data bus, but increases the logic consumption.
Burst Unit Size	The unit size in bytes for burst transfers and controls the smallest burst transmitted, except when it sends an EOP. The POS-PHY level 4 transmitter only inserts a control word when required by the specification. This parameter may have different names for different ASSP devices, and some devices may simply use **MaxBurst2** as the minimum burst size. In general, a larger burst size corresponds to less SPI4.2 overhead, but longer turn-off times, resulting in more latency and lower shut-off thresholds. The valid range for this parameter is from 16 to 1,024 bytes, in 16-byte granularity. When the data path width is equal to 128 bits and the **Lite Transmitter** feature is turned off, the unit size is 32 bytes, up to 1,024 bytes in 32-byte granularity.
Theoretical Maximum Burst Size Efficiency	For reference, the calculator provides the efficiency of the SPI-4.2 link based solely on the burst size, where an infinite burst length corresponds to 100% efficiency.
Training Sequence Interval	The interval at which the training sequence occurs—16 to 65,535 bytes. It is measured from the end of the last training sequence to the beginning of the next training sequence. The training sequence is scheduled to be inserted after the `MaxT` counter expires, but is not actually inserted until the burst that is sent is complete. Therefore, the time between training pattern insertions is no less than the value of the `MaxT` parameter, and no more than the value of the `MaxT` parameter plus the burst unit size. If `MaxT` = 0, periodic training patterns are disabled.
Training Pattern Repetitions	The number of training pattern sequence repetitions. This value only applies to interval training patterns. Setting to zero results in an `ALPHA` of 256. The training sequence includes one `IDLE` word, plus `ALPHA`(?) × 20 training words.
Theoretical Maximum Training Efficiency	For reference, the calculator provides the efficiency of the SPI-4.2 link based solely on the training pattern insertion, where no training pattern insertion (`ALPHA` = 0) corresponds to 100% efficiency.

POS-PHY Level 4 MegaCore Functions

When you use POS-PHY Level 4 MegaCore functions, the calculator provides most of the required data via the fields shown in Table 2 on page 5. When you use ASSP devices, you must enter the ASSP device characteristics, based on their respective datasheets.

For most full-duplex interfaces between an Altera FPGA and an ASSP, the calculator is needed twice. It is needed once for the POS-PHY Level 4 transmitter MegaCore function to ASSP receiver interface, and once for the ASSP transmitter to the POS-PHY Level 4 receiver MegaCore function interface.

For each POS-PHY Level 4 MegaCore function, you specify the data path width, the buffer mode to be used, and for the transmitter, whether a lite or full transmitter is used. For the transmitter, you also specify the burst size.

This application note uses the following default values as examples:

- The SPI-4.2 interface carries 10 Gbps of Ethernet traffic.
- A 64-bit MegaCore function is on each side of the interface.
- The transmitter MegaCore function is a lite transmitter.

Burst Size

The burst size corresponds to the smallest burst that the transmitter can send, except when it sends an EOP. The transmitter may insert a control word at the end of each burst. Figure 4 shows the worst case efficiency of a SPI-4.2 interface in relation to burst size. Figure 4 shows not much is gained by setting the burst size to be greater than 128 bytes.

Figure 4. Efficiency Due to Burst Size

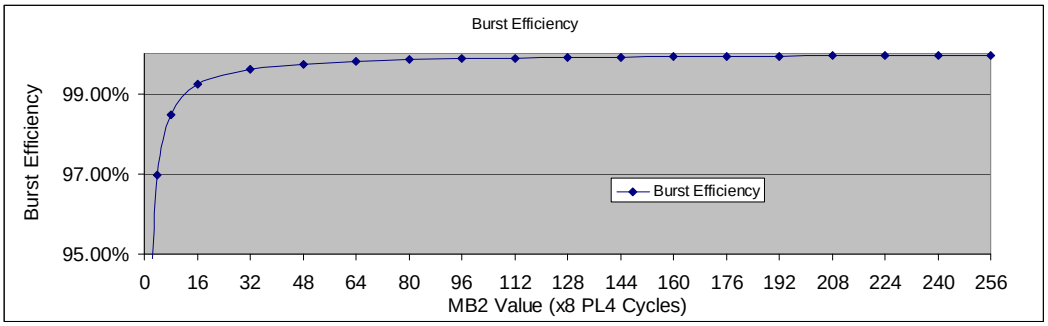

Theoretical Maximum Burst Size Efficiency and Theoretical Maximum Training Efficiency

As an example for this application note, the default values are used. The burst size is set to 128 bytes—corresponding to a 98.46% theoretical burst size efficiency. The training sequence interval (**MAX_T**) is set to 2,048 bytes, and the training pattern repetitions (ALPHA) is set to 8—corresponding to a theoretical training efficiency of 99.42%.

