

the EXTENSION

A Technical Supplement to Control Network

© 2004 Contemporary Control Systems, Inc.

Introduction to Real-Time Ethernet 1

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

INTRODUCTION

Real-time electronic distributed control systems are an important development of the technological evolution. Electronics are employed to control and monitor most safety-critical applications from flight decks to hospital operating rooms. As these real-time systems become increasingly prevalent and advanced, so does the demand to physically distribute the control in strict real-time. Thus, there is a need for control network protocols to support stringent real-time requirements. Real-time networks must provide a guarantee of service so they will consistently operate deterministically and correctly.

Ethernet, as defined in IEEE 802.3, is non-deterministic and thus, is unsuitable for hard real-time applications. The media access control protocol, CSMA/CD with its backoff algorithm, prevents the network from supporting hard real-time communication as due to its random delays and potential transmission failures.

Decreasing costs and increasing demand for a single network type, from boardroom to plant-floor, have led to the development of Industrial Ethernet. The desire to incorporate a real-time element into this increasingly popular single-network solution has led to the development of different real-time Industrial Ethernet strategies. Fieldbus networking standards have failed to deliver an integrated solution. Typically, emerging real-time Industrial Ethernet solutions complement the fieldbus standards—for example, by using common user layers. This article covers an introduction to real-time systems and a study of Ethernet's potential as a real-time network. Part 2 provides a study of the real-time Industrial Ethernet solutions available today.

Real-Time Introduction

Real-Time (RT) systems are becoming increasingly important, as industries focus on distributed computing in automation, see Figure 1. As computing costs decrease, and computing power increases, industry has relied more on distributed computers to deliver efficiency and increased yield to production lines. RT does not automatically mean faster execution, but rather that a process is dependent on the progression of time for valid execution.

RT systems are those that depend not solely on the validity of data but also on its timeliness. A correct RT system will guarantee successful system operation—so

far as its timely execution is concerned. RT systems are generally broken into two main sub-categories: hard and soft.

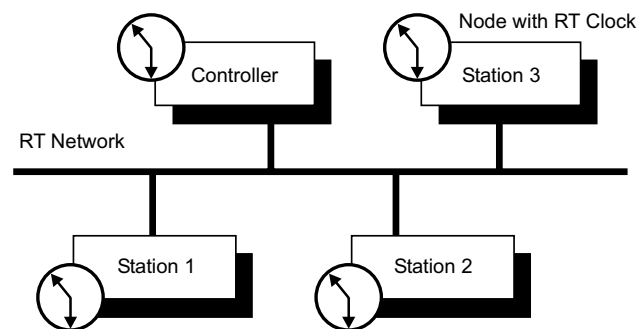


Figure 1—Distributed Real-Time Processing

Hard Real-Time (HRT) systems are those in which incorrect operation can lead to catastrophic events. Errors in HRT systems can cause accidents or even death, such as in flight control or train control.

Soft Real-Time (SRT) systems, on the other hand, are not as brittle. An error in a SRT system, while not encouraged, will not cause loss of property or life. SRT systems are not as safety-critical as HRT systems, and should not be employed in a safety-critical situation. Examples of SRT systems are online reservation systems and streaming multimedia applications where occasional delays are inconvenient but not serious.

Jobs are the RT system's building blocks. Each RT job has certain temporal quantities (Figure 2):

1. Release Time,
2. Ready Time,
3. Execution Time,
4. Response Time,
5. Deadline.

Release Time of a job is when a job becomes available to the system. **Execution Time** is the time for a job to be completely processed. **Response Time** is the interval between release time and completion of the execution. **Ready Time** is the earliest time a job can start executing (never less than the Release Time). The **Deadline** is the time by which execution must be finished, beyond which the job is late. A deadline can be either hard or soft, indicating the job's temporal dependence. As mentioned earlier, a missed hard deadline can have serious consequences. All RT systems have a certain level of **jitter** (a variance on actual timing). In a RT system, jitter should be measurable so

system performance can be guaranteed. For textbook information on RT systems, refer to [1].




