

Notes

By Derek Huang & Upendra Kulkarni

Introduction

The PCI Express Base Specification states that the reference clocks for ports at the two ends of a PCI Express Link must be modulated such that they never exceed a total of 600 ppm (parts per million) difference. PCI Express specifications also permit use of SSC (Spread Spectrum Clocking) with +0 to -5000 ppm modulation of clock frequency at a modulation rate not to exceed 33KHz. The 600 ppm rule applies even when SSC is used. If SSC is desired in PCIe based systems, the implication of the 600 ppm rule to simple system implementations is that if the reference clock used by a port at one end of a Link is SSC then the port at the other end of the link must use the same clock source. For a simple PCIe Switch, this implies that if one or more devices connected to the switch are clocked with SSC, then the switch and all devices connected to the switch must use the same SSC source. This may not be feasible in many system topologies for a variety of reasons. A solution to this problem is to create a PCIe Switch which can isolate the clocking domains on a per port basis such that each Switch port can function properly with a potentially unique SSC source required by and shared with its link partner. The Switch port and its Link partner can be logically thought of as functioning within their own local clocking domain, unconnected to the rest of the Switch device/ports in terms of clocking.

The IDT PCIe Gen2 System Interconnect and Inter-Domain Switch families offer devices which implement local port clocking functionality described above. This application note describes various clocking topologies which can be implemented using these Switches. IDT System Interconnect devices which offer this local port clocking functionality include the 89H64H16[A]G2, 89H48H12[A|B]G2, 89H34H16G2, 89H32H8G2 and 89H22H16G2. All of the IDT PCIe Gen2 Inter-Domain Switches offer local port clocking. Examples of such Inter-Domain Switch are the 89H32NT24[A|B]G2, 89H24NT24[A]G2, 89H32NT8[A|B]G2, 89H16NT16[B]G2, 89H12NT12G2, 89H16NT2G2, etc. For simplicity, this document often uses a representative product name “PES64H16AG2” or similar shorthand to represent all products with this local port clocking capability.

Note that there is a slight difference between the implementations of local port clocking in the System Interconnect switches and the Inter-Domain Switches. In the System Interconnect Switches it is possible to have a unique local port clock for each port, while in the Inter-Domain Switches one local port clock is shared by two to four ports. Please refer to the Inter-Domain Switch User Manuals for additional specifications related to which ports share which local port clock source and the limitations this may pose in some systems. The remainder of this document primarily refers to the one local port clock per port implementation for simplicity.

Clocking Architecture

Figure 1 represents the Clocking Architecture within the IDT Gen2 PCIe System Interconnect Switches. The devices have two differential Global Reference Clock inputs (GCLKP[1:0]/GCLKN[1:0]) and differential Port Reference Clock inputs (PxCLK) for ports that support local port clocking. Note that for the Global Reference Clock inputs, none must be left unconnected, and both must be connected to the same source. Unused Port Reference Clock inputs must be connected to Vss.

