

OPTIMISING NETWORKED DATA ACQUISITION FOR SMALLER CONFIGURATIONS

DAVE BUCKLEY

**ACRA BUSINESS UNIT, CURTISS-WRIGHT CONTROLS AVIONICS &
ELECTRONICS**

ABSTRACT

Network switches are a critical component in any networked FTI data acquisition system in order to allow the forwarding of data from the DAU to the target destination devices such as the network recorder, PCM gateways, or ground station. Larger configurations require one or more switch boxes to handle aggregation, routing, filtering and synchronization via the IEEE 1588 Precision Time Protocol. However, for smaller configurations where space and weight restrictions are more stringent, a separate switch box may not be practical. This paper discusses how all the essential features of an FTI network, such as flexible forwarding and filtering and IEEE 1588 synchronization, can be maintained without the need for the separate switch box thus making significant savings on weight and size and reducing cost.

KEYWORDS

FTI switches, IEEE 1588, Ethernet to PCM gateway

INTRODUCTION

This paper explores the key functions of network switches as they are used in Flight Test Instrumentation (FTI). A standard FTI network consists of several Data Acquisition Units (DAUs), which acquire data from various sensors and avionics busses. Each DAU will have an Ethernet node which packetizes the acquired data for transmission in the network. Within the network this packetized data will move from the DAUs towards the sinks or users of the data. These sinks may include a network recorder to record the data for post flight analysis, an RF transmitter to transfer a subset of the data to the ground for real time analysis or a laptop to allow real time analysis to be carried out on board.

One or more FTI switches will sit between the sources and sinks of data. The FTI switch is responsible for forwarding the data from the DAUs to the sinks, filtering streams of data such

that only the data that each sink requires arrives at that sink, monitoring the health of the network to allow the Flight Test Engineer to determine if any data is being lost during the flight and synchronising every DAU in the system.




Figure 1: Typical FTI network

The acquired data consists of parameters. A parameter can be a measurement on an individual channel from say a strain gauge or a thermocouple or can be a measurement contained within a bus message. To enable time correlation of parameters across all DAUs synchronization is essential.

However, for smaller configurations where space and weight restrictions are more stringent, a separate switch box may not be practical. This paper discusses how all the essential features of an FTI network, can be maintained without the need for the separate switch box. Each function of the switch is briefly discussed and a solution for moving that functionality into the DAU is proposed. Furthermore other functions carried out in a large FTI network such as recording and bridging between Ethernet and IRIG 106 PCM telemetry can also be incorporated in the DAU when there are size and weight are limitations on the configuration.

FUNCTIONS OF AN FTI SWITCH

PACKET FORWARDING

The primary purpose of the FTI switch is to transfer acquired data from its source to its sink. The DAU which is the source of the acquired data encapsulates that data in network packets. These packets are then forwarded by the switch from the ingress port to the appropriate egress port. The

number of sinks and sources and the flow of data in any network can vary so the FTI engineer needs the flexibility to configure the switch to support any topology. A fully configurable crossbar switching fabric allows the FTI engineer to do this.

However the FTI switch must also have other key attributes. The forwarding table needs to be static and live at power up. Essentially this means that there is no learning of the forwarding requirements of the network. The topology of the network is preconfigured into the switch and upon power up the switch begins forwarding according to this topology immediately. The effect of this is that in the event of a power glitch during flight there is no boot up or learning delays during which data could be lost. Also the switch must transfer data with low latency and in a deterministic manner. This requirement is particularly important for data coherency in a PCM telemetry link. Coherency considerations for an Ethernet to PCM bridge will be discussed later.

PACKET FILTERING

In addition to forwarding data from ingress ports to egress ports the network switch has the ability to pass only certain streams of data from the ingress port to the egress port. Data from a particular DAU can be grouped into streams. A stream is a sequence of packets which are defined by one entry in the metadata. The analysis software can use the stream identifier (IENA key in the case of IENA packets) to identify every parameter in the packet. In some cases one stream with a subset of acquired parameters may be destined for telemetry and real time analysis while another stream with all acquired data from a DAU may be destined for a recorder. The FTI switch will have the ability to filter the streams such that only streams destined for a given sink arrive at that sink. Individual streams can be filtered based on application header information such as stream ID or IENA key, or groups of streams can be filtered based on fields such as MAC address.

