PTP Enabled Network for Flight Test Data Acquisition and Recording

Hung Mach
Boeing Commercial Airplane Group
P.O. Box 3707
Seattle, Washington (98124)
hung.k.mach@boeing.com

Øyvind Holmeide
Westermo OnTime AS
Glads Vei 20
N-0489 Oslo, Norway
oeyvind@ontimenet.com

Evan Grim
Southwest Research Institute
6220 Culebra Road
San Antonio, Texas, USA (78238)
evan.grim@swri.org

Chris Calley
Symmetricom, Inc.
3750 Westwind Boulevard
Santa Rosa, California, USA (95403)
ccalley@symmetricom.com

Abstract

Large-scale data acquisition and recording systems have long sought to benefit from the bandwidth, scalability, and low-cost of Ethernet and Internet Protocol (IP). However, these systems' requirement for reliable correlation of data with time is impeded by Ethernet's inherently non-deterministic transit delay. With the advent of Precision Time Protocol (PTP), these challenges can now be overcome by deploying synchronized data sources that timestamp data at the source. Furthermore, data producers and consumers constitute a multicast data distribution model, where a single data source is observable by any interested subscribers. This paper details our work for Boeing's 787 which deployed these technologies to build an innovative system capable of providing gigabit data throughput with sub-microsecond synchronization.

1. Introduction

To support the 787 and future airplane testing platforms Boeing's Commercial Airplanes Flight Test Department set out to develop a next-generation Flight Test Data System that would offer a highly flexible, scalable, and cost-effective platform capable of fulfilling Boeing's immediate and foreseeable flight test verification and certification needs (see Figure 1).

This paper outlines the findings discovered during the on-going development of this system over the past two years, and specifically how the application of IEEE 1588 enabled and shaped the design process. Familiarity with the general system architecture and requirements is crucial to understanding the process detailed in this work. So in order to aide in this, a simplified but realistic system deployment can be seen in Figure 2, depicting data acquisition modules placed throughout the aircraft to acquire data from a variety of data sources. Data sources can range from simple transducers to complete avionics buses, and are fed into the data acquisition modules where they are time stamped using the module's synchronized clock. These modules transmit the data over a mixed 100 Mbps and gigabit Ethernet network to a recorder and real-time data processing systems. Data is transmitted using User Datagram Protocol/Internet Protocol (UDP/IP) multicast and the network fabric can support up to 400 multicast addresses. The system is monitored using custom software employing Simple Network Management Protocol (SNMP) to configure, monitor and control all system components. The recorder captures all system data by subscribing to predefined IP multicast streams and is capable of recording at data rates up to 500 Mbps.