Clock and Data Rates

Under **Clock and Data Rates**, the calculator provides useful information about the application's requirements. Table 3 lists and describes the applicable fields.

Table 3. Clock & Data Rates Fields (Part 1 of 3)

Field	Description
Required Minimum LVDS Rate	This field provides the absolute minimum LVDS data rate at which you can run an ideal SPI-4.2 interface, and still pass the mix of traffic you specified under Traffic Characteristics.This value takes into consideration the traffic mix from the traffic section, and the burst size and training pattern insertion parameter values.
Recommended LVDS Rate	This value is calculated from the required minimum LVDS rate.

Table 3. Clock & Data Rates Fields (Part 2 of 3)

Field	Description
LVDS Data Rate	For a transmitter, the LVDS data rate specifies the data rate out of the FPGA, on each LVDS pair. For a receiver, the LVDS data rate specifies the data rate into the FPGA, on each LVDS pair. This value defaults to the recommended value, however you are not obligated to keep this value. If you change the value in this field, that new value is used for all subsequent calculations. The calculator issues a warning when the LVDS data rate is set to be less than the required minimum. For maximum supported LVDS data rates for each device family, refer to the *POS-PHY Level 4 MegaCore Function User Guide*.
SPI-4.2 Clock Frequency	The SPI-4.2 clock rate is provided for reference purposes only. It is always ½ of the LVDS data rate. Data and control words arrive on the data bus, and are sampled on both edges of SPI-4.2 clock.
Status Clock Frequency	When a receiver MegaCore function is used, the status clock maximum frequency is always ¼ of the SPI-4.2 clock frequency, and cannot be changed. If the receiver is not a POS-PHY Level 4 MegaCore function, the status clock changes to an input field, and you can enter the status clock frequency to be used by the system. The status clock frequency must be equal to, or lesser than, ¼ of the SPI-4.2 clock. The transmitter MegaCore function supports status clock frequencies that are faster, but faster frequencies may occasionally lead to missed status updates.
Transmitter Minimum Atlantic Frequency	The required minimum Atlantic frequency based on the traffic mix, the required throughput, and the data path width. Applies only to transmitter MegaCore functions. Your user logic uses this frequency to provide transmit data, thus you should ensure that the MegaCore function and your logic meet timing at the frequency listed in this field.
Transmitter Reference Frequency (`trefclk`)	The reference clock frequency for the transmitter MegaCore function based on the LVDS data rate and the MegaCore function data path width. The MegaCore function uses this frequency internally, thus you should ensure that the MegaCore function meets timing at the frequency listed in this field. For 128-bits variation, 1/8 of LVDS data rate. For 64-bits variation, 1/4 of LVDS data rate. For 32-bits variation, 1/2 of LVDS data rate.
Transmitter Minimum System Clock Frequency (`txsys_clk`)	The minimum frequency for the transmitter MegaCore function system clock, `txsys_clk`. The `txsys_clk` frequency must be faster than, or equal to, the `tsclk` frequency.
Receiver Minimum Atlantic Frequency	The required minimum Atlantic frequency based on the traffic mix, required throughput, and the data path width. Applies only to receiver MegaCore functions. Your logic uses this frequency to receive data, thus you should ensure that the MegaCore function and your logic meet timing at the frequency listed in this field.

