APPLICATIONS AND OPPORTUNITIES FOR THE IEEE 1588 STANDARD IN MILITARY APPLICATIONS

John D. MacKay
Progeny Systems, Inc.

Abstract

Many government/military data acquisition and signal processing systems are being migrated away from full mil-qualified components in favor of systems that maximize the use of commercial off-the-shelf (COTS) components. Field and ship-board systems provide environmentally conditioned "hotel" space for these components, which are then chosen for performance, low cost, and common use in the commercial world. The use of COTS has its drawbacks, however. Both in the front-end data acquisition and in the devices that bridge the COTS systems to the legacy components that still remain, data and interface synchronization becomes difficult, since many COTS devices replace legacy equipment that has hard-wired deterministic hardware.

IEEE 1588 (Precision Time Protocol) provides both a viable solution to this dilemma, and opportunities to maximize the full performance capabilities of the COTS system. It is a network time protocol that can afford the system a universal notion of time in a network-based architecture, as well as across the legacy boundary. This paper discusses a number of potential use cases for IEEE 1588, addressing the need for a COTS-equivalent deterministic clock source, a common system and intra-system time reference, and potential opportunities for future systems.

THE ROLE OF IEEE 1588 IN MILITARY SYSTEMS

The transition in the past 10 years for many military-technology-based systems from a full mil-qualified development process to a COTS-based one has been largely successful. This model, in fact, is reflected in other industries, such as industrial automation and telecommunications. In all these application areas, the rapid advancement of consumer- and commercial-based technology has been utilized for the development of systems that in turn rapidly advanced the technology base of the technology systems, whereas the previous paradigm has emphasized custom development.

IEEE 1588 is a precision time standard for packet networks, and thus has applicability in many of these newly developed baseline systems. The fundamental leveraged technology for many of these systems is a networked architecture, due solely to the advancement in personal computer technology. IEEE 802, or Ethernet, is the arguably the predominant manifestation of a packet-based data network infrastructure. It is the merging of these two standards that is the focus of this overview, and (more important) is the key to process improvement and advancement in modern military "COTS-based" systems.

As stated before, the shift in military system development in the past 10 years has been to replace mil-grade electronics with mil-qualified "hotel" space for commercial-grade electronics. The system would then be mil-qualified as a Unit. The success of such a system is measured on how well it can leverage rapid advances in commercial processing at much lower cost.

This change in focus has been anecdotally linked to a ship captain in the first Gulf War who, upon assessing the damage caused by enemy fire, noticed that the PC in his office survived the shock event and subsequent environmental conditions as well as the nearby very expensive mil-grade equipment. Whether this is the case, the idea was put into practice by the Navy program offices in 1995 with, among other programs, the Acoustic Rapid COTS insertion program, or ARCI. The development was focused on a single sonar array aboard a Los-Angeles-class submarine, and was very successful in demonstrating the concept. This program enjoys success today as technology insertion for LA-class and VA-class fast attack submarine systems. This is due largely to the recognition by the Navy and its contractors that technology insertion is not a static process. In order to take advantage of the "latest technology," the process of technology insertion must keep up with the advancement cycle. IEEE 1588 is currently being evaluated as a candidate for technology insertion within the context of this cyclic process.

IEEE 1588 aligns with this Mil System Development Concept due to significant development of "1588 on Ethernet." The attraction to Ethernet is clear for many industries: the promise of a networked system allows for scalable systems with a stable infrastructure. Buying into Ethernet not only means low cost for equipment, but also for total system ownership. Code development on open systems, consumer-tested hardware, operating systems, and applications, and common tools and standards significantly impact development, and the consumer demand for a reasonable level of backward compatibility provide the path to system upgrades.

A counterpoint to these advantages resides in the need for real-time processing and controls. The best consumer solution will focus on the cost versus performance, in that order of precedence. Industrial controls and military applications will put first priority on meeting performance requirements. Ethernet typifies this conflict. It is low cost, and performs well for most applications, but as there was no hard real-time requirement levied on it, a cheaper implementation won out, namely the notion of an asynchronous packet structure with an embedded clock. For real-time systems, this is a significant drawback, where synchronous operations were the basis of most of the processing infrastructure. This is the very area, of course, that IEEE 1588 addresses. Specifically, for the technology insertion process employed by the Navy sonar development, it qualifies as a candidate for the following reasons:

- Technology Enabler
    - Data mining and correlation applications can take advantage of a fully synchronous Data Acquisition System
- Problem Solver
    - Systems with disparate time protocols can be improved by the use of a single unified protocol
- COTS Development Gateway
    - Precision time as a network service enables low-cost data-gathering technology insertion and can shorten development efforts.

Adoption into the Tech Refresh Cycle will depend on key adoption milestones in commercial equipment:

- Ethernet silicon manufacturers need to put it onto the wafer
- High sample rate interfaces developed that are "close to the wire." This is needed for high accuracy
- The protocol needs to be included in the stack
- The PTP time service needs to be integrated into the OS
- Consumer applications need to make use of precise time.

