

© 2004 Contemporary Control Systems, Inc.

Introduction to Real-Time Ethernet II

By Paula Doyle, a doctoral researcher with the Circuits and Systems Research Centre at the University of Limerick in Ireland

INTRODUCTION

In “Real-Time Ethernet I”, Part I, we introduced the basic concepts of Ethernet’s capacity to deliver a real-time communication system. “Real-Time Ethernet II” introduces some of the real-time solutions available to industry today: EtherNet/IP, PROFINet, EtherCAT and ETHERNET Powerlink. This article also provides an introduction to a single standard, IEEE 1588, that is growing in popularity amongst real-time Ethernet developers to provide sub-microsecond synchronization accuracy of distributed clocks over Ethernet.

Ethernet/IP

EtherNet/IP (EIP, where the IP stands for Industrial Protocol) is an open application-layer protocol developed and maintained by CI (ControlNet International), ODVA (Open DeviceNet Vendors Association) and IEA (Industrial Ethernet Association). It is built on the existing IEEE 802.3 physical/data layers and TCP/UDP/IP—giving it optimal interoperability with most information-level networks. EIP offers a real-time (RT) solution when using strict guidelines, but is not deterministic. It uses the open object-oriented CIP (Control and Information Protocol) as its application layer—the same layers 5–7 as both DeviceNet and ControlNet - yielding full interoperability with them (Figure 1).

CIP [1] is a flexible and scalable automation protocol well suited to distributed systems with properties such as: object-orientation, Electronic Data Sheets and device profiles. EIP with CIP is not an RT protocol. To

achieve RT for EIP, CIPSync (a high-speed CIP synchronization solution) is employed. CIPSync is based on IEEE 1588. Using 100 Mbps Switched Ethernet, CIPSync can deliver synchronization accuracy of better than 500 nanoseconds

between devices [2] although jitter introduced by the protocol stack will still be an issue.

EIP uses both TCP and UDP with IP for communication. When a connection-oriented exchange is preferred, e.g. at initialization, TCP is used (Explicit Messaging). Explicit Messaging contains protocol information and service information, but does not have strict timing requirements; therefore, it is sufficient to use the slower, yet guaranteed TCP protocol. For RT traffic, EIP uses the unicast and multicast capabilities of UDP to implement the producer/consumer model of communication—popular with control applications. Implicit messages contain no commands, only data. The meaning of this data is configured at initialization—reducing run-time processing in the nodes. [PD1] Typical traffic on an EIP network is cyclic (although CIP also specifies polled, change-of-state and strobed traffic). Network collisions are avoided by switches, and EIP is generally implemented in a star topology.

Unlike other RT solutions, EIP uses UDP/IP for RT communication, adding jitter and non-determinism. If this jitter is quantifiable and does not infract on the system model, the system can still be RT but will be unsuitable for fast and hard RT systems like motion control. Specific advice relating to implementing an RT version of EIP is offered in [3], and although EIP with CIP is not an RT protocol, the currently achievable end-to-end response time in an EIP control system of eight producers and one consumer was determined to be 7 ms—when implemented with the recommendations in the following paragraph.

One recommendation in [3] introduces VLANs and locates all devices sharing time-critical data in the same VLAN so RT multicast EIP traffic will not need to exit the VLAN. Another recommendation is to use the routing functionality of Layer 3 switches and set the TTL of IP multicasts to 1. This keeps RT EIP traffic in its subnet and keeps normal traffic out. Hence, according to [3], it is possible to achieve RT performance between EIP devices (using CIP) if the following guidelines are met:

1. Devices sharing RT information must co-exist in the same subnet.
2. The EIP segment must be isolated from the main network multicast traffic.

EIP use of CIP with commercial off-the-shelf components, including TCP/UDP/IP, for all devices will benefit certain customers while an EIP solution using

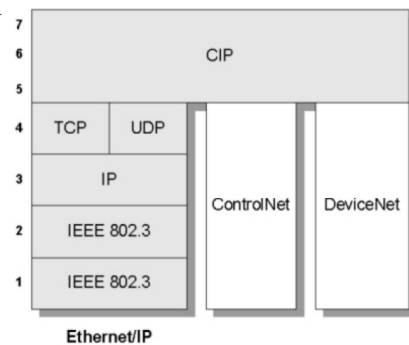


Figure 1—EtherNet/IP Stack

CIPSync will be more beneficial to those requiring sub-microsecond synchronization accuracy.

EIP can cater for many RT systems where the device count is limited, device synchronization is in the order of microseconds and determinism is not required.

