INTRALAB MEMORANDUM

To: Distribution

From: Joe Paradiso

Date: 30-May-1991

Subject: **The Quadrature Sampler**

Memo #: EJC 91-1029

Copies:

Abstract

This memo is essentially an overview and operational guide for the EJ Quadrature Sampler System, which was developed to efficiently digitize bandlimited signals from sonar beamformer arrays (i.e. the WQS1). The functions are first discussed from a "black box" point of view (essentially a users manual), followed by the schematics and a brief analysis. Suggestions for improvements and extensions are provided.

1) Second-Order Quadrature Sampling

The theory behind quadrature sampling was discussed in a previous memo[1], which also showed MATLAB simulations of quadrature-sampled beamformers and gave a strawman sketch of the circuitry that was later realized (as will be described in this report). The paragraphs below are excerpted from Ref. [1], and form a brief introduction to second-order quadrature sampling.

Active sonars often employ highly narrowband signals. In the beamforming application illustrated here, essentially a pure sinusoid is emitted, and the return signal is broadened somewhat by Doppler scattering, dispersion, etc. (but remains narrowband). Any bandpass signal can be represented by a sum of in-phase and quadrature components:

1) $$x(t) = x_I(t) \cos \omega_0 t + x_Q(t) \sin \omega_0 t$$

The in-phase [x_I] and quadrature [x_Q] signals represent a slowly-varying complex envelope which is applied to the high-frequency carrier at ω_0. Under quadrature sampling [2], the signal can be sampled at the Nyquist cutoff of the complex envelope (i.e. at the signal bandwidth) rather than at twice the carrier frequency ($2f_0 = \omega_0/\pi$), yielding an appreciable reduction in sampling rate for narrowband signals.

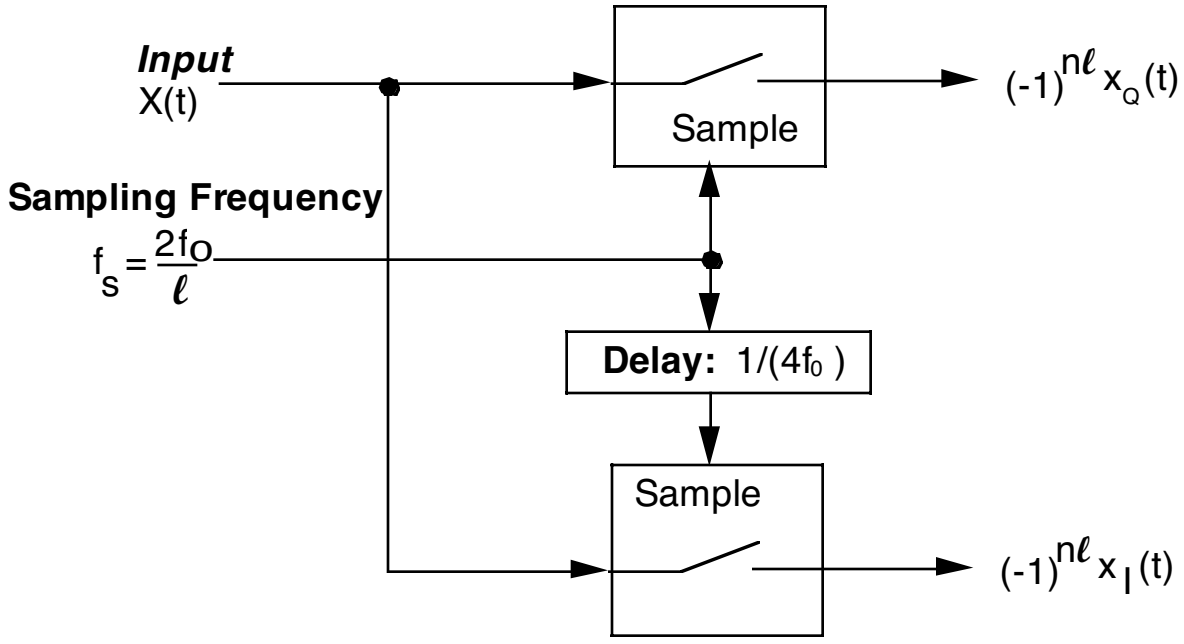

Figure 1: Second-Order Method Approximating Quadrature Sampling

If the input signal is highly narrowband, as it is in the sonar applications considered here, quadrature sampling may be performed via a second-order technique [2]. In-phase signals are sampled at a period of $\Delta = \ell\,(2f_0)^{-1}$, where ℓ is an integer that must be selected such that Δ remains less than W^{-1}, the inverse of the input signal's bandwidth, to avoid aliasing. The quadrature samples are taken after the signal is delayed by α, where $\alpha = (4f_0)^{-1} + K\Delta$, where K is an integer. The simplest case is to pick K to be zero; here the quadrature samples are taken after a quarter-period of the carrier has elapsed, relative to the in-phase samples. The sampling frequency $f_s = \Delta^{-1}$ is again well below the carrier frequency here, since we are sampling the narrowband complex envelope of the signal.

The input signal must be appropriately bandlimited when quadrature sampling at a lower frequency, as in Fig. 1. If any baseband frequencies are present on the input signals (i.e. AC hum, surf noise, fish sounds, etc.), they will alias into the digitized samples. In such a case, the input signal must be bandpass filtered before being input to the second-order sampler.

Other options are available for performing the quadrature sampling. One is to extract the complex envelope by demodulating the signal down to baseband with a quadrature oscillator, analog multipliers, and low-pass filters (as outlined in Ref. [1]). This procedure can involve a degree of complexity in the analog electronics that is avoided by the second-order sampling scheme. Another possibility is to sample the inputs at the carrier Nyquist, then do the quadrature basebanding multiplication and downsampling in a fast DSP. This is indeed possible, but the bandwidth required on the front-end to digitize a multi-element beamformer (i.e. 16 sensors) may present some difficulty. After some consideration of these other options, the second-order technique was adopted for hardware implementation.

In the discussions presented here, the in-phase "I" signal is assumed to be sampled *earlier* than the quadrature "Q" signal. This is the reverse of the standard used in much of the literature, where the Q signal itself is delayed (as opposed to the sampling gate being delayed), thus the convention is for Q be taken before I (see Ref. [1] for detail). Since the hardware has been labeled in the reverse order (I before Q), we will hold to this assumption in this memo. The analytical effect of this change will be to flip the sign on the measured Q component, as has been accommodated by the additional minus sign at the Q output in Fig. 1 (of course, the identities of "I" and "Q" can readily be swapped in postprocessing software to avoid this difficulty).

