



Synchronizing Mechatronic Systems in Real Time Using FPGAs and Industrial Ethernet Communications

Sari Germanos, Technology Marketing Manager, Ethernet POWERLINK Standardization Group
Wolfgang Seiss, Manager Open Automation Solutions, B&R Industrial Automation Corp.
Elias Ahmed, Sr. Strategic Marketing Manager, Altera

WP-01249-1.0

White Paper

Traditional centralized I/O systems are inadequate for building large machines. Running high-speed I/O lines over long distances increases cabling requirements and errors due to signal noise with higher frequencies. Industrial Ethernet solves many of these problems by providing an accelerated and deterministic methodology to manage I/O systems, sensors, and actuators for high-frequency applications such as motion control. The FPGAs provide a protocol for real-time deterministic jitter-free Ethernet communications, as well as reliable deterministic synchronization capabilities over large distances.

This paper will outline some of the major Industrial Ethernet communications protocols in use today, and then present an example implementation of an open source protocol (Ethernet POWERLINK) implemented on an FPGA.

Introduction

Ethernet, or IEEE 802.3, is a Layer 2 protocol in the Open Systems Interconnection (OSI) model of computer networking implemented at the media access controller (MAC) layer. Although Ethernet-based communication has now become increasingly popular for industrial applications, this has not always been the case. For much of the last 30 years, serial fieldbuses (real-time digital data communications protocols used for measurement and control in industrial systems) were the predominant technologies. Some examples of widely used fieldbus implementations are:

- PROFIBUS
- Modbus
- CAN
- FOUNDATION fieldbus
- ControlNet /DeviceNet
- CC-link
- Sercos

Over time, fieldbus technologies have reached their technological limits and are now being replaced by Industrial Ethernet. The main driving forces of the change are the needs for higher bandwidth and performance and the ability to connect to more network nodes at longer distances, and the requirement for interoperability between the factory floor equipment and the enterprise network running traditional Ethernet.



101 Innovation Drive
San Jose, CA 95134
www.altera.com

© 2015 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Furthermore, the ubiquity of IEEE 802.3 Ethernet devices has reduced component costs to the point that Ethernet in the factory is now a viable low-cost alternative to fieldbus systems.

However, standard Ethernet is not without its challenges, since it is based on a “shared” medium concept. Historically, this shared medium was the Ethernet cable, and the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) algorithm resulted in non-deterministic wait times that are unacceptable for many industrial applications such as motion control. Modern Ethernet equipment now employs full-duplex links and near-to-zero delay switches with queuing. Although an improvement, the non-deterministic factor is still present, as the unknown latency is now occurring within the switch queue rather than on the link.

The next sections outline the major Industrial Ethernet protocols in use today and the technologies they leverage to guarantee deterministic behavior. Finally, Ethernet POWERLINK, an open source protocol, is used as a working example of how such technologies are implemented.

How Does Industrial Ethernet Work?

Standard Ethernet is non-deterministic due to the nature of the underlying CSMA/CD algorithm. All industrial Ethernet protocols address this non-determinism by either adding mechanisms at a higher layer of the OSI model or by applying modifications to the Ethernet protocol and frame handling at OSI Layer-2. The Layer-2 modified protocols by definition require dedicated hardware such as ASICs, ASSPs, or FPGAs. Industrial Ethernet protocols working at Layers-3 and above can simply be implemented in software with any standard Ethernet MAC.

The major Industrial Ethernet protocols in use today are Modbus/TCP, PROFINET RT/IRT, EtherNet/IP, Ethernet POWERLINK, EtherCAT, Sercos III, and CC-Link. The next section briefly outlines the features of these protocols.

Modbus/TCP

Modbus was originally developed by Modicon (now Schneider Electric). Modbus is an application layer protocol that establishes master-slave and client-server communication between devices. Modbus/TCP simply takes the Modbus instruction set and wraps TCP/IP around it. It is essentially the Modbus RTU protocol with a TCP interface that runs on Ethernet.

PROFINET

PROFINET is the Industrial Ethernet version of the fieldbus technology referred to as PROFIBUS. Originally developed by Siemens and the member companies of the PROFIBUS user organization, PNO, PROFINET has two main performance classes:

- *PROFINET-RT for soft real-time*—PROFINET-RT uses a standard Ethernet MAC. However, the TCP/IP layers are bypassed for time critical data to provide deterministic performance. This is a software-based solution.
- *PROFINET-IRT for hard real-time performance*—PROFINET-IRT requires dedicated hardware (either ASIC or FPGA) for both, the master and slave. It ensures hard real-time performance by allocating a portion of each cycle to hard-real time.

