1. Introduction

As computers and communication networks become ever faster, there is a need to upgrade the timekeeping technology which timestamps transactions, controls real-time playback of video and audio data and synchrU.S.onizes the network itself. For instance, the switches of the Highball project must be synchronized to within the microsecond. Existing timekeeping technology, such as used in the DARTNET project, is capable only of some three orders of magnitude greater than that.

Among the various methods of synchronization to national standard time are various radio and satellite services designed specifically for time transfer, together with other radio and satellite systems designed primarily for navigation, but providing time transfer as an auxiliary service. While the radio and satellite services operated by the National Institute of Science and Technology provide time-transfer accuracies to the order a millisecond for the WWV/WWVH high-frequency radio service, 50 μs for the WWVB low-frequency radio and GOES satellite services, receivers capable of the ultimate accuracy of these media are not presently manufactured. For instance, the timing accuracy of the Spectracom Netclock/2 WWVB receiver is specified as ±100 μs, while that of the TrueTime 468-DC MK III GOES receiver is specified as ±250 μs. However, the performance of these receivers in service materially degrades due to various propagation irregularities, local noise sources and receiver design deficiencies. For instance, long experience with the Netclock/2 and its predecessors confirm reliable accuracies only to the order of a millisecond.

For accuracies much improved over the available WWVB and GOES receivers, only the Long Range Navigation (LORAN-C) and Global Positioning Service (GPS) navigation systems are practical, although other systems can in principle provide similar accuracies. Both LORAN-C and GPS are controlled from the U.S. Naval Observatory to maintain agreement with Coordinated Universal Time UTC(USNO) to within a small number of nanoseconds. Corrections showing the relative timing offsets of individual LORAN-C chains and GPS satellites published weekly suggest 50-100 ns as the ultimate obtainable accuracies available with these services using simple equipment and averaging periods up to ten minutes and without differential corrections or, in the case of GPS, knowledge of the SA codes. While GPS provides data messages from which UTC can be extracted directly, UTC time can be extracted from LORAN-C transmissions only with the aid of additional data, as described later in this report.

When long term accuracies less than a microsecond are required, even rubidium-vapor and cesium-beam clocks are not suitable, since they drift in time, although slowly, and must be periodically recalibrated using time or space transfer relative to UTC(USNO). The well equipped standards laboratory today uses multiple cesium-beam clocks, together with GPS timing receivers which employ long averaging times and a special feature of the GPS data message which provides differential corrections relative to UTC(GPS) as broadcast.

However, timing receivers for LORAN-C and GPS are very expensive and have limited utility. Presently available receivers are priced in the $10,000 to $20,000 range, making their use possible only in the most demanding environments. While receivers designed for navigation are priced much less, their adaptability as a precision time source is doubtful. Available LORAN-C navigation receivers provide no timecode suitable for use as a primary clock for a computer network, while available GPS navigation receivers require clear-sky view of the horizon and short antenna-receiver cable runs, which can cause siting problems.

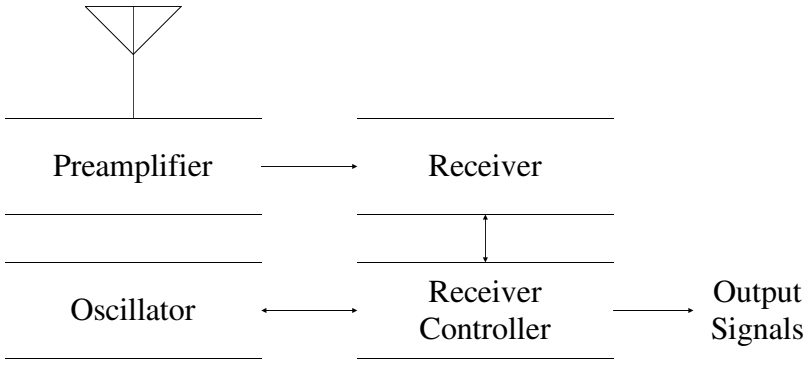

Figure 1. LORAN-C Timing Receiver

With the above in mind the LORAN-C project was initiated to explore the feasibility of constructing an inexpensive LORAN-C timing receiver which could provide accuracies comparable to the ultimate achievable with the medium. The primary objective was to supplant and eventually replace a cesium clock and LORAN-C timing receiver on loan from the U.S. Coast Guard. The ultimate goal was to produce an easily duplicatable, low cost primary clock suitable for computer networks using the Network Time Protocol [MIL91]. A prototype timing receiver has been completed, tested and evaluated. The results of this project, including the hardware and software design and test results, are summarized in this report.

The LORAN-C Timing Receiver described in this report is designed to provide precise frequency and time with respect to one or more selected LORAN-C master or slave stations operating at an assigned group-repetition interval (GRI) in the range 40,000 to 100,000 μs. The design objective is a time-transfer accuracy of ±100 nanoseconds and frequency stability of $\pm 10^{-10}$ relative to the received LORAN-C signal. While intended primarily as a timing source and frequency reference, the design includes the capability to determine geographic position using multiple LORAN-C stations and hyperbolic-coordinate navigation principles.

A block diagram of the receiver is shown in Figure 1. The antenna and optional preamplifier are connected to the receiver itself, which is assembled in a small metal box. The timebase consists of a disciplined frequency standard in the form of a 5-MHz oven-controlled, quartz-crystal oscillator from which all internal and external timing signals are generated. The receiver controller, which is designed for the IBM PC/XT/AT or compatible host computer, contains various counters, registers and converters to generate various timing signals and support the receiver. The reciever outputs include the internal 5-MHz oscillator, together with various other derived signals as determined by software, including a 1-pps signal synchronized to UTC(LORAN).

Software in the host computer supports the acquisition and tracking algorithms, as well as selects the various operating modes. The receiver tracks the selected LORAN-C station(s) in both envelope phase and carrier phase and generates a correction signal to adjust the oscillator relative to the transmitted signal. While not yet realized, a design goal is that the host computer need not be dedicated to the receiver-control function, since for many applications the support program needs only be run infrequently, such as when the computer is not otherwise in use. For the highest accuracy and stability of which the receiver is capable, however, the support program would be run continuously, most conveniently as a transient-stay-resident (TSR) program in the DOS environment.

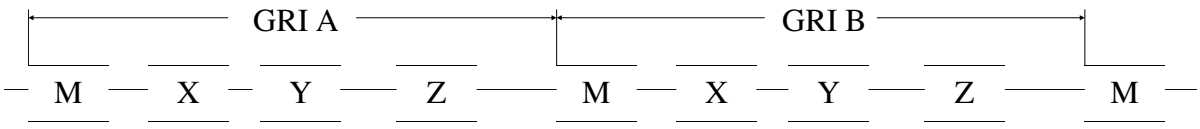

Figure 2. LORAN-C Group Repetition Interval

Master		Slave	
A	B	A	B
+ + − − + - + - +	+ - - + + + + + -	+ + + + + - - +	+ - + - + + - -

Figure 3. LORAN-C Pulse Coding

Since the LORAN-C signal waveform provides no timecode, outside information is necessary to resolve the epoch of the internal receiver time with respect to UTC. The support program can calculate its position in the same manner as an ordinary LORAN-C navigation receiver and determine the propagation delay relative to UTC(LORAN). It is only necessary to resolve which timing interval corresponds to UTC as broadcast. Conventional procedures call for the use of published numbers, called time-of-coincidence (TOC), and an external 1-pps signal, such as provided by an existing WWV, WWVB or GOES timing receiver or cesium clock. However, in the intended application the controlling computer is connected to a network supporting the Network Time Protocol (NTP) [MIL91], which ordinarily provides timing from redundant sources to an accuracy of a few milliseconds. Thus, NTP can be used to resolve the interval ambiguity. For this application the primary purpose of the timing receiver is to reduce the timing jitter ordinarily available with NTP by some three orders of magnitude.

For the highest accuracy and stability, the receiver should be located within groundwave coverage (about 600 km) from the selected LORAN-C station(s). For navigation purposes the accuracy of the LORAN-C system itself is specified as 0.25 nautical mile, which corresponds to a timekeeping accuracy of about 1.5 μs; however, under groundwave propagation conditions between the master and at least two slave stations of a LORAN-C chain, and if the receiver can determine an accurate and precise geographic position, accuracies to the order of 100 ns are achievable. The main accuracy limitations in groundwave conditions are local precipitation, noise levels and propagation conditions; however, the system can operate under skywave conditions, which prevail in some areas at night, with a diminished accuracy of about 50 μs.

2. LORAN-C System Operation

The LORAN-C Timing Receiver is designed to operate using transmissions of the LORAN-C radionavigation system operated by the U.S. Coast Guard [FRA83]. All LORAN-C stations transmit using a biphase-modulated, pulsed carrier at a radio frequency of 100 kHz. The signal consists of a group of eight or nine pulses transmitted at intervals of 1 ms. A LORAN-C chain includes a master station and at least two slave stations, all of which repeat the pulse groups at the same GRI assigned to the chain. For instance, the Northeast U.S. chain with master at Seneca, NY, is assigned a GRI of 99,600-μs (commonly expressed in units of 10 μs or 9960), so that the pulse groups repeat at intervals of slightly less than 100 ms. Figure 2 shows the sequencing of the pulse groups in the repetition interval for a typical master (M) and three slaves (X, Y and Z). The relative timing of the

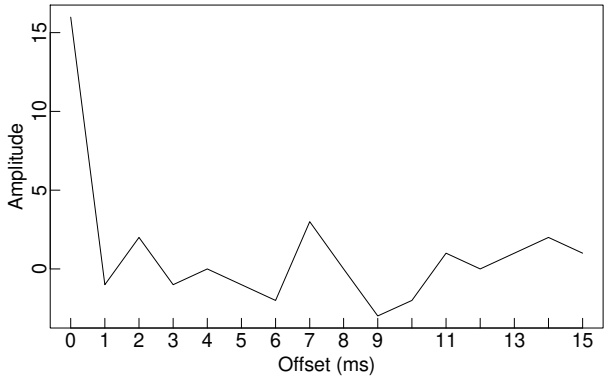

Figure 4. Master Autocorrelation Function

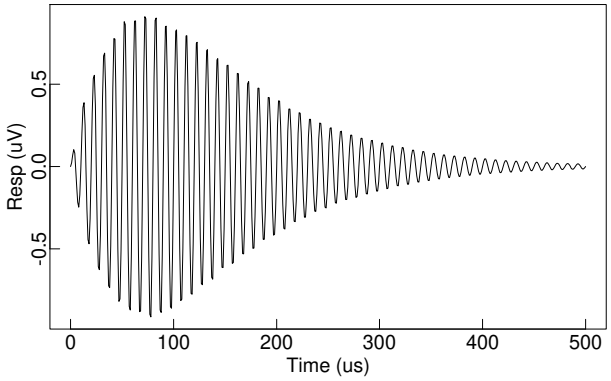

Figure 5. LORAN-C Signal

master and slaves is designed to prevent pulse-group overlap anywhere in the intended coverage area.