2) Operation of the Quadrature Sampler System

Commercially available multi-channel data samplers generally lack the ability to precisely latch a large number of analog inputs simultaneously. The second-order sampling technique discussed in the previous section requires the in-phase and quadrature signals to be sequentially sampled at exact instants, which can be difficult to realize with the multiplexed sample/hold front ends on most commercial units. For example, with a 200 Khz center frequency (f_0), as used in the WQS1, the in-phase and quadrature samples must be taken within an interval of 1.25 μsec. of one another. In order to enable second-order quadrature sampling to be performed on a sonar sensor array, an analog front-end has been developed to latch the in-phase and quadrature samples at the precise instants required. These signals may then be digitized and stored by a commercial sampling unit.

The Quadrature Sampler has been realized as a modular device, consisting of 3 rack-mount elements. The simplest is the power supply, providing a power switch and +5,±12 Volts through rear-panel connectors. A Quadrature Clock Generator provides the phased clocking signals needed to take the quadrature samples (plus generate a delayed range gate, toggle the output multiplexer, and strobe the Data Acquisition System [*DAS*]). The analog samplers are in a crate, configured to hold up to 4 units (all 4 have been constructed), with 8 channels per unit (32 channels per crate). The way in which these modules are interconnected and interfaced with the Analogic DAS (to be used in the upcoming WQS1 sonar tests) is illustrated in Fig. 2. The 8 Input signals are applied to the DB25 connector at the left side of a module. These are balanced signals, and the +, -, ⏚ inputs appear on the DB25 where labeled in the Analogic documentation. There are two output connectors on each sampler module, labeled "I/Mux out" and "Q out". These are single-ended outputs, but adhere to the Analogic DB25 standard by tying the "-" end to the local ground. The signals that will normally be applied to the DAS appear at the "I/Mux" output. In standard operation, this will present the "I"(early) sampled signal to all 8 outputs for half of the low-frequency sampling period (i.e. $\Delta/2$), and present the "Q"(later) sampled signal to all 8 outputs for the other half of the sampling period. In this fashion, both in-phase and quadrature signals are multiplexed onto a single cable, and will appear "interleaved" in the data. This comes with a minor penalty, in that the DAS will have to sample two signals per channel during the sampling period Δ (the benefit is that only one 8-channel DAS unit is needed to quadrature-sample 8 input signals, but the DAS must run twice as fast). If one has twice as many DAS channels than input signals, then one can disable the multiplexing (via the "I/Q Mux" switch on the Clock Generator); now only the I signal will appear at the I/Mux output for the entire sampling period Δ, and likewise, the Q signal is available at the Q output. Both of these outputs can then be then digitized by separate DAS units, using the entire Δ interval. At present, insufficient DAS units are available to digitize I and Q separately, thus the multiplexed operation is assumed, as depicted in Fig. 2.

The sampler channels have a gain of 3, and an available dynamic range of up to ±10 Volts, thus input signals should be limited to within ±3 Volts or so. The noise floor of these circuits appears to be somewhere near 5 mV or so (essentially dominated by clock-related parasitics from the sample/holds), but this should be additionally investigated. There are no bandpass filters on the analog inputs; their response is nearly flat out to 350 Khz or so. If the signal source is not truly narrowband, such a filter must be inserted before the sampler unit to avoid aliasing. If a standard DB25-DB25 computer cable is used to interconnect the sampler output to the DAS, it should be kept as short as possible to avoid potential crosstalk and slew problems.

During the construction of the Sampler modules, the estimated output signal of the WQS1 was revised downward, becoming on the order of 0.4 Volt P-P. As this will only produce a 1.2 Volt P-P signal with the Sampler's gain of 3, it was decided that additional gain would be desirable to increase the signal/noise through the Sampler units. Rather than modify the already-constructed circuit cards, it was decided to introduce additional gain in another to-be-constructed unit. This will be a very simple device based around wideband quad VCA chips made by CEM or SSM, and will enable the gain of several channels to be varied simultaneously, as desired, by adjusting a single voltage. This device can also serve as an adaptor between the BNC cables coming from the WQS1 and the DB25 standard of the Analogic DAS (and Quadrature Sampler modules).

Four BNC cables link the Sampler Crate with the Quadrature Clock Generator. One pair of these cables provides the I & Q gates for the Sample/Holds, and the other pair of cables provides the gates for the I/Q output multiplexers. As these cables are somewhat under-terminated, they should be kept short (i.e. within 2-3 ft. if possible).

If desired for coincident sampling applications, the Sampler Crate can also be used on its own, without the Quadrature Clock Generator. In this case, TTL-level sampler gates can be applied to the I and/or Q S/H inputs, which will simultaneously drive all 32 channels in the crate (providing a ready method of "simultaneously" [within 10 nsec. or so] sampling multiple input signals). The multiplexer operation can be locally disabled by putting the toggle switch at the lower-right of the sampler crate in the up ("free") position. When the system is run with the Clock Generator, this switch should be down ("run"), to enable multiplexer toggling. The narrow module that houses this switch is an input conditioner, which buffers the 4 inputs (from the BNC cables), and drives the backplane of the Sampler Crate.

One clock generator has the potential ability to drive several Sampler Crates. One can stack crates for more channels (at 32 per crate) by splitting the 4 interconnect signals with a simple "T" (but keep the cables short!). As it currently stands, the clock generator should be able to fan-out into 2 crates without any problem. One can reliably add even more crates by building a trivial unit to do a TTL fan-out (use 74LS244 drivers or equivalents) to break these 4 signals into multiple outputs. Try and keep the cable lengths identical to eliminate delay shifts between different crates.

