Using LORAN-C for Time and Frequency Calibration

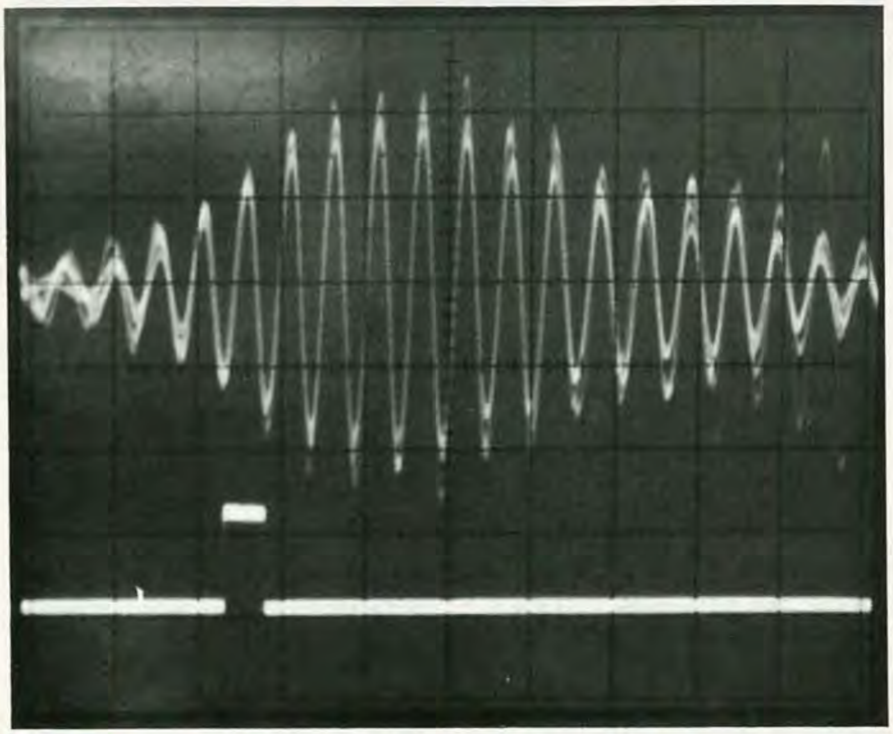

Here's a look at the Loran-C navigation system—what it is and how it works. We will also discover how Loran-C signals can be used as frequency standards for calibrating oscillators.

R.W. BURHANS

THE LORAN NAVIGATION SYSTEM OPerates on two different frequencies, each with a different set of characteristics. Each frequency has its own Loran designation—Loran-A and Loran-C. Although Loran-C's primary purpose is for long-distance navigation, it has another important use. Because Loran stations have to maintain a high level of precision they can—if used properly—serve as extremely accurate frequency standards. In this article we'll be discussing how we can use Loran-C signals for such applications as calibrating a frequency standard or a frequency-counter timebase.

Before we get into the details of the Loran-C system, let's take a brief look at how Loran-C signals can be used for calibration purposes. The Loran-C signals are observed on an oscilloscope that is externally triggered by a special pulse generator. (The pulse generator is quite simple, and we will provide construction details for it shortly.). The pulse generator is driven by the frequency standard you wish to calibrate. (Details for calibrating a 1-MHz frequency standard will be given, but other frequency-standards can also be calibrated using this technique.) The stability of the frequency standard can be obtained by determining the time it takes the display of the Loran-C signal to drift a given distance across the screen. Using this technique you can calibrate a 1-MHz oscillator to better than 0.001 Hz.

The Loran-C navigation system

Loran-C signals are broadcast on a frequency of 100 kHz with a 20-kHz bandwidth (from 90 to 110 kHz). Because of their low frequency, Loran-C signals tend to be ground waves—they follow the earth's curvature. The signals are usually very stable because they are not affected by the ionosphere. But how are they used for navigation?

Loran signals are sent from a *chain* (usually three to five) of stations. One station in each chain is the *master* and the others are *slaves*. The master station transmits groups of pulses that are received by the slave stations. Each slave station transmits similar groups of pulses, and adds a fixed time-delay between the groups of pulses transmitted by the master and its own pulse groups. A Loran-C receiver receives both pulse groups and calculates the time difference between them. It is that time difference that is used to establish a line of position that is used for navigation.

If (at the receiving end) the time difference between the received signals were to equal the original delay added by the slave station, the receiver would be somewhere along a straight line equidistant from both the slave and the master transmitters. If the time difference were to deviate from that fixed time-delay, the receiver would be somewhere along a particular hyperbola. (A hyperbola is a curve where the *difference* of the distances from any point on the curve to two fixed points is a constant.)