2.1. Signalling Characteristics

The biphase modulation of the pulse groups is designed to minimize interference between the chains operating on the same frequency, as well as provide identification of the master and slave stations. All master stations are assigned one encoding schedule and all slave stations are assigned another. Both schedules consist of a frame of two alternating fields A and B in which separate 8-bit Golay codes are used to modulate the pulse group. The period of each field is equal to the GRI assigned the chain, with the A codes transmitted in even fields and the B codes in odd fields, as shown in Figure 3. In the table "+" indicates an in-phase pulse, "–" indicates a 180-degree reversed-phase pulse and transmission begins with the most significant (leftmost) pulse. The schedule for masters includes a ninth pulse for identification and status reporting, but this pulse is not used by the receiver.

In order to assess the operation of the receiver under conditions where the signal level for the selected station at the receiver may vary widely relative to other stations on the same frequency, it is useful to explore the correlation functions for the various codes shown in Figure 3. Figure 4 shows the autocorrelation function for the master codes, assuming the A and B fields follow each other in that order. The function has a peak of 16 at zero offset and no more than three at any other offset, for a processing gain or margin of about 15 dB. As might be expected, the autocorrelation function of a slave is similar, while the crosscorrelation function for a master with a slave is very small, as long as their A and B fields are aligned.

If the A and B fields are not aligned; that is, the A field of the receiver does not correspond to the B field of the transmitter, the value of the various crosscorrelation functions at various offsets is larger than three and, in some cases reaches eight. This reduces the margin and increases the false-alarm rate where the receiver locks onto an incorrect station or at an incorrect offset. If because of signal level mismatch or noise the margin is exceeded, the receiver may synchronize incorrectly. The software design described later in this report deals with this problem by using various linear and nonlinear filters and multiple integration cycles.

2.2. Establishing the Reference Epoch

Each of the eight 1-ms pulses in each field has a controlled waveform in order to minimize errors due to multipath (skywave) interference and minimize spectrum requirements. Figure 5 shows a

typical LORAN-C pulse. Only the first several 10-μs carrier cycles in each pulse before the envelope peak are used to establish cycle numbering. The amplitude of the first seven of these cycles is controlled at the transmitter to conform to strict standards established by the Coast Guard. The standard (positive-going) waveform begins at time zero, with the maximum amplitude of the first carrier cycle occurring at 2.5 μs, the second at 17.5 μs and continuing to the peak of the seventh carrier cycle at 67.5 μs.

The cycle amplitude $f(t)$ of the standard waveform, expressed as a measure of the antenna current, conforms to the equation [CG81]

$$f(t) = A(t-\tau)^2 \exp\left(\frac{-2(t-\tau)}{65}\right)\sin(0.2\pi + PC) \quad \text{for} \quad (\tau \le t \le \tau + 65),$$

where A is a normalizing constant, t is the time in μs, τ is the envelope-to-cycle difference (ECD) in μs and PC is the phase code parameter in radians. The ECD is determined by a procedure which measures the rms error between the actual LORAN-C cycle amplitudes and the cycle amplitudes implicit in this equation, where the normalizing constant A is set so that the peak amplitudes of the two signals are equal. The rms error ε is then calculated as

$$\varepsilon = \left(\frac{\sum_{n=1}^{8}(I_n - S_n)}{8}\right)^{1/2},$$

where I_n and S_n are the actual and standard amplitudes of the nth carrier cycle. The ECD of I relative to S is that value of τ which minimizes the rms error. With an appropriate value for A and $\tau = 0$, the peak cycle amplitudes of the first eight positive (+) and negative (–) cycles according to the standard waveform are:

Cycle	1	2	3	4	5	6	7	8
Peak+	0.6	13.6	38.1	62.6	81.7	93.8	99.4	99.5
Peak–	4.7	25.3	50.8	72.9	88.6	97.3	100.0	97.9

Due to near/far-field effects, there is a phase shift between the antenna current and the far E-field signal beyond some kilometers from the transmitter. This results in a 90-degree rotation (2.5 μs), ordinarily does not affect the utility of the signal for navigation purposes, but must be taken into account for precise timekeeping purposes.

In order to minimize timing errors, it is desirable to use the groundwave signal, which is usable up to about 1200 nautical miles from the transmitter. However, at the greater ranges and especially at night, the skywave signal can considerably distort the LORAN-C pulse, leading to considerable navigation or timing errors. However, since it must travel a longer path via the ionosphere, the skywave signal is delayed at least 30 μs relative to the groundwave signal. Therefore, the reference epoch of the pulse group is established as the positive-going zero crossing between the end of the third carrier cycle and the beginning of the fourth cycle in the first pulse of the eight-pulse group beginning the A field of the LORAN-C frame.

The domestic LORAN-C chains are synchronized to UTC with respect to station timing and time-of-coincidence (TOC) events. Since the GRI of all chains is an integral multiple of 10 μs, the reference epoch of each GRI precesses over the standard second and can be synchronized to it. TOC

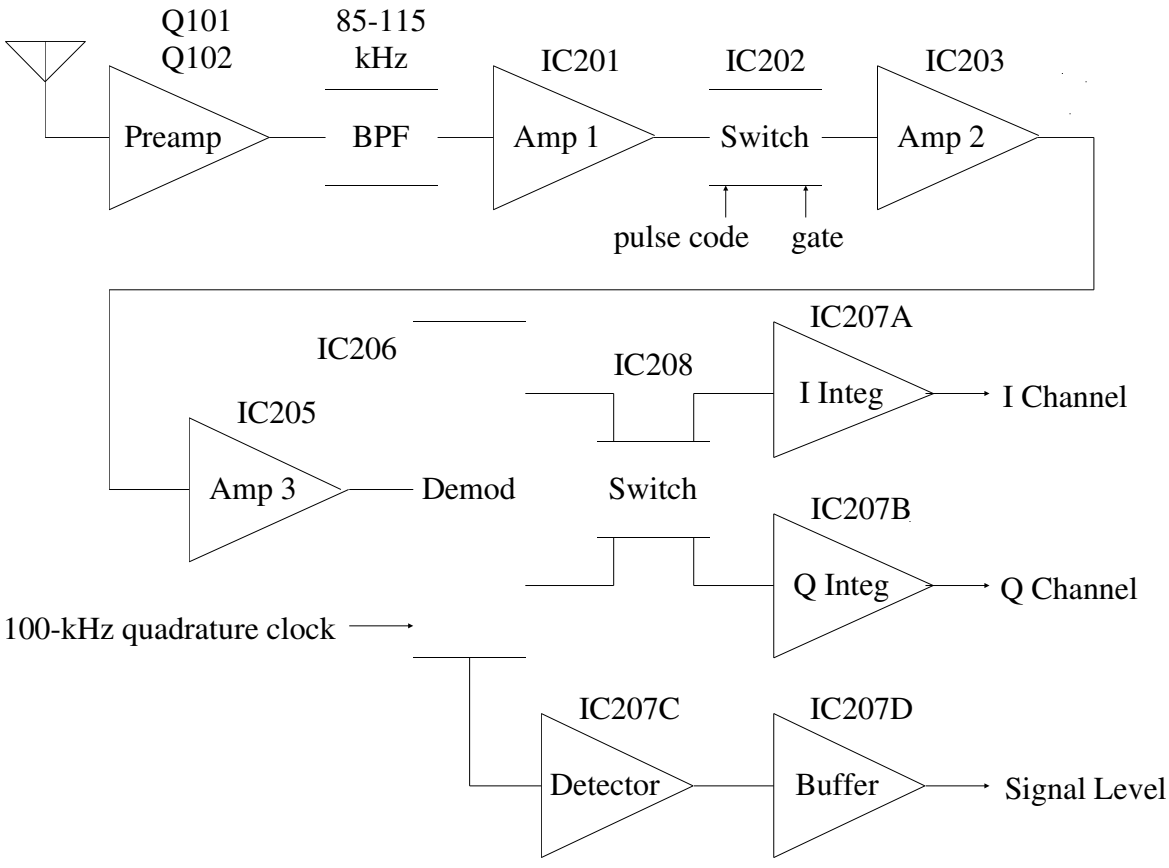

Figure 6. Receiver Block Diagram

events occur at intervals ranging up to a few minutes, 249 s in the case of the Northeast U.S. chain. The conventional procedure is to refer to a list of published TOCs and set a digital clock synchronized to the frequency standard to the time of the next one. The equipment is then initialized to begin counting at the selected TOC and remains synchronized to UTC after that.

3. System Design Description

The receiver design consists of three subassemblies shown in Figure 1: the preamplifier, which is mounted along with a loop antenna at the antenna site; the receiver and oscillator, which are in a shielded, standalone enclosure; and the receiver controller, which is an adapter card designed for the IBM PC/XT/AT host computer or compatible. The functional organization of these three subassemblies is shown in Drawing Sheet 1 in Appendix C of this report. Ordinarily, all components are powered directly from the host computer. For the most demanding applications the receiver, including its oscillator and frequency-controlling circuitry, should be equipped with a backup battery in order to preserve oscillator frequency and phase when the computer is powered down.

3.1. Receiver and Preamplifier

Figure 6 shows a block diagram of the preamplifier and receiver subassemblies. The antenna and preamplifier are connected to the receiver using 50-ohm coaxial cable, which may be up to several hundred feet in length. Power is supplied to the preamplifier using the center conductor of the cable. The antenna consists of a shielded, balanced loop approximately two feet in diameter, which

provides substantial directivity and isolation from local noise sources. The preamplifier uses a pair of dual-gate MOSFETs (Q101, Q102) in a balanced configuration to provide approximately 20 dB of gain. Alternatively, in installations where sufficient signal energy is available, a whip antenna can be connected directly to the matching circuit included in the receiver itself.

The detail design of the preamplifier and receiver-input circuitry is shown in Drawing Sheet 3. The antenna loop is tuned by C101 for a center frequency of 100 kHz. Construction of the loop antenna and preamplifier is not complete at this time; however, while not shown in the drawing, it is likely that the loop will need to be loaded by a swamping resistor in order to lower the Q and obtain adequate bandwidth. The transformer T101 is broadly tuned to 100 kHz by C104. It is implemented using a ferrite toroid and either bifilar or trifilar windings as indicated in the notes. Power to the preamplifier is supplied from an 8-volt regulator (IC204 on Drawing Sheet 2) through T102 and R107 bypassed by C110 and C111. The components are chosen so that no damage will occur if the coaxial cable is opened or shorted.

