

“Modern Industrial Ethernet Switches Can Be Deterministic & Real-Time Responsive”

Steve C. Tang, Sr. Staff Engineer, Victor K. Liang, V.P. Marketing, TC Communications Inc.

Abstract – Ethernet based Local Area Networks (LANs) have become the most prominent technology that promises to revolutionize automation communications. The advancements in Ethernet technology, from collision-prone, shared 10Mbps LAN, to collision-free, duplex 100Mbps switched LAN, and onward to 1000Mbps LAN, Ethernet technology is close to taking over the world of communications.

Up until now, most long haul or inter-remote communications are provided by SONET or SDH communication infrastructure. These legacy systems, although reliable and well understood, are voice-centric, less flexible, difficult to scale, and most of all, are not optimized for packet data transmissions.

With the maturity of Voice over IP technology and products, combined with the advancement in switching technology and long haul optics, Ethernet-based multi-service access platforms have become much better and cheaper alternatives to the SONET/SDH systems for automation and control applications.

However, as the size and range of the automation communication network increases, it is expected to carry a lot more diverse traffic e.g. SCADA information, office data, voice and video packets. As a result, the concern over data integrity and security has increased.

This paper looks at the new capabilities and functionalities of today’s Ethernet switches, and the special features and tools available on the JumboSwitch, to effectively eliminate the risks of impact on real time performances of mission-critical traffic by less-critical data. This paper will also address the common concerns over data reliability issues associated with normal Ethernet technology.

I. Prelude

When Ethernet was standardized in 1985, all communication was half duplex. The nodes connected to the Ethernet bus were restricted to either sending or receiving at any given time, but not both concurrently. All nodes sharing a bus are within a collision domain. That means these nodes will have to compete for access to the same bus. The frames have the potential to collide

with other frames on the bus. Ethernet at this level could not be used for real time applications due to this un-deterministic nature.

With 802.3x standardized in 1998, full duplex communication was finally possible. In essence, full duplex communication enables a node to transmit and receive simultaneously. For a network of 2 nodes, full duplex communication basically eliminated the collision domain completely. But, this didn’t solve the problem if a bigger network were to be built with more than 2 nodes.

The advancement in large scale integrated circuits, and switching chip technology came to the rescue at this time. With switched Ethernet, each link is a domain by itself, and with full duplex communication, the collision phenomenon is totally eliminated. Now, when switches are connected with copper or fiber, for short range or long haul, there is no need to worry about the distances between them any more. And there is no need to calculate the time of arrival for avoiding possible collisions, because the use of full duplex transmission eliminates worries about the collision domains.

With the earlier collision-prone and non-real time technology out of the way, full duplex, switched Ethernet opened up the possibility of using Ethernet for real-time critical applications, such as tele-protections, motion controls, or voice over IP. However, two main concerns remained when deploying Ethernet for real time applications: 1) the un-deterministic nature of Ethernet, which does not guarantee the real time response to a system request; and, 2) the potential corruption of data, or loss of packets, due to large and sudden increase of traffic. Special interest is on the impact of non-critical traffic, such as office intranet or Internet, to more real-time critical data traffic, such as SCADA or control systems.

II. Defining Real-Time & Determinism

In a modern computing and control dictionary, Real-Time System means “a system which responds to an input immediately.” Using this definition, it also implies that “any system that uses a shared transmission medium cannot be a real-time system.” While there may be an arguably quantitative definition for the word “immediately,” general perception is that TDM systems like SONET/SDH are Real-Time and Ethernet systems

are not. Let's not make a final judgment until the following sections on "Determinism" is explained and understood.

