

Latency in Factory Automation

Patrick O'Farrell, Aniruddha Khadye

ABSTRACT

Latency is a critical parameter in Ethernet networks developed for Factory Automation applications. Latency is not a defined value for Ethernet as specified by the IEEE 802.3 standard, nor is Ethernet inherently synchronous or repeatable. This disconnect between the inherent characteristics of Ethernet and the needs of Factory Automation applications must be bridged through careful architecture of Ethernet Physical Layer devices like the DP83867 and DP83869.

The latency through the Ethernet Physical Layer is a key limiting factor of control cycle time in these applications. The capability of the DP83867 and DP83869 to operate with less than 400ns latency in 1000Base-T operation can improve the cycle time. This low level of latency is also comparable to levels seen in 100Base-TX operation, thereby simplifying the process of upgrading a 100Base-TX network to 1000Base-T and increasing the network bandwidth.

Contents

1	Ethernet for Factory Automation	2
2	Factory Automation Background.....	2
3	Impact of Latency on Factory Automation Applications	3
4	1000Base-T Ethernet for Factory Automation	4
5	Design for Latency.....	4
6	Latency Configuration	7
7	Latency Data	8
8	Conclusions.....	9

List of Figures

1	Factory Automation Network	2
2	Ethernet Application Latencies.....	3
3	DP83867E Functional Block Diagram	5
4	DP83869HM Functional Block Diagram	6

List of Tables

1	RGMII Control Register (RGMIICTL), Address 0x0032	7
2	RGMII Control Register 2 (RGMIICTL2), Address 0x0033.....	7
3	100Base-TX Configuration Register (100CR), Address 0x0043	7
4	DP83867 Latency.....	8
5	DP83869 Latency.....	9

⁽¹⁾ All trademarks are the property of their respective owners.

1 Ethernet for Factory Automation

Ethernet is everywhere. The communication protocol defined by the IEEE 802.3 standard has been deployed for decades as a reliable, easy to use technology for Local Area Networks. While Ethernet may be traditionally thought of in an office environment connecting a computer to a printer, it has been adapted to new environments.

The reasons for this proliferation are straightforward. The technology is well understood and easily implemented. It supports different network topologies. Because it is standards based, interoperability amongst components is ensured.

Over time, Ethernet has evolved into new and unexpected applications. These new applications benefit from the many advantages of traditional Ethernet, but they also introduce challenges that are not directly addressed by the IEEE 802.3 standard. These new applications require device characteristics that are above and beyond the needs of a standard Ethernet implementation. This is especially true at the Ethernet Physical Layer, the portion of the application that interfaces with the physical media of communication, often a twisted pair or fiber optic cable.

As Ethernet has been adapted to industrial environments, it is important that Ethernet components adapt to new needs. An Ethernet Physical Layer device, commonly referred to as a PHY, must be rugged and reliable to thrive in the harsh environmental conditions. It must also meet the demands of real-time systems. The PHY should transmit and receive packets in a repeatable manner and minimize the time required to traverse the PHY.