The receiver itself includes a whip matching circuit controlled by switch S201. This circuit may be used instead of the loop antenna in many areas where the LORAN-C signals are relatively strong. The whip antenna can be a CB antenna mounted on a window sill or wire tossed over a nearby tree. Silicon diodes D101 and D102 protect the circuitry in case of static discharges or strong electro-magnetic fields near the antenna. Components C107, L102 and C108 form a lowpass filter and π-section impedance-matching circuit loaded by R105. Source-follower JFET transistor Q103 and resistor R106 isolate the receiver from the effects of the antenna system and provide the required source impedance (about 500 ohms) for the bandpass filter.

The receiver includes three stages of gain using the MC1590G rf-amplifier chip (IC201, IC203, IC205), a bandpass filter and three low noise analog switches using the ADG509A analog-switch chip (IC202, IC206, IC208). The bandpass filter is a three-pole Butterworth type designed for a center frequency of 100 kHz and 3-dB bandwidth of 30 kHz. The three rf-amplifier stages provide up to approximately 110 dB of gain controllable by the host program. One of the analog switches (IC202), called the signal gate, decodes the LORAN-C carrier phase according to a predetermined pulse code, depending on which master or slave station is selected. The second analog switch (IC206), called the demodulator, synchronously demodulates the received signal. The third analog switch (IC208), called the integrator gate, controls the integration time constant and provides a means to dump the charge on the integrator capacitors when necessary. A signal-level voltage is produced from the demodulator input using operational amplifiers IC207C and IC207D.

The demodulated signal consists of two components, an in-phase (I) component and a quadrature (Q) component. The I and Q signals are processed by separate operational-amplifier integrators (IC207A, IC207B), which integrate over an 8-millisecond period, following which their outputs are sampled by an 8-bit analog/digital (A/D) converter in the receiver controller. The controller also includes two buffered 8-bit digital/analog (D/A) converters and associated operational amplifiers (not shown), one of which provides a DC voltage used to control the demodulator clock frequency (VCO) and the another to control the receiver gain (AGC). The gain and offset of the two D/A converters can be set using potentiometers provided for the purpose. Analog switch IC208 is controlled by a strobe signal and latch included in the receiver controller. The integrator gain can be programmed for three different time constants and the integrator capacitors can be either open-circuited or short-circuited to hold or dump the accumulated charge.

The detail design of the receiver is shown in Drawing Sheet 2 in Apendix C of this report. Components C201, L201, C202, L202, C203 and L203 comprise the bandpass filter. Gain at the rf frequency is provided by amplifiers IC201, IC203 and IC205. The amplifier tank circuits are stagger-tuned to provide an overall response characteristic described later in this report. IC202 provides for three modes of operation: normal phase, inverted phase and short-circuited to the power rail. The last mode is used when the signal gate is off in order to avoid integrator drift. The switch is driven by the phase-code (CODE) and envelope-gate (PCI) pulses generated by the receiver controller.

IC205 drives the synchronous-demodulator IC206, which is connected as a DPDT switch driven by the clock signals P0 and P1 generated by the receiver controller. Signal P0 toggles at a 200-kHz rate, while P1 toggles at a 100-kHz rate, with quadrature decoding done within IC206. Operational amplifiers IC207A and IC207B and capacitors C227 and C228 operate as sample-and-hold integrators which accumulate charge for transfer to the A/D converter in the receiver controller. Analog switch IC208 controls the time constant of integration and provides a path to short-circuit the integration capacitors to bleed off charge following an integration interval. Various data pertaining to IC202, IC206 and IC208 are shown on the sheet.

The demodulator input signal appearing at transformer T201 is rectified by D201 and D202 and amplified by IC207C and IC207D, which operate as a “superdiode” with time constant established at about 1 s by R210 and C216. The resulting voltage is one of the inputs to the A/D converter in the receiver controller and can be used as a level indicator. The receiver gain is controlled by the bias voltage applied at pin 2 of IC203 filtered by C211 and C299. This voltage is generated by a D/A converter in the receiver controller.

The receiver includes a precision frequency standard X201 in the form of an oven-controlled, quartz-crystal oscillator (Isotemp OCXO107) with frequency adjustable over a narrow range by a DC voltage. All timing signals used by the receiver and receiver controller are derived from this oscillator, which operates at a frequency of 5 MHz and generates both sinewave and TTL outputs. The oscillator has a rated stability of:

Temperature: $\pm 5 \times 10^{-9}$ from +5 deg C to +55 °C (referenced to +25 °C)
Aging: 10^{-9} after 30 days of operation
Voltage: $\pm 1 \times 10^{-8}$ from 10.8 VDC to +14.0 VDC
Frequency control: ±4 Hz over the range 0 to 8 VDC

In addition to the 5-MHz output, the receiver includes BNC connectors for the envelope gate PCI (J205), Q-integrator output S1 (J204) and demodulator input SIG (J206). These signals are intended to be used with an oscilloscope for receiver alignment and signal monitoring. Ordinarily, J205 is connected to the horizontal trigger input of the oscilloscope, J206 to the channel-1 vertical deflection input and J204 to the channel-2 vertical deflection input. With the oscilloscope controls set for alternate sweep, this arrangement allows the operator to view the leading edge of the despread LORAN-C signal and verify the position of the cycle strobe, which should straddle the positive-going zero crossing following the third carrier cycle of the envelope pulse.

3.2. Receiver Controller

Figure 7 shows a block diagram of the receiver controller, which is designed to operate with an IBM PC/XT/AT computer or compatible. It provides the counters, registers and other logic necessary to generate the LORAN-C pulse code, gating signals and demodulator clock. It includes a program-

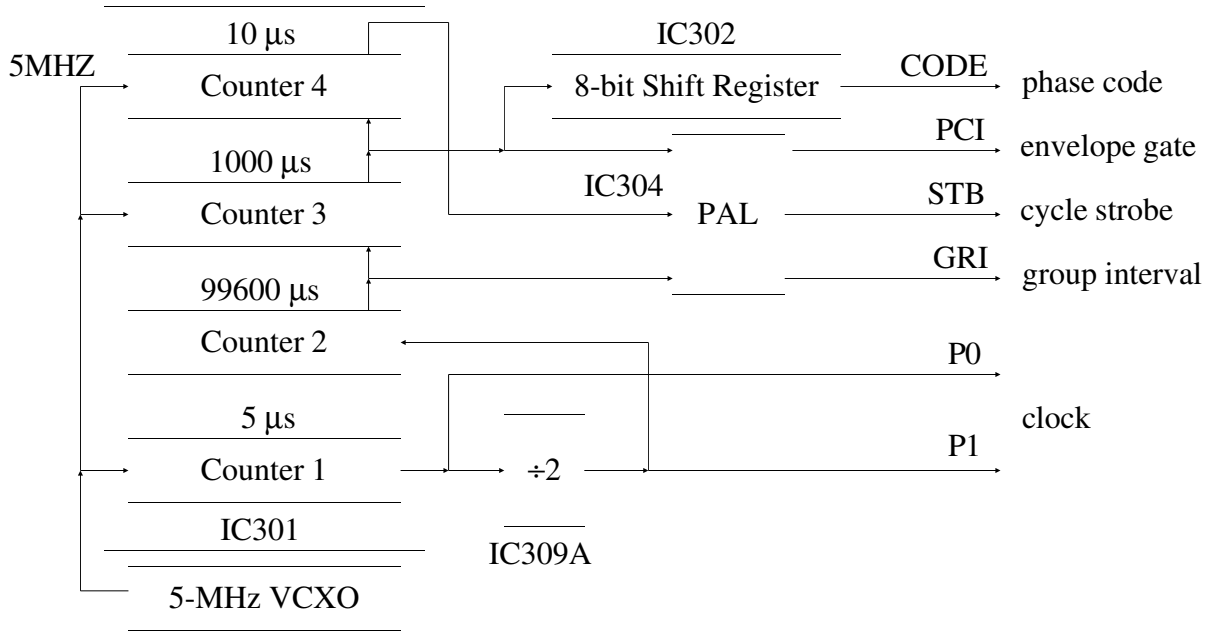

Figure 7. Receiver Controller Block Diagram

mable, multifunction timing generator based on the AMD9513A chip (IC301), a programmable buffer/shift register used to generate the phase code (IC302) and a programmed-array logic PAL (IC304) to generate the receiver gating signals. Not shown are an 8-channel, 8-bit A/D converter based on the ADC0809 chip (IC305), a two-channel D/A converter based on the AD7528 chip (IC307), an 8-bit parameter register (IC303) and miscellaneous supporting logic.

The various signals generated by the timing generator chip are shown on Drawing Sheet 6. Operating from the 5-MHz TTL output from the receiver, the timing generator and flipflop shown produce the 100-kHz and 200-kHz clock signals used by the demodulator, the group interval (GRI) pulse, the envelope gate (PCI) pulse and cycle strobe (STB) pulse. The GRI waveform (Figure 1 on Drawing Sheet 6) consists of an idle interval followed by an 8-ms pulse. This cycle repeats at an interval characteristic of the selected LORAN-C chain, in this case 99,600 µs for the Northeast U.S. chain. The PCI waveform consists af a 300-µs pulse which repeats at 1000-µs intervals as enabled by the GRI pulse, or eight times for each field. The CODE waveform controls the phase of eight PCI pulses as shown. The STB waveform consists of a 10-µs pulse delayed from the rise of the PCI pulse. The amount of delay, which corresponds to a selected cycle of the 100-kHz LORAN-C pulse envelope in the PCI interval, can be controlled by the program. The demodulator gate signal ENB is used during acquisition and tracking phases to gate the I and Q integrators. It is derived from the GRI, PCI, STB and other signals as determined by the PAL and the parameter register (not shown).

Drawing Sheets 4 and 5 show details of the receiver controller. The timing generator IC301 generates outputs OUT1 through OUT5 labelled as shown on Drawing Sheet 4. The selected LORAN-C pulse code is generated by an 8-bit shift register IC302, which can be loaded by the host program. The phase code causes the received signal to be processed normally for a zero bit and phase-flipped (180 degrees) for a one bit. The shift pulse (PLX) for this register is generated by PAL IC311 and the load/shift level (ENG) selected by a bit of the parameter register IC303. The PCI pulse is generated by PAL IC304 and used as the envelope gate to sample the received signal envelope. The PAL logic

equations are shown next to it on the sheet. The falling edge of the GRI pulse, which occurs after eight shifts, triggers the A/D converter IC305. The converter can also be triggered directly by the host program using a bit of the parameter register IC303. When conversion is complete, a host-computer interrupt is requested, following which the host program reads the A/D converter buffer and loads the pulse code and possibly other parameters to be used in the following field.