EtherCAT

EtherCAT (Ethernet for Control Automation Technology) is the motion-control RT solution from Beckhoff. It can process 1000 I/Os in 30 μ s [4], but requires full-duplex. It can use copper or fibre optic cables. EtherCAT is based on the master/slave principal and can interoperate with normal TCP/IP-based networks and other Ethernet-based solutions such as EIP or PROFINet. It also supports any Ethernet topology, including the bus.

The EtherCAT master processes RT data via dedicated hardware and software (Beckhoff currently use their PC-based TwinCAT OS and TwinCAT Y driver). In the future, further variations will be introduced that will also provide the same guarantees. The master prioritizes EtherCAT frames over normal Ethernet traffic, which is transmitted in gaps. The master controls traffic by initiating all transmissions.

The telegrams are standard Ethernet, and the data field encapsulates the EtherCAT frame (an EtherCAT header and one or more EtherCAT commands). Each command contains a header, data and Working Counter (WC) field. Each Ethernet telegram can contain many EtherCAT commands - realising a higher bandwidth and more efficient use of the large Ethernet data field size and header (see Figure 2). The standard Ethernet CRC is used to verify message correctness.

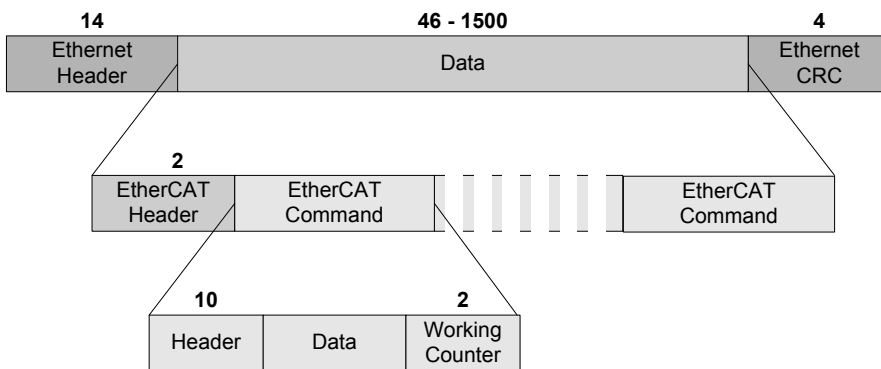


Figure 2—EtherCAT Encapsulation

The EtherCAT master fully controls its slaves. Its commands only elicit responses; slaves do not initiate transmissions. The two EtherCAT communication methods used are “Ether Type” or UDP/IP encapsulation.

The “Ether Type” uses the type field (defined in Ethernet II), which is more commonly known as the length field in IEEE 802.3. The Ether Type implementation does not use IP, thus limiting EtherCAT traffic to the originating subnet. Encapsulating commands using UDP/IP allows EtherCAT frames to traverse subnets, but has drawbacks. The UDP/IP header adds 28 (20: IP, 8:

UDP) bytes to the Ethernet frame and undermines RT performance through its non-deterministic stack.

EtherCAT slaves range from intelligent nodes to 2-bit I/O modules and are networked via 100Base-TX, fiber optic cable or E-bus (depending on distance requirements). E-bus is an EtherCAT physical layer for Ethernet offering a LVDS (Low Voltage Differential Signal) scheme. Slaves are hot pluggable in any topology of branches or stubs. Multiple “slave rings” can exist on a single network if connected by a switch, (see Figure 3).

EtherCAT slaves have integrated memory from 2 bits to 64 Kbytes. They appear to the Ethernet as a single device though actually comprising up to 65,535 devices. They are configured in an open-ring topology, with the Ethernet interface at the open end. Masters transmit commands to the MAC address of the first device. When the signal reaches the Ethernet/slave interface, it is converted to E-bus specifications (if E-bus is employed) and forwarded.

A slave receives a telegram, processes it (in hardware) then forwards it to the next slave on the ring. This processing delays the telegram by an order of nanoseconds. The last slave returns the completed telegram, via the ring, to the master. On the return route, each slave amplifies and regenerates the signal. Each slave has two Tx & Rx interfaces, so bi-directional communication occurs without contention.

In each EtherCAT command, the WC increments when a slave processes a command addressed to it, allowing the master to determine if each addressed slave is exchanging data, although correct data is not guaranteed.

The FMMU (Field Memory Management Unit) of each configurable slave converts a logical address to a physical one, and that information is available to the master at initialization. Thus, each slave needs a special ASIC (Application-Specific Integrated Circuit). On telegram reception, a slave determines if it is addressed, and then passes data to/from the telegram, incurring a delay of

some nanoseconds. EtherCAT is also internally synchronized by a distributed clock algorithm (a simplified version of IEEE 1588) although external synchronization is achievable with IEEE 1588.

