

1. Introduction

The work described in this report continues similar work described in a previous report [3]. That report describes a software program that implements an optimal linear demodulator and decoder for direct-printing radiotelegraph signals commonly used by amateur and commercial stations in the shortwave radio spectrum. The work described in that report and this one started as an exploration of projects useful in laboratory courses for upper division computer engineering and digital communications courses; however, both have taken on lives of their own. While laboratory experiments involving audio signal generation and playback are used in existing courses, the limited processing capability of most available PCs precludes interesting experiments possible with more sophisticated digital signal processing (DSP) hardware. Since DSP chips have become readily available and widely used in embedded systems, emphasis is placed on experiments in which hardware and software designs could be explored in the context of a PC, DSP evaluation board and a shortwave radio.

This report describes a software program that functions as a radio clock using shortwave radio signals transmitted by National Institute of Standards and Technology (NIST) radio stations WWV and WWVH. Operated in conjunction with an inexpensive, fixed-frequency shortwave radio, it has nominal timing errors less than 125 μ s when tracking one of the stations and frequency variations less than 0.5 parts-per-million (PPM) when not tracking either station. The clock produces an ASCII timecode that can be used to set the time of another device, such as a computer, as well as precision reference signals that can be used for other purposes, such as to drive laboratory test equipment.

The primary reason for choosing the DSP radio clock application is that signal propagation is via the ionosphere, which acts as an unreliable, multiple reflector, and that the noise process is very bursty. Thus, the transmission channel model is very much time-varying, sometimes badly distorted by multipath, often contaminated by interfering signals and atmospheric electrical noise, and sprinkled with dropouts due to strong adjacent channel signals. This presents an interesting challenge for the student of digital signalling, one unlikely to be found in textbooks. There are opportunities to devise interesting and insightful approaches to traditional problems such as carrier dropout, synchronization recovery, etc., that occur with ionospheric channels.

The primary motivation for this report is as an example and case study of optimum demodulator and decoder design using a maximum likelihood approach and matched filter, synchronous detection and soft decision principles. The clock discipline is modeled as a Markov process, with probabilistic state transitions corresponding to a conventional time-of-century clock and the probabilities of received decimal digits. The result is a performance level which results in very high accuracy and reliability, even under conditions when the one-minute beep from the WWV/H signal, normally its most prominent feature, cannot be detected by ear with a sensitive communications receiver.

This report proceeds first with a discussion on the design approach followed by a description of the WWV/H transmission format and an overview of the hardware architecture. The detailed design of the DSP clock is described by each functional block. This includes the processing steps at the radio frequency and baseband stages of processing. The discussion proceeds with a detailed mathematical analysis of performance and ends with an evaluation of performance under real

world conditions. Appendix A contains operation notes, including status indicators, command functions and signal inputs and outputs.

2. Design Approach

The DSP clock is designed for use with an inexpensive shortwave radio, although for testing it has been used with a sensitive communications receiver and laboratory equipment, including a signal generator, oscilloscope and calibrated cesium clock. It is intended to work with indoor antennas in most cases, although for testing an elaborate outdoor antenna has been used. The radio IF bandwidth and stability are ordinarily not a factor, since the emissions are amplitude modulated and the DSP-93 program includes all necessary signal conditioning and bandwidth filtering. Most shortwave radios today have high sensitivity for use with telescoping rod antennas. However, shortwave stations normally operate with hundreds of kilowatts and highly directional antennas, while the WWV/H stations operate with only 10 kW and nondirectional antennas. In most situations, something more than a telescoping rod, perhaps a hank of wire strung across the ceiling or tossed out a window may bring better results.

Shortwave transmissions are propagated via single or multiple reflections between the ionosphere and the ground. The path geometry depends on up to three ionospheric layers, which vary in height over the days, seasons and years. One result is the best frequency to use at any hour of the day depends on the height of these layers and the intensity of solar radiation over the ray path. Usually, the higher frequencies work best during the day and the lower at night. Commercial radio clocks that use WWV/H signals, such as the GC-1000 Most Accurate Clock, originally marketed by the Heath Company, and the 1020 Time Source, originally marketed by Precision Standard Time International, use frequency-agile receivers that automatically tune the embedded receiver to the frequency with the strongest signal.

However, to avoid complexity and cost, the DSP clock is designed to operate on only one frequency. The best frequency to use depends on the distance from the transmitter and must be chosen by compromise. Usually, the lower frequencies are better nearer to the transmitter and the higher frequencies are better further away. When more than one frequency is usable, it is usually best to choose the higher one. In some cases it may be best to choose a compromise frequency that works best near sunrise and sundown. In any case, the best choice depends on location, ambient noise level and antenna configuration.

In order to succeed in this approach the DSP clock must have the highest sensitivity and reliability when only marginal signals are available. This is the main reason so much emphasis has been placed in sensitivity and reliability under adverse conditions. In extreme cases, it may be possible to synchronize the clock only near sunrise and sunset and, even then, maybe not every day. A critical consideration in this regard is the frequency variations of the DSP chip clock oscillator due to ambient temperature variations. With the current hardware and ordinary room temperature variations, this limits the accuracy to the order of a millisecond per hour when the radio signals are not available.

3. WWV/H Signal Design

The DSP clock uses signals transmitted by frequency and time stations WWV, located in Ft. Collins, CO, and WWVH, located in Kauai, HI. Both stations transmit amplitude-modulated signals

Sec		Sec		Sec		Sec	
0	no carrier	15	minute 10	30	day 1	45	
1	DST2	16	minute 20	31	day 2	46	
2	leap warning	17	minute 40	32	day 4	47	
3		18		33	day 8	48	
4	year 1	19	P2	34		49	P5
5	year 2	20	hour 1	35	day 10	50	DUT sign
6	year 4	21	hour 2	36	day 20	51	year 10
7	year 8	22	hour 4	37	day 40	52	year 20
8		23	hour 8	38	day 80	53	year 40
9	P1	24		39	P4	54	year 80
10	minute 1	25	hour 10	40	day 100	55	DST1
11	minute 2	26	hour 20	41	day 200	56	DUT .1
12	minute 4	27		42		57	DUT .2
13	minute 8	28		43		58	DUT .4
14		29	P3	44		59	P6
						60	leap second

Figure 1. WWV/H Timecode Format

on standard frequencies of 2.5, 5, 10 and 15 MHz (WWV includes 20 MHz) using a program of voice announcements and audio tones at 440, 500 and 600 Hz. In addition, the stations transmit a time-of-century digital timecode designed for displays and other electronic media. Only the timecode and related pulses at 100, 1000 and 1200 Hz are used by the DSP clock.