IC305 is an eight-input, 8-bit A/D converter with internal analog multiplexor and sample-and-hold features. only three of the inputs are used, one each for the I integrator, Q integrator and level indicator. Potentiometer R303 controls the reference voltage for the A/D converter and normally is set at 5.0 V. The two-channel, 8-bit D/A converter IC307 and operational amplifiers IC308A and IC308B provide signals which control the receiver gain (AGC) and oscillator frequency (VCO). Potentiometers R305 and R307 control the offsets in these two channels, respectively, while potentiometers R306 and R309 control the gains.

Of the remaining components shown on Drawing Sheet 4, flipflop IC309A is used to generate the 100-kHz component of the demodulator clock and IC306 is a bidirectional transceiver for the PC data bus. IC307 is a two-channel D/A converter buffered by operational amplifiers IC308A and IC308B. IC303, which functions as the parameter register, is an 8-bit latch used to hold various parameters controlling receiver operation. The components on Drawing Sheet 5 include various ancillary PAL and logic components, with the PAL logic equations next to them on the sheet. IC310 is used to generate chip-select signals, while IC311 is used to generate read/write select and related signals. Flipflop IC309B, buffers IC312A, IC312C and IC312D and inverters IC314A and IC314B provide interrupt signals for the PC bus upon completion of an A/D conversion. Selection of the interrupt line is by dipswitch S301. Buffer IC312B is used to provide a clean logic-high signal for various purposes. Flipflop IC313A can be used in conjunction with Counter 5 to generate interesting signals for system test and debugging, but is not used in regular operation. Fliflop IC313B generates a signal used to gate the integrators during the period between GRI pulses.

Power for the receiver, preamplifier and analog circuitry in the receiver controller is supplied by an 8-volt regulator IC204 in the receiver. All analog grounds are returned to the common pin 2 on this component. Zener diodes D204 in the receiver and D301 in the receiver controller, together with their associated bypass capacitors, provide a regulated bias source which allows use of a single 12-volt power supply, normally provided by the computer. Fuse F1 in the receiver controller protects the controller traces in case of a cable or receiver short. The regulator has internal protection against shorts on the 8-volt power bus.

3.3. System Timing and Sequencing

A timing diagram for the signals used by the PC bus interface are shown in Figure 4 of Drawing Sheet 6. Only the signals necessary for the bus handshaking operations are shown, including the 12-bit address bus A, address-enable (AEN), read-select (RD_) and write-select (WR_) signals. A bus cycle begins when the CPU places an address on the address bus and lowers AEN. Data are transferred from the device to the CPU on the trailing edge of RD_ and from the CPU to the device on the trailing edge of WR_.

In normal operation the program initializes the counters and registers of the receiver controller to automatically integrate over the next field, measure and convert the charge on the I integrator and interrupt the program. The program reads the result and initiates conversion of the charge on the Q integrator. The program then reads the result and initiates conversion of the signal-level voltage.

Finally, the program reads the result and initializes the counters and registers for the next field. A timing diagram for the detailed waveforms required for these operations is shown in Figure 3 on Drawing Sheet 6.

The GRI pulse repeats continuously during system operation. The rise of this pulse is shown at time *a* on the diagram. The ENG signal is program-controlled as bit 8 of the parameter register IC303. When high, the A/D converter IC305 is started automatically at the fall of the GRI pulse; while, when low, the converter is started by the program. The fall of the STRT signal at pin 6 of IC305 initiates conversion, which is shown on the diagram as the result of the GRI pulse at point *b* and of the program at point *f*. The EOC signal is asserted high at pin 7 of IC305 when conversion is complete, while the OE signal is asserted high at pin 9 of IC305 when the converter is read by the program. The DONE signal originates at flipflop IC309B, which is set by the rise of the EOC signal and cleared by the rise of the OE signal.

The sequence of operations is as follows. Previous to point *a* the program has set the A/D converter multiplexor to channel 0 (I integrator). Upon completion of the field at point *b*, the converter is started automatically. When complete at point *c*, a program interrupt is signalled, which causes the program to read the A/D converter buffer at point *d*. Next, the program lowers the ENG signal at point *e*, sets the multiplexor to channel 1 (Q integrator) and starts a second conversion at point *f*. When complete at point *g*, the program sets the multiplexor to channel 2 (level) and starts a third conversion to read the signal level (not shown). Finally, the program raises the ENG signal and sets up for the next field.

The GAT signal is generated by flipflop IC313B to control the integrator sample-and-hold function. It set by the rise of the ENG signal and cleared by the rise of the GRI pulse. When high, the integrator capacitors are shorted, preventing the accumulation of charge due to leakage or bias currents in the operational amplifiers. When low, the PCI1 signal generated by PAL IC304 gates the integrators to accumulate charge at the rate selected by bits 0 and 1 of the parameter register, including .036 ms, 0.264 ms and 1.0 ms. A fourth selection shorts the capacitors in order to dump the accumulated charge. In addition, bits 2 and 3 of the parameter register select which gating signal to use for the integrators, including the GAT, PCI and STB pulses plus always-open. Ordinarily, the PCI pulse is used during pulse-group search and the STB pulse during cycle search and tracking.

Details of selected counter operations are shown in Figure 2 on Drawing Sheet 6. They show the relationships between the 100-kHz signal counted by counter 2, the 5-MHz signal counted by counter 3 and the GRI, PCX and PCI pulses. Numbering of the pulses begins at the value programmed into the LOAD and HOLD registers and shown on the diagram, then decrements to a phase called terminal-count (TC), which occurs at the rise of what would be pulse number zero of the LOAD interval, and then continues at the value programmed in the HOLD register. The TC falls one clock cycle later, which toggles the output of the counter to produce the high level. This output persists until the next TC, at which time the output is toggled to produce the low level and the cycle repeats.

Note that the PCX (inverted PCI) output of counter 3 can be as much as 400 ns late relative to the GRI output of counter 2, due to way in which the TC operates and the setup times for the counter gates. Ordinarily, this causes no problems, since the sampling intervals established by the STB output of counter 3 used during cycle identification and tracking are toward the middle of the PCI interval.

3.4. Analog System Simulation

Major components of the LORAN-C Timing Receiver have been simulated using PCSPICE. Drawing Sheets 7 and 8 in Appendix C show the simulation models, with results shown in Figures 8 through 10 of this report. Sheet 7 shows the model used to refine the gain and passband characteristics, but includes only a crude model for the demodulator and integrator circuits. Source-follower JFET J1, shown as a 2N3819, has characteristics similar to the MPF102 shown as Q103 on Sheet 3. Voltage-controlled voltage sources E1 through E3 represent the MC1590G amplifiers shown as IC201, IC203 and IC205 on Drawing Sheet 2. The various resistors and gain parameters were chosen to match the characteristics shown on the data sheets for these devices. All inductors and transformers shown include a resistance chosen to represent the typical loss characteristics associated with ferrite-core components operating at 100 kHz.

Figure 5 shows a typical LORAN-C signal pulse as transmitted. For the purposes of simulation, this waveform is approximated by

$$s(t) = at\, \mathrm{e}^{-at}\sin(2\pi 10^5 t)\ ,$$

where $a = 14286$ to match the LORAN-C signal waveform such that the maximum amplitude is reached at the seventh carrier cycle, or 70 μs after the beginning of the pulse. This signal is synthesized by the product of the voltage sources VCAR, VEN1 and VEN2 shown on Sheet 7. VCAR generates the 100-kHz carrier, while VEN1 generates the linear component and VEN2 generates the exponential component. The voltage-controlled voltage source EENV combines the linear and exponential components to produce the pulse envelope, while the voltage-controlled voltage source EEIN combines the carrier and envelope components to produce the actual signal.

Drawing Sheet 8 shows the model used to refine the demodulator and integrator circuits. For the purposes of simulation, the 2N3819 JFETs have characteristics similar to the ADG509A analog switches shown as IC206 and IC208 on Drawing Sheet 2. The LORAN-C signal is produced in the same way as on Sheet 7. The voltage sources V1, V2 and V3, together with the voltage-contolled voltage sources E1 and E2, synthesize the quadrature clock that drives the demodulator switches J1 and J2 and the integrator A1A, here representing IC207A and IC207B shown on Sheet 2.

The simulation is primarily of interest in order to determine the bandwidth and envelope delay of the receiver. In order to maintain the system specification flat within 6 dB over the range 90-110 kHz, the input matching circuit is swamped with resistor R1 and the source impedance of the bandpass filter is controlled at 500 ohms, which is also the output impedance of the source follower. The termination impedance of the bandpass filter is established at 500 ohms by the parallel combination of R3 and the input impedance of the first RF amplifier E1 (IC201 on Drawing Sheet 2). In order to maintain adequate bandwidth it is necessary to stagger-tune the first two RF amplifiers, one E2 (IC203) to 85 kHz and the other E3 (IC205) to 115 kHz, as well as swamp the last RF amplifier tank circuit with R10 (R207).