Notes

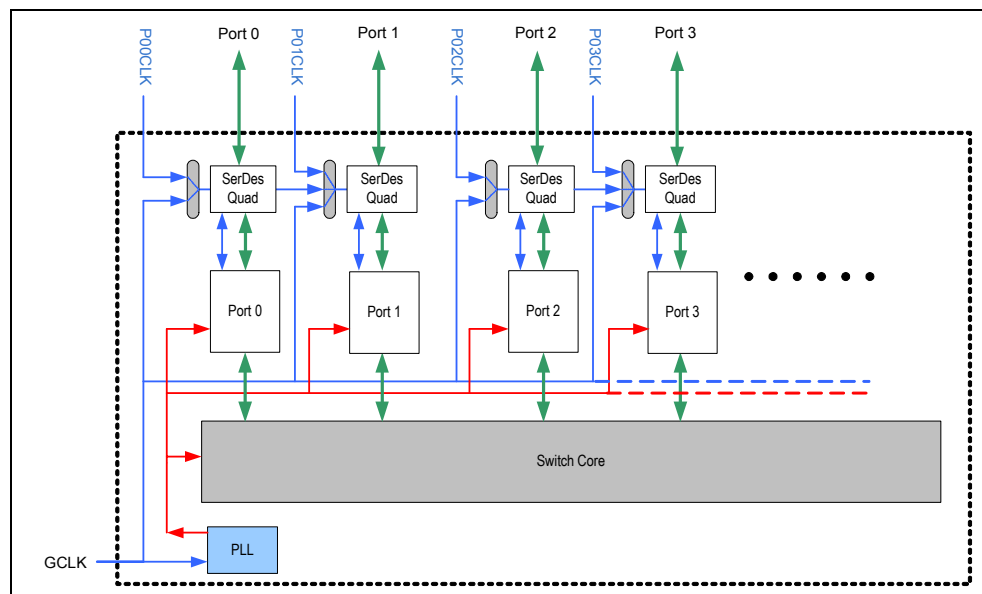


Figure 1 Logical Representation of Clocking Architecture

The Global Reference Clock inputs to PLL are used to generate all of the clocks required by the internal Switch logic and the SerDes. The Local Port Reference Clock input associated with a port (or a group of ports in the Gen2 Inter-Domain Switches) is used by the SerDes only when a port (or a group of ports in the Gen2 Inter-Domain Switches) is configured to function in local port clocked mode. Local Port Clock refers to the clock that a port uses to receive and transmit serial data.

The differential clock inputs require the signal source to drive 0V common-mode and the REFCLK signal must meet the electrical specifications defined in the PCI Express Base 2.0 Specification and the IDT Device Datasheet.

The Port Reference Clock inputs support Spread spectrum clocking (SSC) for reduced EMI and allow isolation of this SSC from the rest of the Switch. The PCI Express Base 2.0 specification permits SSC down-spread technique which only allows from +0% to -0.5% (+0 to -5000 ppm) of the nominal data rate frequency at a modulation rate not to exceed 30 kHz - 33 kHz while still meeting ± 300 ppm requirement (link partners can not exceed a total of 600 ppm difference).

If SSC is used by the global reference clock, the same clock source or generator must be used for all link partners of the switch, both upstream and downstream.

There are no skew requirements between the Global Clock input and a Port Reference Clock input or between any of the Port Reference Clock inputs. A constant phase difference is acceptable.

The switch supports both 100 MHz and 125 MHz reference Global Clock inputs and the frequency selection is made via the Global Clock Frequency Select pin (GCLKFSEL).

Port Clocking Modes

A specific Port Reference Clock input (PxCLK) is associated with only one specific port (or a specific group of ports in the case of Inter-Domain Switches). Depending on the port clocking mode, a differential port reference clock input is driven into the device on the corresponding PxCLKP and PxCLKN pins. The frequency of a PxCLK input must always be 100 MHz.

A Switch port may operate in Global Clocked mode or Local Port Clocked mode. A port (or specific groups of ports in the case of Inter-Domain Switches) may be independently configured to operate in any port clocking mode without regard to the mode of any other port (or groups of ports in the case of Inter-Domain Switches).

Notes

A Port Reference Clock input is associated with each x4 port (or a SerDes quad). When two x4 ports are merged to create a x8 port, the Port Reference Clock input used by the merged port is the one associated with the even port. The Port Reference Clock input associated with the merged odd port is unused and may be left unconnected or terminated to ground.

The initial port clocking mode of a port is determined by the state of the CLKMODE[2:0] pins in the boot configuration vector. The port clocking mode can be modified at any time by programming the Port Clocking Mode (PxCLKMODE) field in the Port Clocking Mode (PCLKMODE) register.

If a port operates in global and common clocked mode with its link partner, SSC is supported as shown in Figure 2. However, if a port is in global and non-common clocked mode with its link partner as shown in Figure 3, both GCLK and the clock source to its link partner must disable SSC. SSC is only supported in common-clocked configurations.

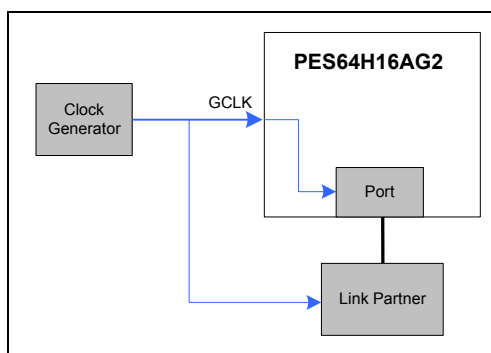


Figure 2 Global, Common Clocked Configuration

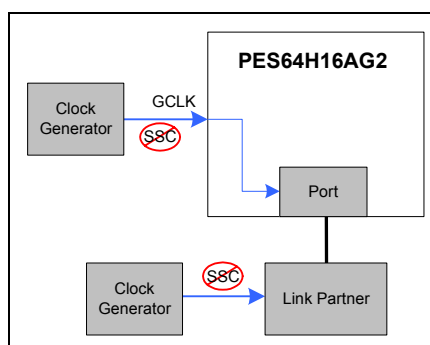


Figure 3 Global, Non-Common Clocked Configuration

SSC can be supported on the Port Reference Clock input (PxCLK) and reference clock to its link partner as shown in Figure 4 and such a port must be configured for common clock mode. In this scenario, the Global Reference Clock inputs must disable SSC. For a port operating in local port clocked and non-common clocked mode with its link partner as shown in Figure 5, all clock sources must disable SSC. There is no reason to use local port clocking in this scenario.