The WWV signal format is described in [5]. It consists of three elements, a 5-ms, 1000-Hz pulse, which occurs at the beginning of each second, a 800-ms, 1000-Hz pulse, which occurs at the beginning of each minute, and a pulse-width modulated 100-Hz subcarrier, which carries the data bits. The WWVH format is identical, except that the 1000-Hz pulses are sent at 1200 Hz. The second and minute pulses are transmitted at 100-percent modulation, while the subcarrier modulation is transmitted 10 dB below the 100-percent level, one pulse per second. In the Inter-Range Instrumentation Group (IRIG) code, from which the WWV/H code is derived, a zero is encoded as a 200-ms pulse, one as a 500-ms pulse and a special position marker as a 800-ms pulse. However, in the WWV/H format, the subcarrier is suppressed for the first 30 ms of the second, in order to make room for the 5-ms sync pulse.

Each minute encodes nine BCD digits for the time plus seven miscellaneous bits, as shown in Figure 1. There is no subcarrier modulation in second zero of the minute, in order to make room for the 800-ms sync pulse. Also, second 60 is present only on the occasion of a leap second. There are six position identifiers, encoded P1 through P6, which are intended to assist the synchronization and decoding functions in some equipment, but are not used by the DSP clock. Seconds not numbered in the figure are not used in the current transmission format and are sent as zeros.

The time-of-century is sent as nine BCD-encoded digits, low-order bit and least-significant digit first. The leap warning bit indicates that a leap second is to be inserted after the last minute of the last day of June or December and is numbered second 60 of that minute. The DST1 bit is set to

one at 0h UTC on the day of transition to daylight time and thereafter and set to zero at 0h UTC on the day of transition to standard time and thereafter. The DST2 bit follows the state of the DST1 bit, but delayed 24 hours later. The DUT sign and magnitude encode the UT1 time correction in tenths of a second. The UT1 correction is used by astronomers and navigators to establish a precise star transit time.

4. Hardware Functional Description

The hardware chosen for the DSP clock is called the DSP-93 and is supplied in kit form by the nonprofit Tucson Amateur Packet Radio (TAPR) organization operating in partnership with the Amateur Satellite (AMSAT) organization. Previously, TAPR developed the packet radio technology commonly used by for-profit firms manufacturing inexpensive VHF and UHF wireless modems. The DSP-93 includes a Texas Instruments (TI) TMS320C25 DSP chip [6] and various interface chips for analog/digital conversion, radio control and serial input/output interface to the PC. The TMS320C25 uses a 16-bit pipelined design and operates with a 40-MHz clock and 10 MIPS instruction rate. It communicates with the PC using a universal asynchronous receiver/transmitter (UART) chip operating at 19,200 baud.

The DSP-93 consists of a digital circuit board, an analog circuit board and an external power supply transformer. The digital board includes the TMS320C25, 256K words of EPROM used for the monitor program, 32K words of program RAM and 32K words of data RAM. The monitor program is used only during program development. The analog board includes a TI TLC32044 14-bit voiceband analog interface chip [7], a 16550 UART and various chips for audio processing and transceiver control.

Figure 2 is a block diagram of the DSP-93 hardware components. The DSP-93 program controls the front panel LED display, the UART used to communicate with the PC program, the analog/digital converter (ADC) and digital/analog converter (DAC) in the analog interface, as well as two audio switches and the transceiver control interface. One of the audio switches selects which of several sources are input to the ADC, including the Audio In input of either transceiver port, the DAC output for test signal generation and monitoring, or the auxiliary (Squelch) input of either transceiver port. The other audio switch is used in the feedback path of a differential amplifier to implement a programmable gain function. Potentiometers are provided to set audio input and output levels for each transceiver port.

The DSP-93 program can individually control either of two transceivers, one connected to each radio port, using the push-to-talk (PTT) output, which serves the same function as on a typical mobile microphone. Most transceivers can connect the Audio Out directly to the microphone audio input; however, some transceivers equipped for direct FSK keying can use the Key output instead. Some transceivers are equipped to shift operating frequency in discrete steps up or down in response to a button press on the microphone. The program can use the Freq Up and Freq Down outputs to perform these same functions.

The analog interface includes, besides the basic ADC and DAC functions, a set of programmable dividers driven from the TMS320C25 clock. As used by the DSP-93 program, these dividers are programmed to provide a basic analog/digital conversion clock rate as close to 8000 Hz as possible. This clock rate was chosen both for compatibility with standard telephone conventions, as well as a compromise between various sources of noise and available processing latencies.

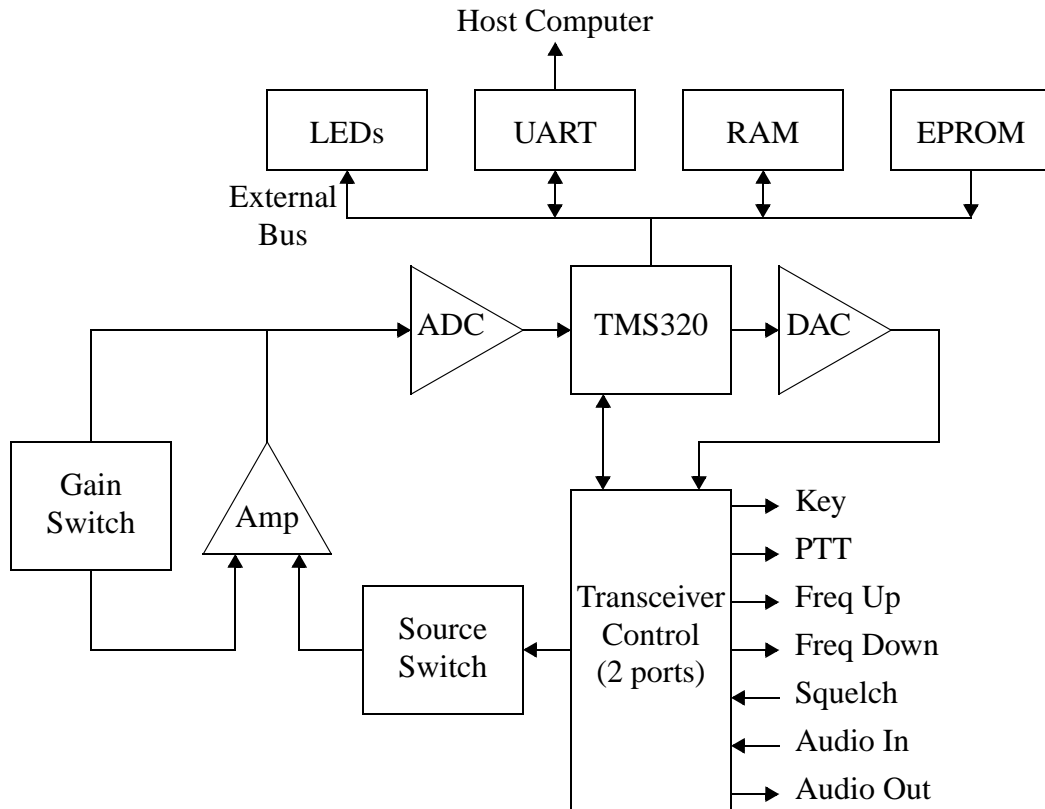


Figure 2. DSP-93 Hardware Architecture

The monitor program resides in EPROM on the digital board. It includes a rudimentary debugger, program loader, a suite of prebuilt modem programs, and a set of utility programs, including a digital oscilloscope and spectrum analyzer. The program loader is used to load DSP-93 programs developed on the PC, but is disabled when the DSP-93 program is in operation. Current program development tools include a rudimentary shareware assembler, which is quite adequate for this application.