IEEE 1588 GRANDMASTER




Figure 2: IEEE 1588 Protocol

Inter DAU time correlation of sampled data has always been straightforward when parameters sampled at the same rate are sampled at exactly the same point in time. In order for DAUs to sample parameters at the same time the DAUs must first be synchronised with each other. IEEE 1588 Precision Time Protocol (PTP) [1] is a network protocol which allows all DAUs in a system to be synchronised to an accuracy of 100 nanoseconds. The protocol is carried out using the same Ethernet wires which are used to transmit the data. Typically there is one grandmaster and multiple slaves. Each DAU in the system is a PTP slave. The grandmaster seeds the time to the slaves and through a handshaking protocol the slave figures how accurate its own time is with regard to the grandmasters time and adjusts its own clock accordingly. While the PTP grandmaster can come in a standalone box typically it is built in to the FTI switch.

IEEE 1588 TRANSPARENCY

The handshaking protocol built into PTP which allows the slave to figure out the error between its own clock and the grandmasters clock relies upon the trip time from the grandmaster to the slave being exactly the same as the trip time from the slave to the grandmaster. This is typically not the case in networks and rarely the case in FTI networks. In an FTI network all the data flows from the DAUs which are the sources of data towards the sinks of data such as the recorder or the telemetry bridge. This means that the delay from the slave to the grandmaster is typically much larger, as the PTP packets in this direction get delayed behind the data packets than the delay from the grandmaster to the slave. This variable and mismatched delay can introduce synchronisation errors of several milliseconds in large FTI networks.

There are two methods of correcting for these errors boundary clocks and transparency. Transparency is typically more accurate than boundary clocks so it is the more commonly used technique in FTI switches. With transparency the time in residence of the PTP packet in the switch is remembered and the switch adjusts the timestamp of the packet to make it look like the delay in the switch never happened. Using this method the delay through the switches is effectively removed and the accuracy of sub 100 nanoseconds can be achieved. PTP transparency is a feature that needs to be built into the fabric of the network switch. Each PTP packet must be hardware time stamped on the way in and the way out of the switch and the contents of the PTP packet need to be adjusted to correct the timestamp.

NETWORK HEALTH MONITORING

Transporting acquired data via Ethernet networks is relatively new to FTI. Traditionally data was carried through an FTI system using IRIG 106 PCM [2]. Now while Ethernet networks offer significant advantages [3] over PCM including throughput, simplicity and flexibility the guaranteed determinism of PCM is not a feature of standard networks. A standard office network will not be concerned with data getting delayed or lost in switches and routers as the upper layer protocols will simply resend the data. In FTI large latency and data loss are unacceptable.

While it is possible to design a deterministic, lossless, low latency FTI network it is also possible to design an FTI network that does not meet all these requirements. The adoption of Ethernet by the FTI industry has allowed the hardware from multiple vendors to be easily integrated into one configuration. However using hardware from multiple vendors can make it more difficult to design a deterministic, lossless, low latency network.

To this end Network Health Monitoring is good way to monitor the data flow through your network to ensure in real time that all the data that was transmitted from the DAUs has passed through the network in a timely manner with no loss. To allow the flight test engineer to monitor the network the FTI switch reports details such as the throughput of each port and whether any packets were dropped due to buffer overflow.

MOVING FUNCTIONALITY TO THE DAU – CONTROLLER WITH IN BUILT IEEE 1588 GRANDMASTER

Next we consider how to move each function of the network into the DAU. In an FTI network every DAU will have an Ethernet node which behaves as an IEEE 1588 slave. In a KAM 500 system the Ethernet node is built into the backplane controller. When attempting to fit a large switch box into the form factor of a small data acquisition card it is necessary to push some of the functionality of the switch onto other cards. Given this requirement the backplane controller which acts as a network node is the perfect home for the PTP grandmaster. The controller will have the grandmaster and slave functionality built in and the FTI engineer will be able to decide whether he wants the backplane controller to act as a PTP slave or grandmaster.