"Determinism," is commonly recognized as "the ability of a system to respond with a consistent and predictable time delay between input and response." This blurs the distinction between SONET/SDH systems, or TDM systems alike vs. Ethernet systems. Let's take a look at the following example:

The Cycle Time between the Master's "Poll" and the "Response" of the 30th RTU in a fiber based multi-dropped RS-485 bus system running at 1Mbps, in the following Fig. 1, can be calculated:




Fig.1 RS-485 Multi-Drop Bus Network

Broadcasting a polling message of 640Bytes from the Master to the Slaves, assuming negligible fiber transmission delays between Master and the 1st Slave, and the 1st Slave to the 2nd, 3rd, all the way to the 30th Slave, will take $640 \times 8 \times 1\mu\text{S} = 5,120\mu\text{S}$. Assuming the processing time for each addressed Slave is $5\mu\text{S}$, and the Response is also 640Bytes. The total Cycle Time for this TDM like set up is $5,120\mu\text{S} + 5\mu\text{S} + 5,120\mu\text{S} = 10,249\mu\text{S}$, roughly equivalent to 10mS.

If the same SCADA communication is to be set up through 10BaseT or 10BaseF Ethernet switches, as in Fig. 2, the cycle time can be calculated in the following.

Since the switches operate on a store-and-forward mechanism, each switch is introducing a $5\mu\text{S}$ processing latency to the cycle time. The data transmission between switches can be calculated as: $640 \times 8 \times 0.1\mu\text{S} + 5\mu\text{S} = 517\mu\text{S}$ per link. For the Master to send a polling message to the 30th switch, it'll take $517\mu\text{S} \times 29$ links = $14,993\mu\text{S}$ one way, and $29,986\mu\text{S}$ round trip.

Now, let's assume the switches are replaced with Fast Ethernet switches operating at 100Mbps, or 100Base, with the same latency of $5\mu\text{S}$ per switch. The cycle time is now: $((640 \times 8 \times 0.01\mu\text{S}) + 5\mu\text{S}) \times 29 \times 2 = 3,260\mu\text{S}$.

If replacing the 10Base switches with GigE switches, the cycle time would be reduced to around $587\mu\text{S}$,

which is 0.59mS.




Fig.2 10Base Ethernet Multi-Drop Bus Network

While 10mS cycle time is considered "Deterministic," in a RS-485 network which is a TDM like system, 0.59mS in a Gigabit Ethernet system is also nonetheless "Deterministic." Of course, faster Ethernet technology has the speed advantage to make "latency" no longer an issue, but it will take other measures or standards to make that "latency" a constant expectation.

III. Real Time Ethernet

There has been a lot of standardization efforts to make Ethernet technology closer to real-time performance, and the control systems running on top more responsive. For example, Ethernet/IP, or EIP (where IP stands for Industrial Protocol), is an open application layer protocol developed by ControlNet International and IEA. It is based on the existing IEEE802.3 physical/data layers and TCP/UDP/IP; PROFINet uses object orientation and available IT standards, TCP/IP, Ethernet, XML, COM. It is also based on IEEE802.3 and is interoperable with TCP/IP. PROFINet V1 has a response time of 10 to 100mS, while PROFINet-SRT allows PROFINet to operate with a factory automation cycle time of 5 to 10mS, achieving this kind of real time response solely by software.

We will focus on the requirements of the Ethernet switch that are essential in supporting such a real-time control system. Over the past few years, it is recognized by the industry that the following key specifications are needed behind the predictable, real-time Ethernet technology:


1. 100Mbps or above port speed
2. Full duplex (IEEE802.3x)
3. Priority queues (IEEE802.1p)
4. Virtual LAN (802.1q)
5. Loss of link management
6. Rapid Spanning Tree Protocol (IEEE802.1w)
7. Back-pressure and flow control (IEEE802.3x)
8. IGMP snooping and multicast filtering
9. Fiber optic port interfaces

10. Remote monitoring, port mirroring, and diagnostics
11. Extended temperature range & EMI hardened

100Mbps port speed and full duplex switches are becoming the basic components of a modern day Ethernet switch, thanks to the fast advancement of chip technology. Some of the chips on the market also come with pre-fabricated fiber optical drivers for interfacing with the optical transceivers directly. The store-and-forward mechanism inside the switch basically keeps all the incoming pipelines open while making it possible to schedule higher priority outgoing traffic. This switching chip design not only created a collision-free environment, but also helped increase the throughput of a switch tremendously. Although the port speed may be 100Mbps, the internal fabrics are designed to support the switching of all ports simultaneously; therefore, the bandwidth of a switching chip could be in the range of multi-gigabits per second.

