

What is PXImc?

*By Chetan Kapoor, PXI Product Manager
National Instruments*

Overview

Modern day test and control systems are growing larger, more complex and more intricate. Most of these intricacies are in order to address new challenges such as the need for higher speed, higher channel-count data acquisition and generation and inline data processing. These new test, measurement and control systems need to address these needs while still having to provide other features such as low cost, scalability, vendor interoperability and flexibility. At the core of these issues there is a fundamental need to transfer and process large amounts of data with minimal delay.

The PXI MultiComputing (PXImc) specification, announced by the PXI System Alliance in November 2009, attempts to address such needs of these complicated systems by using low-cost, off-the-shelf technology. The specification provides the hardware and software requirements to connect two or more intelligent systems together using PCI or PCI Express based interfaces via a PCI(e) Non-Transparent Bridge (NTB). The PCI Express bus offers multi-gigabyte per seconds of practical data throughput and few microsecond latency and is thus well suited to meet these data transfer needs.

This white paper explores the technical details of PXImc and the specific use cases that it helps to address.

Understanding the technology underneath

Figure 1 illustrates the topology of a typical PCI Express based system. A single host bus interface, or root complex, is the interface for the computing system to communicate with PCI devices. The computing system will interface with this root complex to assign the required resources to all PCI(e) devices in the system.

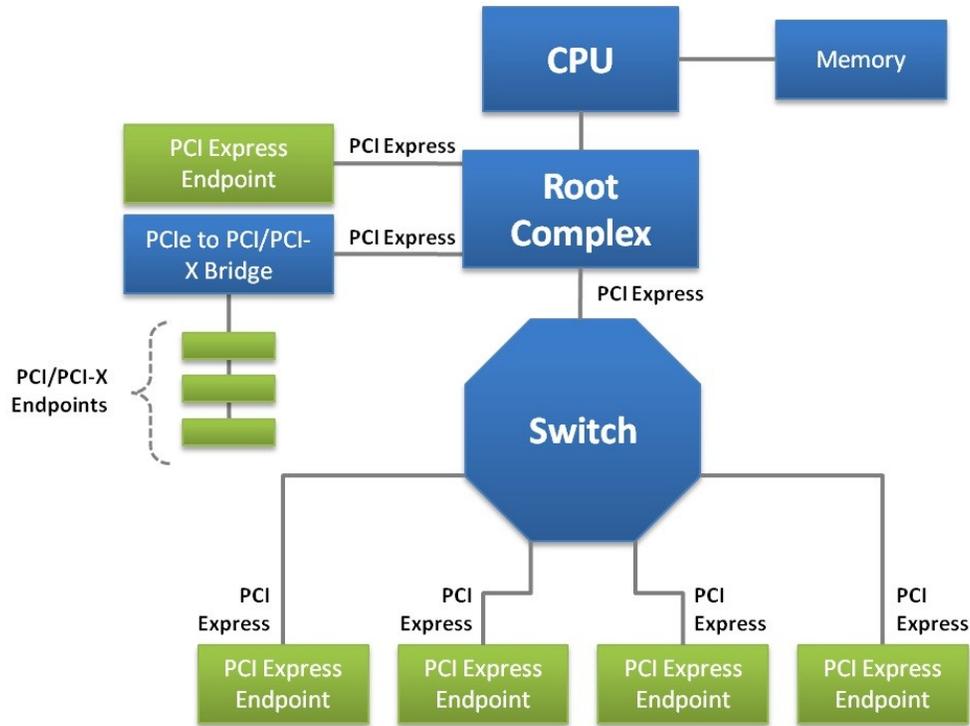


Figure 1: PCI Express System Topology

As mentioned earlier, PXImc provides a high bandwidth and low latency communications model by leveraging PCI or PCI Express as the physical communication layer. Two systems with their own root complexes cannot be directly connected via PCI or PCI Express because of various contentions between the two PCI domains such as bus ownership and endpoint resource allocations.

The use of a Non-Transparent Bridges (NTB) address these contentions by logically separating the two PCI domains while at the same time providing a mechanism for translating certain PCI transactions within one PCI domain into corresponding transactions in another PCI domain.

Figure 2 illustrates this concept. Both System A and System B have complete control of resource allocation within their own domains and the presence of the non-transparent bridge does not affect either's resource allocation algorithm.