In short – we need to see the label "1588 Inside" on commercial network equipment.

TIME SERVICES AND STANDARDS

The typical time standards that will be considered as candidates for technology refresh are grouped into two categories: time synchronization, and frequency syntonization. Time synchronization refers to the accuracy between a clock and a time reference keeping the time of day, or a relative notion of time, such as seconds since midnight. Frequency syntonization refers to the frequency and phase difference between oscillators. There are well defined standards and protocols employed by data acquisition systems for both.

Time synchronization

- Inter Range Interface Group (IRIG) physical transport mechanism is shielded wire – 1 μs accuracy
- Network Time Protocol (NTP) transport mechanism is network PHY – 10 ms accuracy

Frequency syntonization

- "ATS" custom clock/marker schemes – used to sample data
- IRIG master clocks output pulse per second (PPS) or reference signal
- FDDI-based taxi, ATM-based TAIPT.

The significant feature of the IEEE 1588 protocol is that it does both over any packet-based network as a single homogenous standard with significantly higher time synchronization accuracy.

SYNCHRONIZATION ACCURACY AND IMPLEMENTATIONS

The accuracy of implementations of IEEE 1588 depends directly on implementation, as expected. An important distinction to other protocols is that 1588 is independent of these implementations, allowing for both hardware and software implementations.

On a 100baset network, the following has been achieved:

- A software-only implementation that does not optimize the operating system can achieve 10 μs of synchronization accuracy

- A software-only implementation that modifies the O/S to time-tag 1588 messages as soon as they are received can achieve about 1 μs of synchronization accuracy
- Hardware-assisted implementations have achieved from 500 to 4 ns synchronization.

Gigabit Ethernet implementations, coupled with designs that increase sampling granularity, can achieve sub-nanosecond accuracy with careful network design and methods to null asymmetry in path lengths

An important short-term impact to the adoption of this standard in a technology insertion program is the revision of the standard to Version 2. While this is somewhat disruptive in the process, there are several benefits worth the impact. Key components for mil applications are:

- Sub-nanosecond representation
- Varied sync update rates
- Different topologies
- Rapid reconfiguration
- Fault tolerance
- IPv6
- Security
- Alternate time standards such as Coordinated Universal Time (UTC).

APPLICATION EXAMPLES

TEST & MEASUREMENT IN A "HARDWARE IN THE LOOP" SCENARIO

In this scenario, time synchronization is established to the required accuracy by the distribution of an IRIG-B signal to several data acquisition and actuation components. When the system requires upgrades for any IRIG receiver, the component selection is limited by compatibility with IRIG components. The original tasking for this program was in fact to develop an IRIG interface card using standard interfaces, such as PCI. Once the components are selected, a non-trivial development effort is needed to capture the IRIG time and use it in the acquisition application.

Without regard to 1588, network-based solution takes advantage of COTS. The 1588 advantage depends on the following:

- 1588 is part of network solution – built into the NICs and switches
- 1588 is part of the network protocol stack
- 1588 is part of the O/S time service

These all depend on industry adoption of 1588.

When this occurs, networked devices are precisely synchronized simply by being on the network. Therefore, there is no driver development, no integration of timing services, and significantly lower cost of ownership for this implementation.

SONAR INTER-ARRAY SYNCHRONIZATION

The current baseline for the ARCI sonar system operates by the independent data acquisition of each sensor array. The forward array is fully synchronized to itself, but is not typically synchronized to

reference hydrophones or the aft towed array. The objective is to tie all sonar arrays to single time/clock reference. Sub-systems that can take advantage of this are Active Intercept & Ranging (AI&R), which upgrades the Active Intercept ships safety functions on LA Class fast-attack subs, and Common Undersea Picture (CUP)/Universal Data Exchange Manager, which are technology enablers for data exchange on the Global Information Grid (GIG).

The current baseline system is underutilizing the aggregate data that are acquired, and without a common time base, data acquired and recorded cannot be correlated accurately in post-operation reconstruction.

- Universal Data Exchange is impacted by the use of different time standards at different accuracies
- Time standards across systems, both on a single ship and across multiple platforms, should be common
- Data correlation across arrays increases data value and improves situational awareness.
- Larger arrays being developed need synchronization for multiple equipment racks.

For the application mentioned, the following advantages can be realized:

- Active Intercept & Ranging
    - Data correlation across arrays promotes improvements to all AI&R functions
    - Total ship monitoring – better self-noise correlation
    - Ranging – wider sensor geometric base reduces ranging error
    - Beamformer gain from multiple arrays increases sensitivity to transients
- Common Undersea Picture (CUP) and Universal Data Exchange Manager
    - Very low latency situational awareness
    - Enables low-bandwith platforms to synchronize
    - Enables high-fidelity scenario recreation.