5. Software Functional Description

The DSP-93 program consists of interlocking loops driven by hardware interrupts from the ADC, DAC and UART. At the heart of the program is a circular buffer which holds input samples from the ADC before processing by the program and output samples from the program before output to the DAC. The buffering strategy, adapted from [3], is designed to relieve the programmer from intricate loop unrolling and decimating. Most of the interrupt, buffering, command interpretation and storage management functions were incorporated intact from the high performance modem program described in [3] and will not be described further here.

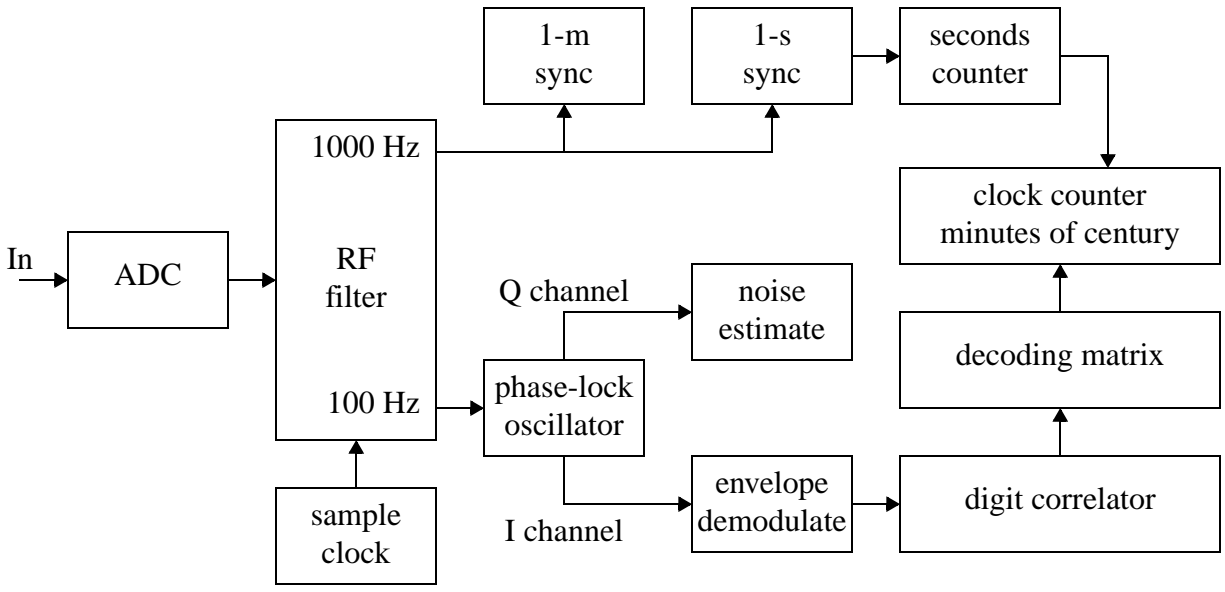


Figure 3. DSP Clock Software Architecture




Figure 4. 400-Hz Bandpass Filter Response

5.1 RF Filtering and Sync Extraction

A block diagram of the DSP processing is shown in Figure 3. The analog audio signal from the radio is sampled at 8000 Hz and converted to digital representation. The 1000/1200-Hz pulses and 100-Hz subcarrier are first separated using two FIR filters, a 400-Hz bandpass filter centered on 1100 Hz and a 150-Hz lowpass filter. Figure 4 shows the response of the bandpass filter, while




Figure 4. 150-Hz Lowpass Filter Response




Figure 5. 1-s Matched Filter Output - WWV Signal

Figure 4 shows the response of the lowpass filter. As evident on both figures, the stopband attenuation is at least 40 dB.

The 1-s sync pulse is extracted using a 5-ms matched filter and exponentially averaged in a comb filter. Figure 5 shows the output of the WWV matched filter with the 5-cycle WWV signal, while Figure 6 shows the output of the same filter with the 6-cycle WWVH signal. The plots show a 6-dB difference in the peak amplitude, which is the only way the DSP clock can differentiate between the two stations. The maximum over all 8000 stages of the comb filter establishes the first sample of the second. The 1-m sync pulse is extracted using a 800-ms noncoherent integrator and exponentially averaged in a comb filter. The maximum over all 60-stages of the comb filter establishes the first second of the minute.




Figure 6. 1-s Matched Filter Output - WWVH Signal

Considerable care is taken in the suppression of errors due to jitter in the 1-s sync signal, which is the master timing source for the DSP clock. A three-stage median filter is used to clean up sync samples produced by the 1-s comb filter. Samples are rejected if the span (maximum minus minimum) value exceeds 1 ms or if the amplitude of the pulse is less than a decision threshold of 1500. A noise gate requires a run of ten identical samples before changing the epoch.

Depending on the location of the radio and time of day, it may happen that reception of WWV and WWVH signals may overlap in time. The matched filter used for the selected station suppresses the unselected station by 6 dB, but this might not be enough to avoid hopping between stations. To deal with this, A jitter gate discards all samples with offset from the current epoch greater than 1 ms, unless a run with length equal to the current averaging interval has occurred. This works on most occasions where the overlap does not last more than fifteen minutes and the difference in propagation delays is greater than 1 ms.

The ADC sample clock is derived from the 40-MHz DSP-93 clock oscillator and a system of hardware counters. The sample clock frequency is actually near 7949 Hz, due to restrictions on the counter ratios; however, samples are stuffed by the DSP-93 program in order to steer the average frequency very near to 8000 Hz. The clock frequency is disciplined by a frequency-lock loop (FLL) which operates independently of the sync and data recovery functions. At averaging intervals determined by the measured jitter, the frequency error is calculated as the difference between the most recent and the current epoch divided by the interval. The sample clock frequency is then corrected by this amount using an exponential average.

Another jitter gate rejects frequency averaging intervals where the difference between the beginning and ending epoch is greater than 500 μ s, unless two adjacent intervals occur with differences greater than 500 μ s, in which case the latest difference is used. The program begins with an averaging interval of 8 s, in order to refine the nominal frequency offset to less than one sample clock tick per averaging interval, or about 16 PPM. After four intervals when this is the case, the program doubles the interval and refines the offset to less than one tick per interval, or about 8 PPM. Under most conditions, the averaging interval doubles in stages from 8 s to over 1000 s, which




Figure 7. 100-Hz Subcarrier Envelope

results in an ultimate frequency precision of 0.125 PPM, or about 11 ms/day. This is consistent with the stability of a TCXO oscillator and somewhat better than the uncompensated clock oscillator used in the DSP-93.