The results of simulation show an overall gain of about 110 dB with 1μV at the preamplifier input. The overall selectivity is shown in Figure 8, from which it is evident that signals an octave above and an octave below the 100-kHz working frequency are attenuated at least 90 dB. Figure 9 shows the bandpass characteristics are uniform within about 5 dB over the range 90-110 kHz. Figure 10 shows the resulting signal at the receiver demodulator. The pulse has been delayed approximately 50 μs and the risetime has been degraded somewhat due to the bandpass characteristic. In practice,

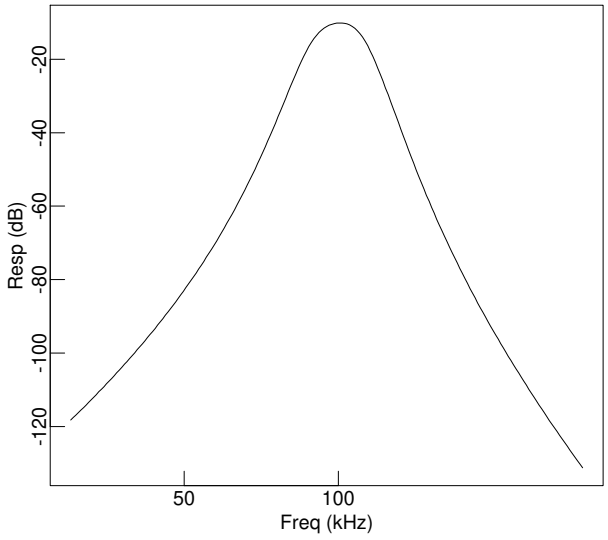

Figure 8. Receiver Selectivity

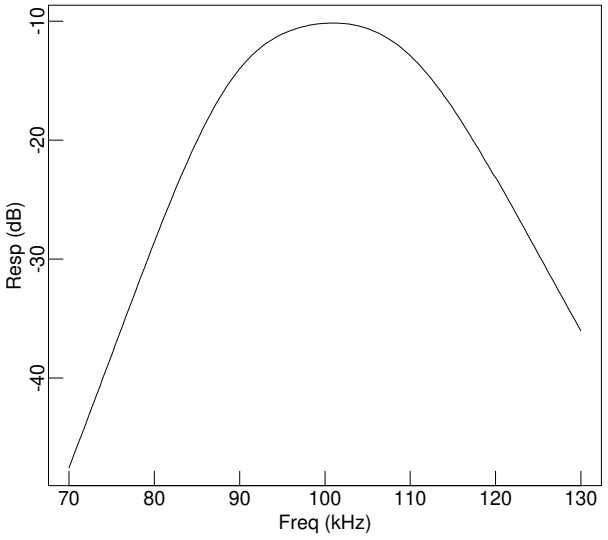

Figure 9. Receiver Bandpass

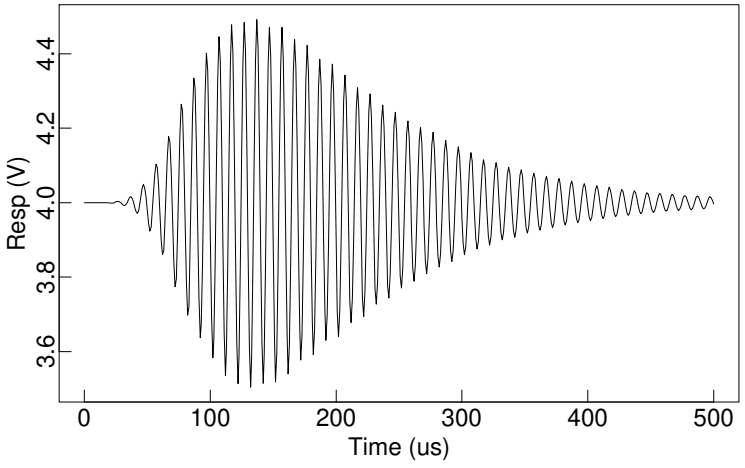

Figure 10. Receiver Output

these characteristics are likely to be modified by component tolerances, final alignment and gain adjustment. System gains and bandwiths of this order are normally realized only with extreme care in circuit layout, shielding and bypassing.

Further study of the signal waveform reveals useful data for the design of the integrators. The demodulator produces a pulse represented by the envelope $t\exp(-t)$ multiplied by the carrier $\sin(\omega t)$. Since the envelope consists of many carrier cycles (about 50), the integral can be approximated by the product of the integral of the envelope multiplied by the factor $\frac{2}{\pi}$ due to the full-wave detected carrier. The integral of the envelope over the complete pulse is calculated to be e times the peak value and both the peak of the pulse and the zero-crossing of the derivative occur at $t = 1$. In one of the tracking modes the signal is integrated only over the rising edge of the pulse, while in another it is integrated only over the third carrier cycle. By numeric integration the scale

factors are found to be 0.264 and .0355 respectively. These factors are reflected in the gain parameters selected by analog switch IC208.

4. Software Operations

The Loran-C Timing Receiver is controlled by a software program designed for the IBM PC/XT/AT or compatible host computer. The control program requires only minimal memory and no disk resources. Since most of the arithmetic is in floating point, use of a floating-point chip is required, at least with the 386SX or slower computers. It is not clear whether the PC or XT, with or without the chip, is fast enough for the operations required.

After initialization the program executes a continuous loop driven by the ready signal of the A/D converter. Since only one converter and one 8-bit pulse-code buffer is provided, it is necessary to timeshare it for the I channel, Q channel and level channel and for both the A and B fields separately. A single frame of two fields thus produces six interrupts, three for the A interval and three for the B interval.

The program operates in two phases, one for the A interval and the other for the B interval. Separate software accumulators are provided for the I, Q and level channels and cleared at the beginning of the A interval. The three A/D converter channels are read at the end of the next A and B intervals and added to the accumulators. At the end of the B interval the program calculates the envelope amplitude and phase from the I and Q accumulators. In order to avoid overrun due to the rather long time to convert and display monitoring information, the display function is executed only during a B interval when otherwise very little is happening.

The program begins by initializing the timing generator and various counters, time-constants and signal information. It reads the GRI, master/slave switch and initial gain and oscillator D/A converter biases from the command line. Normally, it then automatically sequences through six modes of operation in order as it searches for the reference cycle and phase of the LORAN-C signal. For testing and experimentation purposes, the operator can redirect this sequencing and adjust the timing and gain parameters. The program operations in each of these six modes are described in following sections.

4.1. Mode 1: Receiver Offset Calibration

In this mode, which is normally the first in the sequence, the receiver gain is set to minimum and the program waits a few seconds for the receiver circuitry to stabilize. It then reads and averages the I, Q and level channels for a number of fields in order to obtain reliable offset data. In subsequent operations these offsets are subtracted from received signals in order to suppress residual integrator and converter biases. When the average error has decreased below a threshold, the program automatically sequences to mode 2.

4.2. Mode 2: Receiver Gain Calibration

In this mode the receiver gain is set to maximum, and the program waits a few seconds for the receiver circuitry to stabilize. This normally results in amplifier or demodulator saturation due to various signals and noise processes appearing at the receiver antenna. The program then reads and averages the level channel in order to establish the maximum level signal, which is used to calibrate the gain transfer characteristic. When the average error has decreased below a threshold, the program automatically sequences to mode 3.

4.3. Mode 3: Signal Level Calibration

In this mode the receiver gain is set to a value calculated from the minimum and maximum level channel values and the program waits a few seconds for the receiver circuitry to stabilize. It then measures the level channel under nominal operating conditions and adjusts the receiver gain for optimum signal/noise ratio. The level channel itself, from which the receiver gain signal is derived, is sensitive to the peak pulse energy in the 90-110 kHz spectrum. The program executes a search procedure in order to set the gain so that it operates in the linear region for LORAN-C signals, but clips noise pulses more than about twice the pulse peak. When the average error has decreased below a threshold, the program automatically sequences to mode 4.

4.4. Mode 4: Pulse-Group Search

In this mode the receiver gain is fixed at the value calculated in mode 3 and the oscillator frequency is fixed at a nominal value close to 5 MHz. During this and preceding modes the envelope gate has been set to the PCI pulse, which consists of a group of eight 300-μs pulses at frequency of 1 kHz and the group itself repeating each field. The receiver is enabled only for the duration of each pulse, so accumulates signal energy only during that time.

The program then executes a search in order to find the correct position of the LORAN-C pulse group relative to the receiver timing signals. The search is designed to accumulate the I and Q signals for each 100-μs bin over the entire frame interval consisting of both the A and B fields. Each bin thus contains the accumulation of sixteen envelope samples despread by the selected master or slave pulse code. Thus, as many as 2000 bins, depending on the GRI, are produced for a complete search. This can take almost seven minutes at the longest GRI, during which the nominal oscillator offset relative to LORAN-C timing must be held to within about 100 μs. In order to insure this, the oscillator must have an inherent frequency error less than 0.25 ppm.

While the inherent processing gain achieved by the despreading and synchronous envelope integration results in a considerable improvement in the signal/noise ratio, additional linear and nonlinear processing is needed in order to improve the ratio still more. Note that there are three bins within the 300-μs envelope gate interval, each separated by one frame interval. In the case of a pulse group from a selected chain, there will be some offset relative to the LORAN-C frame in which all three bins will contain substantial energy, while the remaining bins will contain the results of integrating noise or pulses from another chain operating at a different GRI.

As long as no other chain in the same area is operating at a GRI within 300 μs of the selected chain, it is not possible that another chain could produce more than one pulse group within the envelope gate. The program implements a noise gate which requires for each window of three envelope samples received that the center one have a relative amplitude less than three times the average of the other two samples. If not, it is replaced by the average of the other two. This technique has been very successful in reducing cross-rate interference and results in a substantial number of cross-rate pulse groups being suppressed before any linear filtering is done.

In addition to the noise gate, the program implements a matched-filter edge detector adjusted for a 300-μs signal interval in a 900-μs noise interval. It is implemented using a nine-stage shift register, one stage for each 100-μs sample from the noise gate. The output of the filter consists of the sum of the three center samples less the average of the remaining six samples.

The maximum filter output found during the entire pulse-group search represents the location of the selected LORAN-C signal relative to the receiver timing. If the calculated signal/noise ratio at the conclusion of the pulse-group search falls below an equivalent unprocessed signal/noise ratio of -5 dB, the program remains in mode 4 and executes another pulse-group search. If the signal/noise ratio exceeds this value, the program initializes for the cycle search and automatically sequences to mode 5.

4.5. Mode 5: Cycle Search

In this mode the receiver gain is fixed at the value calculated in mode 3 and the oscillator frequency is fixed at a nominal value close to 5 MHz. The program executes a search of each 10-μs cycle of the 300-μs envelope gate. For this purpose the signal gate IC202 (see Drawing Sheet 2) is controlled by the pulse-code pulse PCI, while the integrator gate IC208 is controlled by the strobe pulse STB. The strobe signal is a 10-μs pulse that can be delayed from the beginning of the envelope gate pulse in increments of 10 μs. The program operates to rotate the strobe through all 30 cycles in the envelope gate, accumulating certain signals for each cycle separately.

There are three signals of interest, the cycle amplitude, cycle phase and cycle error. The cycle amplitude signal is calculated as the square root of the sum of the squares of the I and Q channels. The cycle phase signal is taken directly from the I channel. The program processes each of these signals using multistage median filters for each cycle of the pulse. For this purpose the program maintains two shift registers for each cycle plus a single temporary array used to sort the samples. The filter output is taken from the center of the sorted array and the jitter determined as the difference between the first and last element of the array.

The median filters reduce errors due to two causes. One cause is the use of dual-rated LORAN-C stations. A station is dual-rated if it can transmit pulses for more than one LORAN-C chain. This reduces the number of stations required to service overlapping areas at a cost of somewhat higher average power requirements. However, a dual-rated station cannot transmit two pulses at the same time, so occasionally deletes a pulse for one chain while transmitting a pulse for another. This leaves holes in the chain pulse schedule which only modestly affects the processing gain at the receiver. However, the median filter makes up for this loss in most cases by synthesizing the deleted pulses.