All data acquisition modules carry a synchronized

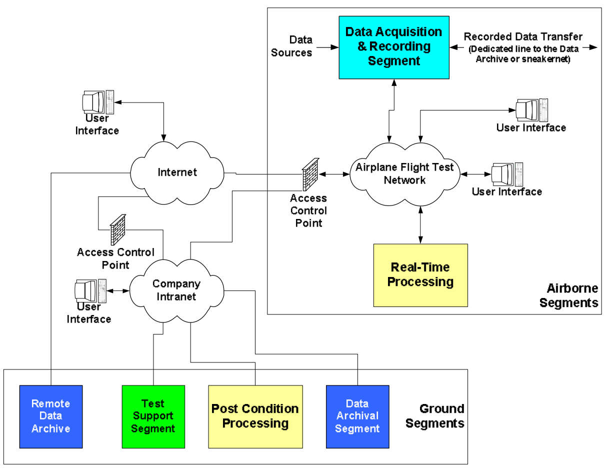

Figure 1 – Flight Test Data System

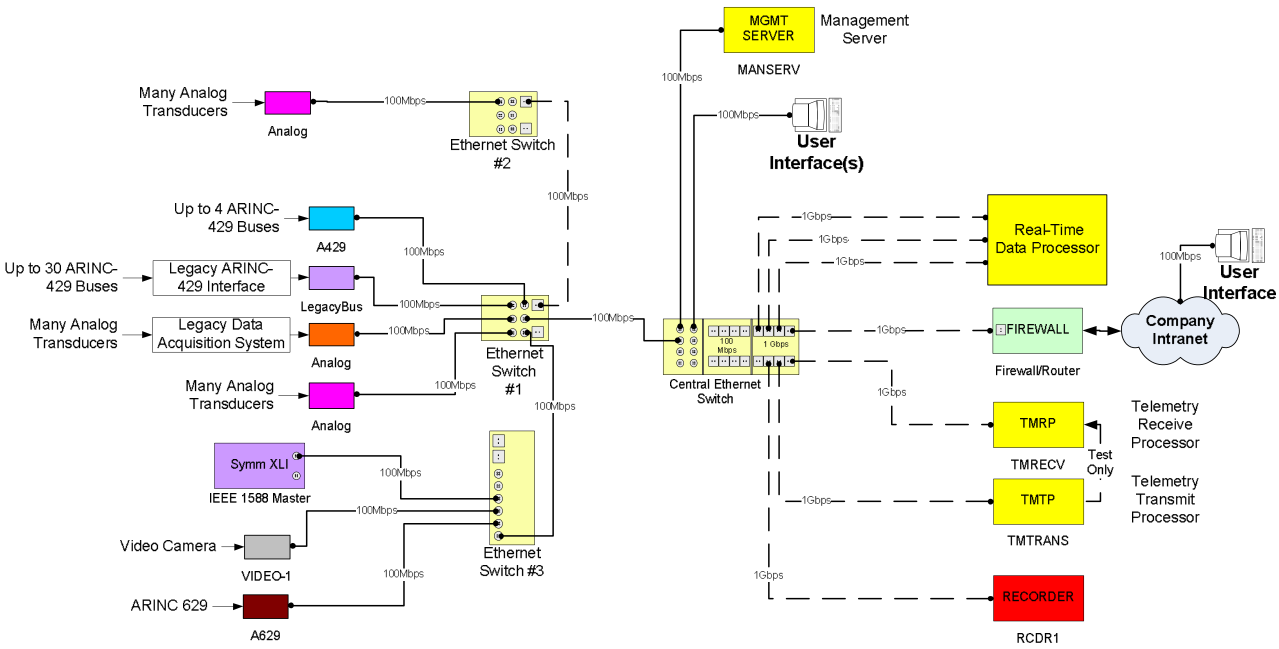

Figure 2 – Flight Test Network

time-of-day clock that is maintained over the network using the IEEE Standard. 1588-2002 Precision Time Protocol (PTP) and PTP aware Ethernet switches that further enhance the system's synchronization precision. Data is tagged at the source, allowing the data to be useful regardless of network transit time. The real-time data processing system selects a few hundred parameters from the multicast messages; time orders them using the time tags applied at the source and does the requested processing using the selected data.

At the completion of a test, the data is transported from the recorder to a data archive by either removing the recording media from the recorder and physically transporting it to the archive or by connecting a gigabit Ethernet network to the airplane (while grounded) and transmitting the data to a ground-based data archival system. In the archive, the data is stored in a RAID and copied to magnetic tape in a tape robot. Data requests from the post-test data processing system are sent to the archive using the company Intranet. The Data Archive then extracts the requested data along with the time tags, places them in a Hierarchical Data Format 5 (HDF5) and sends it back to the data processing system. The time tags are then used to correlate the measurements that were acquired.

2. Motivation

The use of Ethernet and associated network protocols as a transport media is a significant departure for this type of system. While the wide availability and low-cost of Ethernet systems have long been an attractive technology for data acquisition systems, the necessity for time correlation of data from many different data sources with a very high timing precision has posed a significant hurdle standing in the way of its adoption. Time correlation was traditionally accomplished by minimizing the transport time from the data producers to the data consumers. However, with the use of switched Ethernet the transport delays are neither constant nor predictable.

To address this, a logical solution involved moving the time tagging function from a central timing-stamping authority out to the data acquisition elements of the system. Since it is highly undesirable to individually wire each data acquisition element for time, we recognized early on that an optimal time synchronization mechanism must employ the Ethernet network itself to distribute timing information. Furthermore, given the timing properties of the legacy system, any replacement architecture must be capable of providing timing uncertainties no greater than 15 microseconds between distributed clocks. This put the requirements in order of magnitude outside the capabilities of the *de facto* Ethernet-capable clock synchronization mechanism Network Time Protocol (NTP). Instead, it became clear that PTP alone offered an acceptable solution [1].