Priority queuing and VLAN – Because of the increasing expectation and the flexibility of an Ethernet network, it is carrying vastly different types of traffic. This is best illustrated by the picture of an Ethernet network that is carrying mission-critical commands from a Central Control to trip protective relays in a power substation system, regular SCADA type traffic, VoIP traffic, CCTV traffic, and the intranet/Internet office LAN traffic. It is obvious that different messages have different delivery time requirements, making it necessary to separate the incoming traffic into different priority queues.

Priority queuing mechanism is defined in IEEE802.1p standard, and is accomplished by inserting a 4-byte tag into the Ethernet frame header.



A Virtual LAN is a mechanism for partitioning the network into multiple “virtual” domains. Messages will only be sent to devices within the same “sub-net”, therefore, forming different “pipes” within the same network. The boundaries of the pipes are defined by the “sub-net” addresses. For example, same type of incoming traffic can be assigned by a switch port with a pre-defined sub-net, and the messages will be processed, forwarded within that same group of devices, without the possibility of mixing up with the data with a different sub-net or domain.

This mechanism can be thought of as analogous to the Virtual Concatenation or VC feature in SONET or SDH system, whereby different time slots can be allocated to form pipes for processing different traffics,

with different bandwidth. There is a minute difference, however, in that the original motivation for standardizing VCs on SONET and SDH was for bandwidth amplification and management especially designed for handling data traffic; whereas, the VLAN creation was for the segregation of packets due to different characteristics of the different traffic types, and for the different routing and priority treatment of these packets or frames.

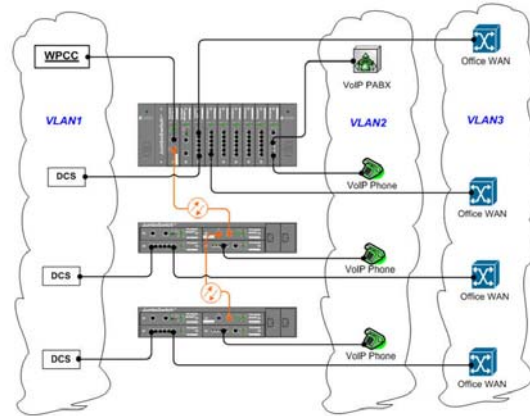


Fig. 3 VLAN Concept

VLAN and QoS have been two of the most important standards for Ethernet technology to become more practical for the deployment in real-time control applications. However, there is still an unsettled concern from most people that are familiar with SONET or SDH systems. That is, with Virtual Concatenation or VC, a network administrator can allocate definitive bandwidth for each data pipe for a certain type of traffic. In other words, the sizes of the data pipes in SONET or SDH systems can be pre-defined and managed by system administrator so that less important traffic will not impact the performance of more important traffic, such as mission-critical messages, no matter how sporadic or busy the less important traffic is. There is always a reserved capacity for more important traffic.

VLAN and QoS, though effective in traffic segregation, prioritization, and thus minimizing latencies, do not provide a mechanism in the standards for defining the “sizes” of the pipes.

IV. Controlling the Pipe Size – JumboSwitch Way

Separation and prioritization of data traffic in modern day layer 2 managed switches has been properly defined and implemented by manufacturers through IEEE802.1p and 1q, namely VLAN and QoS. But there is still no guarantee that, in the case of sudden and drastic increases of non-critical data traffic, the throughput for mission-critical traffic will be maintained or left un-impacted.

Port & Slot Level Rate Limiting - The only reasonable way to control the “sizes” of the data pipes of an Ethernet switch is to design a capability so that the

number of frames per second, or the rate of the throughput of certain type of traffic can be controlled at the ingress level of the connected port. JumboSwitch, a GigE fiber switch from TC Communications, allows network administrators to control this throughput “rates” both at the port level and at the aggregate level of an interface card slot, as shown in Fig. 4.