EtherCAT is a fast RT Ethernet solution and deterministic if not used with UDP/IP or intermediate switches or routers between master and slaves.

ETHERNET Powerlink (EPL)

EPL is a hard RT protocol based on Fast Ethernet. Like EtherCAT, it uses the Ethernet II Frame type field. EPL devices use standard Ethernet hardware with no

special ASICs. EPL can deliver a cycle time of 200 μ s with jitter under 1 μ s. Its frame is encapsulated as illustrated in Figure 4.

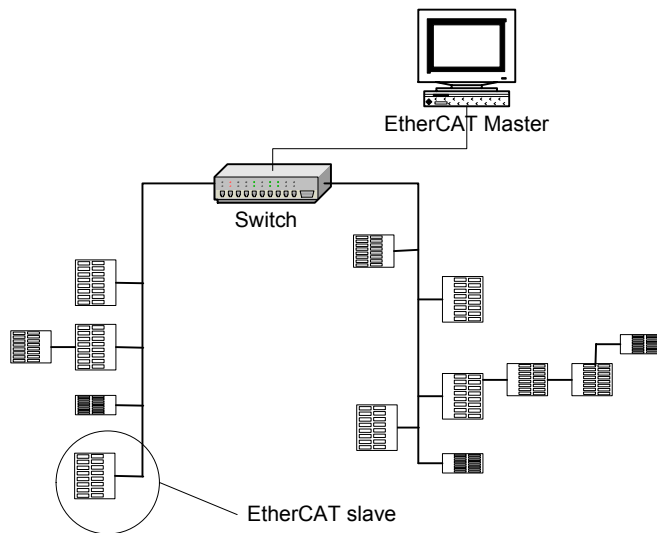


Figure 3—Sample EtherCAT Implementation

EPL uses cyclic communication with time-slot division and the master/slave model. One master (manager) is allowed per network. The master schedules all transmissions and is the only active station—slaves transmitting on request.




Figure 4—Powerlink Encapsulation

The four EPL cycle subdivisions are illustrated in Figure 5.

During the Start Period, the EPL master broadcasts the “Start-of-Cyclic” (SoC) frame which synchronizes the slaves. The timing of this frame provides the only time base for the network synchronization: all other frames are purely event-driven.

After transmitting the SoC frame, the Cyclic Period occurs as the manager polls each station with a “Poll Request” frame. Only then does the slave respond with a “Poll Response” frame containing data—hence, collisions are avoided. The slave broadcasts its response to all devices; thus, inter-slave communication can occur.

After successful polling of all slaves, the master broadcasts the “End-of-Cyclic” (EoC) frame, informing each slave that the cyclic traffic progressed correctly.

The Asynchronous Period allows non-cyclic data transfers under master control. To transmit during this period, a slave must have informed the master in its “Poll Response” during the Cyclic Period. The master builds a list of waiting slaves and employs a scheduler to guarantee that no send request will be delayed indefinitely. During the Asynchronous Period, standard IP datagrams can be transferred.

Unlike PROFINet, EPL does not employ switches to avoid collisions or to provide the network synchronization; the master controls this. EPL networks can be built with standard hubs. It is proposed that each device incorporate a hub for ease of bus implementation. Switches, although not prohibited, are not recommended for EPL because they add jitter and reduce determinism. Since the EPL network avoids collisions via time-controlled bus access, up to 10 hubs can be cascaded (an allowable exception to the 5-4-3 Ethernet rule).

Currently EPL devices demanding RT communication cannot co-exist on the same segment as non-RT Ethernet devices. However, EPL devices can operate as normal Ethernet devices. In Protected Mode, the RT segment must be separated from normal traffic by a bridge or router. In Open Mode, RT traffic shares the segment with normal traffic, but RT communication is compromised. In the next Powerlink version (V3), IEEE

1588 will be used to synchronize traffic across multiple RT segments—providing a more distributed EPL implementation, but true RT segments will still contain only EPL devices. Unlike PROFINet where normal Ethernet and RT devices can co-exist and not affect RT traffic, EPL must be protected from non-RT communication through bridges or routers. Unlike PROFINet or EtherCAT, which need special ASICs, EPL

employs standard Ethernet hardware.

PROFINet

PROFINet is a plant-wide fieldbus standard for distributed automation systems. It uses object-orientation and available IT standards (TCP/IP, Ethernet, XML, COM). PROFINet is also built on IEEE 802.3 and is interoperable with TCP/IP—allowing it to be implemented on existing Ethernets. It is compatible with PROFIBUS-DP.