As shown in Fig. 1., a second pair of transmitters (the same master but a different slave) can be used to construct a second hyperbola. The intersection of the two hyperbolas is the receiving point. To be useful for navigation, at least three Loran stations in a chain must be re-

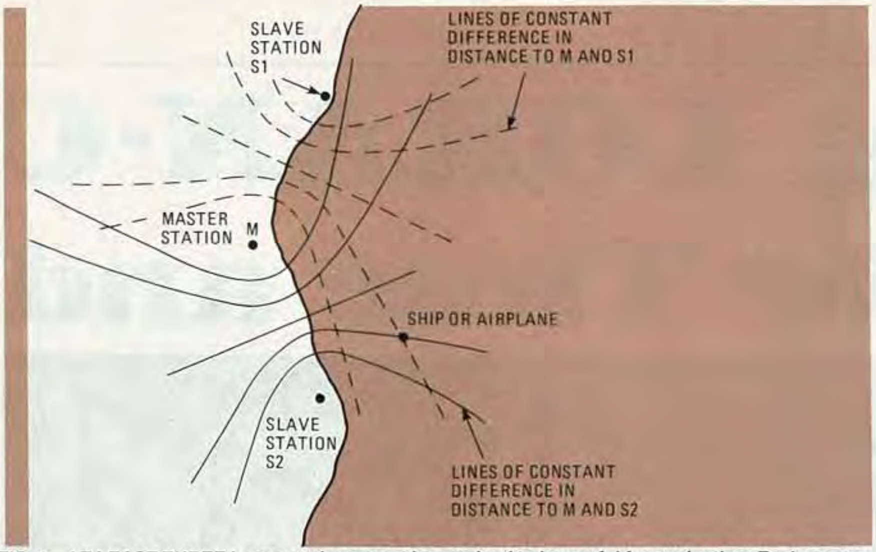

FIG. 1—AT LEAST THREE Loran stations must be received to be useful for navigation. Each master-slave pair is used to establish a particular hyperbola of constant time-differences. The intersection of the hyperbolas gives the receiver location.

ceived. However, for time-and frequency-calibration applications—which is what we are interested in—only one Loran station needs to be received.

One source of more detailed information on the Loran-C navigation system, is the *Loran-C Handbook*. For information on its price and availability, write the Superintendent of Documents, U.S. Government Printing Office, Washington D.C. 20402. The book's stock number is 050-012-00171-5.

Loran-C receivers

The sophistication of modern Loran-C receivers is now very high—some will directly compute your latitude, longitude, range, and bearing. The cost of those receivers is also high, typically \$5000 for marine systems and \$10,000 for airborne receivers. The navigational precision that can be obtained varies with the chain's geometry and distance, but it can be as good as ± 50 feet when differential and propagation corrections are taken into account.

A complete Loran-C navigation receiver usually consists of the main components shown in Fig. 2: an active-antenna coupler, an RF sensor, a senor processor, a navigation processor, and finally a data-display device. The last three are microprocessor controlled.

The Loran-C signal

Before we can understand how a Loran-C signal is used by the receivers, and how we can use it for frequency calibration, we have to look at its characteristics. A Loran-C signal consists of pulses. An ideal pulse is shown in Fig. 3. Each Loran-C transmitter transmits a series of 8-pulse groups with the pulses spaced 1 millisecond apart. An extra (ninth) pulse (sent two milliseconds after the 8-pulse sequence) is used to identify the master station. (The master station is the beginning of the chain's sequence.) The sequence is repeated at a certain GRI (*G*roup *R*epetition *I*nterval) that identifies the chain. Figure 4 shows the relative amplitudes of the pulse envelopes as received from one Loran-C chain (Northeast US: the repetition period is 99,600 μs). Each vertical line in the figure represents one pulse. Note that the relative amplitudes as well as the time separations between the groups will vary depending on the receiver location. (Of course if that weren't true, Loran would not be very useful for navigation.)

An additional characteristic that we should point out, although it is not shown in Fig. 4, is that the Loran signals are phase coded. That allows Loran receivers to automatically identify master and secondary stations; to have an automatic search mode, and to reject multi-hop sky waves.

Sky waves

One problem encountered when using Loran-C is that the ground waves are often contaminated by sky waves. (Remember—one of the reasons that Loran can be used for great accuracy is due to the stability of ground waves.) So that the signal at the receiver will not be contaminated by the arrival of sky waves, a signal with a fast risetime is used. That allows the pulse to build up to its maximum value at the receiver before the sky waves arrive. Also, the tail of one pulse should be low in amplitude when compared to the beginning of the next one so that the trailing sky waves will not contaminate the beginning of the next pulse. The limiting constraint on a signal's risetime is its bandwidth. (For example, a squarewave has a very fast risetime, but its bandwidth would be too large to be used by the Loran system.)