In the event that absolute time is not important and the requirement is to simply have all channels in the network synchronised with each other one of the controllers in the network will be designated as the grandmaster and it will synchronize all the other DAUs to its local free running time. In this scenario all other controllers in the network will be designated as PTP slaves. The fact that each controller can be configured as either a slave or a master gives the flight test engineer the flexibility to easily change which of his DAUs can be setup as time master.

For scenarios such as time correlation between two unconnected systems where absolute time is a requirement, time will be acquired in a time code generation module. This time will then be pushed across the backplane to the controller. The controller will then use this time to synchronise all the other DAUs in the configuration.

MOVING FUNCTIONALITY TO THE DAU – SWITCH MODULE

The other functionality of an FTI switch is intrinsic to the switch core and cannot be easily moved to elsewhere in the DAU. The standard FTI switch has 8 ports and supports speeds of up

to a gigabit per port. A switch module (or KAM switch) which fits into one slot of a compact DAU chassis clearly cannot match the exact functionality of the standalone switch or it would be required to be the same size as the standalone switch. A sensible trade-off is to make the KAM switch a 4 port switch with ports running at up to 100 Mbps. A 4 port switch allows data from the chassis the switch is located in plus 2 other DAUs to be aggregated and routed to a standalone recorder or a built in recorder in a 4th DAU. For a large configuration with 40 plus DAUs the data rates are at about 200 Mbps. Therefore 100 Mbps for a 3 DAU system is more than enough bandwidth. In the event that more than 4 ports is required 2 or more KAM switches can be cascaded to provide as many switch ports as necessary.

Of course with a fully configurable switching fabric the KAM switch could be used in many different topologies. For example it could be used as a two input, two output switch. With the two outputs being a recorder, a telemetry link or a lap top for on board analysis or programming. The switch could also be configured as a network tap for an external network, with two ports being used to tap off the external network and another being used to aggregate data from both directions on the external network. The KAM switch can also be used to locally aggregate data from a small number of chassis which can be forwarded onto a larger network. Equally if the flight test engineer architects a dual redundant network the KAM switch can provide dual outputs from the chassis in which it is housed.

A standard office switch will filter based on MAC or IP address as this is how data is routed through such a network. Many switches used in FTI follow this paradigm. However the packets travelling through FTI networks have specific application header identifiers which allow the packets to be identified in the metadata description of the configuration. Packet flows are split up into streams using these identifiers, stream IDs in INET [4] and keys in IENA. The FTI engineer is offered much more flexibility if the switch allows filtering based on these identifiers. In fact by designing an FTI switch core rather than repackaging off the shelf cores the KAM switch will allow filtering on any field within the entire Ethernet frame.

As previously mentioned network health monitoring is an important requirement of the switch. The KAM switch gathers information on network health such as throughput per port and packet drops due to buffer overflow or transmission errors. Each network health metric is made available to the FTI engineer to place in any network packet or PCM frame for telemetry. This gives the FTI engineer flexibility to monitor the network health to whatever extent he deems appropriate.

MOVING FUNCTIONALITY TO THE DAU – ETHERNET TO PCM BRIDGE

Although IRIG 106 PCM has been replaced in many FTI configurations with Ethernet in general the telemetry link remains at the last stronghold of PCM. This means that for telemetry an Ethernet to PCM gateway or bridge is required. Now in a PCM frame parameters from different channels and different DAUs are grouped together. Each group will contain one instance or sample of each parameter. In order to facilitate correlation of signals across multiple DAUs it is

necessary that for each group the sample from each DAU are sampled at the same time. In a typical PCM configuration the data from each individual DAU (or slave chassis) is transported via PCM to a master chassis where the overall PCM frame is assembled. The inherent determinism or constant delay of IRIG 106 PCM allowed the parameters to be transferred from the slave chassis to the master chassis such the parameter instances grouped together are always guaranteed to be sampled together. From an IRIG 106 PCM standpoint this data is defined as coherent.