2 Factory Automation Background

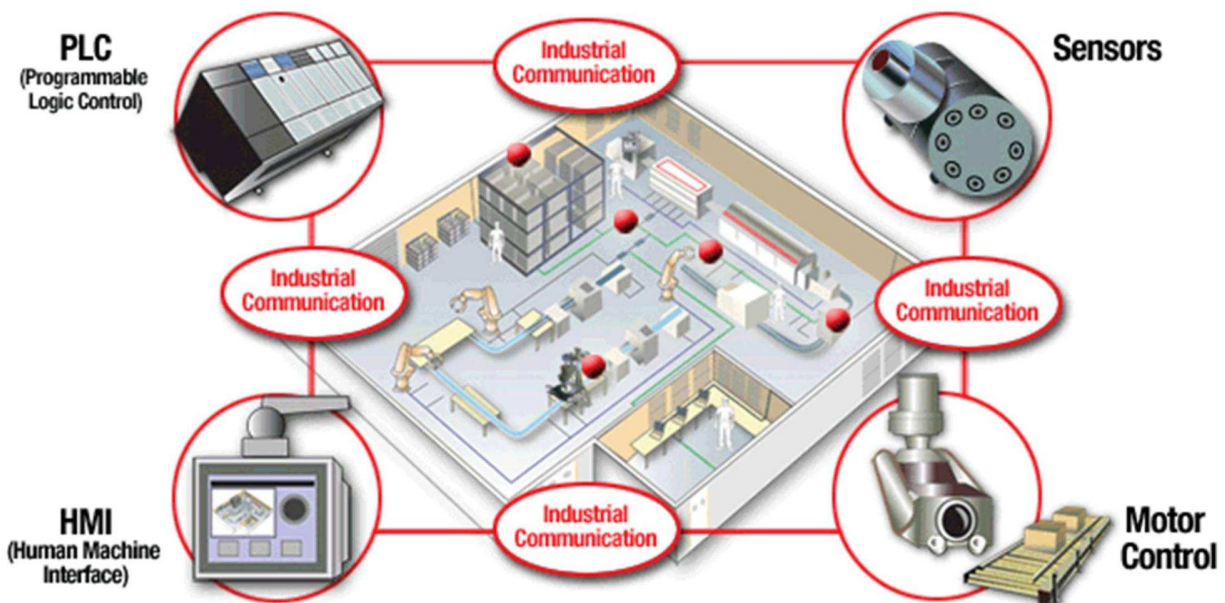


Figure 1. Factory Automation Network

There are multiple industrial protocols used in factory automation applications. Profinet, Ethernet/IP, EtherCAT, SERCOS III, and CC-Link are some prominent examples. Although the structure and function of the individual protocols varies, they are all based on Ethernet.

The topology of the automation network can vary. It may be a ring or a line topology implemented in a daisy chain. It may be a switch based star topology. The flexibility offered by Ethernet supports all of these options.

A factory automation application can include a number of different components. These components can be categorized as Programmable Logic Controllers (PLCs), Human Machine Interfaces (HMIs), Sensors, and Motor Drives. All of these components can be connected via an Ethernet backbone.

The network architecture and configuration for applications will vary based on the protocol. Regardless of the network topology or the industrial protocol, the protocols have a common goal of providing precise control to the different nodes on the factory floor. This can be accomplished by timestamping the transmitted and received packets and using those timestamps to align the network time across the nodes of the network.

The network time is shared by the protocol in the packet data. Timestamp units (TSUs) at each node mark the time. Variation in the timestamps will decrease the accuracy of the system. Therefore variations in the latency must be minimized.

3 Impact of Latency on Factory Automation Applications

The time required for packets to travel from its source to its destination is referred to as the latency. In an Ethernet system, there are many sources of latency. Some latency is due to physical media like cabling and PCB traces. Delays through the PHYs, MACs, switches, and other components in the path also contribute to the latency.

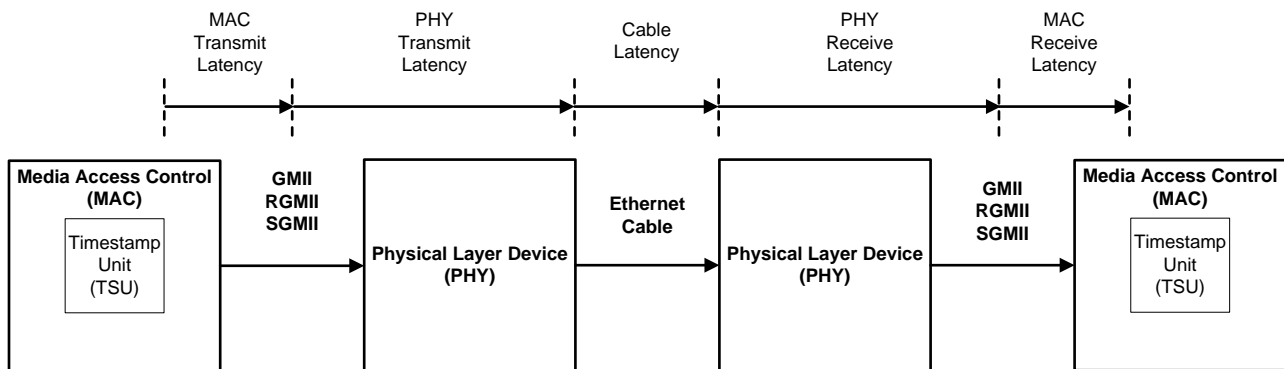


Figure 2. Ethernet Application Latencies

Traditional Ethernet is non-deterministic. The IEEE 802.3 standard does not impose specifications on the time for packets to traverse the PHY. However, latency is an important criterion for real-time Factory Automation applications.

When considering the latency through the PHY, it is important to note that the time to pass through the PHY often differs for transmitting and receiving packets. Therefore, it is necessary to consider both transmit and receive latency. It is also important to consider the variation in the latency through the PHY. Ideally, the latency time through the PHY would be both short and repeatable. This would commonly be known as low, deterministic latency.