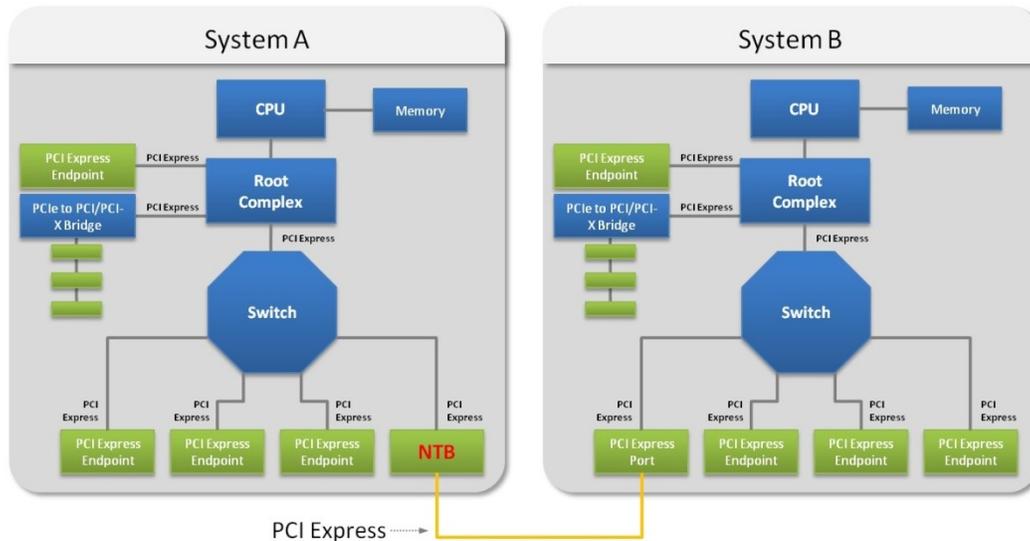


Figure 2: Two Systems connected via PCI Express by leveraging a Non-Transparent Bridge (NTB)

The non-transparent bridge responds to resource requests from the root complex similar to the other PCI endpoints in the system, requesting some amount of physical address space. The system BIOS will assign a specific range of physical address space to the non-transparent bridge. As this resource allocation occurs on both System A and System B, the non-transparent bridge will acquire resources in both PCI domains. As shown in figure 3, the address space acquired by the non-transparent bridge within PCI domain of system A will act as a window into physical address space within PCI domain of system B, and the address space acquired within PCI domain of system B will act as a window into physical address space within PCI domain of system A.

After resource allocation executes on both System A and System B, the non-transparent bridge will contain memory mechanisms to transfer data between the two systems. These include items such as scratch pad registers to send data, doorbell registers to assert interrupts, and large blocks of address space to translate across the NTB into the opposing address space.

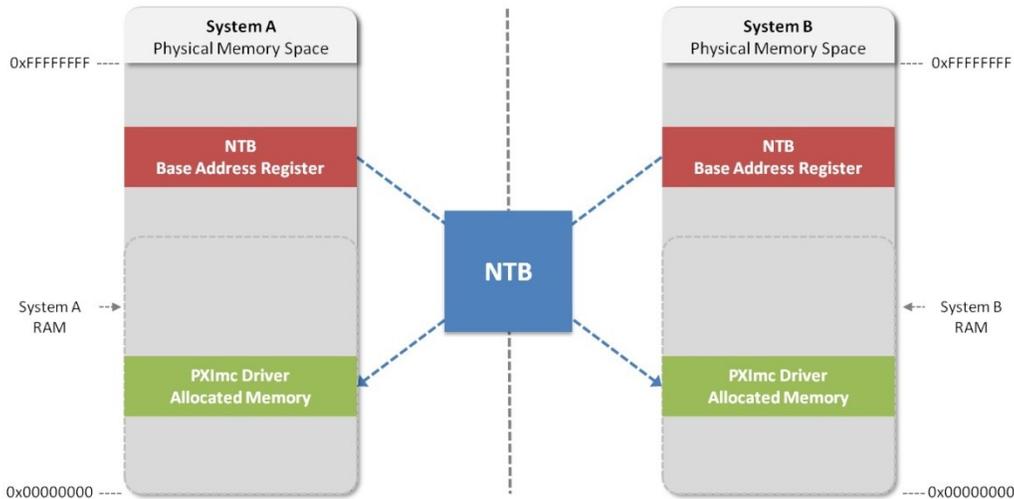


Figure 3: Communication mechanism between two PCI domains by using NTBs.