5.2 Subcarrier Demodulation

The phase of the 100-Hz subcarrier relative to the 1-s sync pulse is fixed at the transmitter; however, the audio highpass filter in most radios affects the phase response at 100 Hz. The phase delay is specific to each radio and cannot be predicted. In order to adapt to each radio, the data pulses are demodulated using a 100-Hz phase-locked oscillator (PLO) and two 170-ms matched filters, one for the I (in-phase) signal, the other for the Q (quadrature-phase) signal. The I signal tracks the subcarrier envelope amplitude, while the Q signal is used to minimize the phase difference between the PLO and received subcarrier. In addition, the absolute value of the Q signal is exponentially averaged for use as a noise estimate. It is important to note that the PLO does not control the 100-Hz frequency, since that is established by decimating the 8000-Hz sample clock frequency, which in turn is disciplined by the FLL and ultimately the 1-s sync pulse. Thus, the PLO controls only the 100-Hz phase relative to the 1-s sync pulse.

Figure 7 shows the subcarrier envelope signal after the I-channel 170-ms matched filter. The data bit probabilities are determined from the envelope using a threshold-corrected slicer. The envelope amplitude 30 ms in the second establishes the minimum (noise floor) value. The envelope rises steadily after that until reaching a peak at 200 ms, which establishes the maximum (signal peak) value. The slice level is midway between these two values. The negative-going envelope transition at the slice level establishes the length of the data pulse, which in turn establishes probability estimates for each of three signals, binary zero (P0), binary one (P1) and a position identifier (P2), which is not used by the program. The probability values are established by linear interpolation between the pulse lengths for P0 (300 ms), P1 (500 ms) and P2 (800 ms) so that the sum of all three values is equal to one.

Since the DSP clock is expected to perform well under conditions when the WWV/H signals are marginal, a series of sanity checks is performed to insure that the data are valid and do not represent fortuitous noise spikes that look like data. The sanity checks include the following:

1. The DSP clock must be synchronized to the second and minute,
2. The signal peak must be greater than a decision threshold of 1000,
3. The ratio of the signal peak to noise floor must be greater than a decision threshold of 8,

4. The pulse length must be in the range 200-1000 ms.

If any of these sanity checks fail, the data bit is considered invalid and all three probabilities are set to zero.

The difference between the P1 and P0 probabilities for each data bit is exponentially averaged in a set of 60 accumulators, one for each second, to determine the semi-static timecode data, such as DST indicator, leap second warning and DUT correction. In this design, an average value larger than a positive threshold (+1000) indicates as one and a less than a negative threshold (−1000) as zero. Values between the two thresholds, which can occur due to signal fades or loss of signal, are interpreted as “don’t care,” and result in no change of indication.

The sanity checks and decision thresholds have been established experimentally for reliable sub-carrier demodulation, while the probability mappings have been chosen for reliable pulse-width demodulation of the data signal itself. Correct PLO phase is also essential for reliable data recovery. The principal hazards in this process are phase discontinuities during times the WWV/H signals are rising from the noise to useful levels and, especially, when signals from both WWV and WWVH are received at the same time and interfere with each other due to different ionospheric ray paths.

5.3 Digit Decoding

The BCD digit in each digit position of the timecode is represented as four data bits, all of which must be valid for the digit itself to be considered valid. If so, the bits are correlated with the bits corresponding to each of the valid decimal digits in this position. If the digit is invalid, the correlated value for all digits in this position is assumed zero. In either case, the values for all digits are exponentially averaged in a likelihood vector associated with this position. The digit associated with the maximum over all of the averaged values then becomes the maximum likelihood selection for this position.

The decoding matrix contains nine row vectors, one for each timecode digit position. Each row vector includes the likelihood vector, along with other related data. The maximum likelihood digit for each of the nine digit positions becomes the maximum likelihood time of the century. A built-in transition function implements a conventional clock with decimal digits that count the minutes, hours, days and years, as corrected for leap seconds and leap years. The counting operation also rotates the likelihood vector corresponding to each decimal digit as it advances. Thus, once the seconds and minutes have been synchronized, the maximum likelihood time in any minute should correspond to the BCD timecode transmitted in that minute.

Each row of the decoding matrix includes the likelihood vector, clock digit, a compare counter and the difference (mod 10) between the current clock digit and most recently determined maximum likelihood digit. If a digit probability exceeds the decision level and the difference is constant for four successive minutes in any row, the maximum likelihood digit replaces the clock digit in that row. When the differences are zero for all rows, the clock is synchronized and delivers correct time to the integral second. The millisecond within the second is derived from the ADC sample clock, which runs at 8000 Hz and drives all system timing functions.

In addition to the synchronization indicator character in the ASCII timecode, four status bits are encoded in a quality character. One bit indicates whether the seconds and minutes are synchro-

nized, another indicates whether any of the likelihood values are less than the decision threshold, a third indicates whether the subcarrier envelope for any second is valid, and the fourth indicates whether any maximum likelihood digit differs from the corresponding clock digit. In addition, a set of control commands can be used to select the operating mode and optional debug message types.

An important feature of the DSP clock design is that the maximum likelihood digit depends only on the correlated digit values, which are exponentially averaged. Valid new data in effect charges the likelihood values, while the absence of new data slowly discharges them. Since the maximum likelihood digit depends only on comparing likelihood vector entries and the integration constant for all entries is the same, the digit is unchanged even under propagation conditions which result in missing timecode bits or BCD digits. If under these conditions the maximum likelihood value for any digit decays below the decision threshold, the compare counter is reset and updating of that digit is suspended.

6. Performance Analysis

Since the DSP clock operates under conditions where received WWV/H signals normally vary from unusable to excellent, the DSP-93 program design places considerable emphasis on avoiding false indications due to marginal data and destructive interference between WWV and WWVH signals when both are received at the same time. A successful synchronization sequence requires each of a series of steps to be completed before the next step is begun and all steps must be completed before any changes are made to the clock digits or miscellaneous bits. This section discusses the performance of the DSP clock with respect to signal to noise ratios and error rates for the synchronization, demodulation and decoding functions.

In the following, the signal to noise ratio (SNR) is defined as symbol energy divided by noise power density E_s/N_0 , where the noise process is assumed Gaussian. The general plan is to first derive the processing gain, defined as the SNR improvement between the input and output of the various filtering and averaging algorithms, then establish the required SNR for specified error rate, and finally compute the minimum acceptable SNR at the DSP-93 audio input.

The first step is second synchronization recovery. From standard analysis [1], the 400-Hz band-pass filter and 5-ms matched filter provide an SNR improvement of $10\log(2T_sB)$, where $B = 2100$ Hz is the input bandwidth and $T_s = 5$ ms is the integration time, or 10.2 dB. It is necessary to distinguish the center cycle of the matched filter signal from the eight others that accompany it; but as shown in Figure 5, this cycle has amplitude only 1/5 more than the two surrounding it, which results in a processing loss of 7 dB and a total processing gain of $10.2 - 7 = 3.2$ dB. The averaging constant of the 1-s comb filter follows the frequency averaging interval at 1/4 its value. At the longest interval, the averaging constant is 256, which results in a processing gain of $10\log(256) = 24.1$ dB and a total processing gain of $3.2 + 24.1 = 27.3$ dB.