Table 3. Clock & Data Rates Fields (Part 3 of 3)

Field	Description
Receiver Reference Frequency (`rdint_clk`)	The reference clock frequency for the receiver MegaCore function based on the LVDS data rate and the MegaCore function data path width. The MegaCore function uses this frequency internally, thus you should ensure that the MegaCore function meets timing at the frequency listed in this field. For 128- and 64-bit variations, 1/8 or 1/4 of the `rdat` data rate. For 32-bit (quarter-rate) variations, 1/2 the `rdat` data rate.
Receiver Maximum System Clock Frequency (`rxsys_clk`)	The reference clock frequency for the POS-PHY level 4 transmitter, based on the LVDS data rate and the data path width. The transmitter uses this clock internally. Ensure the transmitter meets timing at the frequencies listed here. For the worst case frequency requirement of `rxsys_clk` and for `rxsys_clk` frequency restrictions, refer to the *POS-PHY Level 4 MegaCore Function User Guide*.

For guidance on the frequencies that the MegaCore function can be expected to achieve, refer to the *POS-PHY Level 4 MegaCore Function User Guide*.

The information in this section is calculated from the parameters in the previous sections. However, it should also be used to ensure that the selections in the previous sections are reasonable.

For example, this application note uses the following default values:

- The SPI-4.2 interface carries 10 Gbps of Ethernet traffic.
- The transmitter and receiver MegaCore functions, on either side of the interface, are 64 bits wide.

With a 64-bit internal bus width, the required Atlantic clock frequency is 152 MHz, which is quite feasible. However, consider the following modifications:

- If the Atlantic interface width for either of the MegaCore functions is changed to 32 bits, the MegaCore function requires considerably fewer logic resources in the target device; however, the required Atlantic clock rate is now 303 MHz, which makes it difficult, or impossible, to meet timing.
- If the Atlantic interface width for either of the MegaCore functions is changed to 128 bits, the required Atlantic clock rate is now only 75.97 MHz, but the MegaCore function requires many more logic resources in the target device.

In light of these results, a 64-bit data path is the best solution for the given traffic and bandwidth requirements.

Latency Related Transmitter Settings

Under **Latency Related Transmitter Settings**, you provide the information the calculator needs to calculate the latencies inherent in the system, and to recommend values for the almost empty (AE) and almost full (AF) thresholds.Table 4 lists and describes the applicable fields.

Table 4. Latency-Related Transmitter Fields

Field	Description
Number of Ports	The number of ports supported by the system. The calculator does not use the number of ports directly, it uses it only to provide a check for the calendar length (CALLEN) parameter.
Calendar Length (CALLEN) × Calendar Multiplier (CALM)	The **Calendar Length** determines the length of the calendar used by the status message. For POS-PHY level 4 MegaCore functions, the CALLEN parameter can only be changed if the **Programmable Calendar Length Support** parameter is selected during the MegaCore function's generation. The calendar length cannot be greater than the number of ports. if the **Programmable Calendar Length Support** parameter is turned off, the calendar length is equal to the number of ports. The **Calendar Multiplier** determines the number of times the calendar sequence is repeated before the DIP-2 parity and framing is inserted. Choose a value from 1 to 256. CALM and CALLEN affect the length of the status message, which directly affects the time it takes for the transmitter to receive the receiver's status.
MaxBurst1 (MB1) MaxBurst2 (MB2)	The sizes of bursts sent by the transmitter when the receiver is starving or hungry, respectively. The transmit scheduler, included with the shared buffer with embedded address mode, does not use the **MaxBurst1**or **MaxBurst2** parameters; it simply reads data out of the buffer in 'Burst Size' blocks (defined in the Transmitter Characteristics section). The credit-based scheduler, included with individual buffers mode, sends data across the interface in units of **MaxBurst1** and **MaxBurst2**. For the transmitter MegaCore function, **MaxBurst2** must be less than or equal to **MaxBurst1**; **MaxBurst1** and **MaxBurst2** must be greater than or equal to burst unit size. Select a value from burst unit size to 2,032 bytes in 16-byte increments.
Status Interpretation Mode	If you select **Optimistic**, the status information is provided to the user logic as soon as it arrives on the status channel before the DIP2 error checking code is received. If you select **Pessimistic**, the latest status information is captured and is stored inside the status processor block until a DIP-2 status is received. If the DIP-2 is valid, the buffered status is passed on to the scheduler or user logic. The time spent waiting for the DIP2 verification causes the round-trip latency in receiving a valid status message to be calendar multiplier × calendar length tsclk cycles longer than the optimistic mode. If the status interpretation mode for an ASSP device is unknown, you should set this parameter to **Pessimistic** because it results in more conservative results for AE and AF.