Figure 2

To develop a RT distributed system of interconnected computers, it is vital to provide communication among the various computers in a reliable and timely fashion. Distributed processors running RT applications must be able to intercommunicate via a RT protocol, otherwise the temporal quality of work is lost. **Real-Time Communication** networks are like any RT system. They can be hard or soft, depending on requirements and their 'jobs' include message transmission, propagation, and reception. A number of RT control networks are employed in industry, but none have the popularity or bandwidth of Ethernet.

The Demand for Real-Time Ethernet

Demand for Ethernet as a RT control network is increasing as manufacturers realize the benefits of employing a single network technology from boardroom to plant floor. Decreased product costs plus the possibility of overlapping training and maintenance costs for information, field level, control and possibly device networks would greatly reduce manufacturers' expense.

At the RT control level, Ethernet offers many benefits over existing solutions. As a control network, 10 Gbps Ethernet offers bandwidth almost 1000x faster than today's comparable fieldbus networks (such as the 12 Mbps of ProfiBus) and can also support RT communication. Distributed applications in control environments require tight synchronization so to guarantee the delivery of control messages within defined message cycle times (typical times appear in Table 1). Traditional Ethernet and fieldbus systems cannot meet cycle time requirements below a few milliseconds, but emerging RT Industrial Ethernet solutions allow cycle times of a few microseconds.

Along with the increased bandwidth and tight synchronization, RT Ethernet gives manufacturers the security of using a physical and data-link layer technology that has been standardized by both the IEEE and the ISO. Ethernet can provide reduced complexity with all the attributes required of a field, control or device network—in operations having up to 30 different networks installed at this level [2]. Furthermore, Ethernet devices can also support TCP/IP stacks so that Ethernet can easily gate to the Internet. This feature is attractive to users since it allows remote diagnostics,

control, and observation of their plant network from any Internet-connected device around the world with a license-free web browser. Although Ethernet does introduce overhead through its minimum message data size (46 bytes), which is large in comparison to existing control network standards, its increased bandwidth, standardization and integration with existing plant technology should generate good reasons to consider Ethernet as a control network solution.

Ethernet and CSMA/CD

Ethernet, as defined in IEEE 802.3, is unsuitable for strict RT industrial applications because its communication is non-deterministic. This is due to the definition of its media access control (MAC) protocol, based on Carrier Sense Multiple Access/ Collision Detection (CSMA/CD). The implementation described in the standard uses a truncated binary exponential backoff algorithm.

With CSMA/CD, each node can detect if another node is transmitting on the medium (Carrier Sense). When a node detects a carrier, its Carrier Sense is turned on and it will defer transmission until determining the medium is free. If two nodes transmit simultaneously (Multiple Access), a collision occurs and all frames are destroyed. Nodes detect collisions (Collision Detection) by monitoring the **collisionDetect** signal provided by the physical layer. When a collision occurs, the node transmits a **jam sequence**.

When a node begins transmission there is a time interval, called the Collision Window, during which a collision can occur. This window is large enough to allow the signal to propagate around the entire network/segment. When this window is over, all (functioning) nodes should have their Carrier Sense on, and so would not attempt to commence transmission.

When a collision occurs, the truncated binary exponential backoff algorithm is employed at each 'colliding' node. The algorithm works as follows:

Initially: $n=0$, $k=0$, $r=0$.

When a collision occurs, the node enters the algorithm which states:

- It increments **n**, the Transmit Counter, which tracks the number of sequential collisions experienced by a node.
- If **n** > 16, (16 unsuccessful successive transmission attempts), transmission fails and the higher layers should be informed.
- If **n** ≤ 16, select a number from the set **k** = min(**n**, 10) (Truncation).
- A random number, **r**, is selected from the set (0,1,2,4...2^k) (Exponential and Binary).
- The node then waits **r** x slot_time before recommencing a transmission attempt.

One advantage of this backoff algorithm is that it controls the medium load. If the medium is heavily loaded, the likelihood of collisions increases, and the algorithm increases the interval from which the random delay time is chosen. This should lighten the load and reduce further collisions.