An Ethernet packet includes a Start of Frame Delimiter (SFD) which identifies the start of the packet. This SFD can be used as the timestamp point for a packet. By timestamping the packet when it is transmitted and again when it is received, the protocol can accurately share time across the network despite the inherently non-deterministic nature of standard Ethernet. The precision of the timestamping will impact the accuracy of the time shared across the network and will therefore affect how well the equipment at each node of the network is controlled.

Regardless of protocol or topology, as packets are shared and timestamped across the network, the end to end latency can affect the precision of the network. At the system level, variation of the latency will be seen as jitter and will directly impact the precision of the network. To some degree, the fixed latency of the system can be accounted for, but lower overall latency is desirable because longer latencies can impose limits on the frequency with which packets can be timestamped or on the number of nodes allowable in a network.

4 1000Base-T Ethernet for Factory Automation

Standard Ethernet is a “best effort” technology. There is no guaranteed delivery. In a factory automation application, upper layer protocols must ensure delivery via structures like a protected window, allowing higher priority packets on the network during certain cycles and interspersing lower priority packets between these cycles. The high priority cycles may account for 10-20% of the network utilization.

The cycle time will depend on the network and the type of industrial automation components in the network. For a motion control application, where precise control is required, the cycle time could be measured in tens of microseconds. At these levels, it is critical to minimize the latency through each component in the network.

Most industrial Ethernet applications currently operate in 100Base-TX, but applications are beginning to move to 1000Base-T and the trend will continue in the years to come. A motor drive today may have over 1000 parameters describing its characteristics as well as a database and a web server. Often applications target limiting network utilization to 50%. Going above this is considered a high risk, especially for safety critical applications. 1000Base-T Ethernet provides increased bandwidth that is required to meet the needs of these additional services and diagnostics.

The encoding and the signaling for 1000Base-T differs from that of 100Base-TX. Accordingly, the transmit and receive paths for the two speeds inside the PHY differ. At the application level however, there are benefits to maintaining comparable latencies for the two modes of operation. A significantly longer latency in a 1000Base-T PHY would have the disadvantage of requiring longer cycle times relative to a 100Base-T application. By contrast, maintaining comparable latencies for a 1000Base-T PHY relative to a 100Base-T PHY provides the benefit of allowing an easier upgrade path to increased bandwidth without significant changes at the application level.

5 Design for Latency

An Ethernet PHY is a complex mixed signal component with many sub-systems and functional blocks. The DP83867E and DP83869HM functional block diagrams shown below provide a high level indication of this complexity.




Figure 3. DP83867E Functional Block Diagram




Figure 4. DP83869HM Functional Block Diagram

There are many system characteristics that can impact packet latency through a PHY. It is important to use a PHY that is designed for a deterministic delay. Moreover, this delay should be repeatable on multiple power ups and from link to link.

Latency can also change based on the speed of operation and the interface to the MAC. Consider the architectural differences of SGMII, RGMII, and GMII. The signal encoding of the interface, the existence of FIFOs or buffers and the need to cross clock domains can all negatively impact the latency.

Moreover, the Ethernet mode of operation can introduce variability in the latency. The data path may change for 1000Base-T versus 100Base-TX. If this is not considered as part of the PHY architectural design, the latency of a packet through the system can change significantly. The architecture of the DP83867 and DP83869 was designed to minimize variability in transmit and receive latencies through the device.

6 Latency Configuration

The DP83867 and DP83869 are architected to optimize the latency through the device. This process requires scrutinizing every sub-block and each clock cycle.

For the lowest latency applications, GMII operation is recommended. However, given the large number of connections required for GMII, it is not optimal for some applications. There are several configuration options available for RGMII to reduce latency to levels close to those achieved with GMII.

The key configuration bits for reducing latency are in register 0x0032, register 0x0033, and register 0x0043. These registers have different names in DP83867 and DP83869 datasheets. The register address and bits are same for both the PHYs. For DP83867, they are register 0x0032 (RGMIICTL), register 0x0033 (RGMIICTL2), and register 0x0043 (100CR). For DP83869, they are register 0x0032 (RGMII_CTL), register 0x0033 (RGMII_CTL2), and register 0x0043 (G_100BT_REG0). Excerpts of these registers are shown below:

Table 1. RGMII Control Register (RGMIICTL), Address 0x0032

BIT	BIT NAME	DEFAULT	DESCRIPTION
6:5	RGMII_RX_HALF_FULL_THR	10, RW	RGMII Receive FIFO Half Full Threshold: This field controls the RGMII receive FIFO half full threshold.
4:3	RGMII_TX_HALF_FULL_THR	10, RW	RGMII Transmit FIFO Half Full Threshold: This field controls the RGMII transmit FIFO half full threshold.

