

Portable Instrumentation for Time Source Verification and Analysis

Nigel Helsby

Research Department
Time & Frequency Solutions Ltd
Witham, Essex CM8 3AL UK
Email: nigel.helsby@timefreq.com

William Dean

Instrumentation Systems Division
Industrial Computing Ltd
Gillingham, Dorset, SP8 4NG, UK
will.dean@indcomp.co.uk

Abstract— The architecture and principles of operation of a new type of portable, battery-powered measuring instrument for precise time source performance analysis are described. Potential applications include the measurement of common time sources, such as simple pulsed systems, asynchronous serial data streams, modulated and unmodulated time-codes and particularly network time servers using Network Time Protocol (NTP) or Precise Time Protocol (PTP). New techniques for high precision 100Mb/s Ethernet packet timing as well as for modulated carrier time-code extraction are described showing how the accuracy of clocks providing these sources can be fully assessed automatically in the field with laboratory precision. The employment of a GPS receiver and antenna within the same compact package to provide the synchronization and disciplining of an internal rubidium oscillator is presented, together with some results of its performance in a typical time source measurement project. This portable clock can maintain reference time with high accuracy when GPS signals are no longer available during the verification of a time source. The method of time comparison using precision clock signals from the rubidium oscillator in conjunction with interpolation to sub-nanosecond resolution is outlined; the principal advantages of the method described are its seamless and continuous time measurement capability over a full +/-500ms range to a resolution of 100ps directly related to the best averaged information from GPS satellites. Traceability of the instrumentation is established and maintained by the use of GPS Common View calibration.

I. INTRODUCTION

Proving Master Clock timing system accuracy during development, Factory Acceptance Testing or Commissioning soon reveals that the engineer has few tools at his disposal for rapid evaluation and easy recording of results. Verifying the accuracy and the types of timecode in use can become a laborious operation involving scrolling through captured timecode frames. Often data is lost when the time error of the clock source under test moves fractionally from positive to negative, necessitating a change of measurement connections. We describe a multipurpose instrument that facilitates this task. It was decided that the instrument should be capable of measuring and identifying the common amplitude modulated

1kHz carrier based codes, their dc level shift (DCLS) equivalents, standard pulses such as a simple 1pps, serial data and network time protocols.

II. SYSTEM ARCHITECTURE

The reference and measurement part of the instrument was designed in three main blocks, with most of the circuitry installed on a single printed circuit board that carries connectors for subsidiary interfaces and the network measurement system. The three main blocks are time reference acquisition with precision measurement, signal acquisition and analysis and network time measurement. They all connect to the main CPU which provides the user interface to an LCD touch-screen and USB port for data-logging.

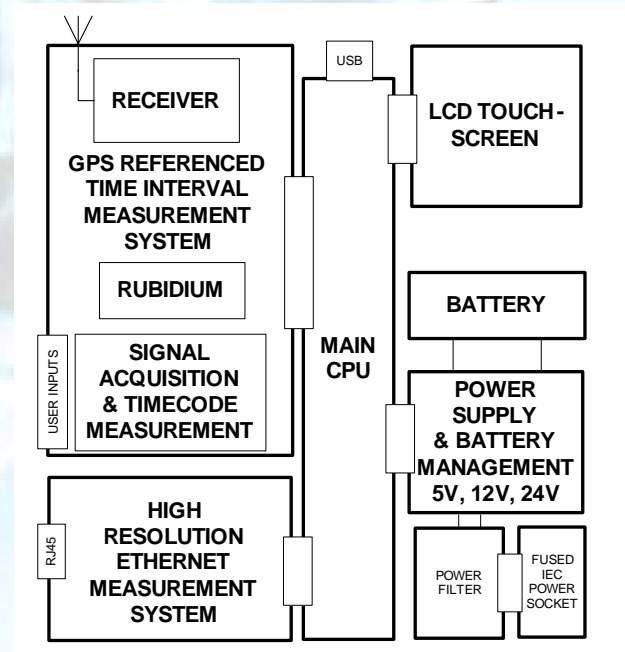


Figure 1. Measurement System Block Diagram

GPS is sometimes not available for synchronization, even via external cable connection; for this situation a rubidium oscillator is fitted, allowing good time transfer after synchronization on an external site. The entire package is powered by rechargeable batteries allowing several hours of operation with no connections other than those to the user's time sources.

