# PCI-SIG ENGINEERING CHANGE NOTICE

| | |
|---|---|
| **TITLE:** | Power-up requirements for PCIe side bands (PERST#, etc.) |
| **DATE:** | September 18, 2014 |
| **AFFECTED DOCUMENT:** | PCI Express M.2 Specification, Revision 1.0 |
| **SPONSOR:** | Dave Landsman, SanDisk; Ramdas Kachare, LSI; David Harriman, Intel |

## Part I

### 1. Summary of the Functional Changes

Section 3.1.3.2.1 is redefined to provide a more realistic timing model for reset. Specifically:

a) Section is pulled out of 3.1.3, because Reset should not be placed under PCIe Side Bands, but separately, at same level.

b) PERST# timing variable ($T_{PVPGL}$) is changed to be “implementation specific”, with a recommended value, instead of the original hard minimum.

c) CLKREQ# timing variable ($T_{PVCRL}$) is removed as it is not needed specifically in relation to reset. CLKREQ# timing rules are referenced elsewhere in the spec.

d) The reset timing diagram and reset variable value tables are modified to reflect (a) and (b), and other changes, such as that “power on” refers to power being valid on the device , not on the system.

e) Language inserted referring to the PCI Express base spec, and BIOS/OS requirements, as they relate to $T_{PVPGL}$, in order to provide clearer guidance to implementers.

### 2. Benefits as a Result of the Changes

This change removes an unrealistic requirement for PERST# de-assertion timing. It also provides a clearer guideline to device and system implementers on what the timing requirement means, and how it relates to the full PCI Express Conventional Reset sequence, so that implementations can be optimized as relevant for the target platform.

### 3. Assessment of the Impact

The designs that meet existing timing requirements wouldn’t have any impact.

### 4. Analysis of the Hardware Implications

This change may help avoid special hardware circuitry that may be required to to meet the current timing parameters.

### 5. Analysis of the Software Implications

None

### 6. Analysis of the C&I Test Implications

None

## Part II

### Detailed Description of the change

#### 3.1.3 PCI Express Auxiliary Signals

The auxiliary signals are provided on the system connector to assist with certain system level functionality or implementation. These signals are not required by the PCI Express architecture, but may be required by specific implementations such as a PCI Express M.2 Device~~Card~~. The high-speed signal voltage levels are compatible with advanced silicon processes. The optional low speed signals are defined to use the +3.3 V supply, as it is the lowest common voltage available. Most ASIC processes have high voltage (thick gate oxide) I/O transistors compatible with +3.3 V. The use of the +3.3 V supply allows PCI Express signaling to be used with existing control bus structures, avoiding a buffered set of signals and bridges between the buses.

The PCI Express M.2 Device ~~Card add-in card~~ and system connectors support the auxiliary signals that are described in the following sections.

#### 3.1.3.1 Reference Clock

The REFCLKp/REFCLKn signals are used to assist the synchronization of the device's~~card's~~ PCI Express interface timing circuits. Availability of the reference clock ~~at the card interface~~ may be gated by the CLKREQ# signal as described in section 3.1.3.2. When the reference clock is not available, it will be in the *parked* state. A parked state is when the clock is not being driven by a clock driver and both REFCLKp and REFCLKn are pulled to ground by the ground termination resistors. Refer to the *PCI Express Card Electromechanical Specification* for more details on the functional and tolerance requirements for the reference clock signals.

#### 3.1.3.2 CLKREQ# Signal

The CLKREQ# signal is an open drain, active low signal that is driven low by the PCI Express M.2 device ~~add-in card function~~ to request that the PCI Express reference clock be available (active clock state) in order to allow the PCI Express interface to send/receive data. Operation of the CLKREQ# signal is determined by the state of the Enable Clock Power Management bit in the Link Control Register (offset 010h). When disabled, the CLKREQ# signal shall be asserted at all times whenever power is applied to the device~~card~~, with the exception that it may be de-asserted during L1 PM Substates. When enabled, the CLKREQ# signal may be de-asserted during the L1 Link state.

The CLKREQ# signal is also used by the L1 PM Substates mechanism. In this case, CLKREQ# can be asserted by either the system or the device ~~add-in card~~ to initiate an L1 exit. See the *PCI Express Base Specification* for details on the functional requirements for the CLKREQ# signal when implementing L1 PM Substates.