EtherNet/IP

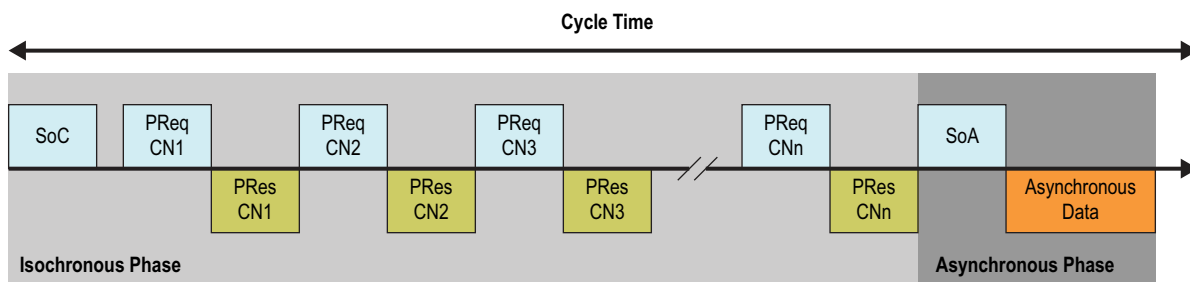
EtherNet/IP was developed by Rockwell Automation and is currently managed by the ODVA group. It uses the standard IEEE 802.3 MAC and runs on the upper layers of the OSI layer model (session, presentation, application layers). EtherNet/IP can therefore be implemented in software. It is based on the producer/consumer model and employs both TCP/IP and UDP/IP for data transfer.

Ethernet POWERLINK

POWERLINK is an Open Source protocol managed by the Ethernet POWERLINK Standardization Group (EPSG) and was originally introduced by B&R. POWERLINK achieves real-time performance via a mixed polling and time-slot procedure that only allows one node at a time to transmit data. The Managing Node (MN) controls the network access where it acts as the arbitrator of the connection to avoid collisions. Controlled Nodes (CN) only respond to requests from the MN. In one cycle, the MN polls each CN using a standard Ethernet frame, and each CN responds appropriately as shown in [Figure 1](#).

POWERLINK can be implemented purely in software as both master and slave. However, to reach maximum performance, time-critical tasks can be moved to dedicated hardware such as FPGAs.

Figure 1. A POWERLINK Cycle



EtherCAT


EtherCAT is an Ethernet-based fieldbus system originally invented by Beckhoff Automation and currently managed by the EtherCAT Technology Group (ETG). EtherCAT is a Layer-2 protocol that requires implementation in dedicated hardware on the slave side (ASIC or FPGA). A master may use any Ethernet MAC, since standard Ethernet frames are transmitted and received.

The EtherCAT frame that is sent from the master contains data for all network devices. Each EtherCAT slave selectively reads the data addressed to it from the frame and inserts data into the frame while it is passing through. As this processing is done "on-the-fly", a specialized Layer-2 implementation in hardware is required.

Sercos III

Sercos III is an open digital interface based on standard Ethernet, developed by Bosch Rexroth Group and currently promoted by the Sercos Group. Sercos works on a producer/consumer model where the master initiates and ends the data transmission during a real-time cycle. Figure 2 shows a typical Sercos III cycle where the time slots are divided into a real-time and a non-real-time channel. This time slot reservation provides bandwidth for critical data to ensure collision-free data transfer. Sercos III uses dedicated hardware for both the master and slave, although a software-only solution is viable for the master.

Figure 2. Sercos III Clock Cycle



CC-Link IE

Originally developed by Mitsubishi and currently managed by the CC-Link Partner Association (CLPA). CC-Link IE ensures real-time behavior through the use of a token-passing technique to minimize collisions. It is currently the only Industrial Ethernet protocol implemented with 1 Gbps line speed. The technology is based on dedicated hardware (ASIC or FPGA) that handles the complete data link layer and transport layer.

Ethernet POWERLINK

Industrial Ethernet adoption is expected to continuously increase in the upcoming years, and as outlined in the previous section, there are many different protocols available to meet the real-time Ethernet requirements for industrial automation. One of these protocols is Ethernet POWERLINK, which is described in more detail in this paper.