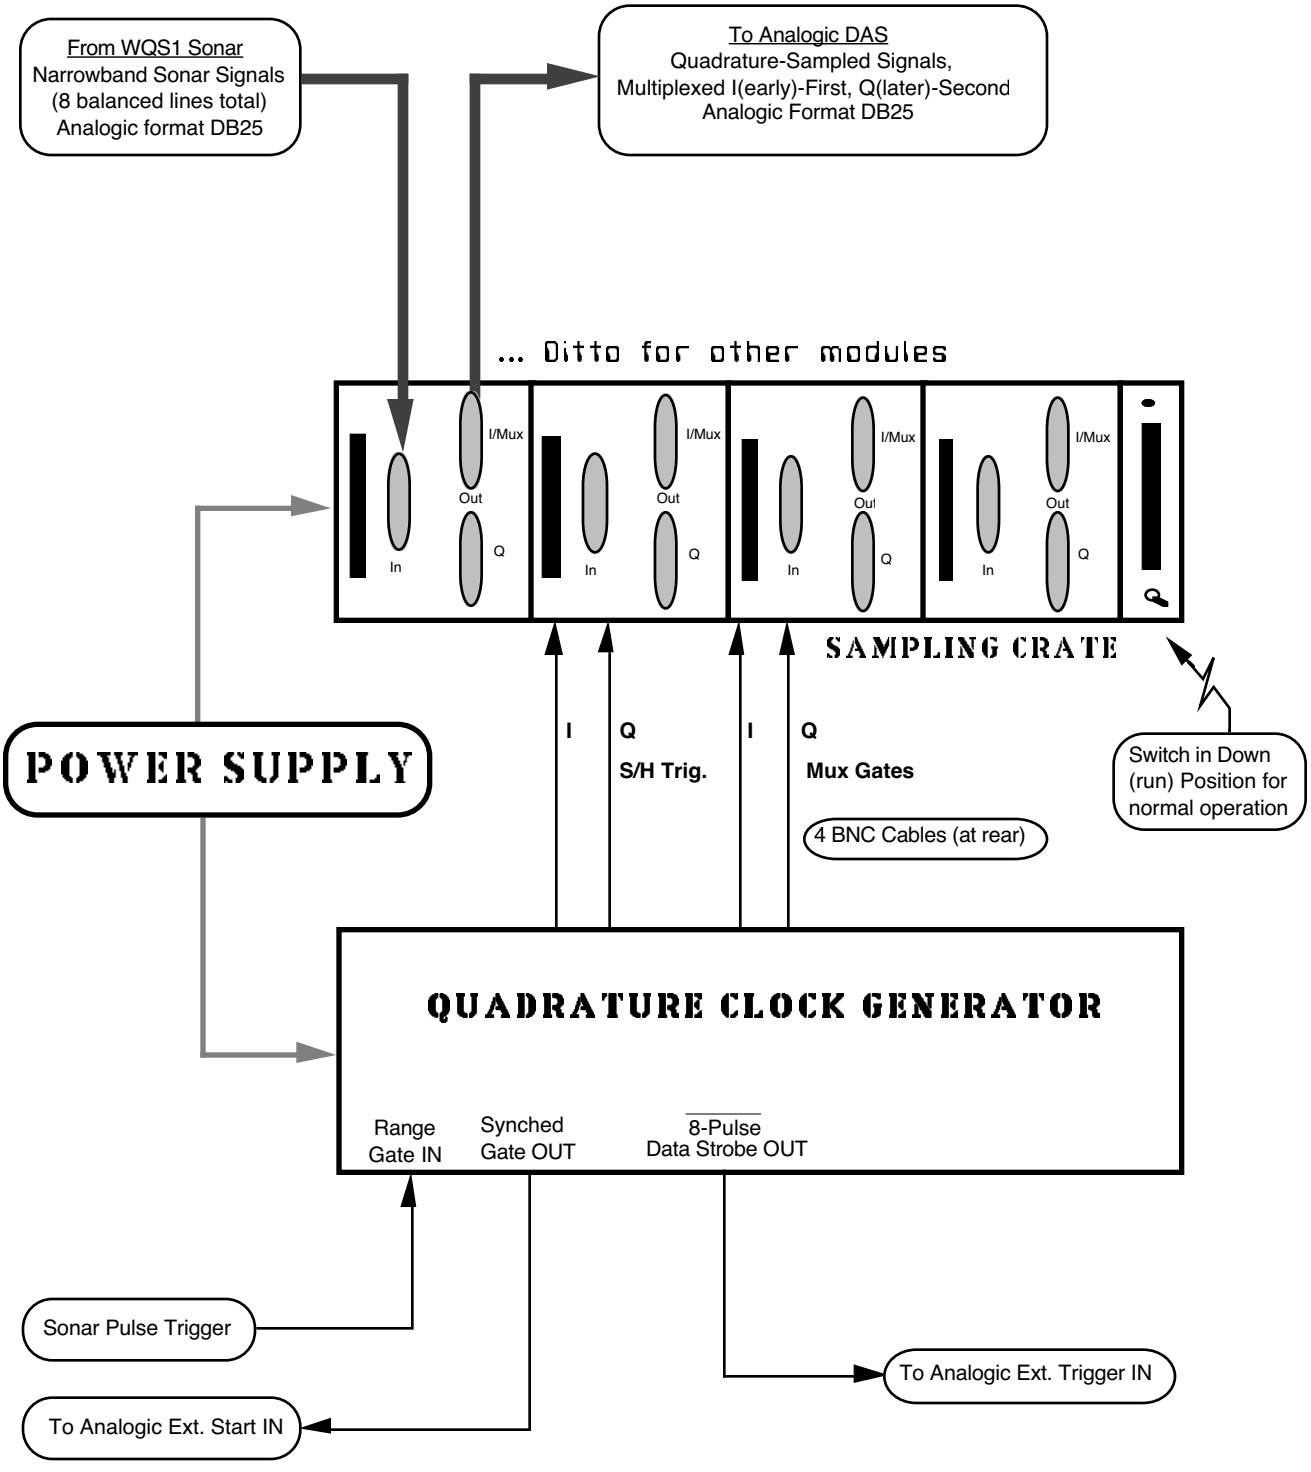

Figure 2: Basic Set-Up for Quadrature Sampler Operation

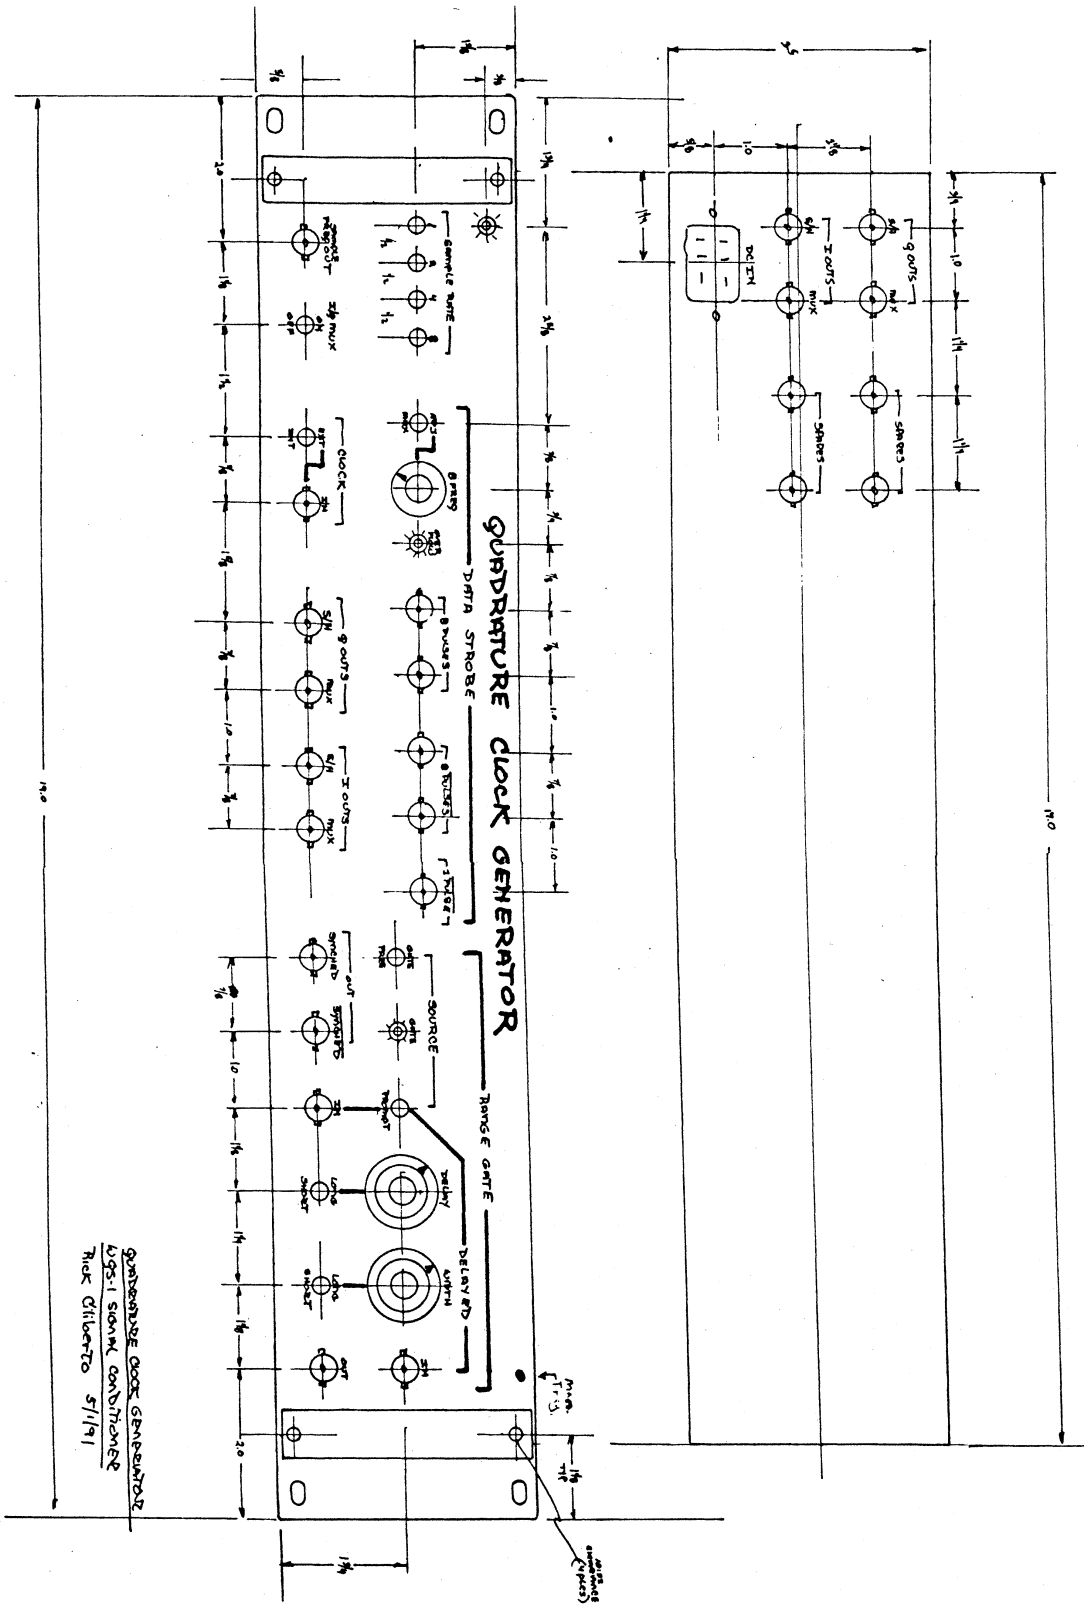

Figure 3: Panel Layout of Quadrature Clock Generator

The clock generator produces the gates that drive the Sampler electronics and strobe data into the DAS in response to a trigger signal received from a data source (i.e. the WQS1 sonar system). It has tremendous capability and flexibility, which I'll delve into below, but, for its essential function, it's connected as shown in Fig. 2. The operational discussion of the Clock Generator will frequently reference its front-panel layout, which is given in Fig. 3.

The Clock Generator is by default locked to a precise internal crystal reference oscillator running at 40 times the assumed center frequency (f_0). An 8 Mhz oscillator is currently installed, which yields an f_0 of 200 Khz. If the "Clock" toggle switch at the lower left of the front panel is set to the "Ext." position (up), an external reference oscillator will be accepted to determine an alternate f_0. This is a standard, unterminated LS TTL input. *The external reference frequency will be at 8 times the desired f_0.* A non-square pulse is admissible, as it is internally divided and squared.