Whenever dynamic clock management is enabled and when a device ~~card~~ stops driving CLKREQ# low, it indicates that the device is ready for the reference clock to transition from the active clock state to a parked (not available) clock state. Reference clocks are not guaranteed to be parked by the host system when CLKREQ# gets de-asserted and devices ~~module designs~~ shall be tolerant of an active reference clock even when CLKREQ# is de-asserted by the device~~module~~.

The device ~~card~~ must drive the CLKREQ# signal low during power up, whenever ~~it~~ the device is reset, and whenever ~~it~~ the device requires the reference clock to be in the active clock state. Whenever PERST# is asserted, including when the device is not in D0, CLKREQ# shall be asserted.

It is important to note that the PCI Express device must delay de-assertion of its CLKREQ# signal until it is ready for its reference clock to be parked. The device must be able to assert its clock request signal, whether or not the reference clock is active or parked, when ~~it~~ the device needs to put its Link back into the L0 Link state. Finally, the device must be able to sense an electrical idle

break on its up-stream-directed receive port and assert its clock request, whether or not the reference clock is active or parked.

The assertion and de-assertion of CLKREQ# are asynchronous with respect to the reference clock.

Devices ~~Add-in cards~~ that do not implement a PCI Express interface shall leave this CLKREQ# output unconnected~~on the card~~.

CLKREQ# has additional electrical requirements over and above standard open drain signals that allow it to be shared between devices that are powered off and other devices that may be powered on. The additional requirements include careful circuit design to ensure that a voltage applied to the CLKREQ# signal network never causes damage to a component even if that particular component's power is not applied.

Additionally, the device must ensure that it does not pull CLKREQ# low unless CLKREQ# is being intentionally asserted in all cases; including when the related function is in D3cold. This means that any component implementing CLKREQ# must be designed such that:

- Unpowered CLKREQ# output circuits are not damaged if a voltage is applied to them from other powered "wire-ORed" sources of CLKREQ#.
- When power is removed from its CLKREQ# generation logic, the unpowered output does not present a low impedance path to ground or any other voltage.

These additional requirements ensure that the CLKREQ# signal network continues to function properly when a mixture of powered and unpowered components have their CLKREQ# outputs wire-ORed together. It is important to note that most commonly available open drain and tri-state buffer circuit designs used "as is" do not satisfy the additional circuit design requirements for CLKREQ#.

### ~~3.1.3.2.1 Power-up Requirements~~

~~CLKREQ# is asserted in response to PERST# assertion. On power up, CLKREQ# must be asserted by a PCI Express device within a delay (TPVCRL) from the power rails achieving specified operating limits and PERST# assertion (see **Error! Reference source not found.**). This delay is to allow adequate time for the power to stabilize on the card and certain system functions to start prior to the card starting up. CLKREQ# may not be de-asserted while PERST# is asserted.~~

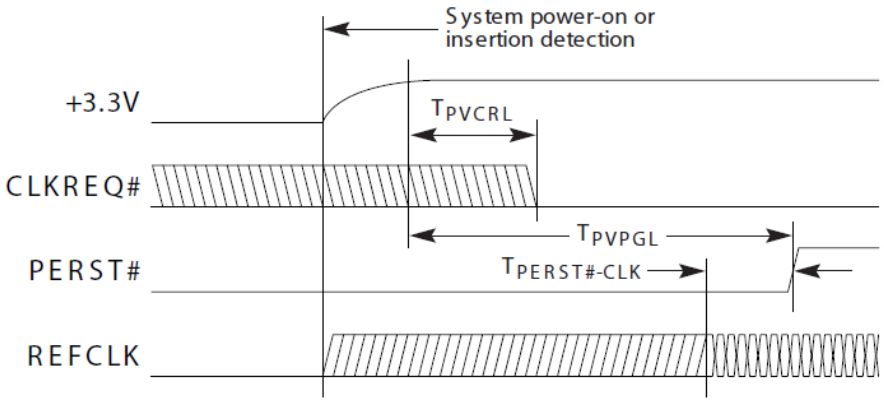


Note: $T_{PVCRL}$ is measured from the later rising edge of +3.3V.