As an example for this application note, the default values are used. The number of ports is equal to four, the calendar length is equal to the number of ports (thus 4), and the message is only repeated once for each status update. The **MaxBurst1** size is 128 bytes, and the **MaxBurst2** size is 64 bytes; both values are multiples of the burst size of 128 bytes. The status interpretation mode is set to pessimistic, meaning that the transmit scheduler is not updated with new credits until the message is confirmed to be DIP-2 error free.

Latency Calculations

The POS-PHY level 4 transmitter MegaCore functions involve two kinds of latency: data latency and status receive latency.

Data latency is defined as the latency from the Atlantic interface that is reading from the buffer to the SPI-4.2 LVDS transmit pins. It does not include the latency through the buffer. For external status, the numbers assume that aN_atxclk is faster than the tsclk thus ensuring that the clock-crossing FIFO buffer is empty.

Status receive latency is defined as the latency from the point at which the last cycle of a valid status message is received (the DIP-2 error code) to the point at which the user logic or the transmit scheduler can use the status information. It does not include the time spent waiting for a complete, error-free status message.

Figure 5 shows a generic picture of the L_{MAX} contributions (transmitter start to transmitter finish gives the transmitter L_{MAX}).

Figure 5. L_{MAX} Top-Level Overview

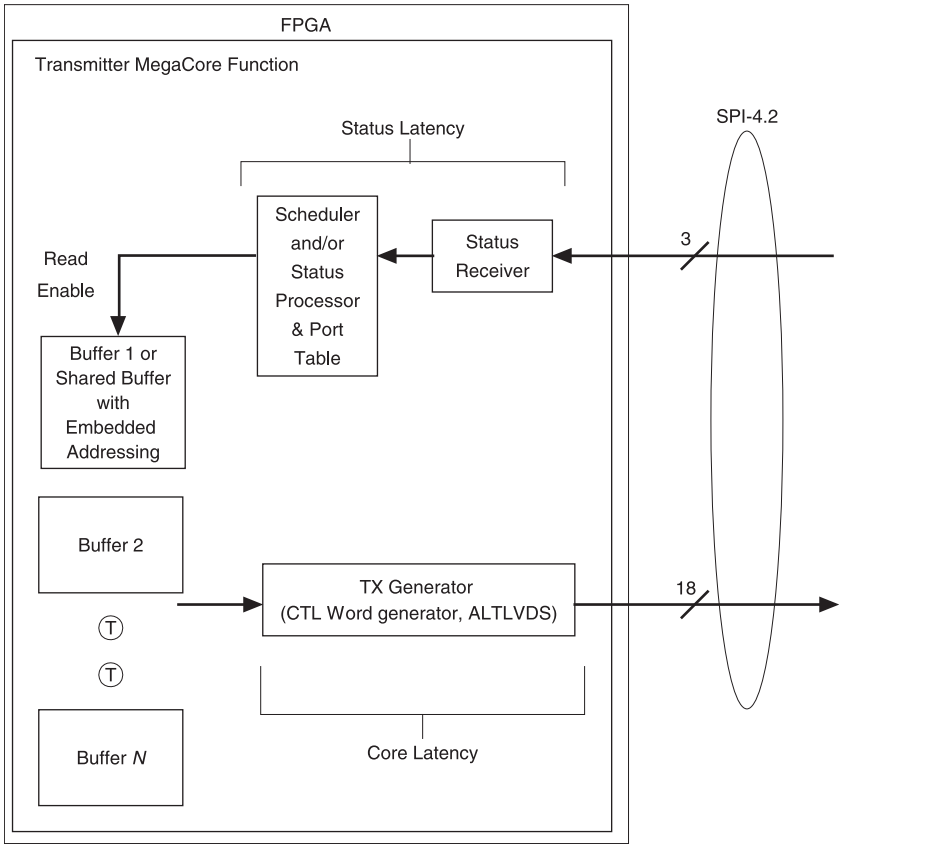

The POS-PHY Level 4 receiver MegaCore functions involve two kinds of latency: data latency and status transmit latency.