The sampling period Δ is set by the row of four toggle switches in the upper left of the Clock Generator's panel. The highest available sampling frequency is specified with all toggles up, and the lowest sampling frequency is selected with all toggles down. If one reads these toggle switches as a binary nibble "ß", with the LSB at left and a logical "1" with a switch in the down position, the selected sampling factor ℓ introduced in the previous section is given by 4(ß + 1). For a 200 Khz f_0, the selected sampling frequency f_s will thus be (100 Khz)/(ß + 1). These switches thus allow us to set a (ß) between 0 and 15, thereby, with the internal reference oscillator, specifying sampling frequencies between 6.25 Khz and 100 Khz (which are rational ratios of f_0 = 200 Khz). For an f_s of 20 Khz (as assumed in the analysis of Ref. [1]), all toggles are up except for toggle #4 (the third from left). The actual sampling clock running at the selected f_s is made available at the BNC connector at the lower left of the Clock Generator panel; this is useful for frequency monitoring and/or triggering applications. The delay factor K of the second-order sample (see Sev. 1) is set to zero; the I and Q samples are sequentially spaced at $\frac{1}{4}$ of the carrier's period.

The I/Q multiplexing operation (discussed earlier) may be disabled by the "I/Q MUX" switch at lower left. The normal position of this switch is up (in the multiplex position), presenting alternate I and Q samples at the multiplexed sampler outputs for periods of nearly $\Delta/2$. If this switch is down, the multiplexing is disabled, and only the I signal is presented to the "I/mux" sampler outputs (the Q signals are always available separately at their output connector). The data strobe pulses (discussed below) know about the setting of this switch. If multiplexing is enabled, strobe signals are created twice during a sampling period Δ, in order to read both I and Q signals sequentially into the DAS. If multiplexing is disabled, a set of strobe signals are only generated once, near the beginning of the sampling period Δ.

The two Sample/Hold gates and multiplexer gates used by the Sampler Crate are output at the back panel of the Clock Generator, as discussed earlier, and seen in Fig. 3. These signals are also made available as separately buffered TTL signals at the front panel (the "I Outs" and "Q Outs" at lower center) for triggering and diagnostic utility. These signals are unaffected by the setting of

the "I/Q Mux" switch (although the Mux gates presented at the rear panel are). The Mux gates are inactive unless enabled by a valid range gate, as will be discussed below.

The Data Strobe section is located at the top-center of the panel. These outputs produce narrow (≈ 180 nsec wide) pulses that can be used to clock data into the DAS. The rightmost output, labeled "1 $\overline{\text{Pulse}}$ " provides a single negative-going pulse when new data is presented at the "I/Mux" sampler output. This pulse appears somewhat early; it takes the data lines a small interval to settle (especially with a long cable attached to the sampler outputs), thus one should wait roughly 2 → 3 μsec after this pulse is issued before latching in the data. This output is capable of driving a 50Ω termination load. Two additional pairs of strobe outputs are also available, termed "8 Pulses " and "8 $\overline{\text{Pulses}}$ ". These are complementary trains of eight narrow pulses (again, ≈ 180 nsec wide) that are issued whenever new data is presented at the "I/Mux" sampler output. These outputs have been explicitly added to this circuit to drive the Analogic DAS, which requires an external clock to advance the input multiplexer (thus needs 8 pulses to advance through all 8 channels). Although the circuit is hardwired for 8 pulses, the pulse quantity can be changed, if desired (i.e. if one uses another digitizer that has a different number of channels and likewise requires multiple clocking pulses) by a simple circuit modification to pick a different tap off an internal counter, as will be discussed in the next section on circuit details. The 8-pulse train is delayed slightly relative to the S/H gates, thus there is no need to insert any additional wait before accepting data (as with the 1-pulse); the digitizer can be directly clocked with this signal. These outputs are also capable of driving a 50Ω termination load. The Analogic DAS is configured to accept the complimented "8 $\overline{\text{Pulses}}$ " signal as its "Ext. Trigger". The "8 Pulses " output is provided for the sake of eventual flexibility. All of the data strobes are 0→5 Volt TTL-compatible signals.

If the switch at the left of the Data Strobe section is down ("max"), the spacing between pulses on the "8" outputs will be fixed at 2.5 μsec (locked to the internal crystal oscillator). This clocks the Analogic DAS at 400 Khz, which is just below its maximum quoted rate of 409.6 Khz. If the switch is up ("Adj."), the spacing between pulses is set by the "8 Freq" potentiometer. This controls the frequency of a synchronized internal oscillator that allows the pulse spacing to range from slightly over 1 μsec down to 10's or 100's of msec. If the pulse spacing is too wide, and a pulse train runs into its successor, the "Overrun" LED is illuminated. In this case, one must back off on the pulse spacing, run at a lower sampling rate, or not use the I/Q multiplexing (whichever is most easily implementable). This LED is driven with a timer that produces roughly a $\frac{1}{4}$ second delay before turning the LED off, thus the response to beneficial action is not entirely immediate.

The rightmost section of the Clock Generator panel is dedicated to the Range Gate circuitry. The toggling of the multiplexers in the Sampler Crate and the production of Data Strobe signals require a range gate signal to be present. If no range gate exists, the Sampler Crate multiplexers are turned off (thus no signal is present at the I/Mux outputs), and no Data Strobes are produced.

The status of the range gate is given by the "Gate" LED. If this LED is illuminated, a valid range gate is present. The range gate may be forced to be always active by placing the leftmost

toggle switch under the "Source" heading in the down ("Free") position. This option is provided for test purposes, when one desires the system to be in a free-running state. If this toggle is up ("Gate"), the Range Gate is derived from the source specified by the rightmost "Source" switch. If this "Source" switch is down, the Range Gate comes directly from the "Prompt" input, located immediately below. This is a standard LS-TTL input, with the active gate at logic high.

If the right "Source" switch is up ("Delayed"), the range gate comes from the "Delayed" input, located at the right of the panel. A negative-going edge applied to this input will trigger a gate pulse that is of adjustable width and arbitrarily delayed (relative to the input trigger). The amount of delay and the gate pulse width are set by a pair of 10-turn potentiometers. The range of these potentiometers is set by toggle switches placed immediately beneath them. The gate pulse width can be adjusted from 0.4 msec → 55 msec ("short" range) or 14 msec → 1.2 sec ("long" range). The delay can be adjusted from 0.5 msec → 58 msec ("short" range) or 11.5 msec → 1.2 sec ("long" range). The delayed pulse is available at the output at lower right of the panel; this is a standard TTL signal provided for test, calibration, and/or external device triggering. It is always present, independent of the gate "Source" switch settings. In sonar operation, one triggers the delayed gate when the outgoing "ping" is sent. The delay is adjusted to match the transit time to-and-from the estimated target, and the pulse width is set to the interval over which data is to be taken. A pushbutton is included at the upper right of the panel that will trigger the Delayed Gate for test purposes.