~~Figure 1. Power-Up CLKREQ# Timing~~

~~The system is required to have the reference clock for a PCI Express device in the parked clock state prior to device power-up. The state of the reference clock is undefined during device power-up, but it must be in the active clock state for a setup time $T_{PERST\#-CLK}$ prior to PERST# de-assertion. Table 1 lists the power-up CLKREQ# timing.~~

## ~~Table 1.  Power-Up CLKREQ# Timings~~

| ~~Symbol~~ | ~~Parameter~~ | ~~Min~~ | ~~Max~~ | ~~Units~~ |
|---|---|---|---|---|
| ~~$T_{PVCRL}$~~ | ~~Power Valid to CLKREQ# Output active~~ | | ~~100~~ | ~~µs~~ |
| ~~$T_{PVPGL}$~~ | ~~Power Valid to PERST# Input inactive~~ | ~~1~~ | | ~~ms~~ |
| ~~$T_{PERST\#-CLK}$~~ | ~~REFCLK stable before PERST# inactive~~ | ~~100~~ | | ~~µs~~ |

### 3.1.3.2.2 Dynamic Clock Control

After a PCI Express device has powered up and whenever its upstream link enters the L1 link state, it shall allow its reference clock to be turned off (put into the parked clock state). To accomplish this, the device de-asserts CLKREQ# (high) and must allow that the reference clock will transition to the parked clock state within a delay ($T_{CRHoff}$). Figure 80 shows the CLKREQ# clock control timing diagram.

To exit L1, the device must assert CLKREQ# (low) to re-enable the reference clock. After the device asserts CLKREQ# (low) it must allow that the reference clock will continue to be in the parked clock state for a delay ($T_{CRLon}$) before transitioning to the active clock state. The time that it takes for the device to assert CLKREQ# and for the system to return the reference clock to the active clock state are serialized with respect to the remainder of L1 recovery. This time must be taken into account when the device is reporting its L1 exit latency.

When the PCI Express device supports, and is enabled for, Latency Tolerance Reporting (LTR), the device must allow that the reference clock transition to the active clock state may be additionally delayed by the system up to a maximum value consistent with requirements for the LTR mechanism. During this delay, the reference clock must remain parked. When exiting the parked state following the delay, the clock must be stable and valid within 400 ns.

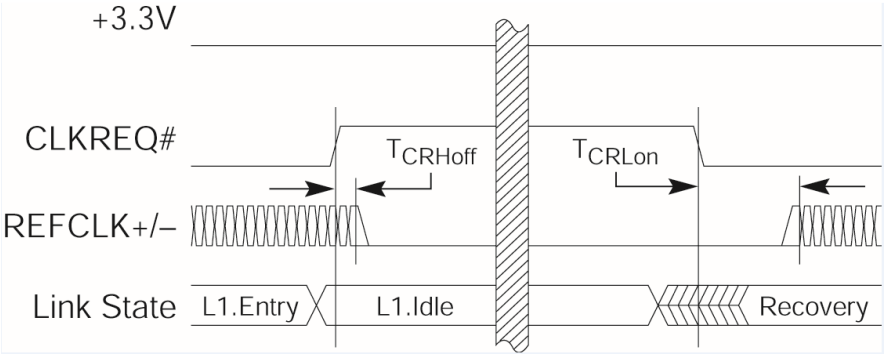


Figure 80.  CLKREQ# Clock Control Timings

All links attached to a PCI Express device must complete a transition to the L1.Idle state before the device can de-assert CLKREQ#. The device must assert CLKREQ# when it detects an electrical idle break on any receiver port. The device must assert CLKREQ# at the same time it breaks electrical idle on any of its transmitter ports in order to minimize L1 exit latency. See Table 17 for CLKREQ# clock control timing.

## Table 17. CLKREQ# Clock Control Timings

| Symbol | Parameter | Min | Max | Units |
|---|---|---|---|---|
| $T_{CRHOFF}$ | CLKREQ# de-asserted high to clock parked | 0 | | ns |
| $T_{CRL}$ | CLKREQ# asserted low to clock active | | 400* | ns |

**Note:** *$T_{CRLon}$ is allowed to exceed this value when LTR is supported and enabled for the device

There is no maximum specification for $T_{CRHoff}$ and no minimum specification for $T_{CRLon}$. This means that the system is not required to implement reference clock parking or that the implementation may not always act on a device de-asserting CLKREQ#. A device should also de-assert CLKREQ# when its link is in L2 or L3, much as it does during L1.