Table 2. RGMII Control Register 2 (RGMIICTL2), Address 0x0033

BIT	BIT NAME	DEFAULT	DESCRIPTION
4	RGMII_AF_BYPASS_EN	0, RW	RGMII Async FIFO Bypass Enable: 1 = Enable RGMII Async FIFO Bypass. 0 = Normal operation.
2	LOW_LATENCY_10_100_EN	0, RW	Low Latency 10/100 Enable: 1 = Enable low latency in 10/100 operation. 0 = Normal operation.

Table 3. 100Base-TX Configuration Register (100CR), Address 0x0043

BIT	BIT NAME	DEFAULT	DESCRIPTION
0	FAST_RX_DV	0, RW	Fast RX_DV Enable: 1 = Enable fast RX_DV. 0 = Normal operation.

Configuration of these bits must be performed after a power cycle, a hardware reset via the RESET_N pin, or a software reset via the Basic Mode Control Register (BMCR), Address 0x0000, RESET (bit 15) or the Control Register (CTRL), Address 0x001F, SW_RESET (bit 15).

6.1 RGMII_RX_HALF_FULL_THR Configuration

In normal operation, RGMII_RX_HALF_FULL_THR[1:0] can be changed from the default of 0x2 to 0x1. This will reduce the latency by one clock period.

6.2 RGMII_TX_HALF_FULL_THR Configuration

If the MAC and the PHY have the same source clock, RGMII_TX_HALF_FULL_THR[1:0] can be changed from the default of 0x2 to 0x1. This will reduce the latency by one clock period.

6.3 RGMII_AF_BYPASS_EN Configuration

Configuration of RGMII_AF_BYPASS_EN is only applicable to 10/100 operation. For systems that can operate in an “MII-like” mode where the PHY provides the receive clock to the MAC, configuration of RGMII_AF_BYPASS_EN bypasses the Receive FIFO, thereby reducing the latency.

6.4 LOW_LATENCY_10_100_EN Configuration

Configuration of LOW_LATENCY_10_100_EN is only applicable to 10/100 operation. When configured, the transmit latency is reduced by one clock cycle.

6.5 FAST_RX_DV Configuration

Fast RXDV mode can be enabled in 100Base-TX operation. This configuration reduces latency by asserting RX_DV when the /J/ symbol of the received packet is detected.

7 Latency Data

The latencies reported in [Table 4](#) reflect the transmit and receive latency as shown in [Figure 2](#). The transmit latency represents the time from assertion of the RGMII TX_CTRL or GMII TX_EN signal on the MAC interface to the time the packet arrives at the Ethernet cable. The receive latency represents the time from arrival of the packet from the Ethernet cable to the time the RGMII RX_CTRL or GMII RX_DV signal asserts on the MAC interface.

7.1 DP83867

Table 4. DP83867 Latency

Speed	MAC Interface	Transmit Latency (ns)	Receive Latency (ns)	Configuration
1000Base-T	RGMII	88	288	
100Base-TX	RGMII	200	272	
100Base-TX	RGMII	200	152	Enable RGMII_AF_BYPASS_EN, FAST_RX_DV, and LOW_LATENCY_10_100_EN
1000Base-T	GMII	72	264	
100Base-TX	MII	88	224	

7.2 DP83869

Table 5. DP83869 Latency

Speed	MAC Interface	Transmit Latency (ns)	Receive Latency (ns)	Configuration
1000Base-T	RGMII	96	288	
1000Base-X	RGMII	151	155	
100Base-TX	RGMII	248	374	
100Base-TX	RGMII	169	192	Enable RGMII_AF_BYPASS_EN, FAST_RX_DV, and LOW_LATENCY_10_100_EN
100Base-FX	RGMII	245	266	
100Base-FX	RGMII	171	185	Enable RGMII_AF_BYPASS_EN, FAST_RX_DV, and LOW_LATENCY_10_100_EN
100Base-TX	MII	64	220	
100Base-FX	MII	80	143	

8 Conclusions

This application note highlights the importance of latency in factory automation and describes some of the challenges faced in the design of an Ethernet Physical layer device to meet these challenges. The DP83867 and DP83869 have been designed to meet these challenges. They provide excellent latency in the default configuration and provide register configurable options for further improving the latency.

Revision History

NOTE: Page numbers for previous revisions may differ from page numbers in the current version.

Changes from Original (October 2015) to A Revision	Page
• Added references to DP83869	1
• Added DP83869 Functional Block Diagram	6
• Added latency data for DP83869.....	9

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL 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 AND 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 skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI 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. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to TI's Terms of Sale (www.ti.com/legal/termsofsale.html) or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2019, Texas Instruments Incorporated