While NTB technology is not new, the industry has lacked a standardized means by which to easily communicate between PCI hierarchies. NTB solutions required custom hardware and software design prior to using them as a communications channel.

The PXI MultiComputing (PXImc) specification developed by the PXI System Alliance (PXISA) provides specific requirement for hardware and software components and hence provides a standardized protocol for communicating between two systems using PCI or PCI Express. From the hardware standpoint, several issues were reconciled to allow two independent systems to communicate directly over PCI or PCI-Express. From the software view, a communications scheme was created to allow each system to discover and configure its own resources in a way to communicate with the other system.

Expected & Benchmarked Bandwidth and Latency Performance

Since PXImc leverages PCI Express as the physical communication layer, the performance of a PXImc link will depend on the type of the PCI express interface used. Table 1 lists the theoretical bandwidth of various PCI Express links.

| PCI Express Link | Generation | Theoretical Unidirectional Transfer Rates | Theoretical Bidirectional Transfer Rates |
|------------------|------------|---|--|
| x4 | Gen 1 | 1 GB/s | 1 GB/s x2 |
| x16 | Gen 1 | 4 GB/s | 4 GB/s x2 |
| x4 | Gen 2 | 2 GB/s | 2 GB/s x2 |
| x16 | Gen 2 | 8 GB/s | 8GB/s x2 |

Table 1: Theoretical bandwidth of various PCI Express links

In order to compare the performance of a typical PXImc link to other alternatives, National Instruments performed some initial benchmarking using prototype hardware and preliminary software stack. For this benchmarking, two PXIe-1082 PXI Express chassis each with their own PXIe-8133 PXI Express embedded controller were connected by using prototype x4 gen1 PCIe based PXImc hardware. Using this setup, we observed a one-way latency of 6 μ s and throughput of 670 MB/s. As compared to Gigabit Ethernet, these performance numbers represent roughly 10 times better bandwidth and 100 times better latency validating PXImc as an ideal interface for building high-performance multicomputer test and control systems. As new improvements are made to the PCI Express standard, the performance offered of a PXImc link will continue to scale.

Using PXImc for Creating High-Performance Multicomputer Test and Control System

Applications that can benefit the most from PXImc can be broken down in to the three categories. First would be hybrid high-performance test and measurement systems, second would be applications that require the use of multiple multicore CPUs for distributed processing and last would be applications that require the use of x86-based CPUs as co-processors in a single system.

Hybrid High Performance T&M Systems

Complex measurement and control systems such as HIL systems typically consist of various discrete subsystems that provide different functionalities. Often times these subsystems are based on different hardware platforms in order to best match the functional requirements of the subsystem to the hardware platform's capabilities.

PXImc specification allows these subsystems based on different hardware platforms to communicate via a high-bandwidth low-latency cabled PCI express link. This ability enables the creation of a hybrid high-performance test and measurement system that can utilize different hardware platforms to offer the ideal combination of performance, flexibility and cost.

By leveraging PXImc as the intersystem communication bus, such systems can achieve faster test times or loop execution rates as compared to using other slower interfaces such as Ethernet and Reflective Memory. This capability thus enables these systems to execute their specific task faster, and in the case of HIL systems, more accurately simulate the real-world conditions.

Figure 4 illustrates a typical configuration for such a system. In this example, the master PXI system communicates via the high-throughput and low-latency PXImc link to the secondary PXI systems and the high-performance boxed instruments to create a hybrid high-performance test and measurement system.