3. PTP and multicasting

PTP and the data acquisition framework described above share in common the fact that they both benefit from the use of a multicast distribution model. This is due to the significant efficiency improvements afforded by the use of IP multicast in the distribution of data as well as time. In this system, the data acquisition end nodes, the multicast producers, send their data as multicast packets. The recorders and real-time processing system components, the multicast consumers, join the multicast groups containing the data streams of interest.

To ensure efficient delivery of the multicast traffic, the switches must be aware of where the data is desired so that they can appropriately forward the traffic. Facilitating this awareness, the Ethernet switches "snoop" Internet Group Management Protocol (IGMP) control packets which consist of various messages to setup and teardown group subscriptions. With this information, the switches can then set up the multicast filters on their ports such that data is only sent out to where it is requested. PTP fits well within such an IGMP-backed multicast system, but PTP applications must also support IGMP for the PTP multicast group in order to properly receive PTP messages. This capability is usually available by default in modern operating systems' built in network stack.

The IGMP Snooping standard, see [2] describes the mechanism by which switches can snoop IGMP control messages off the network. Using this information, switches are able to learn the topology and thus are able to limit the broadcast nature of the multicast traffic so that it only reaches subscribing ports. Non-snooping aware switches simply forward all multicast traffic; including the IGMP control packets, as if it was regular broadcast traffic.

A multicast consumer must send IGMP joins, while a multicast producer is not required to do so unless the multicast producer also wants to receive multicast data sent to the given multicast group. All PTP end nodes in this network, the PTP grand master and the PTP slaves, will both send and receive PTP multicast packets. This means that all these end nodes must join the PTP multicast IP group address: "224.0.1.129".

The Boeing Flight Test Data System is based on using a high number of multicast groups for data acquisition. This is also the trend in modern industrial automation networks. The reason for this is as follows:

1. A multicast group (IP multicast address) is allocated for each type of data source.
2. Data producers always transmit their multicast data to the network.

This means that the multicast data is sent even if no data consumer has explicitly requested the data. Data for a given multicast group will be sent to the switch to which it is connected. If IGMP snooping has revealed that there are no consumers in the network, then the multicast packet for this group will be dropped at this IGMP snooping switch. However, if consumers exist, then they will have notified the network fabric by sending out an IGMP join, and the data will be forwarded appropriately through the network to reach them. The benefit here is that packets are forwarded throughout a network fabric aware of where they are needed and are replicated or dropped as needed to provide a very efficient model for both data and time distribution.

However, this multicast concept will represent a high IGMP control packet load if the end nodes and the switches are based on IGMP version 1 (see [3]) or version 2 (see [4]) since an IGMP Membership Report will be generated for each multicast group each time an IGMP enabled device receives an IGMP Query Report. This becomes a concern because IGMP control packet load can become critical on the snooping switches which potentially can receive a large number of these messages as they traverse the network between the multicast subscribers sending subscription requests to the IGMP queries. However, IGMP version 3 (see [5]) solves this problem since IGMP bulk joins are defined in this standard. Up to 183 multicast groups can be included in single IGMP version 3 Membership Report. This drastically reduces the number of IGMP Membership Reports from a multicast consumer that wants to receive all multicast groups available. For instance, in the Boeing Flight Test Data System, the network fabric is required to support any number of end nodes subscribing to up to 400 multicast groups. If IGMPv1 or IGMPv2 is used, this would represent 400 messages that the switch fabric would have to process, whereas this can be reduced to only 3 messages if IGMPv3 is used (400 groups divided by a max of 183 subscriptions per IGMPv3 report).

To avoid overloading network switches, all multicast end nodes subscribing to large numbers of multicast groups should support IGMP version 3. Furthermore all IGMP snooping switches must support IGMP version 3 as well. This requirement also applies to the PTP multicast group.

