December 2022

BIFURCATED THUNDERBOLT HOST PCLE X16 CARD WITH DISCRETE GPU

HP INC

Bifurcated Thunderbolt Host PCIe x16 Card with Discrete GPU

This disclosure relates to Thunderbolt Type C I/O integration into a host computer architecture using the industry standard PCI Express x16 add-in card with both a discrete graphics processor (GPU) and up to two discrete Thunderbolt host controllers.

An architecture is disclosed that integrates a GPU with up to two Thunderbolt host controllers on shared PCI Express x16 slot add-in card with a bifurcated PCI Express x16 interface. The x16 interface is bifurcated into x8 + x4 + x4 lane device groups to independently enumerate each discrete controller through the shared connector. X8 lanes are reserved for the GPU. X4 lanes are reserved for each Thunderbolt controller, up to two controllers. The result is a solution with two to four Thunderbolt Type C I/O connectors as shown in Figure 1.

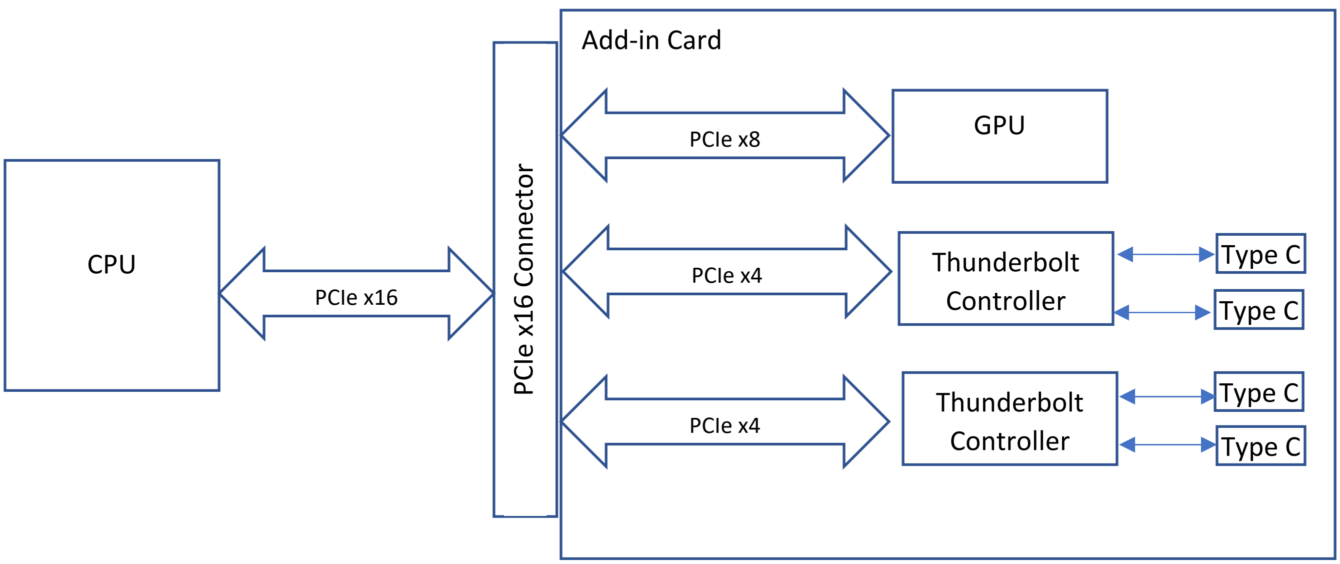

Figure 1 - Bifurcated PCI Express Interface Thunderbolt Add-in Card

PCI Express data lane groups are typically pre-configured when the system is booted and not dynamically reconfigured during run-time. This is challenging for a single add-in card then to bifurcate the connector data lanes into smaller lane groups across multiple devices. The host will detect the first devices but not know to look for the secondary devices on the unused lanes. Solutions exist for switching lanes groups between multiple connectors when multiple cards are detected via the PRSNT2# connector pins, but this physical hardware detection also occurs during boot to configure the lanes for independent devices, not dynamically during run-time. A common example of this solution relates to a x16 connector where the upper x8 lanes are switchable to a second x16 connector to allow two x16 graphics cards to share the same x16 interface between two physical connectors that each receive x8 data lanes. During system boot, this hardware configuration is detected, and the switch is configured as shown in Figure 2.