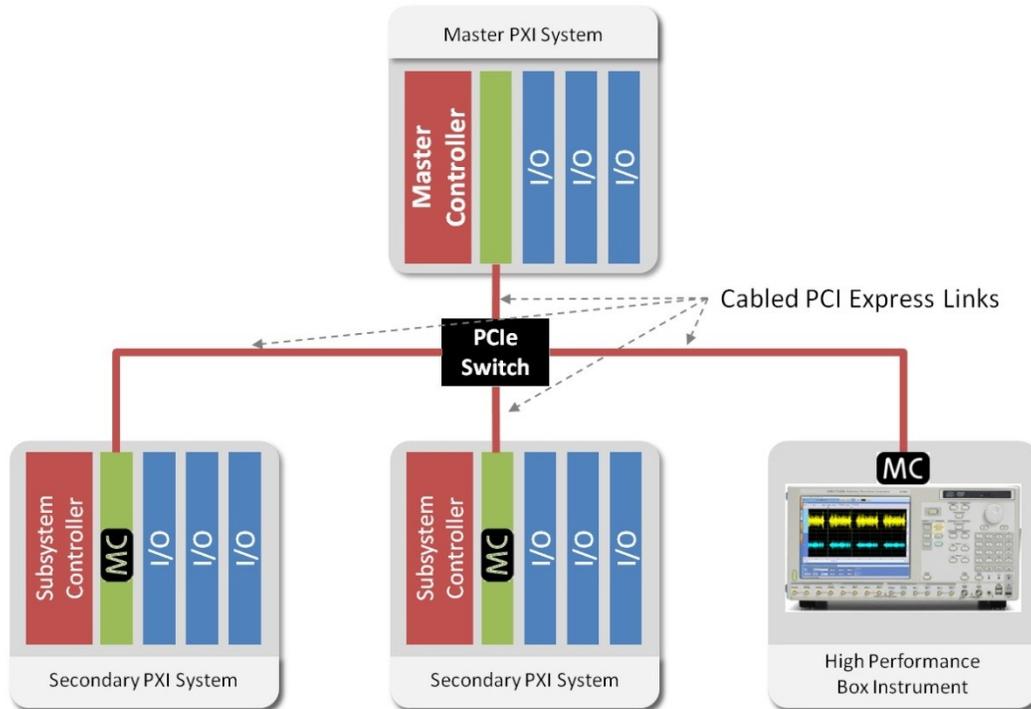


Figure 4: Example configuration for a PXImc based hybrid high-performance test and measurement system.

Using Multicore CPUs for Distributed Processing

Applications such as SIGINT and Real-Time High Performance Computing (RTHPC) that require real-time or in-line data processing typically need intensive data processing capabilities to keep up with the data flow. These applications need the ability to distribute the processing load between multiple discrete processing nodes using a communication interface that offers high-bandwidth and low-latency.

For some distributed processing system, the use of FPGAs and discrete DSPs can address this need. However certain applications require either the use of existing x86 based software IP or they need to perform the computations in floating-point instead of fixed-point. In such cases PXImc provides the ability to create distributed processing systems that utilize the latest multicore CPU based PCs as external compute nodes.

This not only addresses the aforementioned needs but also provides a system with faster development and debugs times compared to systems leveraging FPGA and custom DSPs.

Figure 5 illustrates one of the feasible distributed computing system configurations that can be created using a PXI Express system and PXImc interface cards.

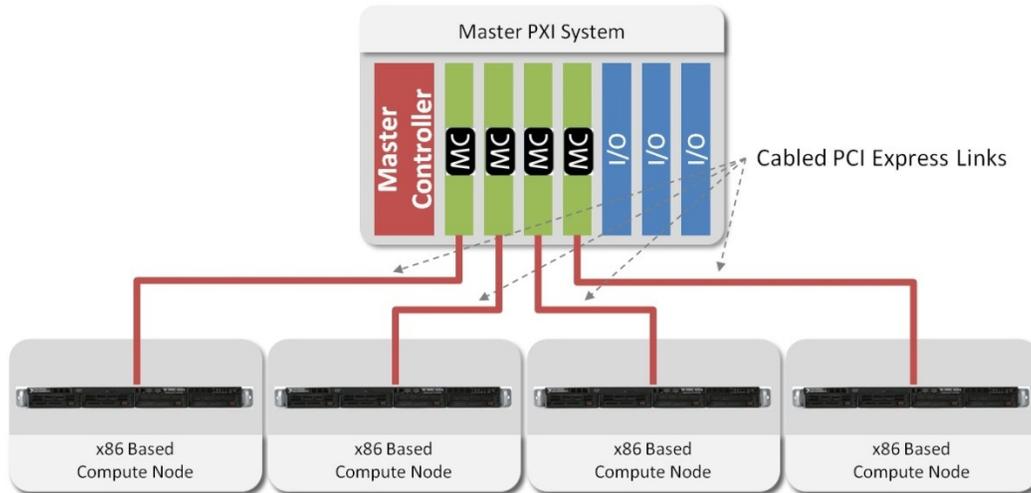


Figure 5: Example of distributed computing system based on a PXI Express system and PXImc interface cards.

In this example, the master controller is responsible for aggregating data from various I/O modules and then distributing that data over PXImc links to four x86 based compute nodes. The compute nodes can be regular PCs or high-end workstations depending upon the required processing performance.