Figure3: PCM Coherency

When data is transferred between chassis using a networked backbone this constant delay no longer exists. The variable delay of a networked system means that parameters from some DAUs will arrive at the Ethernet to PCM bridge later than parameters from other DAUs. In fact which parameters arrive late and which arrive on time will not be constant meaning it will be impossible to be sure of the relationship of data between different DAUs. This effectively leads to incoherent data. This raises the challenge of how to get PCM coherency using a network with variable delays. Designing a small network solution such that the Ethernet to PCM bridge is housed in the DAU gives you the flexibility of using multiple approaches which can be chosen by the addition of one or two cards to one of your DAUs. Some of these approaches are discussed below.

USING PCM AS A PIPE TO CARRY NETWORK PACKETS

The first potential solution to this problem is to effectively do away with the PCM frame. In this paradigm the parameters are packetized in the DAUs passed through the network as packets and passed through the PCM telemetry link as packets. The packets are inserted directly into the PCM frame with any gaps between packets being filled with a known pattern to indicate that there is gap. The network switch would ensure the bandwidth of the Ethernet traffic arriving at

the bridge was less than the PCM bandwidth. The ground station would then remove the packets from the PCM frame and pass the network packets to the analysis software. This solution removes PCM from the picture completely. Because each packet will have a timestamp which is the time of acquisition of the first parameter in the packet it will be possible to do correlation across all parameters in all DAUs even when some packets are delayed significantly relative to others.

However this solution has potential drawbacks. If all the ground station and analysis software is based on PCM frames then this software can no longer be used and new software tools will need to be purchased or developed. This can be a significant added expense and could potentially delay the flight test campaign. Also the variable latency in the network will mean that the analysis software has to wait until the worst case delay through the network has elapsed before displaying the real time parameters to ensure it can correlate all samples for each parameter. It is unlikely that the analysis software will be able to accurately estimate the worst case delay leaving two possible approaches. It can choose a large safe value for the worst case delay. However this value may be too large for latency critical measurements. The other option is for the analysis software to choose a smaller value for the worst case delay. This will allow the display of latency critical measurements and most of the parameters will arrive within this time. However you will see occasional spikes and drop outs due to some packets arriving after this window.

On top of this the passing of network packets over the PCM link may not be very bandwidth efficient. Network packets when transmitted over UDP/IP with some additional application layer headers can have up to 60 bytes header in total. To efficiently use network bandwidth it is advisable to keep the size of the Ethernet packet as close as possible to the maximum transfer unit of 1500 bytes. In this case less than 5 % of the frame is overhead. In order to do this multiple samples from a given parameter are often put in a packet to ensure that the packet is close to its optimal size. However for latency critical parameters the optimal strategy is to packetize on sample of a parameter with one sample of all other parameters of the same sample rate and send to telemetry immediately. Waiting for several samples of the same parameter can lead to unacceptable delays. These latency critical packets tend to be small and in this case over half of the packet is header. Sending these packets over PCM will mean that over 50% of the already tight bandwidth will be used up with packet header. This will mean the total number of parameters that can be telemetered will be halved.

In some configurations this Ethernet to PCM bridge is housed in a separate box. However in a small network one Ethernet to PCM card can be used to do the bridging. The KAM switch or the PCM to Ethernet card could filter the data such that it does not exceed the maximum throughput of the PCM link and the Ethernet to PCM card can buffer the Ethernet traffic and insert into the PCM frame with fill data between the packets. Existing de-commutation hardware can then de-commutate the frames with the addition of some ground station software the packets can be extracted and played back by off the shelf analysis software such as GS Works.