Data latency is defined as the latency from the SPI-4.2 LVDS receive pins to the internal Atlantic interface that is writing into the buffer(s). For the shared buffer with embedded addressing mode, it does not include the time the data spends in the buffer.

Status transmit latency is the number of clock cycles from when the status is provided from the user logic or the Atlantic buffer until it is transmitted to the adjacent device, assuming that the status channel is not disabled. It does not include the latency involved in waiting for the previous transmit message to complete, or in waiting for the status for other ports to be sent.

Figure 6 shows a picture of the L_{MAX} contributions (receiver start to receiver finish gives the receiver L_{MAX}) for a receiver using the individual buffers mode.

Figure 6. L_{MAX} Individual Buffers Mode Overview

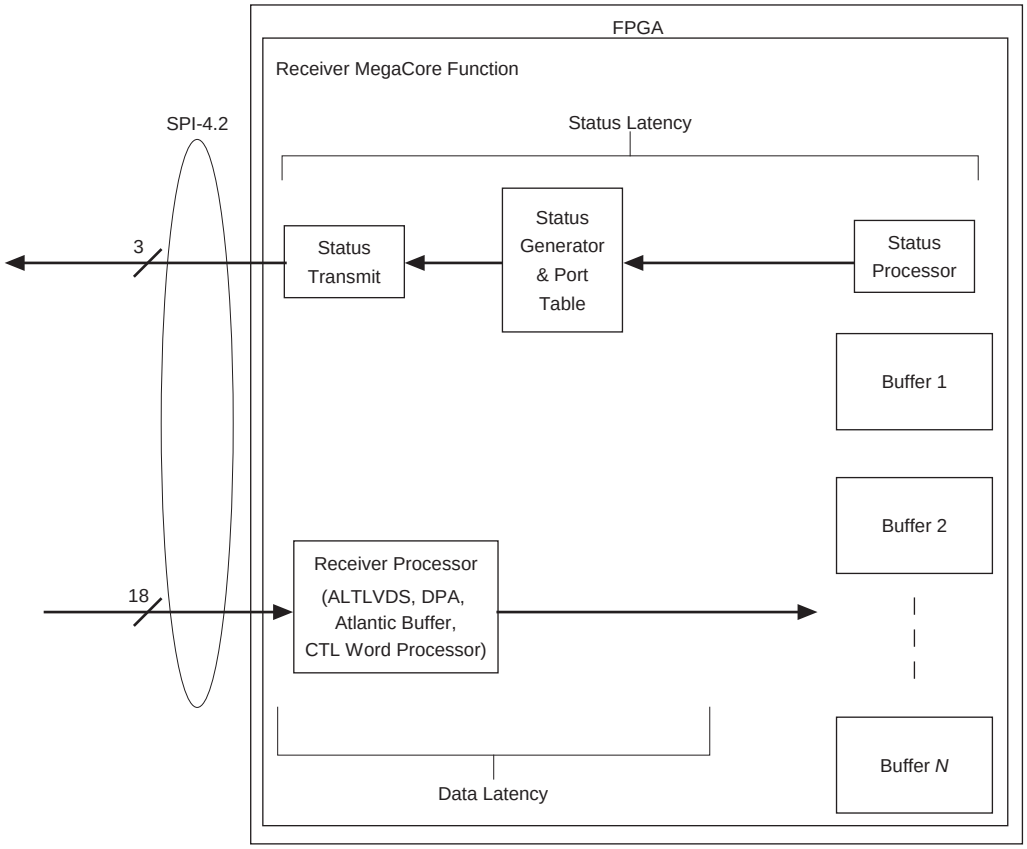

In this section, the calculator uses all of the information from previous sections to determine the round-trip latency (shut-off times) for the system, when the receiver is in the hungry or starving state. This round-trip latency is used to determine appropriate values for the AF and AE thresholds in the receiver. Figure 7 shows the various components that contribute to the round-trip latency.