The second cause of error is due to the fact that subsequent cycle integrations are from different fields, so matching of the envelope signal from one cycle to the next can be uneven. This is primarily due to signal level fluctuations from one LORAN-C frame to the next caused by cross-rate interference, impulse noise and emmisions from stations in other services adjacent to the LORAN-C band. The median filters are quite effective in dealing with impulse noise, particularly with the atmospheric noise bursts prevalent at the 100-kHz LORAN-C radio frequency, since a burst that affects one frame will seldom affect the next.

The cycle error signal is calculated from the (normalized) differences between each carrier cycle in a seven-cycle window, beginning at a selected cycle number, and the corresponding cycle of the LORAN-C model envelope. The cycle error signal, computed as the square root of the sum of the squared differences, is carried out for each cycle separately.

As samples of the cycle amplitude and phase are computed, they are integrated into exponentially-weighted averages separately for each of the 30 cycles of the envelope gate. The time constant controlling the integration process varies depending on the sample variance and history of the process and varies from one cycle scan to several hundred cycle scans, which corresponds to

averaging times up to several minutes. The time constants for the cycle amplitude and phase, as well as the size of the median filters are varied together.

4.6. Mode 6: Phase Tracking

In the mode-5 cycle search the program continuously scans all or some part of the envelope gate interval called the cycle-search window. The minimum cycle error over all cycles in the window establishes the reference cycle, while the zero crossing of this cycle establishes the LORAN-C timing relative to the receiver timing. However, especially during the initial cycle scans when the window is largest and time constant of integration is relatively small, the cycle number of the minimum cycle contains considerable jitter. A multistage cycle-number median filter is used to reliably establish the correct reference cycle.

At intervals roughly equal to the time constant of integration, the program checks the reference-cycle jitter calculated during the window scan. If it is zero or one, there is a high probability the correct reference cycle number has been found. If so, the number is checked against the limits of the envelope gate. If too close to the beginning or end of the gate interval, the envelope gate is slipped an integral number of carrier cycles to align the reference cycle near the center of the window and the program operates as at the beginning of mode 5.

If the reference cycle is safely inside the envelope gate and the jitter is satisfactory, the program automatically sequences to mode 6, in which the receiver is tracking the correct cycle and cycle phase. In this mode the cycle-search window extends only one cycle leading the reference model envelope to one cycle lagging it for a total of nine cycles. Since the window scan produces the index of cycle one in the model envelope, the correct phase-sampling point is cycle three and the correct amplitude-sampling point is cycle seven. It is these values that are actually used to control the oscillator frequency and receiver gain.

Since the demodulator operates over the full 10-μs carrier cycle, it averages both the positive-going and negative-going zero crossings of the reference cycle. While this design does not strictly conform to the LORAN-C waveform specification, the timing errors introduced should be negligible. The polarity of the phase signal determines whether the receiver tracks first the positive-going and then the negative-going zero crossing or vice-versa. Since the cycle amplitudes in the LORAN-C standard waveform are so similar for the positive-going and negative-going envelopes, the sign ambiguity must be resolved externally.

4.7. Auxiliary Functions

The program maintains a monitor display which shows the cycle number, amplitude, phase and error of each cycle scanned in the cycle-search window. The display rotates each cycle in turn upon completion of the window scan, so that a complete picture of the window is produced for every 9 or 30 window scans. Occasional messages indicate the status of the receiver and its operating mode. Additional, related information can be produced on demand using operator commands.

While the LORAN-C Timing Receiver normally operates unattended and automatically, it can be controlled manually for testing and experimenting. There is a suite of operator commands in the form of individual keystrokes. Following is a brief list of those commands presently implemented.

The following commands control the phase of the receiver frame relative to the received LORAN-C signal. They are normally needed only for manual signal acquisition.

+ shift oscillator frequency offset up 10 μs offset per GRI

- shift oscillator frequency offset down 10 μs offset per GRI

0 restore oscillator frequency to zero offset

] advance oscillator phase +10 μs or +100 μs (depending on mode)

[retard oscillator phase -10 μs or -100 μs (depending on mode)

The following commands adjust the receiver A/D and D/A gains and D/A biases. The exact values are determined during initial receiver alignment and compiled in the source code and normally need not be changed in regular operation.

} increase program gain (A/D) by a factor of 1.1

{ decrease receiver gain (A/D) by a factor of 0.9

) increase receiver gain (D/A) by +1

(decrease receiver gain (D/A) by -1

] increase oscillator bias (D/A) by +1

[decrease oscillator bias (D/A) by -1

The following commands select which pulse codes are used (master or slave) and determine which set (A or B) to use. They are normally needed only for manual signal acquisition.

m use master pulse codes

s use slave pulse codes

x exchange A and B pulse codes

The following commands select the envelope gate and integrator gain. These are normally needed only for manual signal acquisition.

u switch to ungated mode (for testing only)

g use GRI pulse to control both the signal and integrator gates (for testing only)

p use PCI pulse to control both the signal and integrator gates (modes 1 through 4)

e use the PCI pulse to control the signal gate and the STB pulse to control the integrator gate (modes 5 and 6)

l lock the D/A converters at their present value (used for receiver alignment and initial calibration only)

The following commands establish the receiver mode. Normally, the acquisition process sequences through these modes automatically through all the modes in the order below. These commands can be used to restart the acquisition and tracking processes at any point.

1 switch to mode 1 (receiver offset calibration)

2 switch to mode 2 (receiver gain calibration)

3 switch to mode 3 (signal level calibration)

4 switch to mode 4 (pulse group search)

5 switch to mode 5 (cycle search)

5. Performance Analysis

Since very long integration times are involved, it is necessary to look closely at the stability of the phase-lock loop (PLL) which controls the receiver oscillator. The loop consists of the phase detector, which is represented by the I-channel signal, the oscillator itself, which is voltage controlled, and a loop filter, which determines the static and dynamic behavior of the PLL. The following sections discuss the small-signal analysis, which determines the static and dynamic behavior of the PLL, and the large-signal analysis, which determines the capture and tracking ranges.

5.1. Small-Signal Analysis

In this application the PLL is designed as a type-I, second-order system, which is characterized by the (closed loop) transfer function [SMI86]

$$\frac{\theta_o}{\theta_r} = \left(\frac{s^2}{\omega_n^2} + 2\frac{\zeta}{\omega_n} + 1 \right)^{-1},$$

where θ_o is the output phase from the VCO, θ_r is the reference phase from the LORAN-C signal and ω_n and ζ are related to the loop gain K_v and loop-filter corner frequency ω_L:

$$\omega_n^2 = K_v \omega_L \quad \text{and} \quad \zeta = \frac{1}{2}\left(\frac{\omega_L}{K_v} \right)^{1/2}.$$

For a first-order filter function $\omega_L = \frac{1}{\tau}$, where τ is the time constant of integration. For a critically damped system ζ should be equal to $2^{-1/2}$, which implies that the product $\frac{1}{2K_v \tau}$ should be unity. This means that, as the time constant of integration is increased, the loop gain must be decreased proportionally in order to maintain the same dynamic characteristics.

The loop gain K_v, sometimes called the velocity factor, is calculated as follows. The program maintains the gain of the receiver to deliver a certain value, in this case 100 program units, at the peak of the pulse at the maximum (seventh) carrier cycle. At this level the amplitude of the reference (third) cycle is about 50 units. Since the slope of the phase characteristic at the zero crossing is unity, the transfer function of the phase detector, including all the gain provided by the hardware and program, is 50 unit/rad. The D/A converters operate over a range of 256 units and deliver an output voltage of 6 V over this range, for a transfer function of .0234 V/unit. The VCO has a rated sensitivity of 1 Hz/V, which is reduced by a factor of ten by a resistor network, for a transfer function of 0.628 rad/V-s. The ratio of the 100-kHz demodulator clock to the 5-MHz VCO frequency is 1/50. Therefore, the overall PLL gain K_v is the product of these factors: .0147 rad/s.

The corner frequency ω_L is the reciprocal of the time constant of integration τ; so, for critical damping the largest τ that can be used without changing K_v is 34, at which the PLL bandwidth is 4.68 mHz (milli-, not mega-). For τ 34 the program operates in an underdamped mode, while for

τ ~34 the program reduces K_V in inverse proportion to τ. In addition, the program controls the size of the median filters so that the delay introduced by the filters is not more than a fraction (0.125) of the time constant and so does not materially affect the dynamic behavior.

The signal/noise performance of the receiver can be determined by the processing gain and signal waveform. Performance must be established for the pulse-group search (mode 4) and cycle-search (mode 5) and phase-tracking (mode 6) modes of operation. In mode 4 the receiver is gated with a gate time of 300 μs and a repetition interval of about 10 ms, which yields a processing gain of 25.2 dB. The synchronous demodulator adds 3 dB to this figure; however, the signal energy is accumulated over only 100 μs, which subtracts 4.7 dB. Over the complete LORAN-C frame 48 coherent pulses are accumulated for each 300-μs integration bin, which adds 16.8 dB for a total of 40.3 dB processing gain in mode 4. In modes 5 and 6 the gate time is reduced to 10 μs, which yields a processing gain of 40 dB. In addition to the 3 dB contributed by the synchronous demodulator, each envelope cycle includes eight signal samples, which adds another 12 dB for a total of 55 dB in modes 5 and 6.

Assuming for the moment that the signal pulse has energy E_b and the noise is Gaussian with rms value N_0 over the 300-μs envelope gate, the signal/noise ratio is

$$\text{SNR} = \frac{E_b}{N_0} .$$

For binary coherent detection and an error rate of 10^{-3}, which seems reasonable here, the SNR should be at least 7 dB or a ratio of 5. In mode 4 the processing gain within the envelope gate includes 3 dB for the demodulator, -4.7 dB for the signal waveform and 16.8 dB for coherent integration, for a total of 15.7 dB. Thus, for an error rate of 10^{-3} or better, the SNR at the receiver demodulator should be -8.3 dB or greater, which is similar to figures quoted for other receivers [FRA83]. In practice the noise is much more likely to be atmospheric noise bursts or cross-rate interference, which has a character distinctly non-Gaussian; however, this model is useful in comparing performance between different designs.

In modes 5 and 6 the processing gains of the median filter and first-order filter can be estimated from the time constant of integration τ, or 23.9 dB for $\tau = 250$ s. In mode 5 the processing gain within the integrator gate adds 3 dB for the demodulator and 16.8 dB for coherent integration for a total of 40.7 dB. In mode 6 the amplitude of the reference cycle is about half that of the peak, so the processing gain is reduced by 6 dB for a total of 34.7 dB. These figures apply for each cycle separately in the envelope gate or cycle window.