4. PTP transparency

The concept of a network device, such as an Ethernet switch, forwarding PTP traffic in such a way as to compensate for the delay it introduces is known as PTP transparency. PTP version 1, IEEE Standard. 1588-2002 [6], is used in the Boeing Flight Test Data System and does not include the concept of PTP transparency as part of the standard (although it is found in the proposal for Version 2). However, the 100 Mbps Ethernet switches used in our network conform to the standard's requirements but also provided the additional feature of transparent operation. These switches' PTP transparency implementation is based on the principles described in [7]. Specifically, the residence times[1] of SYNC and DELAY_REQ packets passing through the PTP Transparency switch are used to correct the precise timestamps of the FOLLOW_UP and DELAY_RESP packets respectively. This concept is similar to the proposed "end-to-end transparent clock" of the PTP version 2 draft except that instead of providing the measured residence time in a separate correction field, the timestamps (and any necessary checksums, etc.) are directly modified in the packets that are being forwarded to their PTP destinations. Direct modification of the time stamps are required in order to avoid any PTP version 1 interoperability problems, and provide for a switch-based correction that is truly transparent to the PTP end-nodes.

PTPv1 specifies that propagation delay measurements are based on round-trip delay measurements where DELAY_REQ/DELAY_RESP and SYNC/FOLLOW_UP are used. By using the special PTPv1 transparency only propagation delays are measured by this process since all switch queuing delays have already been removed. Due to technical difficulties in applying these same transparency techniques to gigabit switches, transparency is not available on the higher throughput (gigabit) aggregation Ethernet switches. This was corrected by adding redundant timing-only links that provide for a purely transparent timing path while still allowing data to flow through the higher throughput links. Details of this temporary workaround are beyond the scope of this paper, but are discussed in more detail in [8].

5. Migration to PTP version 2

Network topology changes must be handled for PTP version 1 slaves since such devices can suffer network topology changes that represent a different propagation delay between the slaves and the PTP grand master.

This can be handled if the PTP transparent switch supports both PTP version 1 and PTP version 2 peer-to-peer capable ports[2] based on PDELAY_REQ and PDELAY_RESP measurements. The switch shall then modify the precise time stamps of both FOLLOW_UP and DELAY_RESP version 1 packets that are forwarded on the switch ports. This technique is also applicable for PTP version 2 slaves that do not have "peer-to-peer" support since such PTP slaves will also suffer from network topology changes.

Figure 3 shows the residence calculation of a PTP version 1 transparent switch and/or a PTP version 2 transparent clock switch with end-to-end support, while Figure 4 shows the residence calculation for a PTP version 2 transparent switch.

If the peer-to-peer residence time calculation principle is used on the peer-to-peer capable ports as proposed above, then the PTP version 1 slaves or PTP version 2 slaves with end-to-end support will only measure the propagation delay on the link that the PTP slave is connected to and the PTP grand master link (if the PTP grand master also has no peer-to-peer support) as shown in Figure 5.

Utilizing the peer-to-peer transparent clock property of a switch for PTP version 2 slaves with only end-to-end support gives us the opportunity to handle various PTP implementations in the same network. However, such a hybrid PTP transparent clock implementation may be interpreted as non-complying to the version 2 draft of the PTP standard which states in section 10:

- "An end to-end transparent clock shall not implement peer delay mechanism [defined in this standard]"
- "It is recommended that all PTP version 2 and higher DELAY_REQ and DELAY_RESP messages be discarded [by a peer-to-peer transparent clock]"

The statements above are, however, meant for the case where either an end-to-end or peer-to-peer transparent clock operational capability is included in a specific clock. Our suggestion would be to include a third hybrid transparent clock implementation where both clock types are supported.

6. Legacy systems

Rarely does any replacement system come about that can completely throw off the vestiges of its predecessor, and the Boeing Flight Test Data System is no exception. As previously stated, the 'network-centric' system outlined in this paper signals a huge shift away from the traditional underlying technologies that have been used in the past. Yet many of these legacy system components represent extensive development efforts and cannot be redesigned to work with the new network and its protocols. Because of this, a large amount of effort was spent developing adapter elements that translate from legacy devices to the new network interfaces, allowing for their seamless integration with the new system.

On one side, a legion of legacy data producers existed that still provided data only in an abundance of data formats. All of these individual source types required an adapter device that could:

- time stamp incoming data (with high precision),
- translate the legacy input to a uniform, network-ready format, and

---

[1] residence time = egress timestamp – ingress timestamp of the PTP Transparency switch.