Figure 7. SPI-4.2 System Latency

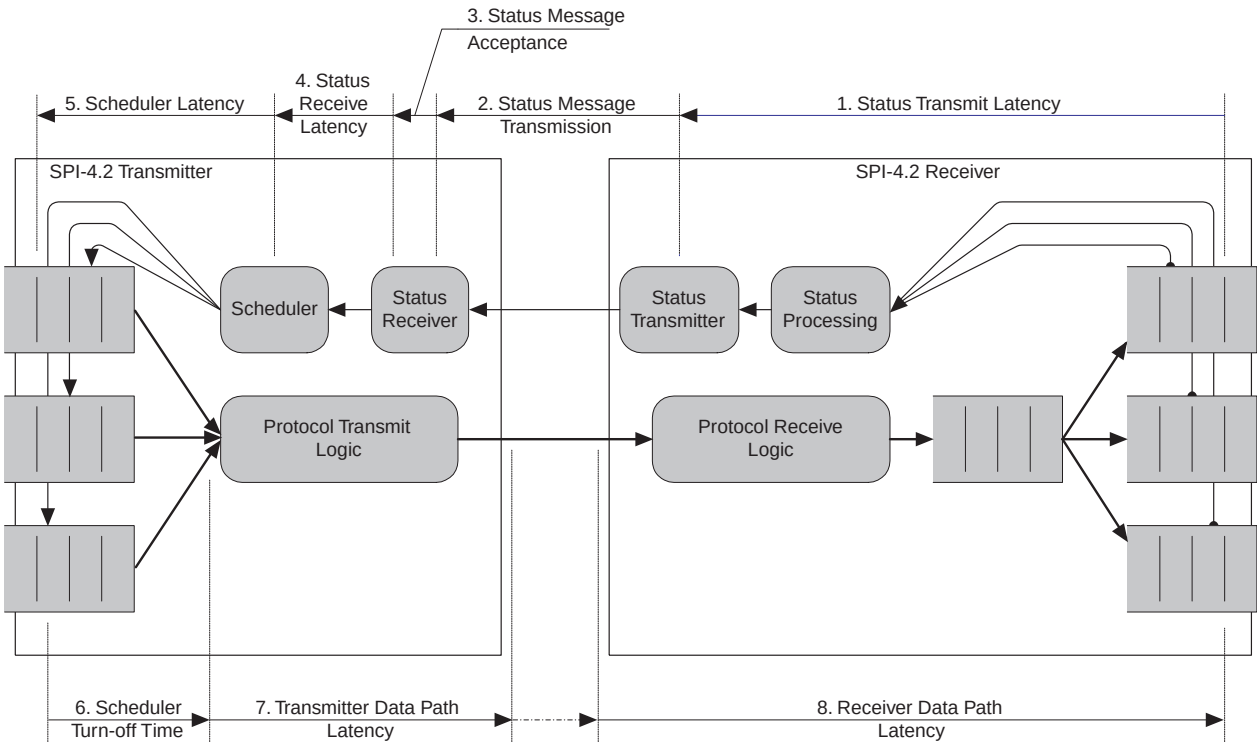

Table 5 lists and describes the applicable latency fields. Because this latency calculation focuses on the number of bytes that are sent and received on the SPI-4.2 interface, from the time the receive buffer crosses the AE or AF threshold to the time the traffic for that buffer stops, the latency fields are all expressed in units of bytes.

Table 5. Latency Fields (Part 1 of 2)

Field	Description
Status Transmit Latency	The number of bytes sent across the SPI-4.2 interface from when the fill level in a buffer crosses the AE or AF threshold to when the receiver is able to send the new status to the transmitter. For the receiver MegaCore function, the calculator inserts this value automatically, based on information from the user guide. The values do not include waiting for the appropriate time slot in the status channel for the status to be transmitted. The values do not include the time spent waiting for a complete error-free status message. The resultant value also reflects the status channel mode-either optimistic or pessimistic.
Status Message Transmission	The number of bytes sent across the SPI-4.2 interface during the time it takes for the status message to be sent. This is equivalent to (CALM × CALLEN) + 2 status clock cycles.
Status Message Acceptance	The number of bytes sent across the SPI-4.2 interface during the time it takes for the status message to be accepted by the transmitter. When the status interpretation mode is optimistic, this value is zero—the transmitter accepts the status as it arrives. When the status interpretation mode is pessimistic, this value is (CALM × CALLEN) + 1 status clock cycles, as the transmitter waits for the DIP-2 error code to arrive and be verified.
Status Receive Latency	The time it takes for the status receiver to process the incoming status message, and update the scheduler with the new information. For a transmitter MegaCore function, the calculator inserts this value automatically, based on the information from the user guide.