A. Reference Acquisition and Precision Measurement

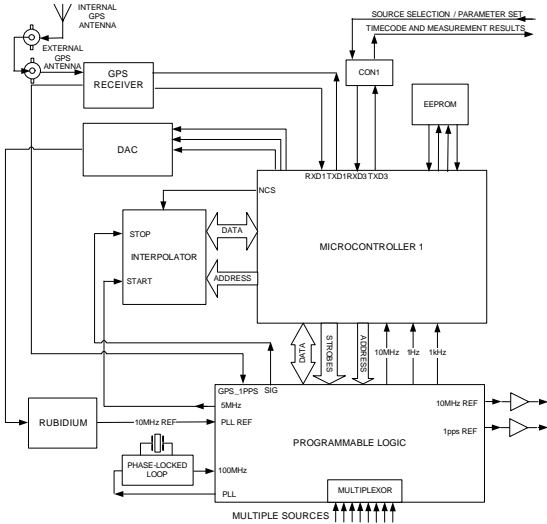


Figure 2. Oscillator Disciplining and Time Interval Measurement

A GPS engine supplies a reference time string and 1pps. For continuous frequency correction of the reference the time interval between the GPS receiver 1pps and 1pps derived from the 10MHz output is measured, averaged and integrated. The continuous measurement and adjustment strategy results in final corrections being applied at the smallest frequency step that can be made (3E-14). A second order phase-locked loop is implemented with variable time constants as described in an earlier paper [1]. The derived 1pps reference is arranged to open a measurement gate once per second. 5MHz from the same synchronous counter system starts and restarts a potential measurement interpolation each 200ns. Signals to be measured are processed such that a single rising edge pulse represents the 'on-time' point. The chosen pulse closes the measurement gate and triggers the interpolator enabling time measurement to within a fraction of the 200ns. After each measurement the same 5MHz is used to carry out a calibration of the interpolator, updating offset and scaling constants for each new sample. The precision of the measurement and interpolation process were explored by slewing a 1pps at 100ps per second into the measurement input and plotting the result with the first order slope removed.




Figure 3. Time Measurement Error

At least 2000 samples were used ensuring that more than one complete cycle was explored. The result shows that there is 'noise' as the phase of the measurement instant traverses transitions of the 5MHz clock but the measurement variation was inside our ± 1 ns target with useful resolution of better than 0.2ns.

Each broad category of input has an interface designed to capture the input signal and locate the timing information with best precision. Provision is made for five broad types of input

- Pulsed Input such as 1 pulse per second.
- DC Timecode pulse-width modulated codes
- AC Timecode Inputs that use amplitude modulated carrier
- A single 100base-T network connection for network time measurement with client and server capability and the facility of remote instrument control
- RS232/RS422 Inputs for serial time-code and miscellaneous measurements such as timing of relay contact changeover.

Each physical input is buffered by appropriate hardware and connected to multiplexer circuitry for selection by microcontroller 1. A command from the main CPU designates the input source. Microcontroller 1 returns measurement data as well as status of the GPS disciplining process for display. Measurements of Ethernet packet timing are carried out by separate dedicated circuitry.

B. Signal Acquisition and Time-code Measurement

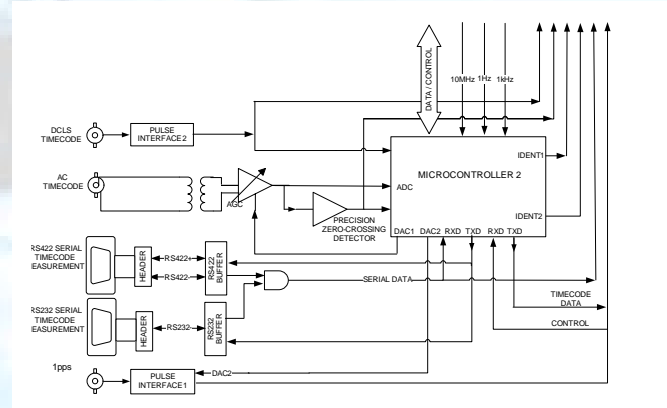


Figure 4. Input Interface and Time-code Reader

The signals to be analysed are buffered and fed to the time-code reader section. Particular attention is paid to the processing of amplitude modulated carrier time-codes such as IRIG B, AFNOR and NASA codes. A precision zero-crossing detector is used to extract the carrier timing using progressive filtering as described by the designers of a detector for frequency stability measurement [2]. This approach allowed the 1kHz carrier to be measured to better than 1µs resolution over more than 10dB carrier amplitude change. All of the hardware signals are routed via the common programmable

logic block, this time to microcontroller 2. The microcontroller anticipates the start bit of each code ‘frame’ to enable an ‘on-time’ 1pps to be forwarded to the time interval measurement system. It also attempts to identify the code type as one of a number of industry standard codes; in some cases a few seconds are required for this process to be completed. When the code type is known it is displayed, providing invaluable information to a user who has the need to confirm it and understand the data content.