The edges of the selected range gate are synchronized with the sampling frequency f_s. This "synchronized" range gate thus goes high before the first set of "I" signals are presented to the the Sampler outputs (and the first Data Strobes are issued), and goes low after the last "Q" signals are presented to the Sampler outputs (and the last Data Strobes have all been issued). Thus the synchronized gates "enclose" the data, and ensure the I-first, Q-second presentation order (beware of my definition of I and Q; see the notes at the close of the first section!). These synchronized gates are made available at the outputs at the lower left of the "Range Gate" section of the panel. They are a complimentary pair labeled "Synched" and "$\overline{\text{Synched}}$", and are 0→5 Volt TTL-compatable. The "$\overline{\text{Synched}}$" output can be used by the Analogic DAS as an "Ext. Start" gate. It is able to drive a load terminated at 50 Ω. The first of the "8 $\overline{\text{Pulses}}$" signals is sent after 4 μsec elapses following the falling edge of "$\overline{\text{Synched}}$", which is more than ample for the Analogic lower limit of 600 nsec.

The Clock Generator and Sampling Crate are powered by another rack mount unit that provides +5 Volt and ±12 Volts. The only control here is the power switch, which is self-explanatory. All Voltages are available for testing & utility at the front banana jacks. The grounds on the +5 and ±12 supplies float when the Clock Generator or Sampling Crate are unconnected.

3) Circuit Details

The circuitry in this system was conceptually laid out in the previous memo[1]. The actual constructed circuitry essentially implements this overview fairly closely, thus Ref. [1] can be consulted for a broader overview with less cluttered drawings (much of the text below was adapted from this source, so if you've already read Ref. [1], this may all sound familiar). The purpose of this section is to present the detailed schematics, together with a fast running dialog of the essential concepts employed. There was insufficient time to load the final schematics into a CAD package for nice printout, so pardon my hand-drawn versions included here. Suffice to say, everything seems to work fairly well....

Fig. 4 shows the Clock Generator schematic. In order to guarantee a stable, precise center frequency, all timing originates from a crystal-locked clock. The master frequency of 8 Mhz was determined in order to enable easy generation of the 800 Khz 4'th harmonic of the desired 200 Khz center frequency. This is accomplished by a division by 5, followed by a division by 2 (to square the pulse shape); both of these divisions are accomplished in a single chip (i.e. the 74LS90). An external clock may optionally be input at this point in the circuit if the sampler is to reference other (adjustable) center frequencies; the clock signal is Schmidt-conditioned, and the ÷2 stage will square this signal as well. To simplify the discussion below, the quoted frequencies and periods assume that an f_0 of 200 Khz has been specified.

The 800 Khz is then divided by 4 in a 74LS93, to yield the 200 Khz fundamental (the f_0). This signal is then additionally divided by a presettable integer (the "sampling rate" setting) via a ÷N counter made up of a 74LS161 and 74LS157 (which passes the clock through when all ÷ switches are off, to yield the ÷1). The resulting signal is again divided by two (using the 74LS93) to regain a square waveform. This signal is essentially our sampling clock running at f_s. A buffered output of f_s is provided for testing and other utility.

The falling edge of the sampling pulse will clock a true state into the lead flip-flop of package C5 (74LS74). This then will be clocked into the other flip-flip of this package one-quarter of the carrier period (1.25 μsec @ f_0 = 200 Khz) later by the next falling edge of the 800 Khz square wave, whereupon the first flip-flop is reset until the next sampling period starts. After another quarter carrier (1.25 μsec) elapses, a zero is clocked into the second flip-flop, which also remains low until the next sampling period starts. Pulling this argument together, this pair of flip-flops will deliver consecutive 1.25 μsec pulses at the beginning of each sampling period. The trailing edges of these pulses are spaced exactly 1.25 μsec apart (barring the <30 nsec. propagation delay from the reset input of the lead flip-flop, which shouldn't be terribly significant). They are output as the I and Q S/H gates, and drive the sample/holds in the Sampler Crate.