Ethernet introduces the possibility of complete transmission failure and the possibility of random transmission time, hence IEEE 802.3's non-determinism and unsuitability for RT communications—especially on heavily-loaded networks. Re-definition of the media access protocol could solve the problem but would leave the new nodes unable to interoperate with legacy Ethernet nodes.

Control Application	Typical Cycle Time
Low speed sensors (e.g. temp. pressure)	Tens of milliseconds
Drive control systems	Milliseconds
Motion control (e.g. robotics)	Hundreds of microseconds
Precision motion control	Tens of microseconds
High-speed devices	Microseconds
Electronic ranging (e.g. fault detection)	Microseconds

Table 1—Typical Cycle Times for Control Applications

Ethernet is non-deterministic only if collisions can occur. To implement a completely deterministic Ethernet, all collisions must be avoided. A collision domain is a CSMA/CD segment where simultaneous transmissions can produce a collision and where collision probability increases with the number of nodes transmitting. Completely avoiding collisions, therefore, gives Ethernet—with a relatively large bandwidth and popularity—an opportunity to be developed for the RT domain. The most common way of implementing a collision-free Ethernet is through the use of hardware.

A situation where two or more nodes compete for medium access to a network segment is called Shared Ethernet.

Full & Half Duplex Ethernet

When Ethernet was standardized in 1985, all communication was half duplex. Half duplex restricts a node to either transmit or receive at a time but not perform both actions concurrently.

Nodes sharing a half duplex connection are in the same collision domain. This means these nodes will compete for bus access, and their frames may collide with other frames on the network. Unless access to the bus is controlled at a higher level and highly synchronized across all nodes on the collision domain, collisions can occur and RT communication is not guaranteed.

Full duplex communication was standardized for Ethernet in the 1997 edition of 802.3 as IEEE 802.3x. With full duplex, a node can transmit and receive simultaneously, but only two devices can be connected on a single full duplex link—typically a node-to-switch or switch-to-switch configuration. Theoretically, full duplex links can double available network bandwidth from 10 to 20 Mbps or 100 to 200 Mbps, but in practice it is limited by the internal processing capability of each node. Full duplex provides every node with a

unique collision domain—completely avoiding collisions and even ignoring the traditional CSMA/CD protocol.

Since full duplex links have a maximum of two nodes per link, such technology is not viable as an industrial RT solution without the use of fast, intelligent switches that can form a network with single collision domains for each node - i.e., Switched Ethernet.

Full Duplex, Switched Ethernet

With Shared Ethernet, nodes compete for access to the media in their shared collision domains. The use of a non-deterministic bus access scheme, such as CSMA/CD, makes shared Ethernet unviable as a RT network solution. The most popular method of collision-avoidance, which produces a deterministic Ethernet, introduces single collision domains for each node guaranteeing the node exclusive use of the medium and eliminating access contention. This is achieved by implementing full duplex links and hardware devices such as switches and bridges. These devices can isolate collision domains through the segmentation of the network because each device port is configured as a single collision domain. Although, both switches and bridges will suffice, switches are more flexible because they allow more segments.

Switches

Switches are data-link layer hardware devices that permit single-collision domains through network segmentation. While a bridge operates like a switch, it only contains two ports compared to switches that have more than two—with each port connected to a collision domain. Switches can operate in half duplex or full duplex mode. When full duplex switches are used with full duplex capable nodes, no segment will have collisions. Today's switches are more intelligent and faster and with careful design and implementation could be used to achieve a hard RT communication network using IEEE 802.3.

Although switches are data-link layer devices, they can perform switching functions based on data from layers 3 and 4. Layer 3 switches can operate on information provided by IP - such as IP version, source/destination address or type of service. Layer 4 devices can switch by source/destination port or even information from the higher-level application.

Further refinements to the IEEE 802 standards, specifically for switch operations, are 802.1p and 802.1Q. IEEE 802.1p (incorporated into IEEE 802.1D [3]) brings Quality of Service (QoS) to the MAC level and defines how these switches deal with prioritization — priority determination, queue management, etc. This is achieved by adding a 3-bit priority field to the MAC header, giving 8 (0-7) different priority levels for use by switches or hubs. As defined, 802.1p supports priorities on topologies compatible with its prioritization service,

but for Ethernet, which has no prioritization field in its frame format, it uses 802.1Q.