LOW-COST SENSOR TELEMETRY SOLUTIONS

Remote shipboard sensors and arrays have been designed with the philosophy that digitization close to the signal source (such as the hydrophone element) provides the greatest fidelity of the signal and reduces design issues, such as noise and signal level. The tradeoff for these obvious advantages manifests in the design of the telemetry used to transport the signal. The design issues specifically have to do with the life cycle of the communication protocol, and the ability to synchronize this signal to the rest of the system.

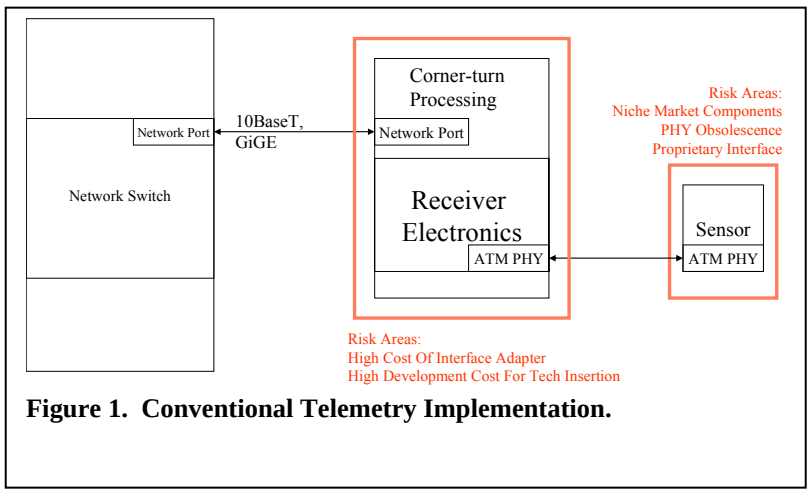

Figure 1. Conventional Telemetry Implementation.

The drivers for the choice of a telemetry interface are link speed and determinism. The link speed, of course, needs to exceed the data rate. Determinism refers to the expected arrival of the data packet after transmission. Within this parameter is the amount of latency, and the random delays imposed by the protocol. As a result, most data links are designed using the transport layer of higher level protocols. This approach has several advantages. For one, the chipsets used for these protocols are commercial off-the- shelf components. Another is the availability of resources for the design of the links – reference designs, lower level software, and firmware.

The use of these links has advantages for data synchronization. Typically, the clock and data are on the same signal, using an encoding scheme such as Manchester encoding. The clock from this interface is used to drive the sampling clock of the sensor electronics. In this way, all sensors on telemetry can be synchronized.

Obsolescence of the chipsets is a significant design issue. Many of these sensor interfaces are designed into tight-fitting enclosures, and are, therefore, custom designs which are costly to redesign and retest. In the past 10 years, every protocol used for this telemetry link has faced obsolescence.

Ethernet implemented with the full stack has been disqualified because of the asynchronous nature of the packet transfer. Often data packets are queued, so a continuous deterministic timeline is not possible. As well, the chipsets for Ethernet do not provide clock recovery, so this link is inherently asynchronous. A lower layer transport mechanism would eventually face the same fate as the other choices. The strategy to employ provides the greatest flexibility for the sensor implementation, and takes advantage of the low-cost Ethernet interface on the receiver.

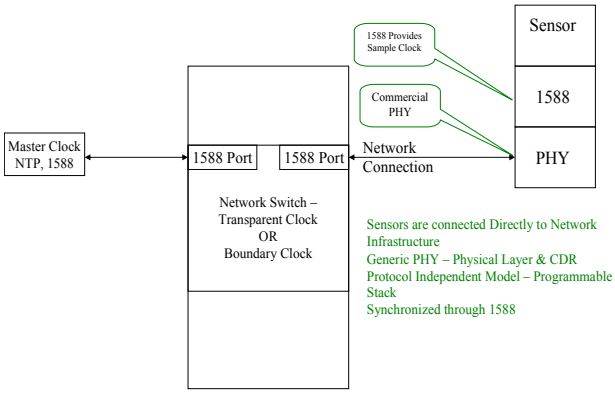

Figure 2. IEEE 1588 and Protocol Independent Telemetry.

In general, this involves replacing proprietary data links with Protocol Independent Interfaces at the sensor end. Using generic building blocks and a highly programmable logic infrastructure allow for a wide range of protocols to be implemented in firmware and software on the sensor electronics.

IEEE 1588 is a significant part of this solution. With an off-the-shelf 1588-aware Ethernet switch interface as the receiver, and a firmware-based Ethernet interface implemented as a 1588 slave, the sensor interface provides commonality, future-proofing from data link evolution, and precise data synchronization.

SUMMARY

IEEE 1588 (Precision Time Protocol) provides both a viable solution to the use of COTS to replace mil-grade hard-wired timing systems, and opportunities to maximize the full performance capabilities of the COTS system. It is a network time protocol that can afford the system a universal notion of time in a

network-based architecture, as well as across the legacy boundary. Systems that are designed as a distributed network can take advantage of the synchronization and syntonization features of 1588.