Ethernet POWERLINK is a patent-free, Open Source protocol that is very efficient and effective at handling both time-critical real-time data and high-throughput best-effort traffic at the same time. It combines the requirements for minimum system jitter as requested by high-precision motion control with demands for bandwidth used in industrial vision applications. All data is transferred on a single wire, typically following the Cat 5e or Cat7 cable standards.

How Ethernet POWERLINK Works

As shown earlier in Figure 1, the Ethernet POWERLINK cycle consists of the real-time isochronous phase and the non-real-time asynchronous phase. The cycle is initiated by the MN issuing a Start-of-Cycle (SoC) frame. This frame is used for network synchronization and indicates to each CN to sample its input data as well as to set active the latest received output data. In the isochronous phase, the MN sends

individual PollRequest (PReq) frames to each CN on the network. The addressed CN then responds with a PollResponse (PRes) frame. The response is sent as a multicast message, thus making the data available to all nodes in the network. This allows direct cross-communication between the CNs within a single cycle, which guarantees the fastest possible synchronization between nodes.

After the isochronous phase, the MN grants permission to send a standard Ethernet frame to any station on the network. This asynchronous phase is used for non-time-critical communication, such as network management, configuration, diagnose or web-traffic.

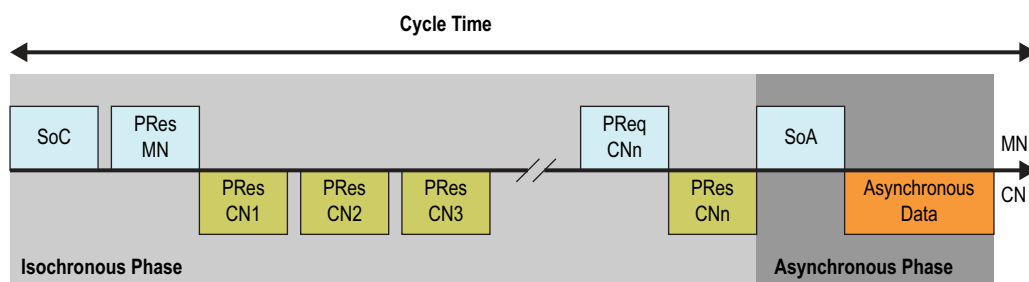
Multiplexing and PollResponse Chaining

An automation network typically contains a variety of components requiring high-frequency data updates such as drives and encoders, and devices requiring lower frequency updates such as sensors. Also, in some cases a centralized control approach is preferred, while in others a decentralized one fits best.

Ethernet POWERLINK offers a feature (Multiplexed Slot Assignment) allowing the master to poll slow stations every n th cycle. Another feature (Poll Response Chaining) combines all outgoing data from the MN into a single-frame request. The CNs will then send their responses at a specified point of time with fixed predetermined time delays. The latter mechanism is shown in [Figure 3](#).

PollResponse Chaining increases the network performance when several nodes with small amounts of process data are connected, because all CNs receive a single PRes frame from the MN. This frame is sent as multicast instead of all the individual PReq frames in standard mode. PollResponse Chaining is ideal for situations where centralized control loops are utilized in applications such as robotics or Computer Numerical Control (CNC).