Notes

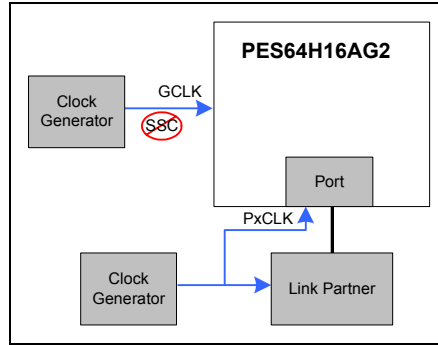


Figure 4 Local Port Clocked, Common Clocked Configuration

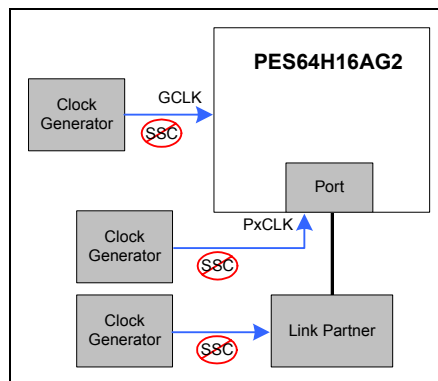


Figure 5 Local Port Clocked Non-Common Clocked Configuration (not a meaningful set up)

Notes

System Clock Usage and Applications

The majority of enterprise class server and storage applications share buffered clock sources from the system or the root complex (CPU and I/O chipset complex) and pass it down to the switch and the endpoints as shown in Figure 6. The Switch ports and their link partners are in common clock mode.

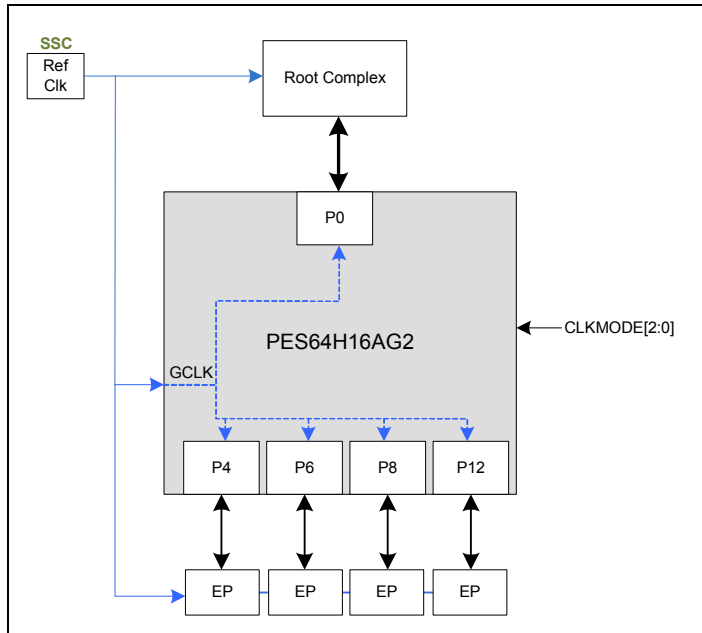


Figure 6 Typical Server/Desktop System Clock Usage Model

In certain applications the Root Complex upstream from the Switch and the endpoints downstream from the Switch are required to work with different clock sources. It is likely that the Root Complex requires SSC. This can be addressed by ensuring a CFC (non-SSC) clock source to the Global Clock inputs and to the endpoints as shown in Figure 7. One or more of the endpoints (typically FPGAs) may require a 125 Mhz CFC reference clock, and this can be supported by the switch.