Table 5. Latency Fields (Part 2 of 2)

Field	Description
Scheduler Latency	This value accounts for the pipeline stages inside the scheduler itself. For the POS-PHY level 4 transmitter, the calculator inserts a value corresponding to two internal cycles at the appropriate data path width.
Scheduler Turn-Off Time	The time it takes the scheduler to stop sending data for this port, after it has an updated status. For a transmitter with the shared buffer with embedded addressing mode, the scheduler at most finishes the current burst, thus the **Burst Size** parameter is used. For a transmitter with the individual buffers mode, the scheduler finishes the current **MaxBurst1**- or **MaxBurst2**-sized burst.
Transmitter Data Path Latency	The time it takes data to travel from the output of the transmit buffer to the LVDS pins. For a transmitter, the calculator inserts this value automatically, based on the information of the user guide. The values in the user guide do not include the latency through the user-side buffers.
Receiver Data Path Latency	The time it takes data to travel from the receiver's LVDS input pins to the receive buffer, where it generates the buffer's fill level. For a receiver, the calculator inserts this value automatically, based on the information from the user guide. The values in the user guide do not include the latency through the user-side buffers. For 64- and 128-bit data path width variations, the values assume that the clock-crossing buffer is empty.
Round-trip Turn-Off Latency	The sum of all of the other latencies, corresponding to the total worst case number of bytes that could be pushed into the receive buffer after it crosses the AE or AF threshold.

If you are using a POS-PHY level 4 receiver with the shared buffer with embedded addressing mode, and the **Status Source** is set to **User Controlled** in IP Toolbench, you need to account for the time it takes your logic to determine the receiver's status to the time provided by the calculator. If you are using a buffering scheme beyond the shared buffer with embedded addressing (usual case), you need to include the latency to read out of the shared buffer and to fill your own buffer. Use the **Additional Receiver Latency** field to add this value to the calculation.

Likewise, if you are using a POS-PHY level 4 transmitter with the shared buffer with embedded addressing mode, and have your own scheduler to read data out of your own buffering structure, you can add the scheduling and data transmission delay to the calculation by using the **Additional Transmitter Latency** field.

For this example, the round-trip **MaxBurst1** and **MaxBurst2** turn-off latencies are 1,304 and 1,176 bytes, respectively.

Latency Related Receiver Settings

In this section, you select a buffer size and, based on this value, the calculator determines the maximum and recommended AE and AF values for the system. Table 6 lists and describes the applicable fields.

Table 6. Latency-Related Receiver Fields

Field	Description
Buffer Size	For the shared buffer with embedded addressing, the **Buffer Size** defines the size of the shared embedded address buffer. For the individual buffers, the **Buffer Size** defines the size of each buffer.
Maximum AF	The maximum value at which AF can be set to pass error-free data. This value is determined by subtracting the round-trip **MaxBurst2** turn-off latency from the buffer size.
Recommended AF	The recommended AF value for the system. This value is determined by subtracting 128 bytes from the maximum AF value, and rounding down to the nearest multiple of 64 bytes.
Maximum AE	The maximum value at which AE can be set to pass error-free data. This value is determined by subtracting the round-trip **MaxBurst1** turn-off latency from the buffer size.
Recommended AE	The recommended AE value for the system. This value is determined by subtracting 128 bytes from the maximum AE value, and rounding down to the nearest multiple of 64 bytes.

For a 10-Gbps Ethernet system, for a buffer size of 2,048 bytes per port, the recommended values for AF and AE are 704 and 576 bytes, respectively.

Conclusion

This application note and accompanying calculator can help you determine the required parameters for your SPI-4.2 interface.

Document Revision History

Table 7 shows the revision history for this application note.

Table 7. Document Revision History

Date and Document Version	Changes Made	Summary of Changes
September 2008 v2.0	Updated parameter names.	—
June 2004 v1.0	First release.	—

101 Innovation Drive
San Jose, CA 95134
www.altera.com
Technical Support
www.altera.com/support