[2] A peer-to-peer capable port is a switch port where both the switch port and its link partner have peer-to-peer support

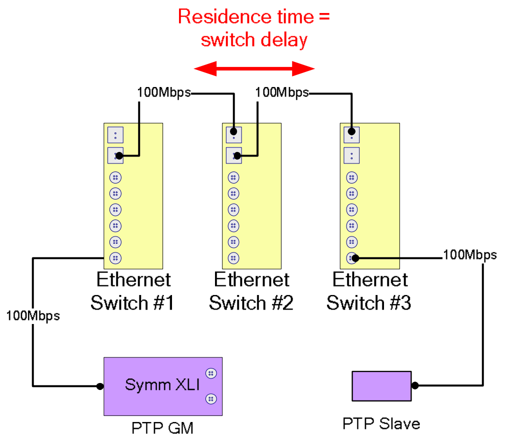

Figure 3 – Residence time: end-to-end transparency

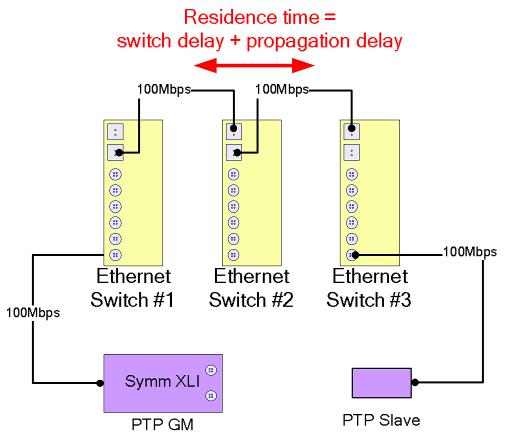

Figure 4 – Residence time: peer-to-peer transparency

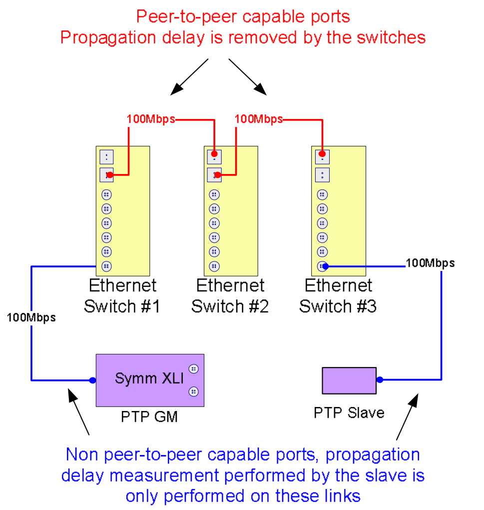

Figure 5 – Alternative method for propagation delay measurement for V1 and V2 end-to-end transparent clock implementations

- transmit the data as multicast packets on the network.

To address this, a data acquisition node was created for each possible data input type. These nodes performed all these steps (facilitated by a built-in hardware PTP implementation). On the flip-side of the system, a similar adaptor layer was created for legacy data consumers. These devices subscribe to the appropriate multicast streams, strip out the data from the network data format and then provide it to the data consumer in the format it expects.

In addition to supporting these legacy systems, the standards-based modular capabilities of the new Flight Test Data System allows for the much quicker and cost-effective development of new system components because they can use more off-the-shelf, standard equipment as a platform for deployment. Indeed, the data recorder, management, and telemetry elements interface directly with the network switches to receive their data.

7. Accuracy

Time correlation accuracy is essential to understanding the temporal relationship between events, and to distinguish between events measured at disparate places on the airplane. It is critical to ensure that every event's timing can be accurately reconstructed within the context of all other events in the system. To do this, the time uncertainty needs to be small enough to ensure that the relationship between events can be reliably determined within a reasonable window. In the Boeing Flight Test Data System, the bar was set at a distributed clock system that carries no more than 15 microseconds of discrepancy between measurement timestamps and a system-wide time source.

Deploying and testing the system described in this paper resulted in a system with extremely satisfactory results. Hardware PTP slaves, synchronized to a Global Positioning System (GPS) linked PTP master clock, routinely achieve stable synchronization comfortably under 100 nanoseconds of deviation from their PTP master. This came in well beneath the application's synchronization requirements (15 microseconds) and proved the viability of network-centric Flight Test Data Systems when deploying PTP. Less time-critical components can be serviced by a cheaper software only solution using PTPd, an open source software implementation of PTP Version 1. These devices can use standard commodity off-the-shelf hardware and are capable of synchronizing to within about 1 millisecond of the PTP master using the same infrastructure as the hardware slaves.