Fig. 4 JumboSwitch on Rate Control

Let’s look at a case where there are VLANs established for DCS Control traffic, VoIP traffic, Office LAN traffic, and CCTV traffic. The network administrator is able to limit the total throughput of an Ethernet switch card at the slot level, let’s say for Office LAN traffic to be at 20Mbps, VoIP traffic at 20Mbps, and CCTV traffic at 100Mbps. This will effectively limit the traffic of each type to the specific rate when sending the packets/frames to the central switching fabrics on the Main card. If finer granularity is required, the network administrator is also able to limit the individual ingress traffic at the port level of the Ethernet card for 1Mbps, 2Mbps, etc.

If all the non-critical traffic is being set to operate within certain bandwidth limits, this essentially leaves the remaining bandwidth in the Gigabit main pipe for the most critical messages: the DCS Control traffic. This network design will almost certainly guarantee the immediate transmission and processing of higher priority data packets when they are received by the JumboSwitch.

Virtual TDM Tunnel – Another special JumboSwitch feature is the capability of forming a virtual pipe, for instance sending real-time critical serial data, through the maze of layer 2 managed switches for point-to-point connections. In many power utilities, the communication channels to connect the protective relays are specified with stringent real time response requirements, e.g. half of the AC cycle time, which is 8.7mS for a 60Hz power system. None of the existing conventional serial server technology products will be able to meet this requirement. Special considerations need to be done to form a connectionless virtual tunnel inside the switching mechanism to guarantee the response performance.

V. Data Reliability

Whenever Ethernet/IP technology is used to transport isochronous signals, such as T1, E1 or T3, E3; and analogue signals, such as voice or video, there are always issues like jitters and wanders that need to be investigated carefully. Jitters and wanders are the results of packet loss, arrival delays and/or packets out of sequence. In the context of our discussion for the real time automation control systems, arrival delays and packet out-of-sequences impact the system with minor real time penalties. But the possibility of packet loss has the most impact on the data reliability aspect of the system.

Exploring deeper into the cause of packet loss or drop is really a necessary self-protection mechanism of the Ethernet switch when stretched to its limit or over its capacity. The solution to this potential mishap is really in the traffic engineering and management of the switched Ethernet network through the tools that are available today. The JumboSwitch brings good news to this situation. Empirical data has shown that it is possible to load a network of 40 JumboSwitches with 95% simulated packet traffic, no jitters nor wanders, that equates to no packet loss when observed on a virtual channel for the TDM traffic continuously for 120 hours.

V. Conclusion

Real time Ethernet and control systems have made tremendous strides over the past few years. Standards organizations and manufacturers alike have spent countless money and effort to make Ethernet/IP technology closer to real time automation applications than ever. This is all because of the indisputable money saving economics demonstrated by Ethernet technology in all industries.

Although the fundamental bus sharing principle of Ethernet has not changed, the advancement in fast and faster Ethernet, 100Mbps, 1Gbps, or 10Gbps, and the duplex operation have helped resolve a significant portion of the problems caused by bus-sharing. Further improvements in standards on the layer 2, 3 and 4, have made Ethernet deterministic. Individual manufacturers may have different implementation methods to add to the progress for Ethernet to become truly a part of the real time system.

Today’s layer 2 and layer 3 switches possess enough capability and functionality to be used in almost any mission-critical applications. With proper engineering and management, switched Ethernet networks can be deterministic and real-time responsive.

VI. References

- [1] Liu, J. Real-time Systems; Prentice Hall Inc., 2000
- [2] Skendzic, V., Moore, R. Extending the Substation LAN Beyond Substation Boundaries
- [3] Chow, F., Priority & QoS – Switch Level Operations
- [4] Tang, S., JumboSwitch Real Time Capabilities
- [5] Taylor, P., Real Time Determinism and Ethernet