5.2. Large-Signal Analysis

The small-signal acquisition and tracking behavior of the LORAN-C Timing Receiver is determined by the loop gain and time constant of integration τ, as discussed previously. The large-signal behavior is dominated by the gain of the VCO D/A converter and VCO itself, together with the maximum voltage limits of these devices. The *capture range* of a PLL is the maximum frequency range over which the PLL will achieve phase-lock, having previously been not in lock. The *tracking range* is the maximum fractional frequency range over which the PLL will remain locked, once phase-lock has been achieved.

A good PLL design has a capture range large enough to handle the maximum open-loop VCO frequency error. Ultimately, the PLL capture range is limited by the maximum sampling rate, which

is itself limited by the LORAN-C frame rate to about 5 Hz. This corresponds to the Nyquist frequency for a PLL bandwidth of 2.5 Hz at the VCO frequency of 5 MHz, or a capture range of 5×10^{-7}. In the prototype receiver the maximum voltage swing produced by the D/A converter is 6 V, which is reduced by a factor of ten by a resistor network. The reason for this network is to reduce the converter quantization error, as described later. Since the VCO has a rated sensitivity of 1 Hz/V, the PLL capture range is limited to 0.6 Hz at 5 MHz, or 1.2×10^{-7}, which is considerably less than the ultimate and about half the 2.5×10^{-7} required for reliable pulse-group acquisition. The capture range could be improved considerably if a 12-bit D/A converter were used, rather than the resistor network and 8-bit converter in the present design.

In order to achieve a capture range of 0.6 Hz at 5 MHz, the sampling rate must be at least the Nyquist rate or 1.2 Hz; however, in the present design the maximum sampling rate is limited to one sample every nine fields, or 0.556 Hz. This limits the actual capture range to 5.56×10^{-8}, which is uncomfortably small. One way to achieve a higher capture range is to change the design so that, immediately upon entering the cycle-search mode, the program continuously samples the reference-cycle phase, at least until the PLL locks and the phase error decreases to a manageable value. However, there are other ways which involve combined frequency-phase loops, which may be a topic for future study.

The ultimate lower bound on stability is the intrinsic jitter and wander of the VCO itself, plus the quantization error of the D/A converter. The VCO used in the prototype receiver has an open-loop, short-term stability relative to GPS of about 5×10^{-10}, but this measurement may not be reliable, since the stability of the GPS signal itself is by policy intentionally degraded. With the 8-bit D/A converter and resistor network the quantization error is also about 5×10^{-10}, which is comparable to the intrinsic stability of the VCO and consistent with the values actually measured under typical operating conditions.

In order to achieve the smallest timing error, together with the highest stability, it is necessary to use a large τ; but, as τ is increased the loop gain K_v must be proportionally reduced to maintain a constant damping factor ζ, as described previously. However, as K_v is reduced, the tracking range is reduced as well. For a type-I PLL the tracking range must not fall below the maximum open-loop VCO frequency error or the PLL may break lock, which in the prototype receiver limits τ to values not greater than 34. For a type-II PLL and assuming an open-loop VCO stability of 5×10^{-10}, the loop gain can be reduced relative to the maximum VCO tuning range of 1.2×10^{-7} by a factor of 240. In the present design this corresponds to $\tau = 8153$, which is probably unrealistic. The values near 250 used in the present design allow for a comfortable margin of error; however, this consideration suggests that a more optimal choice of PLL design may be a type-II system.

5.3. The Cycle Identification Problem

Perhaps the most vexing problem in a LORAN-C receiver is the correct identification of the reference cycle, which is designated as the third carrier cycle of the LORAN-C standard waveform described previously. This is complicated by the fact that individual carrier cycles can be badly contaminated by noise and cross-rate interference and further, in the case of this receiver, by the fact that individual cycle amplitudes are accumulated over different frames. The conventional approach to this problem is to develop a synthetic signal constructed as the sum of the envelope amplitude less some fraction of its first derivative. Properly constructed, this synthetic signal, called the *derived envelope* [FRA83], has a positive-going zero crossing at the third carrier cycle. The

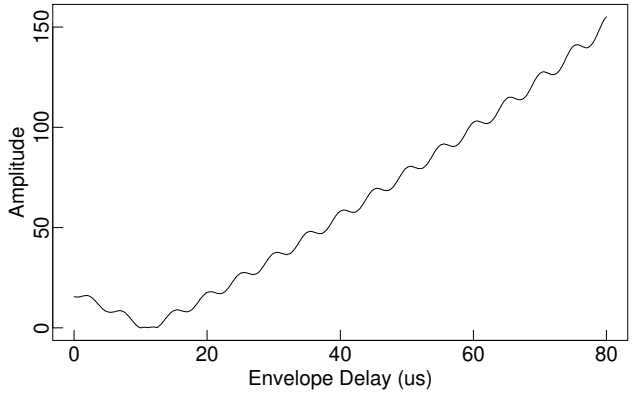

Figure 11. RMS Error Amplitude

control program can then selectively extract the phase signal at the first zero crossing following that as the control signal for the PLL.

Attempts to utilize this approach in the LORAN-C Timing Receiver have not been encouraging. The primary reason for this is the relatively large cycle-to-cycle variations due to the effects cited previously. A more successful approach has been to first normalize the envelope using as metric the maximum cycle amplitude in the envelope gate, then to compute the rms cycle amplitude error between the standard waveform and the actual waveform beginning at each cycle separately in the cycle window. The minimum over all cycles in the cycle window establishes the position of the standard envelope relative to the window and thus identifies the reference cycle.

Figure 11 shows the rms errors between the standard waveform and another shifted relative to it, where the peak amplitudes of the pulses are normalized at 100. In order to faithfully reproduce the receiver operations, the simulated demodulator is an integrating type, not a peak-detector type, which explains the 5-μs wavy artifacts apparent in the figure. All nine cycles in the cycle window are shown, while the null near 10-μs represents the case when the two waveforms exactly overlap.

While Figure 11 illustrates a continuous range of shifts, in actual operation the phase detector operates to line up the waveforms to an integral number of 10-μs cycles starting at zero in the figure; therefore, only the nine values shown at 0, 10, ..., 80 μs are significant. In the figure the rms error one cycle either side of the correct cycle is at least 15.6, which means that correct cycle identification can be reliably determined if the rms normalized envelope error is less than 0.156 relative to the standard waveform.

6. Operational Experience and Performance Evaluation

The receiver itself is constructed on a two-sided printed-circuit board and mounted together with its oscillator on a 12" x 7" x 3" aluminum chassis. It is connected to the receiver controller with a 25-pin DB connector and shielded cable. For test and evaluation purposes an oscilloscope is connected as shown in Drawing Sheet 1 of Appendix C. The following sections describe the initial setup and operation of the system, as well as operational experience and lessons learned.

6.1. Initial Receiver Setup and Adjustment

Winding and alignment instructions for the various tuned coils and transformers are shown on the drawing sheets. For the bandpass filter the simplest way to adjust the coils is to temporarily

disconnect the input, intermediate and output L-C circuits and adjust each of the three for series-resonance at the geometric mean of the high and low bandpass limits or 98.9 kHz.

The receiver controller includes five potentiometers which control the gain and bias of the operational amplifiers IC308A and IC308B and the reference voltage for the A/D converter IC305. The first four potentiometers R304, R306, R307 and R309, are adjusted to match the particular characteristics of the rf amplifier IC203 and oscillator X201, respectively. The goal in the adjustment process is to use the minimum gain necessary, together with an offset that brackets the operational range of the control voltages. The last potentiometer R303 is adjusted for a voltage of 5.0 at pin 12 of IC305.

The receiver gain and oscillator D/A converter bias adjustments are determined from the command line when the control program is started or by default from compilation parameters. The settings for these quantities require a certain degree of experimentation, although ways to do this automatically may be developed in future. The receiver gain depends on the antenna, its siting and various conductors in the near vicinity. can be determined with an oscilloscope while in pulse-group search (mode 4). It should be set so that the maximum signal due to noise plus LORAN-C signals is about 400 mV peak-to-peak, or about 6 dB below the hard-clipping level of the receiver.

It is necessary in setting the oscillator bias that its free-running frequency is within about 0.25 ppm. A simple way to do this is using a communications receiver and the signals broadcast from station WWV or WWVH. For the most precise adjustment, the oscillator harmonic at 10, 15 or 20 MHz can be compared with the broadcast signal. An accuracy of 0.25 ppm is achieved when the period of the beat note at 5 MHz or higher is greater than one second.

Measured receiver gain and passband characteristics after alignment agree with the simulation results within a decibel or two. The passband is flat within 6 dB over the range 90-110 kHz, while the gain can be varied over about a 40-dB range up to about 105 dB. In actual practice the receiver has much more gain than needed, since the ambient electrical noise at the antenna is much greater than the internal receiver noise by at least 20 dB, even with a wire antenna. In spite of the high gain, the receiver is completely stable and has no tendency to oscillate or “motorboat.”

6.2. Receiver Operation

Signal strengths for LORAN-C stations as received in the eastern regions of the country are sufficient to allow use of a simple whip antenna, rather than the loop antenna and preamplifier shown on the drawing sheets. Most of the hardware testing and program development was done using a longwire antenna in a noise-quiet, suburban location. Some of the testing, including an extended period of comparisons between the receiver time and GPS time was conducted using a loop antenna without preamplifier ordinarily used with an Austron 2000 LORAN-C Timing Receiver. The difference between the two antennas is primarily evident in the reduced noise levels with the loop.

Once the program is started, the only interaction between it and the operator is with the commands listed previously in this report. Ordinarily, operation is completely automatic once the GRI, master/slave selection and initial gain and frequency offset are set, either by the command line or as a compiled parameter. The program takes up to several minutes for the pulse-group search to complete, then a minute or two for the cycle search and possibly up to an hour for the initial amplitude and phase transients to die out.