IEEE 802.1Q [4] defines an architecture for virtual bridged LANs, their services and the protocols and algorithms used by those services. 802.1Q allows Ethernet frames to support VLANs (Virtual Local Area Networks)—limiting broadcast domains and thereby reducing broadcast traffic on the entire LAN. This is achieved by inserting 4 bytes between the source address and length/type fields in the frame header, which among other identifiers, includes that of the originating VLAN.

For a RT Industrial Ethernet application, an 802.1p/Q implementation has certain advantages: it introduces standardized prioritization, allowing control engineers up to eight different user-defined priority levels for their traffic. But these standards also have drawbacks including the extra hardware costs for the increased frame length (1522 bytes)—which introduces compatibility issues with legacy Ethernet networks. A RT implementation using 802.1 p/Q requires full duplex, switched Ethernet. IEEE 802.1p/Q are acceptable for certain applications of RT Ethernet in industry when switch ‘through’ time is predictable and an overload situation will not result in hard deadlines being missed.

Although switches can certainly provide RT deterministic Ethernet communication and are the backbone of the Industrial Ethernet solutions available today, they have drawbacks. They are costly—a major influence on cost-conscious industries. They are powered devices capable of failure (a major factor for hard RT control operations). And sometimes the operational predictability is not guaranteed by the manufacturer. A study on switches for RT applications is available at [5].

TCP/UDP/IP for Real-Time Ethernet

With Industrial Ethernet, the trend is to define an application-layer environment along with the TCP/IP protocol, to realize an industrial automation networking solution. Some RT Ethernet solutions (e.g., EtherNet/IP) perform all their communication, RT included, through the TCP/UDP/IP stack. But most solutions, while providing TCP/IP compatibility, do not employ this protocol for RT communication. In a system like EtherNet/IP, TCP is used for initialization and configuration of explicit messages while UDP, with its reduced overhead, is used for RT I/O (implicit messaging).

Typically, RT Industrial Ethernet applications are compatible with TCP/IP, but the protocol suite is bypassed for all RT communication. The ability of a RT

Ethernet solution to intercommunicate with an office-based system is paramount to achieve the Ethernet-technology plant of the future.

CONCLUSION

This article covered a broad introduction to Ethernet for RT. It described concepts to be developed in the follow-up article: Real-Time Ethernet II. The follow-up article will provide detail on existing RT Ethernet solutions such as PROFINet V3, ETHERNET Powerlink and EtherNet/IP plus IEEE 1588—an important supporting technology for clock synchronization.

References

- [1] Liu, J. Real-time Systems; Prentice Hall Inc., 2000; ISBN: 0130996513.
- [2] Crockett, N. “Industrial Ethernet is Ready to Revolutionize the Factory - Connecting the Factory Floor”, Manufacturing Engineer, Volume: 82, Issue: 3, June 2003, pages 41-44.
- [3] ISO/IEC 15802-3: 1998, ANSI/IEEE Std 802.1D, 1998 Edition. Information technology - telecommunications and information exchange between systems - local and metropolitan area networks - common specifications. Part 3: Media Access Control (MAC) bridges.
- [4] IEEE Std 802.1Q-1998 IEEE standards for local and metropolitan area networks Virtual bridged local area networks.
- [5] Georges, J.-P.; Rondeau, E.; Divoux, T “Evaluation of Switched Ethernet in an Industrial Context by Using the Network Calculus”, Factory Communication Systems, 2002. 4th IEEE International Workshop on 28-30 Aug. 2002, pages 19-26.

Look for Paula Doyle’s complete article as Elective IE401 in the Curriculum of Industrial Ethernet University at www.industrialethernetu.com.

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.



CONTEMPORARY CONTROLS
www.ccontrols.com
www.ctrlink.com

Past issues of the Extension are available. If you would like a copy, please send your request to info@ccontrols.com