C. High Resolution Ethernet Measurement System

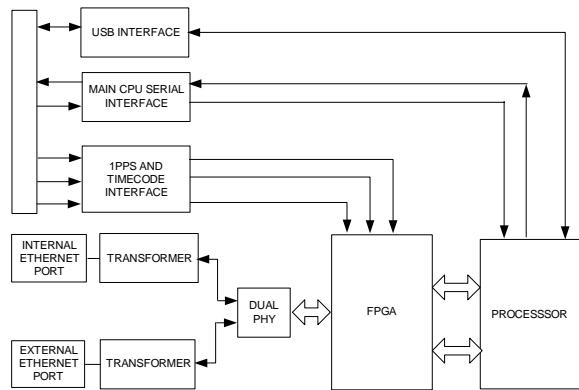


Figure 5. Ethernet Measurement System

The Ethernet measurement system is a combination of custom measurement hardware along with application code executing on the main CPU. The main CPU runs a commercial operating system which does not offer deterministic real-time performance, and is therefore not suitable for making reliable high-precision time measurements. However, the operating system does provide a complete TCP/IP stack and a productive development environment which allows for straightforward implementation of network timing protocols (e.g. NTP).

The custom hardware is capable of generic packet classification and high-precision time-stamping – it is ignorant of the specific network protocol in use. The combination allows easy protocol development and analysis, along with reliable time-stamping.

The measurement hardware uses an FPGA and dual Ethernet transceiver (PHY) to provide two 100MBps Ethernet ports. One of these ports is connected within the instrument to the main CPU; the other port is externally available for connection to the equipment under test. An important, though secondary, function of the FPGA is to provide packet forwarding between these two ports, so that the main CPU can communicate on the external network – such forwarding is transparent to the CPU. The Ethernet ports are of the ‘auto-crossover’ type, allowing direct connection to the equipment under test without special cables or an intervening hub/switch.

A PLL in the FPGA is used to generate a local 50MHz clock from the reference 10MHz clock, and a synthesised UART + ASCII timecode reader is used to qualify the 1PPS signal. Using this combination of 10MHz, 1PPS and ASCII

timecode, the FPGA continuously maintains current time as a 64-bit value, with a resolution of 20ns.

A microprocessor is included on the measurement hardware to provide FPGA configuration and communications to the user-interface processor. It does not participate in network traffic or time-keeping.

Packet capture is performed on the ‘external’ Ethernet port. Independent transmit and receive filters are configured for this port – these filters consist of ‘match’ and ‘mask’ tables, which allow the unit to be set to only capture packets for a specific protocol. Captured packets are stored, along with the time of their capture, in a buffer within the FPGA. Captured packets and their timestamps can subsequently be read by the main CPU, using its RS232 connection to the Ethernet hardware.

The direct capture and time-stamping of packets at the Ethernet transceiver eliminates a number of possible sources of measurement error, including interrupt service latency, Ethernet collision back-off, and MAC processing delays. Remaining sources of error include PHY delays and Ethernet switching or routing delays. The former are typically short enough to be neglected in current applications, and the latter can be avoided by direct connection of the instrument to the equipment under test. The FPGA is capable of filtering packets at ‘wire speed’; its performance is unaffected by the density of packets on the network.

NTP Measurement Example:

To measure performance of an NTP server, the main CPU primes the capture hardware and then generates an SNTP client request. The Ethernet hardware captures and time-stamps both the outgoing request and the server response packet.

The instrument then calculates the difference between its local reference clock and the remote server clock using the standard NTP offset calculation:

$$\text{Offset} = ((T_2 - T_1) + (T_3 - T_4)) / 2$$

Where:

- T_1 = Client transmit time (from capture hardware)
- T_2 = Server receive time (from response packet)
- T_3 = Server transmit time (from response packet)
- T_4 = Client receive time (from capture hardware)

As with all such calculations, the assumption is made that transmit and receive delays on the network are symmetrical.

In a similar fashion, the instrument can calculate the network delay (‘round trip time’) to the device. This uses another standard NTP calculation:

$$\text{Delay} = (T_4 - T_1) - (T_3 - T_2)$$

Where $T_{1,2,3,4}$ are as above.