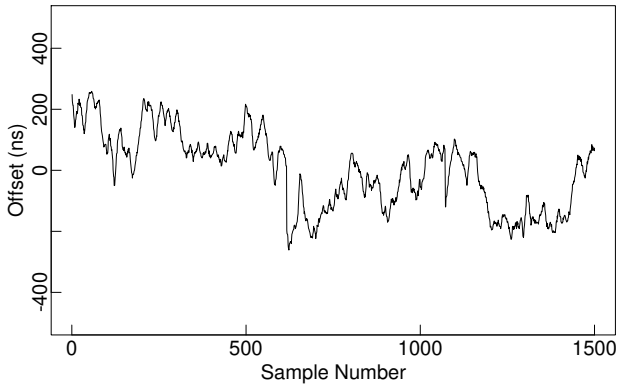

Figure 12. Prototype Timing Errors

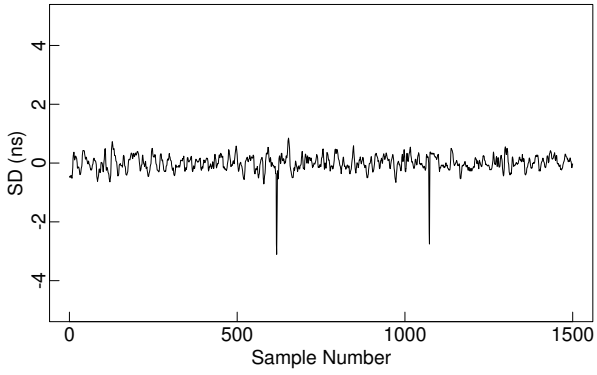

Figure 13. Prototype Frequency Errors

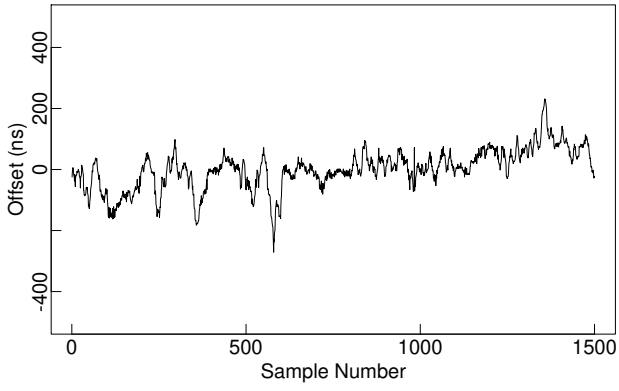

Figure 14. Austron Timing Errors

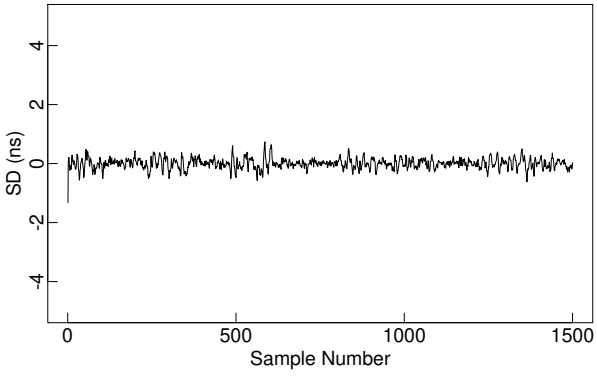

Figure 15. Austron Frequency Errors

The performance of the prototype LORAN-C Timing Receiver in actual service is compared with that of the Austron 2000 LORAN-C Timing Receiver in the series of figures 12 through 15. Note that the Austron receiver must be aligned with the transmitted signal using a tedious manual procedure, while the prototype receiver is completely automatic. The data from which these figures were compiled were produced at 30-second intervals by a time comparator which is an integral feature of the Austron 2201 GPS Receiver. The comparator measures the time difference between the 1-pps signal produced by either LORAN-C receiver and the corresponding GPS time. The LORAN-C chain being tracked during the experiment was the Seneca, NY, master station for the Northeast U.S. chain, which operates at a GRI of 9960.

Figures 12 and 14 show the timing offsets relative to GPS of the prototype receiver and Austron receiver, respectively, over a typical period of about twelve hours. In both cases the offsets are due to timing errors in the GPS receiver, the GPS system itself, the disciplined oscillators in the LORAN-C receivers and the LORAN-C system itself. Previous measurements with the GPS receiver and an HP 5061 Cesium Clock established the noise floor due to the GPS receiver and GPS system itself at a few tens of nanoseconds, which is comparable to the fine-grain behavior of the Austron receiver. The coarse-grain behavior is believed due to the nature of the 100-kHz propagation path and local weather and noise conditions.

As apparent from the figures, the prototype receiver displays timing errors exceeding the Austron receiver. While the peak errors are not very different in the two receivers, the standard deviation of

errors in the prototype receiver, 125 ns, is about twice that of the Austron receiver, 67 ns. The small periodicities apparent in Figure 12 are due to normal PLL hunting as the VCO D/A converter alternates between adjacent discrete values.

In the case of the Austron receiver, a 1-MHz signal phase-locked to GPS is used as the master reference, so that the analog PLL has to operate only over about one-tenth of the carrier cycle, which considerably reduces the incidental noise. In the case of the prototype receiver, the master reference is onboard the prototype receiver itself and the analog PLL has to operate over the entire carrier cycle. While the difference in master reference configurations may explain a portion of the prototype receiver errors, The principal reason for the errors is due to a degraded signal/noise environment. In these tests the receiver signal/noise environment was degraded to the point it was quite difficult for the operator to verify the correct zero crossing using an oscilloscope. The degradation is believed due to a ground loop in the antenna-receiver-computer signal path. Experiments now in progress suggest better performance is possible and that the performance of the Austron receiver can be equalled or surpassed.

Figures 13 and 15 show the frequency errors relative to GPS of the prototype receiver and Austron receiver, respectively, over a typical period of about twelve hours. The frequency stability is estimated as the standard deviation of the first-order differences in a sliding interval of 100 samples spaced at 1 s between samples. As expected, the increase in error of the prototype receiver over the Austron receiver is due to the quantization error in the D/A converter of the prototype receiver. The Austron receiver signal path is analog and has no D/A converter. As mentioned previously, the most obvious way to reduce the frequency error is the use of a 12-bit D/A converter.

Due to the rather aggressive nature of the search-and-integrate techniques, this receiver can dig rather deep into the noise and cross-rate interference to lock onto a selected master or slave station, which results in a moderately high false-alarm rate. If the pulse-group search happens to land on the wrong A or B field or the wrong master/slave station of the selected chain, the mode-5 cycle search times out and the receiver drops back to the mode-4 pulse-group search and tries again. However, under some conditions where the receiver is looking for a selected master station and a slave station from the same chain is much closer to the receiver than the master, the receiver sometimes locks repeatedly onto the first four or last four pulses of the slave station instead.

Inspection of the LORAN-C pulse codes shown in Figure 3 reveals why this happens. The first four phases of the A interval for master stations are identical to the last four phases of the B interval for slave stations, while the last four phases of the A interval for master stations are identical to the first four phases of the B interval for slave stations. If the signal strengths at the receiver between a wanted distant station and an unwanted nearby station differ by more than the system margin, which is 6 dB in the present design, the receiver may operate unreliably. In a navigation receiver designed to track both master and slave stations and equipped with geographic coordinates of each station, it is not possible that the A pulse group of one station can overlap the B pulse group of another station, so such an ambiguity is not possible.

Performance of the receiver with different LORAN-C chains at various locations in the country was assessed by simply attempting to synchronize to the master and slave stations of the LORAN-C chains serving the Southeast U.S. (7980), Northeast U.S. (9960), Great Lakes (8970) and Canadian East-Coast (5930). Synchronization was readily achieved for at least one slave station in each chain, although it was not certain which slave it was, presumably the one with the strongest signal.

However, it was not possible to reliably synchronize to the masters of any chain except the Northeast U.S. master at Seneca, NY, because of the problem mentioned above.

7. Future Development

The LORAN-C Timing Receiver is intended as a prototype demonstration to establish the suitability of LORAN-C technology as a precision synchronization source for a distributed, high speed computer network. While the initial objectives in accuracy and stability have been met in the prototype receiver, there is considerable room for improvement, both in hardware design and algorithm design. For instance, the prototype receiver does not have the software necessary to align the 1 pulse-per-second output precisely to UTC time. It is intended that this capability be provided using the Network Time Protocol and other time servers on the Internet [MIL90]. Also, since the control program is not interrupt driven and monopolizes the resources of the machine in which it runs. The program should be rewritten to use interrupts and operate as a transient-stay-resident (TSR) so that other programs, in particular the NTP code and its network drivers, can operate at the same time.

There is a version of NTP adapted by staff at the MIT Laboratory for Computer Science that runs on IBM-compatible machines. Future work may integrate the control program with this code to produce a standalone, turnkey system that could be deployed in many Internet-accessible places to provide precise time and frequency throughout a building or campus.

The current control program synchronizes to only a single LORAN-C station, so that position determination and propagation delays must be determined from other sources. The program could be enhanced to provide the capability to synchronize to other stations of the same or even different chains. Assuming a sufficient number of stations could be acquired (the master and at least two slaves), it would be possible to calculate the receiver position and determine the propagation delay automatically.

The problem with falsely synchronizing to a slave station when the master of the chain is selected can almost certainly be solved by adroit pulse-code management. One obvious approach is to extend the pulse-group search beyond the single frame interval now used; however, this requires either a higher level of oscillator stability or the ability to integrate more than one bin per frame interval. The latter approach also has the advantage to reduce the pulse-group search time in cases where the difference between master and slave signal levels are small.

The most practical approach may be to simultaneously search for both the master and slaves at the same time by staggering the envelope gate and using two sets of accumulaters and filters. Once the strongest slave or master has been found, the receiver can lock to that station and maintain oscillator stability based on that station. Then, the remaining bins of the frame interval can be searched by timesharing the integrators.

8. Acknowledgements

Neil Corman designed and constructed the receiver printed-circuit board and wired the receiver controller. Grateful acknowledgment is made to the U.S. Coast Guard, who provided valuable equipment, including a Hewlett Packard 5061A Cesium-Beam Clock and Austron 2000 LORAN-C Timing Receiver which materially assisted the progress of this project.

9. References

[CG81] Department of Transportation U.S. Coast Guard. Specifications of the Transmitted LORAN-C Signal. COMDTINST M16562.4, July 1981. This document may be obtained from COMMANDANT (G-NRN-3/TP14) U.S. Coast Guard, Washington, DC 20593.

[FRA83] Frank, R.L. Current developments in LORAN-C. *Proc. IEEE 71, 10* (October 1983), 1127-1139.

[MIL90] Mills, D.L. Network Time Protocol (Version 3) specification, implementation and analysis. Electrical Engineering Department Report 90-6-1, University of Delaware, June 1990.

[MIL91] Mills, D.L. Internet time synchronization: the Network Time Protocol. *IEEE Trans. Communications 39, 10* (October 1991), 1482-1493.

[SMI86] Smith, J. *Modern Communications Circuits*. McGraw-Hill, New York, NY, 1986.