If P_e is the probability of a single error (wrong choice of center cycle), then $P_e(1 - P_e)^n$ is the probability of a run of n samples ending in an error. If half of these runs have ten samples or more, then P_e must be less than .075. This is a conservative bound, since the median filter corrects a single error. Selection of the center sample amounts to detection of an on-off keyed (OOK) wave-

form, which has bit error rate $P_e = Q(E_s/N_0) = \frac{1}{2}\text{erfc}\left(\sqrt{\frac{\text{SNR}}{2}}\right)$ [1]. For the assumed $P_e = .075$, the SNR must be greater than 3.2 dB. Since the 1-s sync pulse is transmitted at 100 percent carrier modulation, which is 3 dB below the carrier power, the minimum SNR at $P_e = .075$ is $3 + 3.2 - 27.3 = -21.1$ dB. At this SNR, it is very unlikely that the 1-s sync pulse can be detected by ear, even with a sensitive communication receiver.

Once the second has been synchronized, the 100-Hz PLO operates continuously to establish the correct subcarrier phase, as described above. The processing gain for the PLO demodulation using matched filters is $10\log(2T_sB)$, where $B = 2100$ Hz and $T_s = 170$ ms, or 25.5 dB. Invoking the same argument as for the 1-s sync signal, but accounting for the subcarrier modulation 10 dB below 100 percent, the minimum SNR at $P_e = .075$ is $13 + 3.2 - 25.5 = -9.3$ dB. Note that the distinction between a P0 and P1 pulse is equivalent to detecting an OOK pulse occupying the interval 200-500 ms in the second. The 170-ms matched filter is suboptimal for this, but not enough to significantly affect the conclusions.

The next step is minute synchronization recovery. The noncoherent integrator sums the I-channel and Q-channel squares for 800 ms in the first second of the minute. The processing gain is $10\log(800/2) = 26$ dB. The 1-m comb filter has an averaging constant of 8, which provides an additional processing gain of $10\log(8) = 9$ dB, for a total of 35 dB. Again invoking the analysis used for the 1-s sync pulse, the minimum SNR at $P_e = .075$ is $3 + 3.2 - 35 = -28.8$ dB. At this level, it is very unlikely that the 1-m sync pulse can be detected by ear, even with a sensitive communication receiver.

Once the minute has been synchronized, miscellaneous timecode bits are demodulated and exponentially averaged as described above. The bit error rate, even at $P_e = .075$, is very low due to the hysteresis thresholds at ± 1000 and the averaging constant of 16, which is equivalent to a time constant of 16 minutes. In effect, this adds another 12 dB to the processing gain and results in a minimum SNR at $P_e = .075$ of $-9.3 - 12 = -21.3$ dB. Note that the minimum SNR for both second and minute synchronization are well below the minimum SNR for subcarrier demodulation, so that in the demodulation analysis the effects of synchronization errors can be neglected.

The next step is to determine the digit in the minute units digit position. This must be done before any other digit in the timecode, since only when the carry from the minute units digit has been determined can the higher order carries be determined. When four minute digits with maximum likelihood values exceeding a decision threshold of 1000 have been successfully compared, the minute units clock digit is set and the remaining likelihood vectors allowed to accumulate values. No matter what the phase of the remaining digits happen to be, there is at least ten minutes to accumulate values before any of these digits can change. Normally, this is sufficient time for the likelihood values to rise above the decision threshold and the compare process to begin.

To set a clock digit, a run of four, 4-bit BCD digits must be correctly decoded. The probability that a run of 16 or more correct bits will occur at least half the time requires P_e less than .045, which in turn requires a minimum SNR of 4.5 dB. But, the likelihood values are integrated with an averaging constant of 16, which provides a processing gain of 12 dB. Together with the subcarrier processing gain derived above, this results in a minimum SNR of $4.5 - 9.3 - 12 = -16.2$ dB.

The clock can be set wrong only if four or more successive maximum likelihood digits are wrong. The probability of a digit decoding error at the $P_e=.045$ is $1 - (1 - P_e)^4 = 0.169$. A clock error will occur if four or more such decoding errors occur in a row, which occurs with probability 9.8×10^{-4} , or about 1.4 errors per day. It is important to note that such periods when operating at the minimum SNR assumed here are short, generally lasting only a few minutes as propagation conditions open and close between WWV/H and the radio location. At other times, either the signal is completely absent or well above the minimum SNR.

7. Measured Performance

An important goal in the DSP clock design is that the accuracy and stability of the indicated time be limited only by the characteristics of the propagation medium. Conventional wisdom is that time synchronization via the shortwave medium is good only to a millisecond under the best propagation conditions. The performance of the DSP clock tracking the WWV/H signals is clearly better than this, even under marginal conditions. Ordinarily, with marginal to good signals and a frequency averaging interval of 1024 s, the frequency is stabilized within 0.1 PPM and the time within 125 μ s. The frequency stability characteristic is highly important, since the clock may have to free-run for several hours when the WWV/H signals are unavailable.

The time accuracy over a typical day has been measured using an oscilloscope and cesium oscillator calibrated with a GPS timing receiver. With marginal signals and allowing 15 minutes for initial synchronization and frequency correction, the time accuracy determined from the 1-s sync pulse is reliably within 125 μ s. In the particular DSP-93 used for program development, the uncorrected clock frequency offset is 45.8 ± 0.1 PPM. Over the first hour after initial synchronization, the clock frequency varies about 1 PPM as the frequency averaging interval increases to the maximum 1024 s.

Figures 8 through 13 show the DSP clock behavior measured under typical conditions over three days and nights in late July, 1997. For these tests the radio was a sensitive communications receiver and the antenna a 135-foot dipole 12 meters above the ground. The tests were designed to establish SNR margins that might be expected using less dramatic radios and antennas. Reception during daylight hours was generally good, with signals ranging from S3 to S9, as measured by the communications receiver signal strength meter. The data from which the figures were made were collected in real time by the running DSP-93 program and saved in a file on a Unix workstation. These files were later processed using the Math Works MatLab [4] and the Signal Processing Toolbox [2] to produce the figures themselves.

At the time of year the tests were performed, the propagation path between the WWV transmitter at Boulder, CO, and the radio at Newark, DE, is open from middle morning to late evening in Newark. Figure 8 shows the 1-s sync pulse amplitude during the measurement period in decibels. The ratio of the maximum daytime amplitude to the maximum nighttime amplitude represents an SNR of about 40 dB during the daytime. The 1-s sync decision threshold is set at 1500, which corresponds to 63.5 dB on the plot. The SNR margin is defined as the excess SNR over the decision threshold, in this case about 17 dB during the daytime. With another radio and antenna, the performance relative to the test setup could be up to 17 dB worse and still successfully synchronize the second. However, the decision threshold is itself set about 24 dB above the noise floor, far above




Figure 8. 1-s Sync Pulse Amplitude




Figure 9. 1-m Sync Pulse Amplitude

the required minimum SNR of 3.2 dB, and could probably be reduced without increasing the false-alarm rate. Note that, near sunrise and sunset at the midpoint of the path, it is not uncommon for the signal to fade in and out and, while not evident from the figure, for both WWV and WWVH stations to be heard at the same time. These are the occasions of maximum stress on the engineered algorithms to avoid false indications.