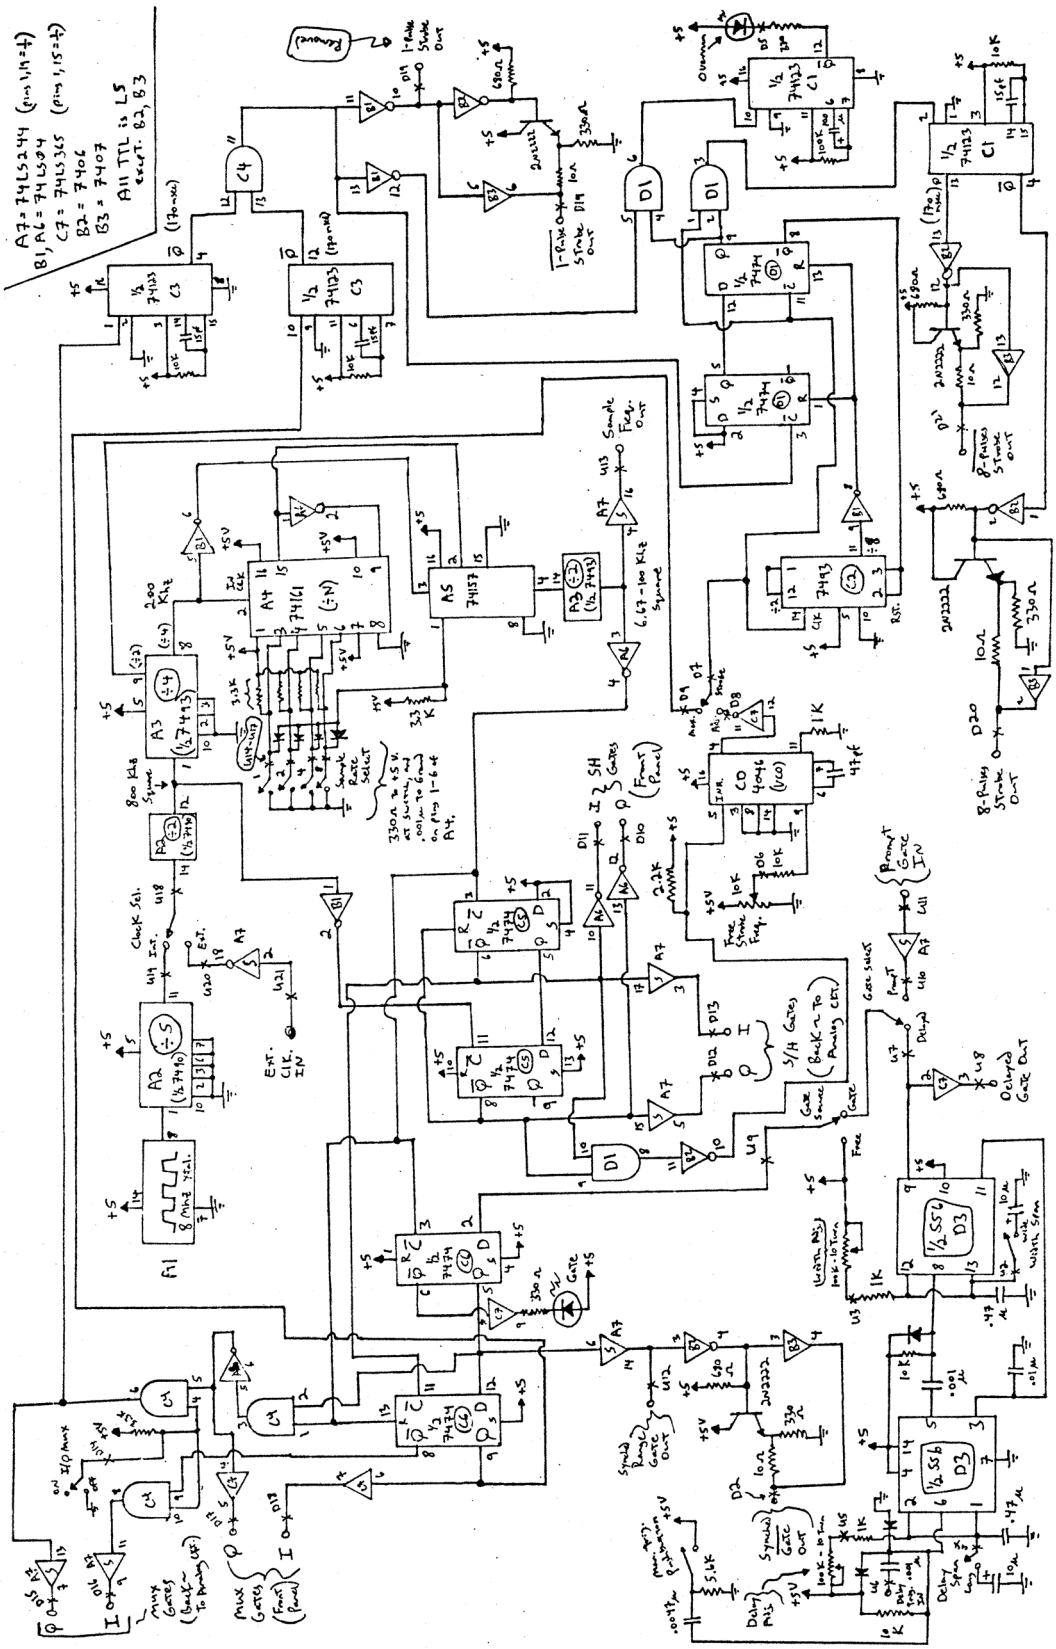

Figure 4: Quadrature Clock Generator Schematic

Samples are presented to the I/Mux Sampler outputs during the receipt of a valid "range gate" signal. An input gate is synchronized to our clock via the lead flip-flop in package C6, which is allowed to go high or low after the falling edge of the sampling signal (at f_s). When it is high, the output of the following flip-flop will go high upon the next falling edge of the I S/H gate; by this time the in-phase signal will be locked into the sample-hold. This signal is thus used as a "valid" multiplexer gate for the I data, and switches the Sampling Crate multiplexers such that they present the sampled I signals at their outputs. One of the monostables in the 74LS123 at the upper right of Fig. 4 is also triggered upon the rising edge of this I gate, which produces a narrow (≈ 180 nsec) "data trigger" pulse that alerts an external digitizer unit to record the multiplexer outputs, and thus capture the I signals. The output of this flip-flop is reset one-half of the sampling period later by the rising edge of the sampling clock (f_s). The sampling clock itself (gated by the synchronized range gate) then provides the Q multiplexer gate, which switches the Sampling Crate multiplexers to provide the Q signals at their outputs and triggers the other monostable in the 74LS123 to produce another narrow data trigger to alert the external digitizer to record the Q signals. In this fashion, the data is organized such that the I samples are presented first (upon the receipt of a range gate) and remain for slightly less than half of the f_s clock (24 μsec at f_s = 20 Khz), followed by the Q samples (which remain for half of the f_s clock; i.e. 25 μsec), then the next I samples, etc. These multiplexer gates are independently buffered and appear at both the front panel (for general utility) and the rear panel (for connection to the Sampler Crate). The rear-panel gates are disabled by the "I/Q Mux" switch, which will always select the I signal when turned "Off".

The outputs of the 74LS123 discussed above are OR'ed together and buffered by a push-pull 50 Ω driver before being made available to the outside world as the "1-Pulse" strobe. This pulse will also clock a logic "1" into the lead flip-flop of package D1. This, in-turn, clocks a "1" into the other flip-flop in D1 upon the next rising edge of the 8-Pulse clock source (as set by the adj/max switch). This enables the 74LS93 counter in C2 to begin advancing, and allows this clock source to trigger one of the monostables in the 74LS123 of C1, which will produce narrow (≈ 180 nsec) pulses on each rising edge. These pulses are generated until the appropriate stage of the C2 counter is asserted (currently set to the ÷8), at which point the flip-flops in D1 are reset (which also resets the C2 counter). This prevents further clock pulses from triggering the monostable in C1 and producing more strobe outputs until another trigger pulse arrives from the 1-Pulse monostable in C3. If less than 8 pulses are desired to be produced (i.e. to interface to a DAS with fewer than 8 channels), the tap off the C2 counter can be correspondingly changed (from its present location at pin 11). The normal and complimented outputs of the C1 monostable are 50 Ω buffered, and made available to the external world as the 8-Pulse strobes. If another 1-Pulse trigger arrives while the D1 flip-flops are high, the other monostable in C1 is triggered, which produces roughly a $\frac{1}{4}$ second pulse to illuminate the "Overrun" LED.