As can be seen, at no point does the main CPU require real-time performance – it merely analyses timestamps generated elsewhere.

III. DATALOGGING

One of the key features of the measurement system is its data-logging capability. After making the source selection applicable to his measurement requirement the user can simply insert a USB memory into a front-panel connector and log any number of samples. Once started, the logging can be stopped at any time or allowed to continue until the specified samples are collected. If the 'engineering log' option is checked then data on the reference performance are also collected on a second-by-second basis.

The files are arranged as ASCII text with comma separated fields for easy integration into either a spread-sheet or into the user's own data processing program for analysis and graph-plotting.

IV. MEASUREMENT RESULTS

Two NTP servers are available and connected to a local network. An engineer can rapidly configure our portable measurement system to specify the I/P address of the server of interest for measurement. The 'client' in the Measurement System queries the server time, time-stamps the packets with precise measurement system time and makes the standard NTP calculation of time offset. The engineer can see immediately from its display how the offset varies and can log the data to memory for analysis together with the times of the client and server requests. In this example NTP Server 1 is polled by the measurement system once per second for about 100s; the graph reveals that frequent outliers larger than 10ms are found when the time offset is computed by the measurement system as shown in Fig. 6.

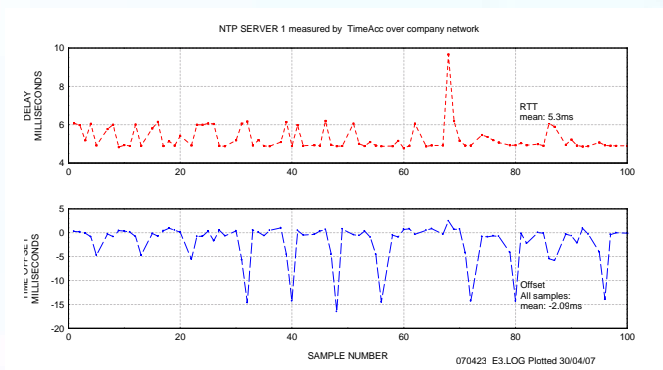


Figure 6. NTP Server 1 Measurement

The second server is connected to the same network and polled by the measurement system. This server contains

hardware which is dedicated to the task of time determination and delivering its responses to the clients. The time offset is much more stable. It is necessary to zoom in (Fig. 8) to observe about 5 microseconds systematic 'jitter' in the time determination. Clearly this server has the potential to deliver time over a simple network to much higher precision, even though tests showed that a standard client would not necessarily chose it in preference to the first server shown above. Within a few minutes the difference in quality between the master clocks under test is revealed in detail.

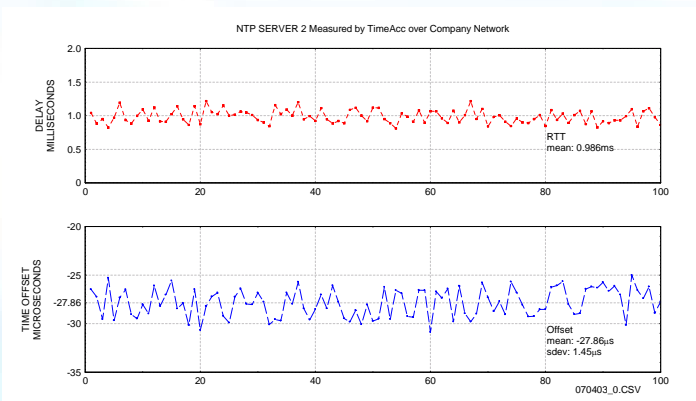


Figure 7. NTP Server 2 Measurement

V. CONCLUSIONS

The instrument we describe has proved extremely easy to use and powerful in its analysis of the performance of many types of master clock and their associated distribution systems. It overcomes the usual limitations of standard laboratory instrumentation in which the measurement of events close to 'on-time' becomes problematical because the first signal to arrive at the start of a measurement may or may not be the reference. As well as identifying, measuring and authenticating standard time-code types it makes available powerful analysis of NTP time-sources, previously unavailable on any other type of instrumentation.

REFERENCES

- [1] Nigel Helsby "GPS Disciplined Offset Frequency Quartz Oscillator", 1992, Proceedings of the 2003 IEEE Frequency Control Symposium, pp. 435-439.
- [2] L.Sojdr, J.Cermak, and G. Brida "Comparison of High-Precision Frequency-Stability Measurement Systems," Proceedings of the IEEE 2003 Frequency Control Symposium, pp. 317-325.