The pulse shape shown in Fig. 3 is used to reduce the problem of one pulse affecting the beginning of the next—the tails of the pulses are greatly attenuated (and the 20-kHz bandwidth constraint is still met.) The third cycle of the pulse (that's the one which is tracked by the receiver) will not be contaminated by sky-wave (or reflected) signals.

Because the transmitted signals have a relatively wide (20-kHz) bandwidth, ordinary communications receivers cannot do a good job of detecting Loran-C pulses, although some receivers with a 12-kHz bandwidth can do a reasonable job for long-term frequency calibration (where the local clock is kept running 24 hours a day).

As we mentioned previously, a Loran-C receiver is designed to detect a point on the signal (the third cycle) before the stronger sky waves have a chance to contaminate the envelope. That task is not easy to perform and there is still argument over the best way to detect the earlier, weak ground wave at long range when it has been contaminated with sky waves (which often have a peak level 20 dB greater than the ground wave). Because

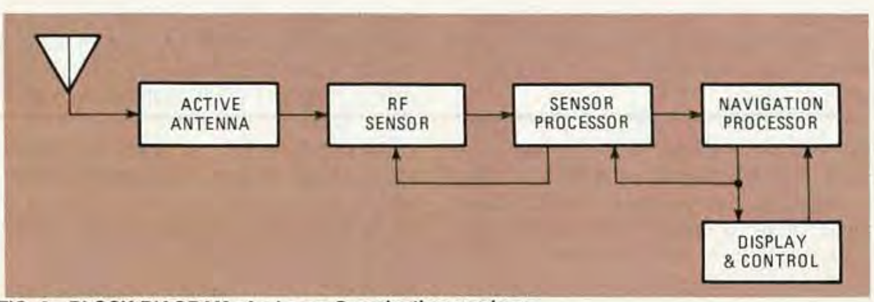

FIG. 2—BLOCK DIAGRAM of a Loran-C navigation receiver.

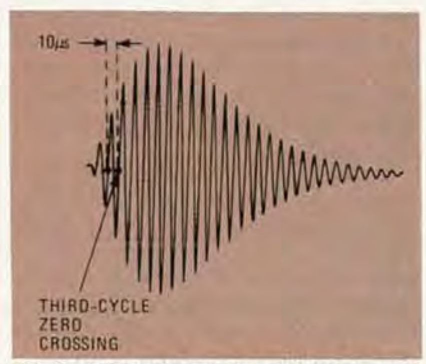

FIG. 3—IDEAL SHAPE of a 100-kHz Loran-C transmitted pulse. Note how the trailing edge is attenuated so that the trailing sky waves will not interfere with the next pulse.

FIG. 4—RELATIVE ENVELOPE AMPLITUDES of a Loran-C chain as received in the midwestern US. Note the "extra" pulse that identifies the master.

the sky-wave signals are not stable—they vary considerably in amplitude and risetime—large errors will be produced if the receiver tracks them.

Another problem with Loran receivers is that any filtering will delay the risetime of the signal so that, with many receivers, a point later than the ideal third cycle is tracked. However, as long as that point is the same for all signals, strong or weak, and is chosen to precede the sky wave's reaching an appreciable amplitude, then the receiver can still operate satisfactorily. Let's look at how the third cycle of the Loran-C pulse is detected.

Third-cycle detection

A theoretical way of detecting a point on the pulse shown in Fig. 3 is to generate the second derivative of the pulse's envelope shape. The resultant envelope has a zero crossing at about 35 μs. The problem in this case is that the envelope generator (with the differentiators) ends up with an extremely wide bandwidth—that adds a lot of noise to the system.

Another way to detect the third cycle is to produce a delay-and-add circuit. Such a circuit delays the signal 5 μs (180°) and then algebraically adds the delayed signal to the original (with a multiplying constant) to produce a phase reversal at about the 30-μs point. A simple implementation of a delay-and-add network—and the effect it has on the Loran-C signal—is shown in Fig. 5.

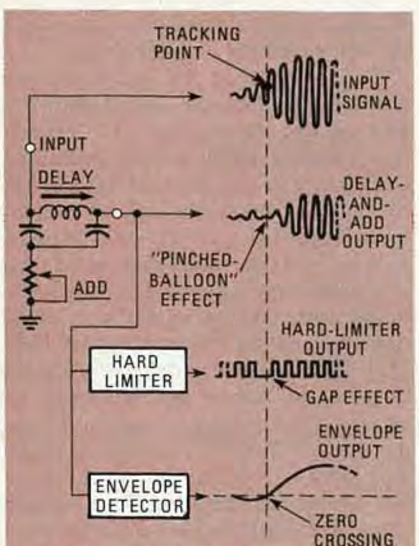

FIG. 5—A DELAY-AND-ADD network can be used to produce a pinched-balloon effect output, which is then used to produce an envelope with a zero crossing at the third cycle.