BUILDING COHERENT NETWORKS

Another approach to this problem is to constrain the Ethernet network such that it is inherently deterministic. FTI networks have a known static topology. Network nodes are not added or removed during a flight. The data rates from each DAU and through each switch are known in advance. The setup software can use this information to model the traffic and the delays through the network. The task of modelling the delays through the network is made easier by the design of simple state machine based switches which offer guaranteed bandwidth to each port and known latency through the switch core. Once the setup software can model the delay through the network it can schedule the packet transmission from each DAU such that the packets that are destined for the Ethernet to PCM bridge have enough time to travel through the network so that the parameters are guaranteed to arrive in time to be placed coherently in the PCM frame. This solution has the benefit of using existing ground station software. The PCM bandwidth is efficiently used as what is transmitted is a traditional PCM frame. Also the setup software schedules the transfer of data such that there are no outliers which can cause spikes or drop outs outside the reasonable time window constraint.

Again this coherent Ethernet to PCM solution can be realised in a standard DAU. The solution requires the use of a standard Ethernet bus monitor and a standard PCM encoder. The flight test engineer will use the setup software to define his PCM frame on the Encoder card. Every parameter in the system will be available to him. Once his frame is defined the setup software under the hood will generate the packets from each of the DAUs and schedule them through the network. It will also setup the switch and the Ethernet bus monitor such that the parameters are guaranteed to be transmitted to the Ethernet to PCM bridge in time for placement in the PCM frame.

MOVING FUNCTIONALITY TO THE DAU – RECORDING

After data is packetized and forwarded the next task is to record the data reliably and in a format that analysis software can easily handle. An important aspect of recording data is choosing a suitable file format that facilitates fast read and write functions in real-time on high speed links. The simplest solution to this is to record the data in its packetized form, that is, timestamp and record the network packets as they arrive and write them quickly and reliably to the storage media. The Packet Capture (PCAP) file format, defines an open standard format to record network packets to file. The PCAP file format is a widely used file format in the networking domain and is most notably supported in the Wireshark application. Wireshark is the most popular network analysis tool that has millions of users and is used in many networking applications not limited to data acquisition systems. In a typical FTI network a PCAP recorder will be connected to the final switch in the topology to record all acquired parameters. The PCAP recorder may have removable Compact Flash or SATA media depending on the size of the configuration and the storage requirement. For smaller networks a compact flash media which can offer 64 Gigabytes of storage will suffice to capture all the data from one or several flights. This compact flash recorder can also be made available in form factor to fit into the standard DAU. This allows the flight test engineer to record all data in the configuration without a separate box. If the configuration contains both a recorder and a PCM telemetry link then it would be advantageous to put both into the same chassis. The Ethernet bus monitor for receiving

data from the network can then be shared by both functions. In a larger configuration a KAM recorder also gives the ability to have redundant local storage of the parameters in one or more DAUs in the event that the larger recorder fails or that the connection to the larger recorder fails.

CONCLUSION




Figure 4: Small FTI network

The FTI switch is perfectly tailored for the needs of networked data acquisition during flight test. However for smaller configurations its size and weight may be prohibitive. All functions of the FTI switch including, forwarding, filtering, IEEE 1588 time synchronisation and network health monitoring can be carried out within the DAU that also acquires the data. This allows a small number of DAUs to be connected up without the need for the separate switch box. On top of this the flexibility of a modular DAU allows several strategies for Ethernet to PCM bridging to be adopted within the same DAU. To complete the picture recording of all data within the network can also be achieved by using one memory card integrated into the DAU.

REFERENCES

- [1] IEEE Standards Committee, “*Precision clock synchronization protocol for networked measurement and control systems*”, IEEE Std. 1588, 2004
- [2] Inter Range Instrumentation Group “*Telemetry Standard RCC Document 106-09*”, Chapter 4, April 2009
- [3] N Cranley, D Buckley “*The Implications for DAU Design in a networked data acquisition system*” ITC 2011
- [4] T Grace, J Kenny, M Moodie, B Abbott “*KEY COMPONENTS OF THE INET TEST ARTICLE STANDARD*”, ITC 2009