Figure 9 shows the 1-m sync pulse amplitude in decibels. While the ratio of the maximum daytime amplitude to the maximum nighttime amplitude represents an SNR well over 60 dB, this figure is misleading, since the 1-m signal is enabled only if the 1-s signal has successfully synchronized the second. The 1-m sync decision threshold is set at 500, which corresponds to 54 dB on the plot and an SNR margin of about 26 dB. As in the 1-s sync pulse case, the threshold could probably be reduced without increasing the false-alarm rate. The plot shows the results of




Figure 10. 100-Hz Subcarrier SNR

almost 16 s of integration and suggests that minute synchronization may never be a problem, even on far weaker signals.

Figure 10 shows the 100-Hz subcarrier SNR, measured as the ratio of the peak I-channel level to the average Q-channel level in decibels. The ratio of the maximum daytime amplitude to the maximum nighttime amplitude represents an SNR of about 25 dB during the daytime. The SNR decision threshold is set at 8, which corresponds to 18.1 dB, and an SNR margin of about 12 dB. As the only function of this threshold is to confirm that the subcarrier phase is in the correct quadrant, it could probably be reduced by a few decibels without increasing the false-alarm rate. While a margin of 18.1 dB is somewhat less than the other margins, there is still much more processing gain before a digit is decoded or a miscellaneous bit is set. One thing to note about the plot is the very wide variation over almost 40 dB during the day as the signals fade in and out.

Figure 11 shows the likelihood values in decibels. While the ratio of the maximum daytime amplitude to the maximum nighttime amplitude represents an SNR well over 60 dB, this figure is misleading, since only valid data bits are considered for likelihood and correlation computations. The likelihood decision threshold is set at 1000, which corresponds to 60 dB on the plot or an SNR margin of about 17 dB. As in the other cases, this threshold could probably be reduced without increasing the false-alarm rate.

With the above data in hand, it is possible to predict how the DSP clock would function with weaker signals than observed during the measurement period. Informal observations confirm that the WWV/H signals produce reliable indications, even when the voice announcements and audio tones cannot be heard. The data show a 1-s sync pulse margin of 17 dB, a 1-m sync pulse margin of 26 dB, a subcarrier SNR margin of 12 dB and a likelihood margin of 17 dB. Clearly, the factor limiting sensitivity is the subcarrier SNR threshold. As mentioned earlier, the SNR threshold could probably be reduced by a few decibels without increasing the error rate, which would raise the overall DSP clock margin to 17 dB. The subcarrier PLO design has not been optimized and the SNR determination is rather crude. However, the present design performs far better than the best




Figure 11. Digit Likelihood Values




Figure 12. Clock Oscillator Frequency

known commercial radio clocks, so it may not be worthwhile to squeeze out the last ounce of processing.

Figure 12 shows the clock oscillator frequency variations over the period. During intervals when WWV/H signals are available, the trace varies over almost 1 PPM, but at other times appears as a straight line. It is not clear whether the variations are due to doppler as the ionospheric layers rise and fall, to the inherent instability of the DSP-93 clock oscillator, or some combination of both. The characteristics shown in the figure suggest that much of the variation may be due to ionospheric motion. To the extent that this motion cannot be reliably predicted and the frequency varies typically over 0.5 PPM, a phase adjustment as high as 18 ms may be necessary after ten hours with no signal. Figure 13 clearly shows these phase adjustments as discontinuities when the WWV/H signals resume after some hours when they are unavailable.




Figure 13. Clock Phase Adjustments

The propagation delay from the WWV transmitter at Boulder, CO, to a communications receiver and DSP clock at Newark, DE, is 23.5 ± 0.1 ms, as measured by a calibrated cesium clock. The delay is measured to the peak of the pulse after the 1-s sync comb filter and includes components due to the ionospheric propagation delay, receiver delay and DSP-93 delay. The measured DSP-93 delay is 5.5 ms, most of which is due to the 400-Hz bandpass filter and 5-ms matched filter.

8. Conclusions

It is tempting to consider the ultimate DSP clock performance capability. The most critical limiting factor in performance is the PLO, which is necessary only because the phase delay characteristics of the particular radio are not known and must be measured during operation. It should be noted that the DSP-93 clock frequency depends only on the 1-s sync pulse and the PLO is used only to adjust the phase. The subcarrier SNR decision threshold is used only to confirm the PLO has indeed locked on the correct phase. Thus, a simple solution would be to reduce the PLO loop gain, effectively increasing the time constant. This would require a better DSP-93 clock oscillator such as a TCXO. Once reliable phase has been determined, which may take some time, the PLO feedback loop can be opened, the SNR gate removed and the 100-Hz signal left on dead reckoning.

Once the subcarrier SNR threshold has been removed, the likelihood margin of 17 dB becomes a limiting factor. As long as correct second and minute synchronization can be maintained, the integration time, currently 16 m, can be increased indefinitely. As a practical matter, an increase of four times (6 dB) coupled with reducing the likelihood threshold by 3 dB would raise the margin to 26 dB. At this point, the 1-s sync margin would be the limiting factor. Increasing the integration time by 6 dB plus lowering the threshold by 3 dB would raise the margin to 26 dB, making the overall margin for the DSP clock to 26 dB. This might even allow a paperclip to be used as an antenna or, with a good antenna, a reduction in WWV/H transmitter power from 10 kW to a pipsqueak 25 W.

A number of improvements other than the above might be considered for further refinement of the DSP clock. Among these might be the following:

1. The clock should provide optional local timezone correction, standard/daylight time adjustment and DUT correction.
2. The clock should provide optional formats emulating other radio clocks on the market.
3. The clock should provide port selection, gain control and possibly other features useful for system integration.
4. The clock should provide automatic UART baud rate control using the AT command.
5. The clock should be rewritten in portable C for use on a fast workstation with integrated audio codec.
6. The clock should support the Canadian time/frequency station CHU. This would require provision for a 300-bps modem and not much else.