National Instruments demonstrated a system based on this configuration at NIWeek 2010 that showed a Real-Time High Performance Computing system in action.

[Click here to view video on YouTube.](#)

x86 based Co-processing

The PXImc specification allows for processing system to have onboard PCI(e) NTB. This capability allows such processing modules to reside in a peripheral slot of a PXI express chassis and use the PCI express signaling in the chassis backplane to communicate with the master controller. Having in-chassis peripheral processing modules greatly increases the total processing capability of a single system while providing a compact and rugged solution. This feature is ideal for processing intensive test and measurement applications that have system footprint constraints and/or portability requirements.

The PXImc specification also allows for direct peer-to-peer communication between PXImc devices and other PCI(e) devices. Using this capability PXImc devices, such as peripheral processing modules, can directly communicate with PCI(e) devices without engaging the master controller. This feature can be used to create subsystems within a single PXI chassis and thus utilizes the chassis resources more efficiently.

Figure 6 illustrates this option to have peripheral processing modules within a single PXI chassis and also the two possible communication methods.

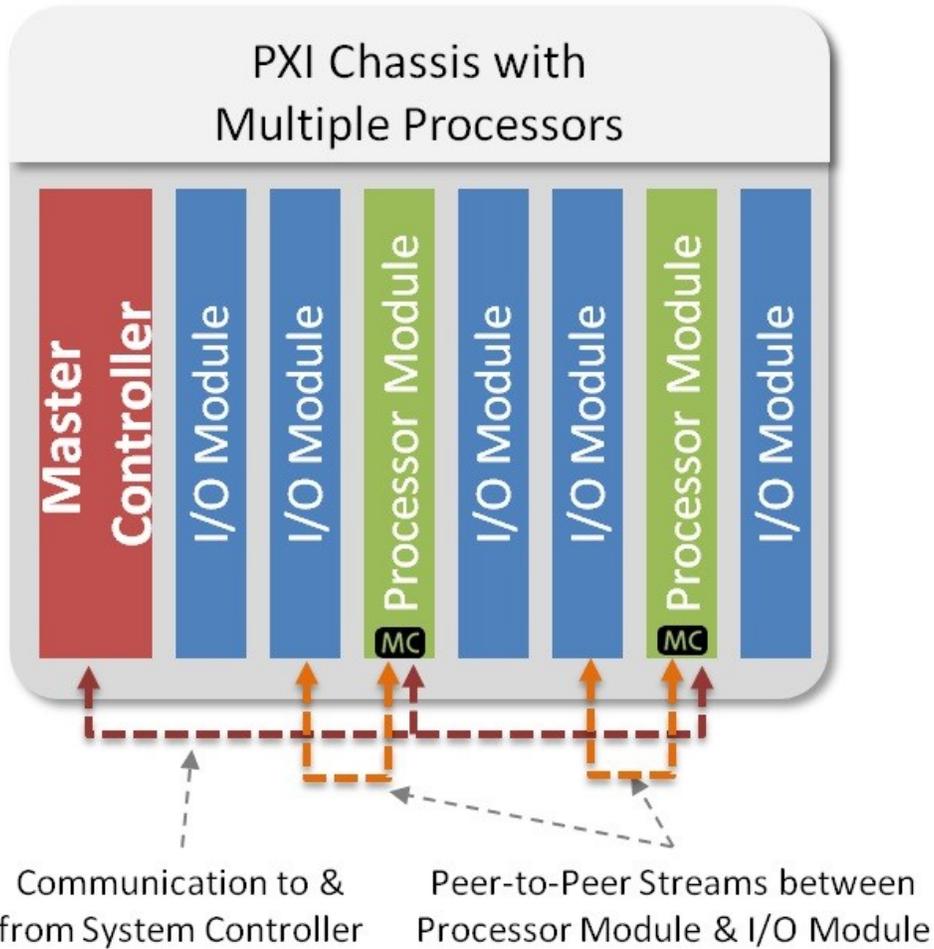


Figure 6: PXImc specification enable PXI systems to have multiple peripheral processing modules in a single chassis.

Conclusion

In conclusion, the PXImc specification, developed by the PXI Systems Alliance (PXISA), provides a vendor interoperable solution for connecting multiple intelligent systems together via a high-throughput and low-latency communication interface to support growing demands of application such as HIL and SIGINT.

PXImc expands the capability of the PXI platform and can be used create complex test and measurement systems as it supports various configurations such as hybrid systems, distributed processing and in-chassis co-processing.