If the output *pinched-balloon* shape of the delay-and-add network is fed to a *hard limiter*, then the result, as shown in Fig. 5, is a rectangular waveform where the phase-reversal point (the third-cycle point) has a gap. All the receiver designer then has to do is to devise some machine-language software to track that gap for the stations of a given chain. That is, of course, no easy task. But that's what is done in the sensor processor of many Loran-C receivers.

Another problem with Loran reception and third-cycle detection is that the conductivities of the earth and of seawater are different. That can cause the group velocity and the phase velocity of the signal to differ, producing an envelope-to-cycle difference (ECD) error up to several microseconds (depending on the terrain and the distance). (Some very precise Loran-C receivers can use that to an advantage. Precision measurements of the amplitude and phase of Loran-C signals made while flying at low altitudes can yield information on the ground's contours.)

Still another problem associated with Loran-C receivers is the fact that strong interference on frequencies like 88 kHz or 116 kHz can produce errors in the navigation data. Fixed, tuned traps—that are designed for particular coverage areas where there are interferring signals—are often found in Loran-C receivers.

RF filtering

A signal takes a finite time to pass through a filter. That time, called a delay, is a function of the signal's velocity. The signal's velocity is a function of its frequency. (For instance, the signal's velocity is lowest at the lower band-edge.) Thus, a filter can cause EDD (*E*nvelope *D*elay *D*istortion).

The effect that a filter with narrow skirts has on the received signal is shown in Fig. 6. Here we have assumed a worst-case sky wave rising at about 30 μs after the start of the ground wave, and we have assumed that the peak amplitude of the sky wave is 20 dB greater than that of the ground wave.

The filter delays the Loran ground-wave signal (the envelope's zero crossing occurs at a later time), but coincidentally, the strong sky-wave signal is also delayed. Unfortunately, however, the ground wave's third cycle (which we want to detect) is at −60 dB. Fortunately, though, even at the 50-μs point the desired pulse is some 30 dB greater than the sky-wave contamination.

A −30-dB contamination of the tracking point at 50 μs would result in a small error in the data. The data will contain additional error because the results are not as precise when the fifth cycle (50-μs point) is tracked instead of the third cycle. However, as long as the sky-wave contamination is low enough, and the same point is tracked for all signals, then the error can be kept down to perhaps 0.1 μs, even for a weak signal.

The effect of filtering and AGC in conventional communications receivers destroys most of the information in the Loran-C signal. However with a receiver such as the Yaesu *FRG7700* in the WIDE-BAND AM (12-kHz) mode, it is just possible to track the fifth cycle (which will be about the start of the pulse envelope as observed at the receiver's line-level output terminal). Noise blankers in communications receivers also destroy the pulse information because they are inherently timed to blank a pulse of the Loran-C shape and duration.

The envelope detector and simple GRI generator that we will discuss next are useful for experimental observations but are not well suited to precision navigation-receiver applications. The minimal equipment required for Loran-C observation is: a triggered-sweep oscilloscope, a frequency standard, a GRI generator (we'll discuss one), and a reasonably wideband AM receiver. If you do not have a suitable receiver available, you can use the Loran-C front end (or envelope detector) that we'll discuss next.

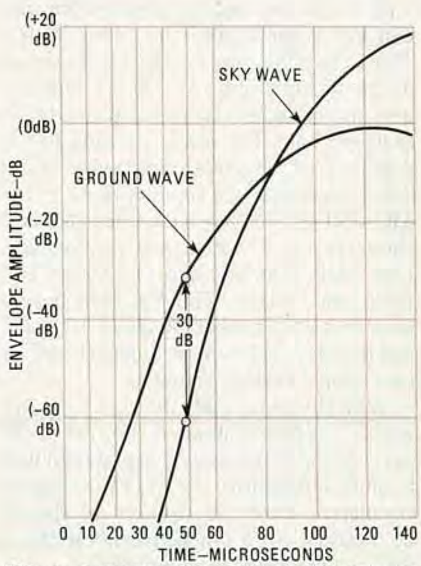

FIG. 6—A FILTER DELAYS both the Loran-C sky-wave and ground-wave signals, resulting in an zero crossing at about the fifth cycle.