### 3.1.3.3 Clock Request Support Reporting and Enabling

Support for the CLKREQ# dynamic clock protocol should be reported using bit 18 in the PCI Express link capabilities register (offset 0C4h). To enable dynamic clock management, bit 8 of the Link Control register (offset 010h) is provided. By default, the device ~~card~~ shall enable CLKREQ# dynamic clock protocol upon initial power up and in response to any warm reset by the host system. System software may subsequently disable this feature as needed. Refer to the *PCI Express Base Specification*, Revision 1.1 (or later) for more information regarding these bits.

### 3.1.3.4 PERST# Signal

- The PERST# signal is de-asserted to indicate when the system power sources are within their specified voltage tolerance and are stable.
- PERST# should be used to initialize the card functions once power sources stabilize.
- PERST# is asserted when power is switched off and also can be used by the system to force a hardware reset on the card.
- System may use PERST# to cause a warm reset of the add-in card.
- PERST# is asserted in advance of the power being switched off in a power-managed state like S3.
- PERST# is asserted when the power supply is powered down, but without the advanced warning of the transition.

~~Refer to the *PCI Express Card Electromechanical Specification* for more details on the functional requirements for the PERST# signal.~~

**Note:** The $T_{PVPGL}$ referenced in Table 16 that is directly related to PERST# de-assertion is defined as Minimum 1 ms in the PCIe Mini CEM Spec and as Minimum 100 ms in the PCIe CEM Spec.

## 3.1.x Power-up Timing (after existing 3.1.3 and before existing 3.1.4)

Figure 79 provides an overview of the M.2 device power up sequence. Table 16 lists the power-up timing variable values.

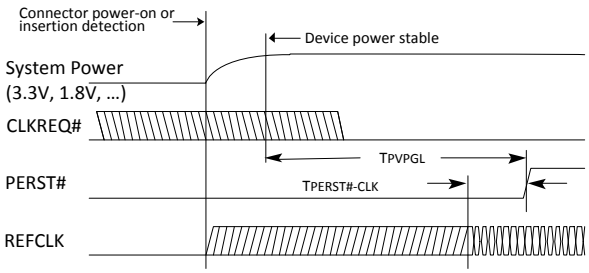


Figure 79. Power-Up Timing Sequence

| Symbol | Parameter | Min | Max | Units |
|---|---|---|---|---|
| $T_{PVPGL}$ | Power Valid to PERST# Input inactive | Implementation specific; recommended 50 ms | | ms |
| $T_{PERST\#-CLK}$ | REFCLK stable before PERST# inactive | 100 | | µs |

Table 16. Power-Up Timing Variables

### 3.1.x.1 PERST# Power-up Timing

The host shall delay de-assertion of PERST# for a period of time ($T_{PVPGL}$) after power is stable on the device (see Figure 79). See section 3.1.3.4 for further details on PERST#.

The PCI Express Specification (see Revision 3.0 or later, Conventional Reset) requires that a PCIe Express device must be in the LTSSM Detect state within 20 ms of PERST# being de-asserted and ready for Configuration Requests within 120 ms of PERST# being de-asserted.

The value of $T_{PVPGL}$ is left as implementation specific, with a recommended value as a guideline. In considering the value of $T_{PVPGL}$:

- Device and host implementers should consult PCIe Express Reset Rules and platform BIOS and OS requirements governing device readiness timing requirements following the de-assertion of PERST#.
- Host implementers should consult device vendors for their $T_{PVPGL}$ values, based on vendor specific device startup requirements.

### 3.1.x.2 REFCLK Power-up Timing

The host shall ensure that the reference clock is in the active clock state for at least a period of time specified by $T_{PERST\#-CLK}$, prior to PERST# de-assertion. See section 3.1.3.1 for further details on REFCLK.

### 3.1.x.2 CLKREQ# Power-up Timing

See section 3.1.3.2 for details on CLKREQ#.