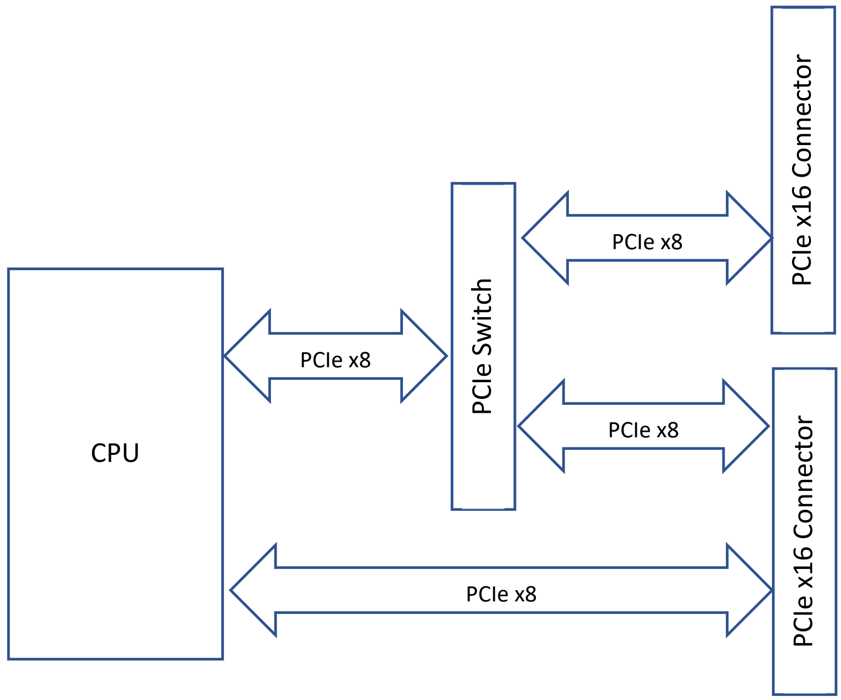

Figure 2 - PCI Express Physical Lane Switch

When multiple independent devices are integrated on a common add-in card, the host initially assumes that the card includes a single device. A PCI Express interface scales the link down to the number of active data lanes based on the device data lane width that responds during training and enumeration. A x16 interface can support x8, x4, x2, and x1 links in addition to the full x16 link and the unused lanes are left idle and not automatically reallocated for use by other devices. The host must be instructed to reallocate these lanes to additional devices. This can be done manually through the system BIOS setup to instruct the host to bifurcate the lanes in smaller groups to accept multiple devices as implemented on the installed add-in card.   For this disclosure add-in card, the BIOS would select a x8 + x4 + x4 device configuration instead of the full x16 lanes.

A more complicated method requires the host to dynamically activate new PCI Express devices when a link is established with less than all available lanes. In this case of this disclosure, the slot connector is originally configured as a x16 interface. The GPU would enumerate first as a x8 link.   The host would acknowledge the first x8 link and then enable a new device on the unused x8 data lanes. The first Thunderbolt host controller would then enumerate with only x4 data lanes. The host would accept the first x4 device and then enable a new device on the remaining unused x4 lanes. The second Thunderbolt host controller would then enumerate on the final x4 data lanes. Refer to Figure 3.