9. References

1. Couch, L.W. II, Digital and Analog Communication Systems. Macmillan, New York, 1993.
2. Little, J., and L. Shure. Signal processing toolbox for use with MATLAB. MathWorks, Inc., 1988.
3. Mills, D.L. An optimal linear receiver and codec for a class of radiotelegraph signals. Electrical Engineering Technical Report TR 95-8-1, University of Delaware, August 1995.
4. Moler, C., J. Little, and S. Bangert. Pro-Matlab User's Guide. MathWorks, Inc., 1990.
5. NIST Time and Frequency Services. NIST Special Publication 432 (Revised 1990), Washington, DC, 1991.
6. Texas Instruments, Inc. Second Generation TMS320 User's Guide. Document Number SPRU0014. Texas Instruments, Inc., 1987.
7. Texas Instruments, Inc. Linear Circuits, Volume 2 - Data Conversion, DSP Analog Interface and Video Interface. Texas Instruments, Inc., 1987.

10. Appendix A. Operation Notes

The DSP-93 program begins operation immediately upon startup. It first searches for 1-s sync, then 1-m sync, which can take up to several minutes, depending on signal quality. When 1-s sync has been acquired, LED4 on the DSP-93 front panel begins to flash. When 1-m sync has been acquired, LED3 begins to flash as well, with the duration of the flash following the pulse-width modulated data pulse. At this time, the 100-Hz phase-lock loop adjusts the PLO demodulator to the data pulse phase, which may take a minute or so. LED5 is on when the amplitude of the data pulse exceeds the threshold; however, under marginal conditions or loss of signal, LED5 may be on only occasionally or not at all.

The program then begins accumulating likelihood values for each of the nine digits of the clock, plus the seven miscellaneous bits included in the WWV/H transmission format. LED6 flashes when the likelihood value is below the decision threshold for a digit or miscellaneous bit. LED7 flashes when the current clock digit disagrees with the maximum likelihood digit. The minute units digit is decoded first and, when four repetitions have compared correctly, the remaining eight digits are decoded. When four repetitions of all nine digits have decoded correctly, which normally takes 15 minutes with good signals and up to an hour when buried in noise, the clock is set within 125 μ s and can be read within 500 μ s with the ASCII timecode formats described later.

At the same time, the clock discipline algorithm refines the frequency offset correction for use during times when the WWV/H signal is unavailable. The algorithm refines the offset using increasingly longer averaging intervals to 1024 s, where the precision is about 0.1 PPM. With good signals, it takes well over two hours to reach this degree of precision; however, it can take many more hours than this in case of marginal signals.

To work well, this DSP clock needs a communications receiver with good audio response at 100 Hz. Most shortwave and communications receivers roll off the audio response below 250 Hz, so this can be a problem, especially with receivers using DSP technology, since DSP filters can have very fast rolloff outside the passband. Some DSP transceivers, in particular the ICOM 775, have a programmable low frequency cutoff which can be set as low as 80 Hz. However, this particular radio has a strong low frequency buzz at about 10 Hz which appears in the audio output and can affect data recovery under marginal conditions. Although not tested, it would seem very likely that a cheap shortwave receiver could function just as well as an expensive communications receiver.

10.1 Front Panel Indicators

The status of the program is shown by the DSP-93 front panel LEDs, which are interpreted as follows:

- LED1 Signal level indicator. Adjust the input signal level until this LED flashes on occasional peaks.
- LED2 Activity indicator. Intensity corresponds to the fraction of idle processor time; i.e., not doing signal processing functions.
- LED3 Data indicator. Follows the 100-Hz subcarrier modulation that carries the data bits. This is active only when in 1-s sync.

- LED4 Sync indicator. Blinks to indicate 1-s and 1-m sync. Short blinks follow the 1-s sync pulse, while long blinks follow the 1-m sync pulse. This is active when in 1-s sync and indefinitely after the clock has been synchronized. If the clock has not been in both 1-s sync and 1-m sync during the minute, status bit 0 is set in the following minute.
- LED5 Data present. On for one second only when in 1-s sync and following a data bit with amplitude above the decision threshold. A data bit with amplitude below the decision threshold has assumed bit probability of zero, but does not necessarily indicate a clock error. Note that there is no data bit in second zero of the minute. If LED5 has not been on during the minute, status bit 1 is set in the following minute.
- LED6 Low decoded signal probability. On only in 1-s sync and when a correlated bit or digit probability is below the decision threshold. The associated bit probability is zero, but does not necessarily indicate a clock error. LED6 when on causes status bit 2 to be set in the following minute.
- LED7 Decoding error. On only in 1-s sync and when the maximum likelihood digit differs from the current clock digit. The associated digit is not used to set the clock, but does not necessarily indicate a clock error. LED7 when on causes status bit 3 to be set in the following minute.
- LED8 DSP-93 Power indicator

10.2 Program Commands

The DSP-93 program responds to a number of commands, which consist of a single case-insensitive letter followed in some cases by one or more digits. Control characters, including CR and LF, are ignored, except in the C and H commands. The station to be monitored is selected using the C (WWV) or H (WWVH) commands. The only affect these commands have is to select which 5-ms matched filter to use for the 1-s sync, which tone (1000 Hz or 1200 Hz) to use for the 1-m sync integrator, and which propagation delay value to use. These commands can be followed by a decimal number, which specifies the propagation delay in 125-us units, and then a CR to enter the value. A C or H command followed immediately by a CR selects the station, but does not change the propagation delay.

B{0-9}

This command selects the UART baud rate. The digit argument is interpreted as follows:

0	300
1	600
2	1200
3	2400
4	4800
5	9600
6	19200
7	38400
8	76800
9	153600

The default is set by the program loader, which is 19200 with the TAPR firmware monitor.

P{0-2}

This command enables or disables either the RTS or DTR line from the UART to follow LED4. Note that internally the UART RTS output is connected to the CTS pin (8) on the DB9 connector. The digit argument is interpreted as follows:

- 0 RTS and DTR always high
- 1 RTS follows LED4, DTR always high
- 2 DTR follows LED4, RTS always high

R

This command resets the DSP clock to the initial state at the point control is transferred by the TAPR firmware monitor to the DSP-93 program.

S{0-9}

This command sets the nine digits of the clock from low to high order. This command is useful in order to verify proper rollover of the leap second and leap year, for example.

T

This command returns the timecode string in the following format preceded by a CR/LF sequence. This format conforms to Format 2 as used by the Spectracom WWVB and GPS radio clocks and supported by the current version of the Network Time Protocol (NTP) software for Unix and Windows.

```
sqyy ddd hh:mm:ss.fff ld
```

- s sync indicator ('?' or '')
- q quality character ('')
- yy year of century
- ddd day of year
- hh hour of day
- mm minute of hour
- ss second of minute
- fff millisecond of second
- l leap second warning '' or 'L'
- d DST state 'S', 'D', 'I', or 'O'

The on-time reference is the start bit of the first (CR) character. The sync indicator s is '?' until the clock is set by valid received data and '' after that. The quality character q is always '' in this format.

The leap second warning l is 'L' if a leap second is expected at the end of the month and '' if not. By international agreement, the leap second is inserted following the last minute of the last day of June or December and is numbered second 60 of that minute. The WWV/H timecode progression just before, during and after the leap on 30 June 1997 is shown below, where the first column is the year, followed by the day, time, leap second warning, DST state and DUT correction (see U command):

```

97 181 23:59:59 LD -5
97 181 23:59:60 LD -5
97 182 00:00:00 D +5

```

Note the seconds progression includes second 60 of the last minute and the day number advances at the next second. Also, the leap warning disappears and the DUT correction -5 changes to +5 at that time; however, in actual use these changes are delayed a few minutes as the result of the integration functions.