It is also important to note that due to the use of a PTP transparent network fabric, the efficacy of the timing system is virtually unaffected by network conditions such as heavy throughput loads. In preparation for this paper, a test was conducted to demonstrate the system's capabilities. Figure 6 shows the timing system only (data nodes are excluded) of a segment of the testbed in our laboratories. This setup demonstrates the timing system running under real-world conditions of a loaded network and using hardware representative of a real Flight Test Data System instantiation. An oscilloscope screen capture spanning 15 minutes and comparing pulse-per-second (PPS) outputs from both the PTP master and slave can be seen in Figure 7 (also, reference footnote 3).

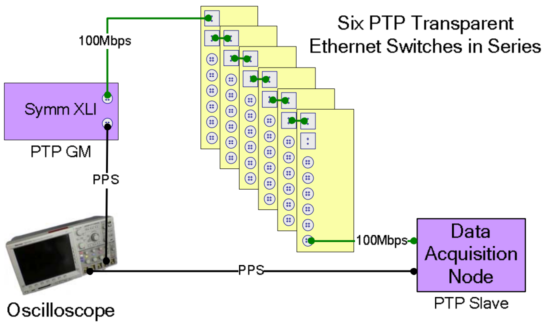

Figure 6 – Timing Accuracy Testbed

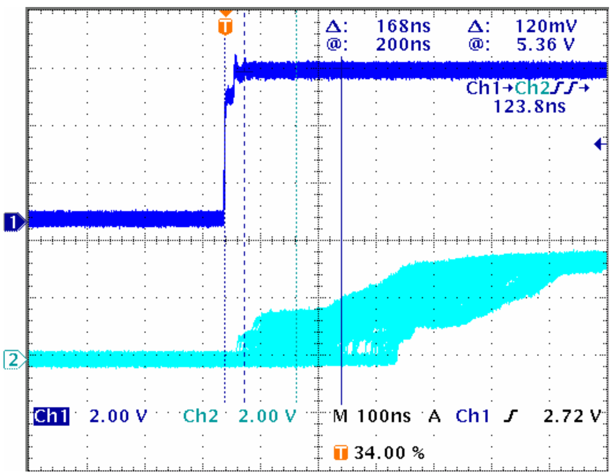

Figure 7 – Oscilloscope Capture[3]

8. Summary

As is described in this paper, making the leap from traditional data acquisition systems to the next generation of network-centric deployments is not without its challenges. Foremost among these is the need for a high-precision timing system to overcome Ethernet's non-deterministic transit delays. PTP provides precisely the required clock synchronization mechanism and with this challenge successfully overcome, the door is opened to the promise of much more flexible, scalable, and cost-effective systems that are possible with the use of popular standards-based networking technologies.

Recognition of PTP's enabling role for these systems and the benefits they bring is increasingly taking hold among the data acquisition and telemetering markets. This is due in no small part to the evidence Boeing's new Flight Test Data System provides as to the efficacy of its use, and in that sense Boeing is truly a pioneer in this field.

References

[1] E.T. Grim, "Achieving High-Accuracy Time Distribution in Network-Centric Data Acquisition and Telemetry Systems with IEEE 1588", International Telemetering Conference USA Proceedings, 2006.

[2] IGMP snooping, RFC 4541, http://www.ietf.org/rfc/rfc4541.txt.

[3] IGMP version 1, RFC 1112, http://www.ietf.org/rfc/rfc1112.txt.

[4] IGMP version 2, RFC 2236, http://www.ietf.org/rfc/rfc2236.txt.

[5] IGMP version 3, RFC 3376, http://www.ietf.org/rfc/rfc3376.txt.

[6] IEEE Instrumentation and Measurement Society, IEEE-Std 1588-2002: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

[7] S. Nylund and Ø. Holmeide, "IEEE1588 Ethernet switch transparency – No need for Boundary Clocks!", IEEE1588 workshop, 2004.

[8] S. Nylund and Ø. Holmeide, "Migration towards IEEE1588 time synchronization over gigabit Ethernet", IEEE1588 workshop, 2005.

[9] PTPd, Sourceforge Project Page available at: http://ptpd.sourceforge.net/.

[3] Note that this capture includes a measured fixed offset of 86 nanoseconds that can be removed by calibration. This calibration was not pursued because it was not necessary to meet the timing requirements for this system.