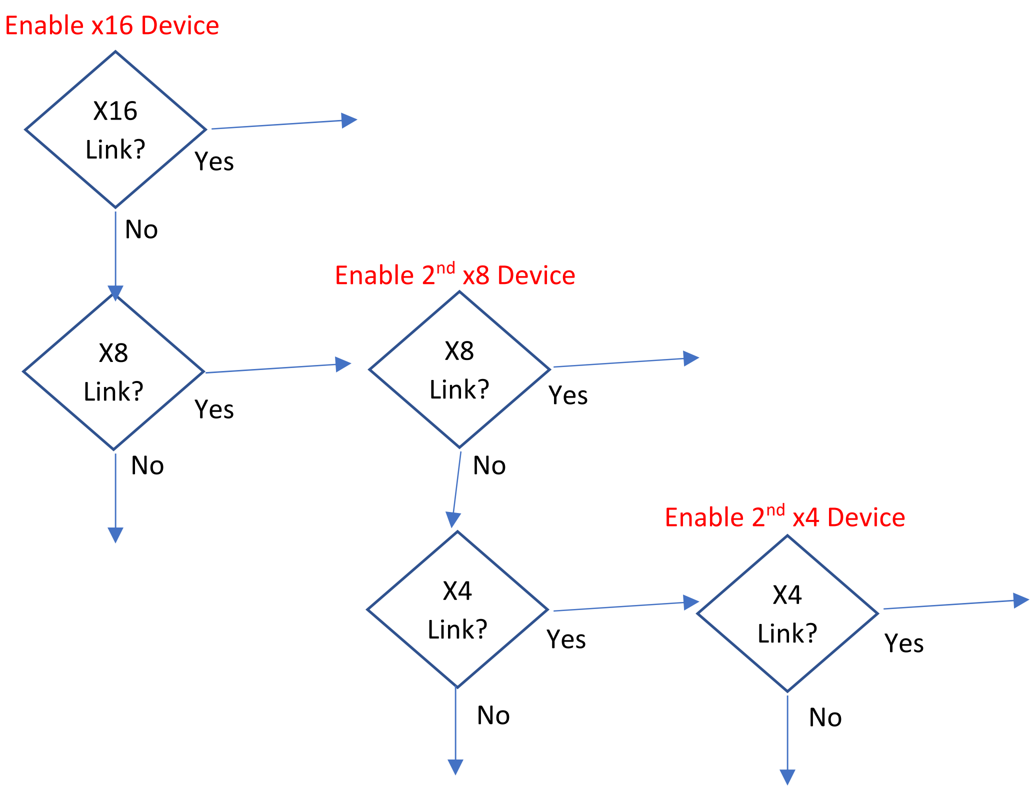

Figure 3 - Automatic Bifurcation and Device Detection

Thunderbolt requires PCI Express and DisplayPort interfaces. USB 3.x and 2.0 XHCI host functions are integrated into the Thunderbolt controller. PCIe Express and power are available via the motherboard slot connector. DisplayPort loopback cables are typically routed from the CPU motherboard to the add-in card to redirect DisplayPort sources from the CPU integrated display adapter. These cables may be internal or external. External cables occupy space on the add-in card backplate for DisplayPort input connectors reducing the available space for thunderbolt Type C output connectors.

The integrated discrete GPU on the add-in card generates the DisplayPort sources locally without cabling. A typical GPU will support four displays but can vary. A Thunderbolt host controller may accept up to three DisplayPort sources and support two Thunderbolt Type C outputs. A discrete GPU supporting four displays may therefore support a single Thunderbolt controller or dual Thunderbolt controllers. Thunderbolt controllers also support a separate DisplayPort output in addition to the two Thunderbolt outputs to redirect the display to a non-Thunderbolt port. If the Thunderbolt controller supports 3 DisplayPort inputs, this extra DisplayPort output may use to cross-couple a third DisplayPort input to support up to three displays at each controller but a maximum of only four displays across both Thunderbolt controllers as shown in Figure 4.

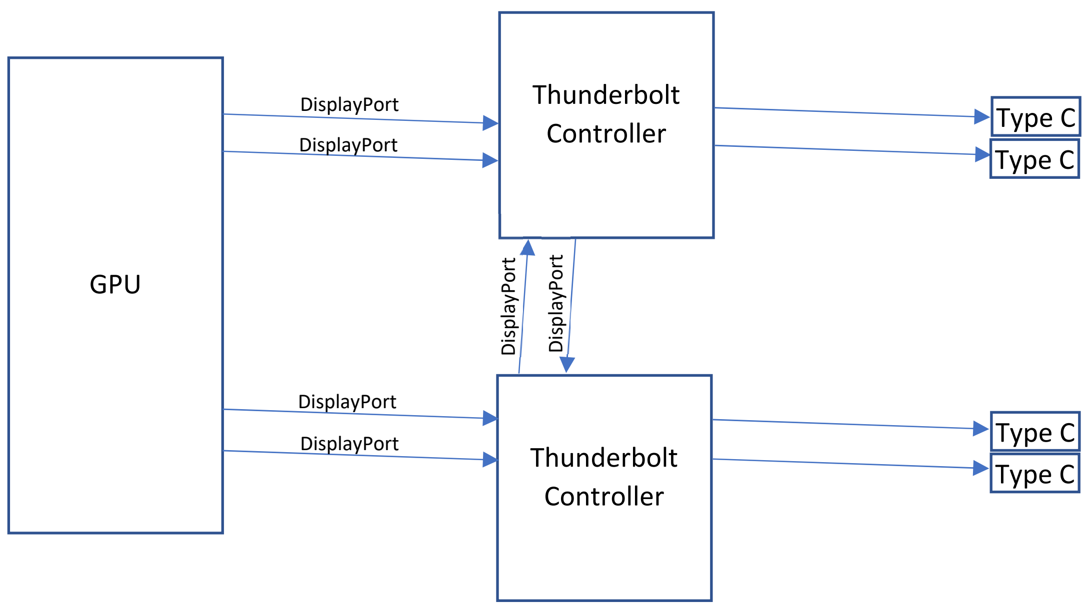

Figure 4 - Cross-Coupled DisplayPort Routing

Disclosed by Patrick Ferguson, HP Inc.