RF envelope detector

For the experimenter who wishes only to study Loran-C or to use the signals as a frequency-calibration source, the envelope method (which detects the envelope zero-crossing as shown in Fig. 5) is simple and effective. A schematic of an envelope detector is shown in Fig. 7. The active-antenna preamplifier has been described previously (see the February, March, and April 1983 issues of **Radio-Electronics**). It uses a 1-meter whip with a 150-kHz lowpass filter at the antenna input. The "receiver" coupler is built as part of the Loran-C front end (Fig. 7) and consists of an 80-kHz highpass filter (instead of the 10-kHz filter used with the previous active-antenna circuits). The output of that coupler is fed to a network consisting of a 100-kHz impedance-matching transformer; to series T-notch traps tuned to 88 kHz and 115.3 kHz; and finally to an MC1350 RF amplifier. The sequence of matching and filtering at the input is an impedance step-up to the MC1350 RF stage such that there is a net voltage gain of 10 dB or more between the active antenna and the RF stage, which provides an added 60 dB of gain. The RF stage is operated with a manual control-voltage for the RF gain (or an external AGC system could be devised). The amplifier drives a transformer and a delay-and-add network. The delayed output and the undelayed (RF-amplifier) output are fed to an MC1357 FM detector. There, the hard-limited RF carrier is mixed or multiplied with the delay signal. That results in an envelope with the same amplitude and phase as that from the delay-and-add network. The zero crossing is adjusted for 40 to 50 μs, and the ADD trimmer potentiometer is adjusted for a good null at the zero-crossing point. The actual pinched-balloon effect can be observed on a scope at pin 11 of the MC1357.

To align the RF transformers and traps with the active antenna preamplifier you can use a signal generator along with a 10-pF coupling capacitor to inject a signal at the antenna input. (By doing that you are simulating the response that will be obtained from the active antenna in the presence of an electromagnetic field.) The response at the input (pin 4) of the MC1350 should look something like that shown in Fig. 8, with a peak at about 103 kHz, nulls at 88 kHz and 115.3 kHz, and fairly wide skirts. The MC1350 output transformer is peaked at about 100 kHz and the delay network is adjusted using an on-the-air Loran-C signal.

With this front end, the zero crossing will actually be at about the fourth or fifth cycle of a strong Loran-C signal. But that is quite satisfactory. If a Loran-C signal simulator is available, the network should be adjusted for a pinched-balloon effect (or zero crossing) at about 40 μs.

The output from the limiter/detector drives a 3-pole 33-kHz lowpass filter (IC3 and its associated components) to generate a DC envelope that can be observed on an oscilloscope. The comparators and flip-flop provide a synchronized 10-μs pulse at the envelope's zero-crossing point, but it is fairly noisy compared to observing the envelope at the analog lowpass-filter output. However, this output pulse for each Loran-C envelope is useful for further experimenting with microcomputer tracking-loops (where the bandwidth of the noise can be narrowed with memory-aided numerical techniques in software).

West-coast and foreign experimenters may wish to align the interference traps at some other frequencies. It is a good idea to keep the traps outside the 90 to 110 kHz region because they are not high enough in Q and could attenuate too much of the desired Loran-C signal. A lot of retuning—going back and forth over the adjustments of all of the tuned circuits—pays off in arriving at a reasonably wide bandwidth with a sharp phase-reversal at the delay-and-add network output.

The adjustable RF inductors and transformers used the circuit shown in Fig. 7 are normally 455-kHz IF transformers padded with additional capacitance to tune them to the 100-kHz region. The particular transformers in this experimental circuit are Mouser 42IF303's, a type sometimes called a 3rd-IF transformer. Other types will sometimes work, except in the case of the transformer at the output of the MC1350. That's because the tap on the transformer is not a true center tap. However, the loading effect of the 0.01-mF capacitor at pin 4 of the MC1357 makes the output look like a balanced load for the RF stage. Some older transformers, from different manufacturers, may have different winding phases, so you may have to reverse the secondary connections to get everything operating properly. The phase of the smaller output-winding has to be such