Figure 3. POWERLINK PollResponse Chaining



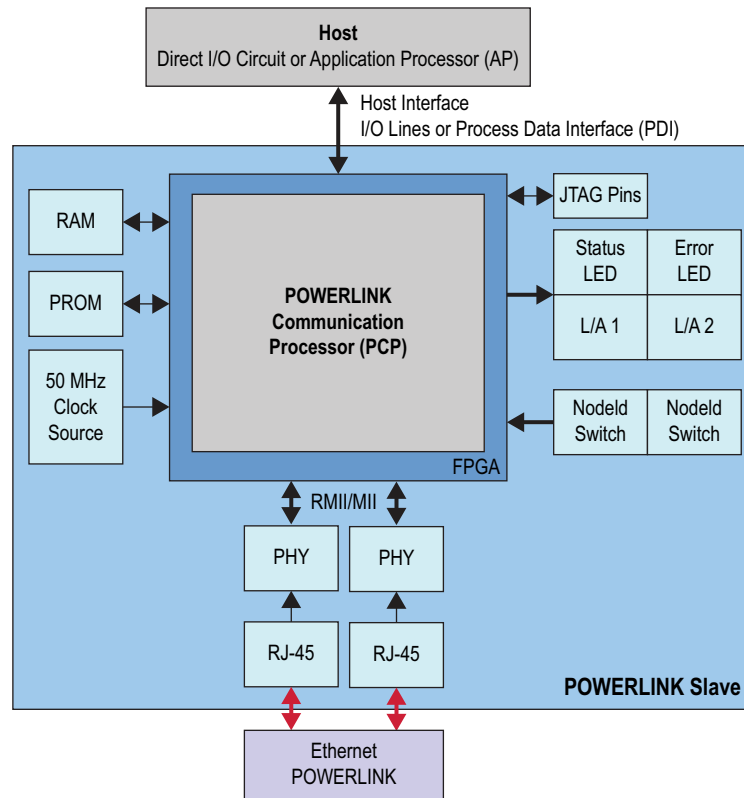
FPGA Implementation of POWERLINK

A typical POWERLINK slave implementation on an Altera® Cyclone®-family device is illustrated in [Figure 4](#). In this configuration, the POWERLINK slave consists of the following hardware components:

- Altera FPGA (e.g., Cyclone IV device)
- Clock source (quartz crystal oscillator)
- Non-volatile memory (PROM)
- Volatile memory (RAM)

- Ethernet physicals (PHYs)
- Ethernet jacks (RJ45) + transformer
- Diagnostic LEDs
- Node ID switches
- Debugging interface (JTAG)

Figure 4. Ethernet POWERLINK FPGA Block Architecture



The Ethernet POWERLINK protocol is implemented in the POWERLINK Communication Processor (PCP). The PCP internally parts into a Soft-IP-Core (POWERLINK IP-Core) and a software stack running on a standard Nios® II embedded microcontroller within the FPGA.

On the left hand side of the block diagram, the required hardware components for booting and running the PCP are listed. They consist of volatile memory (RAM) to run the software stack, non-volatile memory (PROM) to store the firmware, and a clock source.

On the right, hardware parts are shown that directly interact with the user such as the Node ID switches, the diagnostic LEDs (STATUS, ERROR, and PHY LINK/ACTIVITY) and optionally (for debugging purposes) a JTAG interface. At the bottom the connection with the physical network is shown, while the top represents the interface to the host application.

Ethernet POWERLINK Design Options

Figure 5 shows the following three different configurations for Ethernet POWERLINK slave:

- Direct I/O
- External host processor
- Internal host processor configurations

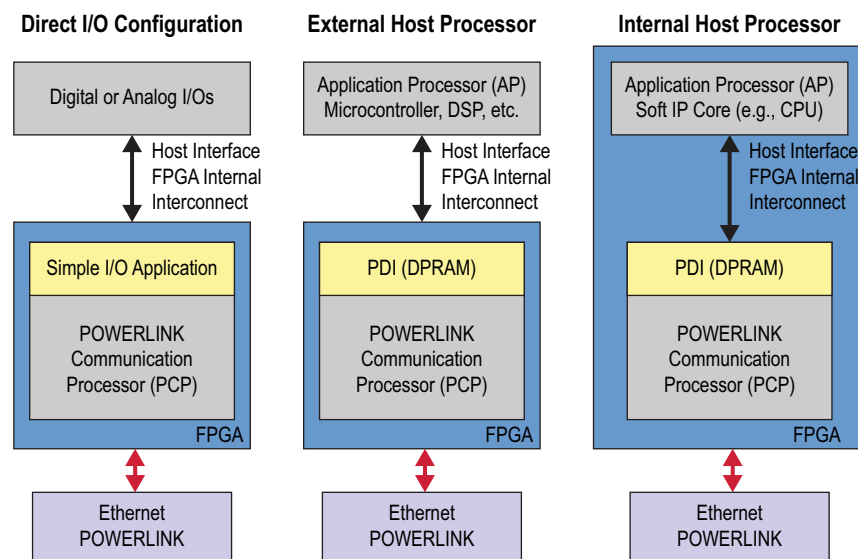
The direct I/O version consists of the PCP solely. It can be directly connected to up to 32 I/O lines. This configuration is used mainly for simple remote sensors interfacing with analog or digital signals.

Typical Ethernet POWERLINK devices require their own application processor on which the main device firmware is run. This application processor may be an external chip (micro controller, DSP) or reside inside the same chip (Dual Nios configuration, or the ARM® cores of a Cyclone V SoC).