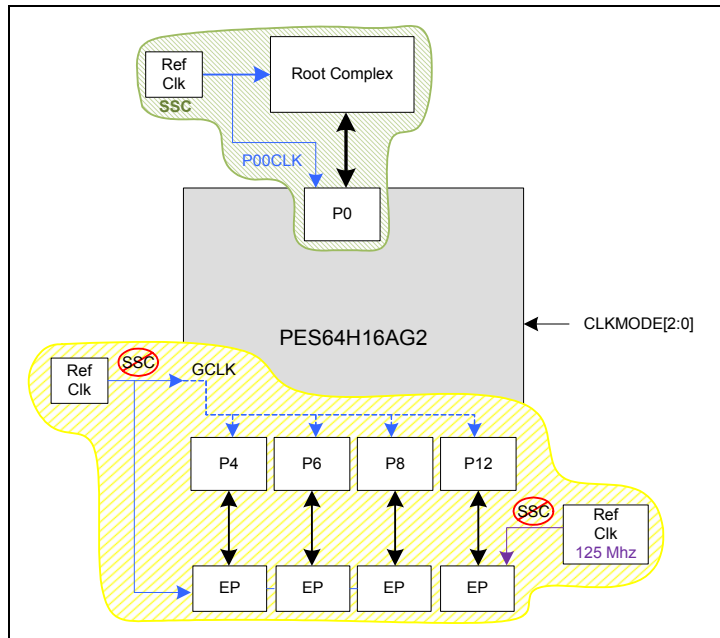


Figure 7 System Clock Configuration #1

Notes

When GCLK is required to use SSC, all link partners must operate in common-clocked mode and from the same clock source as the GCLK. Figure 8 shows an unsupported configuration where the link partners (Root Complex and the endpoints) operate in different clock domains even when the GCLK is SSC.




Figure 8 System Clock Configuration #2, NOT SUPPORTED

In multi-root applications, the switch global clock and the link partners can have different clock sources. Ports must be in local port clock mode if SSC is required (refer to ports 2, 4, and 6 in Figure 9).




Figure 9 System Clock Configuration #3

Notes

Figure 10 shows the switch configured in multi-partition mode with one upstream and two downstream ports in each partition. Each partition operates independently as an independent logical or virtual switch. The upstream ports are configured in local port clocked mode with each clock generator providing SSC to a pair of Root Complex and PxCLK input of corresponding upstream port. The clock source to Global Clock inputs must disable SSC in this configuration. The downstream port link partners that do not share the same clock source as GCLK can use SSC and configure the ports in local port clocked mode.

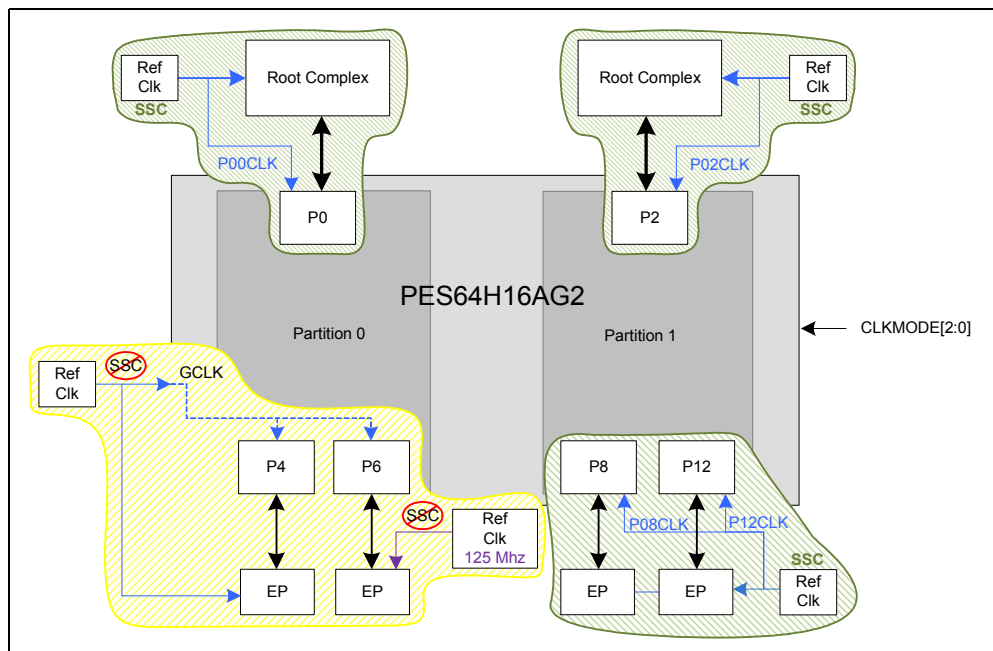


Figure 10 System Clock Configuration #4

Conclusion

The local port clocking feature in IDT PCIe Gen2 System Interconnect and Inter-Domain switches enables system topologies which require multiple PCIe devices to be connected to the switch, each with their own unique SSC or CFC. A large variety of server, storage, communications, and other embedded applications can be realized using this unique feature offered only by IDT PCIe Switches.

References

- [1] IDT 89HPES64H16AG2 & IDT 89HPES32NT24AG2 PCIe Switch User Manuals.
- [2] PCI Express Base Specification 2.0.

Revision History

- June 28, 2009:** Initial publication.
- October 26, 2009:** Updated with Inter-Domain Switch references.

IMPORTANT NOTICE AND DISCLAIMER

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD-PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers who are designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only to develop an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third-party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising from your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Disclaimer Rev.1.01 Jan 2024)

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu,
Koto-ku, Tokyo 135-0061, Japan
www.renesas.com

Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit www.renesas.com/contact-us/.