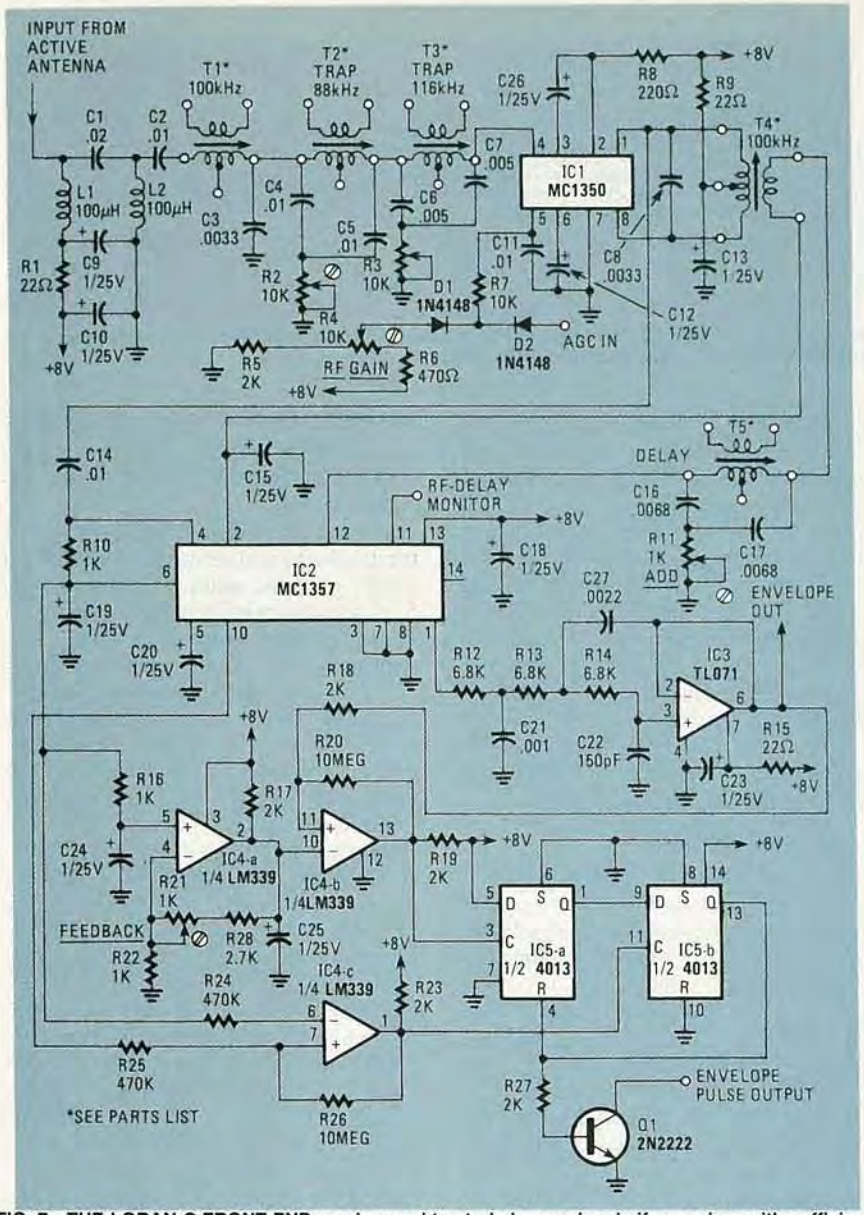

FIG. 7—THE LORAN-C FRONT END can be used to study Loran signals if a receiver with sufficient bandwidth is not available.

PARTS LIST—ENVELOPE DETECTOR

All resistors ¼ watt, 5% unless otherwise specified
R1, R9, R15—22 ohms
R2, R3, R4—10,000 ohms, trimmer potentiometer
R5, R17, R18, R19, R23, R27—2000 ohms
R6—470 ohms
R7—10,000 ohms
R8—220 ohms
R10, R16, R22—1000 ohms
R11, R21—1000 ohms, trimmer potentiometer
R12, R13, R14—6800 ohms
R20, R26—10 megohms
R24, R25—470,000 ohms
R28—2700 ohms
Capacitors
C1—0.02 μF, polystyrene
C2—0.01 μF, polystyrene
C3, C8—0.0033 μF, ceramic disc
C4, C5—0.01 μF polystyrene
C6, C7—0.005 μF, polystyrene
C9, C10, C12, C13, C15, C18, C19, C20, C23–C26—1 μF, 25 volts, tantalum
C11, C14—0.01 μF, ceramic disc
C16, C17—0.0068 μF, polystyrene
C21—0.001 μF, ceramic disc
C22—150 pF, ceramic disc
C27—0.0022 μF
Semiconductors
IC1—MC1350 video-IF amplifier
IC2—MC1357 sound-IF amplifier and quadrature detector
IC3—TL071 JFET-input op-amp
IC4—LM339 quad comparator
IC5—4013 dual D-type flip-flop
D1, D2—1N4148
Q1—2N2222 or similar NPN-type
T1–T5—455-kHz IF transformer, Mouser 42IF303 or equivalent
L1, L2—100-μH RF choke

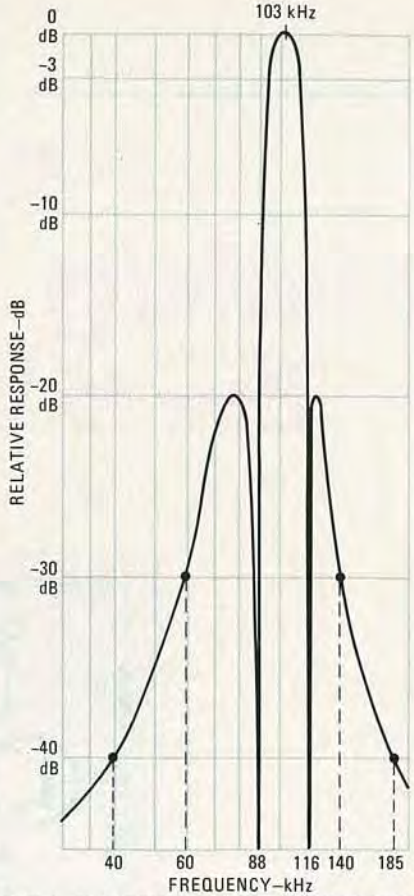

FIG. 8—THE RESPONSE AT THE INPUT to the RF stage (due to the filters) should look like this.