The PCP is connected to the application processor via the process data interface (PDI). Three variants are available for the user to choose from:

- Serial peripheral interface (SPI) for simple applications
- An 8 or 16 bit parallel memory interface for faster throughput
- Altera's on-chip Avalon® memory mapped interface for single chip or SoC solutions

Figure 5. Ethernet POWERLINK Slave Configurations

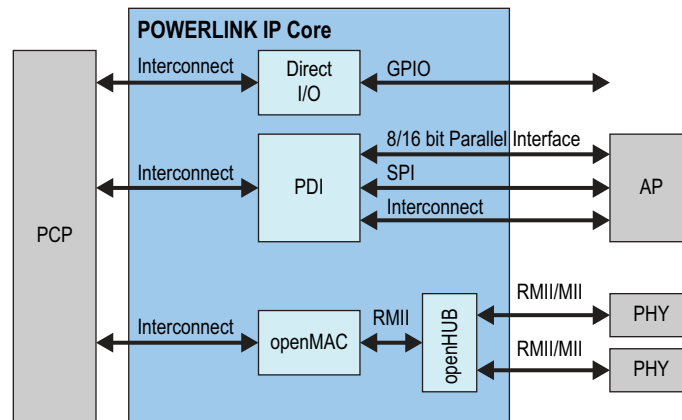


Benefits of the FPGA-Based Implementation

Although there are software implementations of Ethernet POWERLINK that allow running the protocol on standard microcontroller units (MCUs) and processors, the hardware implementation of Ethernet POWERLINK produces the optimal performance. The main reason for it is the Ethernet POWERLINK-aware MAC IP-Core called openMAC, which offers hardware acceleration functionality. An additional benefit of the FPGA implementation is the included openHUB, an Industrial Ethernet hub, that enables placing two (or more) external Ethernet jacks on a device to set up line or ring topologies in the POWERLINK network without a need for further infrastructure components. The PDI is part of the FPGA IP-core.

Note: the PDI and Direct I/O are mutually exclusive

Figure 6. POWERLINK IP-Core Block Diagram



Hardware Acceleration of the openMAC

To understand why hardware acceleration at the Ethernet MAC layer is key, a deeper look into the architecture is necessary.

A main factor for speeding up the automation processes is reducing the overall reaction time of the nodes in the network. In Ethernet POWERLINK, the reaction time of a node is determined by the gap between receiving the PReq from the MN and the start of transmission of the PRes frame: The shorter the gap, the faster the slave.

openMAC now allows to automatically transmit a pre-assembled PRes frame as soon as the PReq is received and recognized. As this is all done in the FPGA logic, the minimum possible gap of 960 ns (the so-called Inter Frame Gap as given by the Ethernet standard) is reached. This auto-response mechanism together with extended filtering on the MAC layer optimally accelerates POWERLINK slave implementations.

In addition, frame transmissions can be scheduled by openMAC (time-triggered sending), which allows Ethernet POWERLINK slaves to operate in PollResponse Chaining mode and Ethernet POWERLINK masters to reduce the system jitter to below 1 microsecond.

With these optimizations inside the FPGA, fastest cycle times and lowest jitter are possible, as shown in [Table 1](#).

Table 1. POWERLINK Performance Specifications in Altera FPGAs

Feature	Master	Slave
Interfaces	Parallel (16/32 bit) PCIe Avalon	SPI Parallel (16 bit) Avalon
Cycle time	250 μ s	250 μ s
Jitter	< 1 μ s	< 1 μ s
Response time	1 μ s	1 μ s
Supported modes	Standard Poll response chaining	Standard Multiplex Poll response chaining
Virtual Ethernet interface	Yes (for tunneling IP traffic)	Yes (for tunneling IP traffic)
Remote update	-	Yes (via POWERLINK master)
FPGA resource utilization for standard configuration	9K LEs 286 KB RAM	7K LEs 286 KB RAM
IP access details	openpowerlink.sourceforge.net	openpowerlink.sourceforge.net

Conclusion

The Ethernet POWERLINK implementation showed combined with FPGA technology make Industrial Ethernet an attractive solution for modern automation systems. The Ethernet POWERLINK protocol implemented on FPGAs allows maximum system flexibility and performance at a cost-effective solution point. With an FPGA, there are various options of implementing Ethernet POWERLINK. This is achieved as either a one-chip solution with the application processor embedded inside the FPGA, or the application processor may be implemented externally. Furthermore, by using FPGAs, it becomes possible to embed additional system functionality such as motor control or programmable logic controller (PLC) functionality into the FPGA to reduce the overall system bill of material (BOM) costs.

Document Revision History

[Table 2](#) shows the revision history for this document.

Table 2. Document Revision History

Date	Version	Changes
April 2015	1.0	Initial release.