PROFINet V1, has a response time of 10-100 ms. PROFINet-SRT (Soft Real-Time) allowed PROFINet to work with a factory automation cycle time of 5-10 ms, achieving RT solely in software. It uses TCP/IP and a dedicated software channel for RT communications. PROFINet-IRT brings a hard-RT element to the

PROFINet protocols. The three PROFINet protocols allow differing degrees of RT. PROFINet for hard RT is PROFINet-IRT.

incorporated into EIP, EPL, EtherCAT and PROFINet—making it a popular standard for delivering RT over Ethernet.




Figure 5—Powerlink Cycle

PROFINet-IRT

PROFINet IRT (Isochronous RT) was developed for systems requiring sub-microsecond synchronisation, typically high-performance motion control systems. The benchmark for such a system is 1 ms cycle time, 1 μ s jitter accuracy, and guaranteed determinism [6]—which IRT fulfills.

Since software introduces jitter above 1 μ s, IRT (unlike SRT) is a hardware solution with highly synchronized Ethernet nodes. Using full-duplex switched Fast Ethernet, it divides the communication cycle into a standard TCP/IP open channel and a deterministic RT channel (see Figure 6). The channel ratio is system-dependent and is chosen by the system engineer.




Figure 6—PROFINet-IRT Channel Division

Each PROFINet-IRT device has a special ASIC for handling node synchronization and cycle subdivision and incorporates an intelligent 2 or 4 port switch.

The PROFINet switch in every node is highly synchronized, contains a schedule of bus access and can deal with RT and non-RT traffic. It prioritizes RT traffic and provides full-duplex links for all ports. Contemporary switches (even cut-through) add jitter that would impact on determinism. PROFINet switches minimise jitter to where it has a negligible effect. The PROFINet communication model allows both RT and non-RT traffic to co-exist on one network without additional precautions.

By 2005, PROFINet-IRT and SRT will incorporate PROFISafe, the PROFINet safety solution for manufacturing and processing industries.

PROFINet, of all the solutions discussed here offers the greatest determinism—and since this is built into the PROFINet-IRT device, the systems engineer is spared from the burden of configuration to guarantee RT communication.

IEEE 1588

IEEE 1588 [7] specifies “A protocol to synchronize independent clocks running on separate nodes of a distributed measurement or control system to a high accuracy and precision.” IEEE 1588 is, or will be,

In IEEE 1588, all network nodes down to the transducer level contain an IEEE 1588 clock, synchronised with all network peers (see Figure 7) using Precision Time Protocol (PTP).

At device level, sensors can timestamp their data locally and actuators can operate at a precise time, avoiding stack and application delays between transducer and controller. The accuracy of the system depends on the synchronisation of local RT clocks.

IEEE 1588 defines two separate types of clocks: ordinary and boundary. Boundary clocks (BC) are employed in devices such as hubs or switches—where more than one PTP communication path (port) exists. Ordinary clocks exist in devices having a single port—e.g., normal network devices. Each BC port can act as a master or ordinary clock in its own segment.

PTP is for networks that support multicasting but keep multicasts within a subnet and where each local clock fulfils exacting requirements. The grandmaster clock (GMC) is the best clock in the system—with the best inherent stability, accuracy, resolution, etc. defined by the standard [8]. The Best Master Clock Algorithm (BMC), run by every live node, determines clock quality. Within each subnet, the

BMC determines the master clock; in a single-subnet system the master is the GMC.

The GMC determines system synchronisation; system clocks synchronise their subnet clocks to the system. There is only one GMC per system, and only one master clock per subnet.

Synchronization is performed as follows. All masters periodically broadcast “Sync” messages containing an estimate of the time the message will physically leave the master. The precise receipt time of these messages is noted at the slaves. The precise sending time of the message is noted at the grandmaster. All precise timing measurements are performed as close to the physical layer as possible—to eliminate the delays from the network stack and operating system—while the estimated times are calculated by the IEEE 1588 code at the Application Layer (see Figure 8). Following the Sync message, the master transmits a related “Follow_Up” message containing the precise sending time of the Sync message. A slave uses the transmission and reception times to calculate its offset and can initiate synchronization with the delay measurement, which is not periodic and not performed as often as the protocol synchronisation. Sync messages do not propagate beyond their originating subnet.