PARTS LIST—GRI PULSE SOURCE

All resistors ¼-watt, 5% unless otherwise specified
R1—47,000 ohms
R2, R3—10,000 ohms
Capacitors
C1, C2—0.001 μF
Semiconductors
IC1—4518 dual decade divider
IC2—12-stage binary ripple counter
IC3—quad 2-input nor gate
D1-D12—1N4148
S1-S3—Hexadecimal thumbwheel switch, 16 position BCD, Unimax SF-54 or equiv.

that a positive-going envelope is created at pin 1 of the MC1357 detector IC. The transformers come with a small 150-pF built-in capacitor across the main primary winding, but it can be ignored since it is very small compared to the 3300-pF capacitor required for resonance at 100 kHz.

The LM339 comparator (for the envelope pulse signal) generates a reference voltage from pin 6 of the MC1357. The small FEEDBACK trimmer potentiometer across the LM339 reference source is used to adjust the DC level of the pulse edge that is fed to the flip-flop. The reason for doing that is that the DC level is controlled by the MC1357, so that drift in the lowpass filter DC-level or in the comparator reference DC-level is self-compensating. The envelope 10-μs-pulse output is intended to drive external logic, usually at a five-volt level. The power source for the whole RF front-end should be from an eight-volt regulated source (using a regulator such as an LM7808).

GRI pulse source

Practically every Loran-C experimenter needs a GRI generator that is capable of producing pulse-repetition intervals of 50,000 to 100,000 μs. We'll discuss a simple generator that consists of a 4040 programmable ripple-counter and a dual decade-divider that is driven from a 1-MHz crystal-controlled frequency standard. There are other methods for producing GRI pulses, but the three-IC circuit of Fig. 10 is about as simple as they come. (You could, for example, use only programmable decade-dividers. And you could use something other than a 1-MHz standard to drive the GRI generator and obtain the same results.) The circuit is programmed in hexadecimal notation for the GRI intervals as indicated in Table 1. That table lists the common designation for the GRI in 4-digit numbers (μs/10), as is done on Loran charts and in United States Coast Guard data. The pulse generator that's shown in Fig. 9 produces a 10-μs pulse that is used to synchronize your oscilloscope for observing Loran-C signals.

The programmable GRI-source allows the experimenter to stop the Loran-C signals on the scope trace and examine them in minute detail. That is also the basic

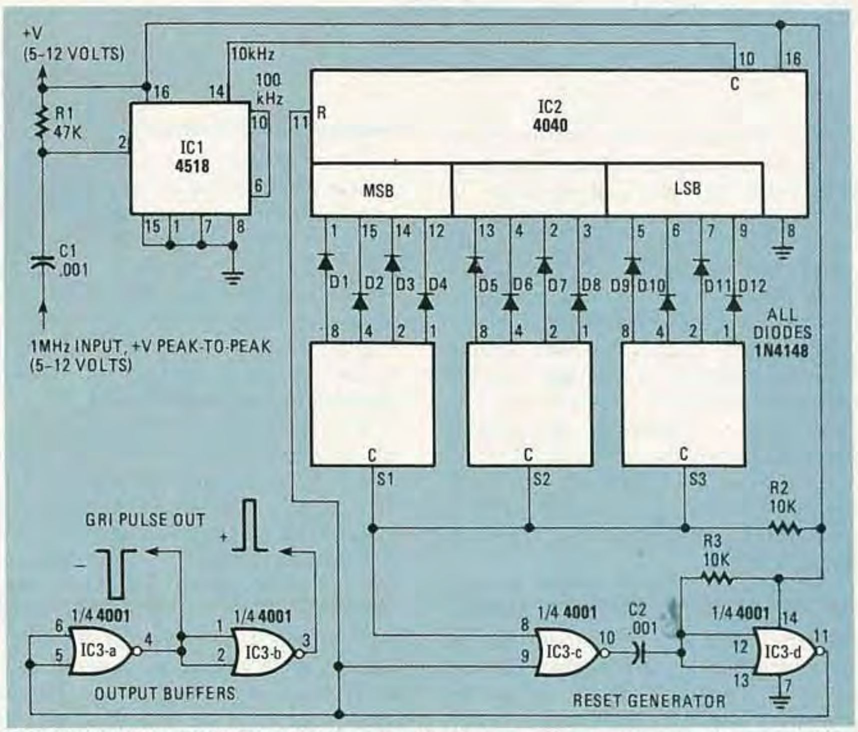

FIG. 9—THE GRI PULSE GENERATOR is used to stop the Loran signals on the scope trace so that they can be examined, and their drift measured. The drift indicates the stability of your frequency standard.