The DST state is determined from the DST1 and DST2 bits of the WWV/H timecode. The DST1 bit is set to one at 0h UTC on the day of transition to daylight time and thereafter. The bit is set to zero at 0h UTC on the day of transition to standard time and thereafter. The DST2 bit follows the state of the DST1 bit, but delayed 24 hours later. The values are decoded as follows:

```

S      standard time
D      daylight time
I      between 0h on the day to daylight time until 0h the next day
O      between 0h on the day to standard time until 0h the next day
U

```

This command returns the timecode string in the following format followed by a CR/LF sequence. This format is intended for experimentation and evaluation and may be changed in future.

```
sq a yy ddd hh:mm:ss.fff ld dut ttttt fffff ggggg uuuuu
```

where the elements of the timecode are as described for the T command, with the addition of the following:

```

q      quality character_ astation select C (WWV) or H (WWVH)
dut    DUT sign and magnitude in deciseconds
ttttt  median filter time offset in 125-us units
fffff  frequency offset in units of about 1.9 ns/s
ggggg  frequency averaging interval in seconds_ uuuuu interval since last clock update in minutes

```

The on-time reference is the start bit of the first (s) character.

The quality character encodes four bits of status formatted in hex. Reading from high to low order, the bits are:

```

0      low 1-s or 1-m sync pulse amplitude
1      low 100-Hz subcarrier amplitude
2      low decoded digit probability_3timecode digit decoding error

```

The DUT sign and magnitude encode the UT1 time correction in tenths of a second. The UT1 correction is used by astronomers and navigators to establish a precise star transit time.

10.3 Debugging and Monitoring Commands

The DSP clock includes a primitive toolkit useful for debugging and evaluation. Debugging mode can be turned on with one of the D commands and off with the D0 command. There are four

debug formats selected by the D1, D2, D3 and D4 commands, respectively. The output is inclusive; that is, one, two or all three formats can be enabled at the same time, but all are disabled by the D0 command.

D1

This command produces one line at second 59 of each minute in the same format as the U command described above.

D2

This command produces one line for each decoded timecode digit in the following format:

```
c m d n q sssss aaaaa mmmm ddddd nnnnn eeeee
```

```
c      clock digit_mmaximum likelihood digit
d      difference (mod 10) between clock and maximum likelihood digits
n      compare counter (max 3)
q      quality character_sssss 1-s sync epoch
aaaaa  1-s sync amplitude (threshold 1500)
mmmmm  1-m sync amplitude(threshold 500)
dddddd 100-Hz subcarrier amplitude (threshold 1000)
nnnnn  noise estimate
eeeee  digit correlator amplitude (threshold 1000)
```

D3

This command produces one line each second in the following format:

```
aaaaa eeeee ppppp fffff ccccc jjjjj vvvvv
```

```
aaaaa  1-s sync amplitude
eeeee  median filter epoch of 1-s sync
ppppp  median filter time offset
fffff  time offset since last frequency update
cccccc compare counter (max 10)
jjjjj  jitter counter (max equal to averaging cycle)
vvvvv  averaging cycle counter
```

D4

This command produces one line at second 58 of each minute in the following format:

```
dddddd eeeee fffff ggggg hhhhh iiiii
```

```
dddddd DST1 integrator value
eeee   DST2 integrator value
ffff   leap warning integrator value
ggggg  DUT± integrator value
hhhhh  DUT1 integrator value
iiii   DUT2 integrator value
jjjj   DUT4 integrator value
```

10.4 Signal Generation and Monitoring

O{1-8}

This command selects certain signals to be monitoring in real time with an oscilloscope connected to the DAC. The synthetic signals are synchronized with respect to the WWV/WWVH transmitter, so may be useful as precision time and frequency sources for other devices, such as laboratory test equipment. Note that the stuffing cycles necessary to produce the standard frequencies from the DSP-93 hardware sample clock produce some phase modulation in the signals, which appear as a low frequency buzz on the 1000/1200-Hz signal. The signals are selected with the O{1-8} command, where the digit is:

- 1 input signal (after peak clipping, but before other processing)
- 2 output of 1-s sync comb filter_3output of data bit matched filter
- 4 synthesized one-second ramp in 125-us increments
- 5 synthesized one-minute ramp in 1-s increments
- 6 synthesized 1000-Hz (WWV) or 1200-Hz (WWVH) sinewave
- 7 synthesized 100-Hz sinewave
- 8 data bit probability - positive (1), negative (0)

For the DC-coupled signals, the oscilloscope should be connected to JP210, pin 2; that is, the pin closest to the power regulators. In the case of the ramps and bit probability, the lowpass filter in the DSP-93 D/A limits the risetime and introduces a moderate degree of ringing.

A speaker can be connected to the DSP-93 output to monitor the operation of the program. The O1 signal is useful for general watchkeeping. The O2 signal is probably the best indication that the WWV/H signal is present. If a faint tick can be heard, the DSP clock can probably synchronize to the signal. Normally, at the decision threshold where the tick can barely be heard, the WWV/H signal cannot be heard in the unprocessed receiver output, although a faint beep may be heard marking the first second of the minute. The O3 signal is best viewed with an oscilloscope, where the waveform should appear as a trapezoid rising at the second and decaying either 200, 500 or 800 ms later. When viewed on an oscilloscope with a calibrated sweep delay and triggered by a PPS signal from a source calibrated to UTC, the O4 signal becomes a precision proof-of-performance indicator.

10.5 Operation Notes

The particular algorithms implemented in the DSP clock have not been thoroughly tested over many months and seasons, nor have the various thresholds been experimentally verified under all conditions. This applies especially to the decision thresholds, called bottom fishers, which determine whether a particular signal is usable or not. Since the DSP-93 program is designed to be very aggressive in sifting signals from noise, if these thresholds are set too low, the program may decode noise as UTC. On the other hand, if these thresholds are set too high, usable signals may be ignored and the clock drift away from UTC. The current thresholds are defined in the output formats described above and in the source code. They are probably set too low in an effort to scrape the weakest signals from the noise. In normal operation, the values are several times the thresholds.

There is a necessary design compromise between weak-signal sensitivity and DSP-93 clock oscillator frequency tolerance. When first started, the program must discipline the sample clock to within 12.5 PPM, in order to reliably identify the 1-s sync signal epoch. At the smallest averaging interval (8 s) and largest jitter threshold (500 μ s) which can reliably extract the epoch, the tolerance must be less than 62.5 PPM. As the frequency offset of the DSP-93 used for program development is over 45 PPM, this could be a problem in some units. While not tested, the program should cope with tolerances greater than 62.5 PPM, but it will take significantly longer to synchronize.

7.