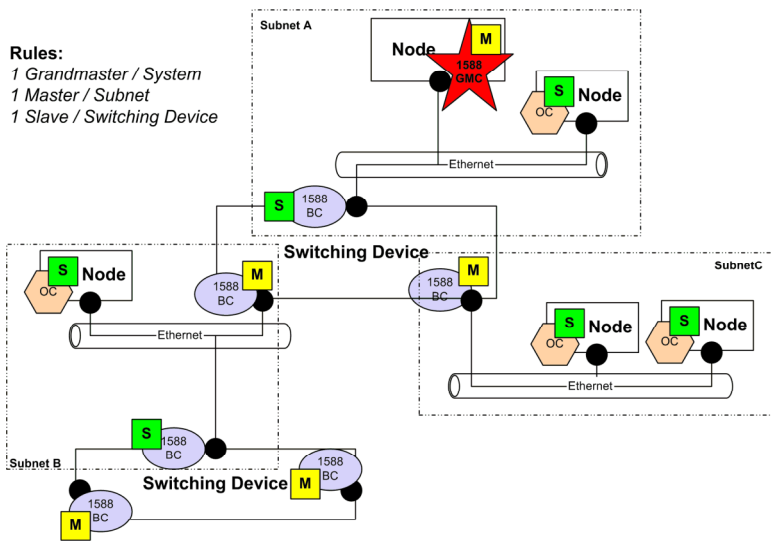


Figure 7—IEEE 1588 Configuration

synchronization accuracy, with IEEE 1588, along with an RT protocol can provide an Ethernet capable of delivering hard, fast and deterministic RT for applications such as motion control, while other solutions cater for softer applications. Therefore, when choosing RT-Ethernet, it is vital to consider the real-time, interoperability and flexibility requirements of your system along with all possible solutions before making a commitment.

The resolution of the system clock is the resolution of the GMC. If required, the GMC can be synchronised to an external source such as GPS.

IEEE 1588 is a highly precise system for synchronizing distributed nodes for applications such as motion control and robotics. It was designed for multicasting networks but with the popularity of Industrial Ethernet, Annex D was included for an Ethernet implementation of PTP. Although IEEE 1588 does not alter Ethernet or make it more deterministic or reliable, it does provide a method for other protocols to do so. A highly synchronized system of distributed nodes—coupled with an application for handling resolution and controlling traffic—could deliver hard, deterministic RT over Ethernet.




Figure 8—IEEE 1588 Node Timing

Real-Time Ethernet is a fast-growing, exciting development of the Ethernet protocol. The ability to have RT control segments running on the same network as office applications brings many new possibilities for industrial applications. With different protocols offering different levels of RT service, it is vital to understand the RT requirement of the system before choosing a solution. Sub-microsecond

REFERENCES

1. Schiffer, V.; "The CIP Family of Fieldbus Protocols and its Newest Member - Ethernet/IP", Emerging Technologies and Factory Automation, 2001. Proceedings. 2001 8th IEEE International Conference on, 15-18 Oct. 2001, Pages: 377-384 vol.1.
2. "ODVA Adds Time Synchronization Services for Real-Time Control to EtherNet/IP™ and DeviceNet™" - available at <http://www.odva.org/index.htm>.
3. Moldovansky, A., "Utilization of Modern Switching Technology in EtherNet/IP Networks," Proc. 1st International Workshop on Real-Time LANs in the Internet Age (RTLIA'2002) in conjunction with the 14th Euromicro Conference on Real-Time Systems, Vienna, Austria, June 18, 2002.
4. Beckhoff EtherCAT available at <http://www.beckhoff.com/english/default.htm?ethercat/default.htm>.
5. Popp, M.; Wenzel, P. "PROFINet-linking worlds," Emerging Technologies and Factory Automation, 2001. Proceedings. 2001 8th IEEE International Conference on , Volume: 2 , 15-18 Oct. 2001, Page(s): 519-522 vol.2.
6. Baumann, G. "Solving the Compatible Real-Time Ethernet Conundrum," The OnLine Industrial Ethernet Book, Issue 16, September 2003.
7. IEEE Std. 1588 -2002, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems."
8. Eidson, J. et al. "Synchronizing Measurement and Control Systems." Sensors, November 2002.

Look for Paula Doyle's complete article as Elective IE402 in the Curriculum of the virtual Industrial Ethernet University at www.industrialethernetu.com. If you wish to contact Paula, her e-mail address is: Paula.Doyle@ul.ie.

Paula Doyle is a doctoral researcher with the Circuits and Systems Research Centre at the University of Limerick in Ireland. She has a first class honors B. Engineering degree in Computer Engineering from the University of Limerick. Her research interests include embedded real-time systems, with emphasis on control networks for smart transducers with applications in the field of industrial automation. Paula is currently in the final stages of her Ph.D. research, based on the time-triggered transducer clusters in an Ethernet networking environment.