continued on page 92

USING LORAN-C ETC.

continued from page 67

idea of using Loran-C for frequency-standard checking/calibration. The GRI source, driven from a local frequency standard, is used to trigger the scope externally. The Loran-C signal, either from a suitable receiver or from the RF front end (the circuit shown in Fig. 7) is displayed, and the sweep time is adjusted to observe some part of a particular signal. The Loran-C signal at a given GRI can be momentarily speeded up or slowed down by "bumping" the thumbwheel switches of the GRI generator to place the Loran-C trace at some point on the scope so that a very small part of the leading edge of a pulse can be observed. Then, by determining the length of time it takes the Loran-C signal to move a given distance across the screen of the scope, the fractional frequency-stability of the local 1-MHz standard can be determined.

How is that time measurement used to obtain a measurement of the relative frequency stability? The two measurements can be related by calculus; we will only give the result: $\Delta f/f = -\Delta t/T$.

Here, Δt is the change in time (the drift) read over a measurement time, T. If, for example, the result was 1.5×10^{-6}, it would indicate that the oscillator's actual frequency was its nominal frequency (in this case 1 MHz) multiplied by 1.5×10^{-6}. In this case, the actual frequency of the standard would be 1.0000015 MHz. Typically, with a good (proportional oven controlled) standard, the Loran-C signal will only move to the right or left about 10 μs/hour. That implies an offset of the order of 3×10^{-9}.

The expanded scope-trace (about 10 μs/division) is useful for examining the actual RF-carrier output from the pin-11 test point of the MC1357 detector in Fig. 7. The envelope of that signal at the output of the op-amp-filter can also be observed, and the movement of the inflection point or zero crossing can be recorded at hourly intervals to check the clock-stability.

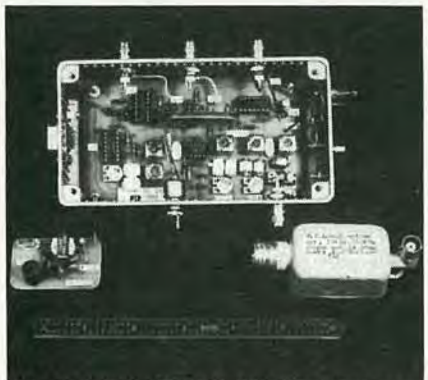

FIG. 10—EXPERIMENTAL MODEL of the Loran-C envelope detector with active antenna preamp and housing.

TABLE 1—GRI RATES

GRI (μs/10)	CHAIN	COUNTER CODE
4990	Suez Canal	1F3
4990	Central Pacific	1F3
5930	Canada East	251
5990	Canada West	257
7930	North Atlantic	319
7960	Gulf of Alaska	31C
7970	Norwegian Sea	31D
7980	US Southeast	31E
7990	Mediterranean Sea	31F
8970	US Great Lakes	381
9940	US Westcoast	3E2
9960	US Northeast	3E4
9970	NW Pacific	3E5
9990	North Pacific	3E7
10000	(HF Radar 0.1s)	3E8
8000	(USSR West)	320
5000	(USSR East)	1F4
1000	(0.01s)	064
100	(0.001s)	00A
10	(0.0001s)	001

When using receivers like the *FRG7700* for Loran-C, the envelope rise-time is smeared out over 150 μs or so. But by observing the change of that signal at daily time intervals, with the clock and GRI source operating continuously, the frequency stability of the clock can be determined without any special Loran-C front-end hardware.

For Loran-C DX hunting on late winter evenings when the noise level is low and DX is coming in from other stations such as Allouis in France on 164 kHz, the GRI rate can be set to try to find some chain not normally observed in the USA. That is done by examining the whole GRI frame in detail (with the oscilloscope in the expanded-sweep mode) and slowly bumping the GRI rate a few tens of milliseconds at a time to find weaker sky-wave pulses standing still. Loran-C signals from the USSR using GRI rates of 8000 and 5000 can sometimes (although rarely) be observed that way. The main problem in looking for weak signals is the cross-rate interference from other chains drifting by the desired small-amplitude signal on the scope trace.

A photograph of an experimental model of a Loran-C RF-envelope detector is shown in Fig. 10. Circuit-board layouts have been prepared for the Loran-C RF detector and the GRI generator. Contact: R. W. Burhans, 161 Grosvenor St., Athens, Ohio 45701. Include a SASE for information on the availability of these boards for experimental use. **R-E**