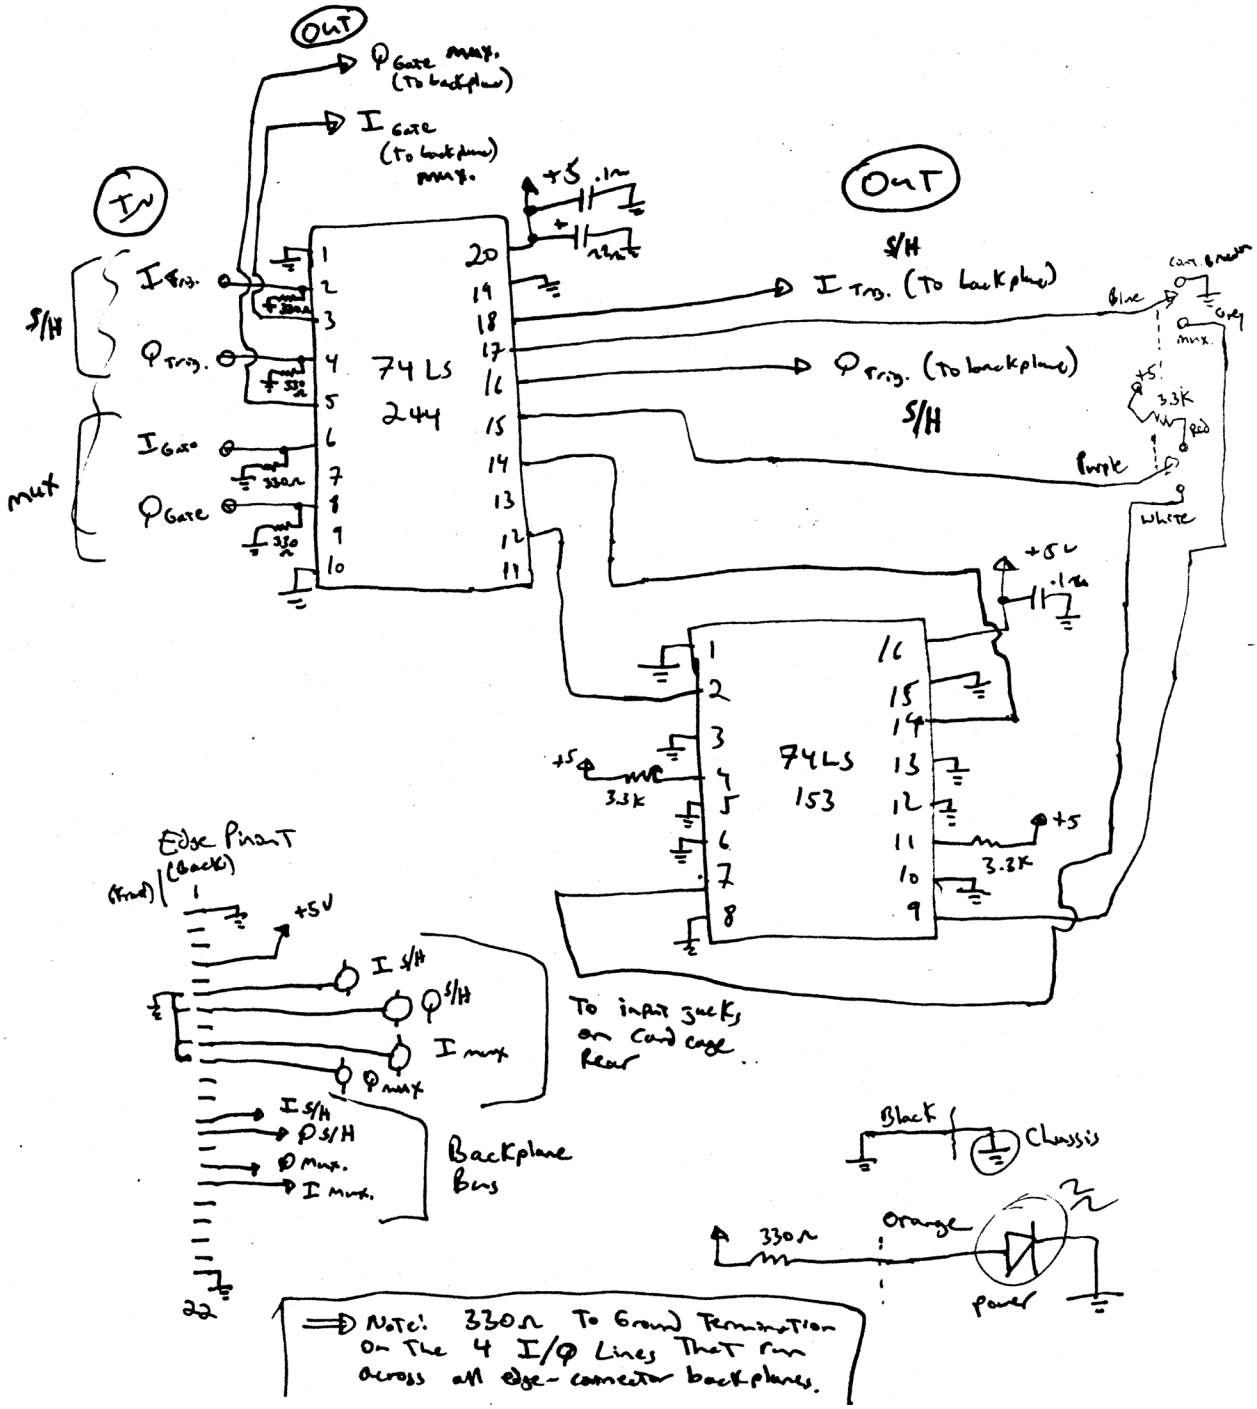

Figure 5: Sampling Crate Clock Driver Schematic

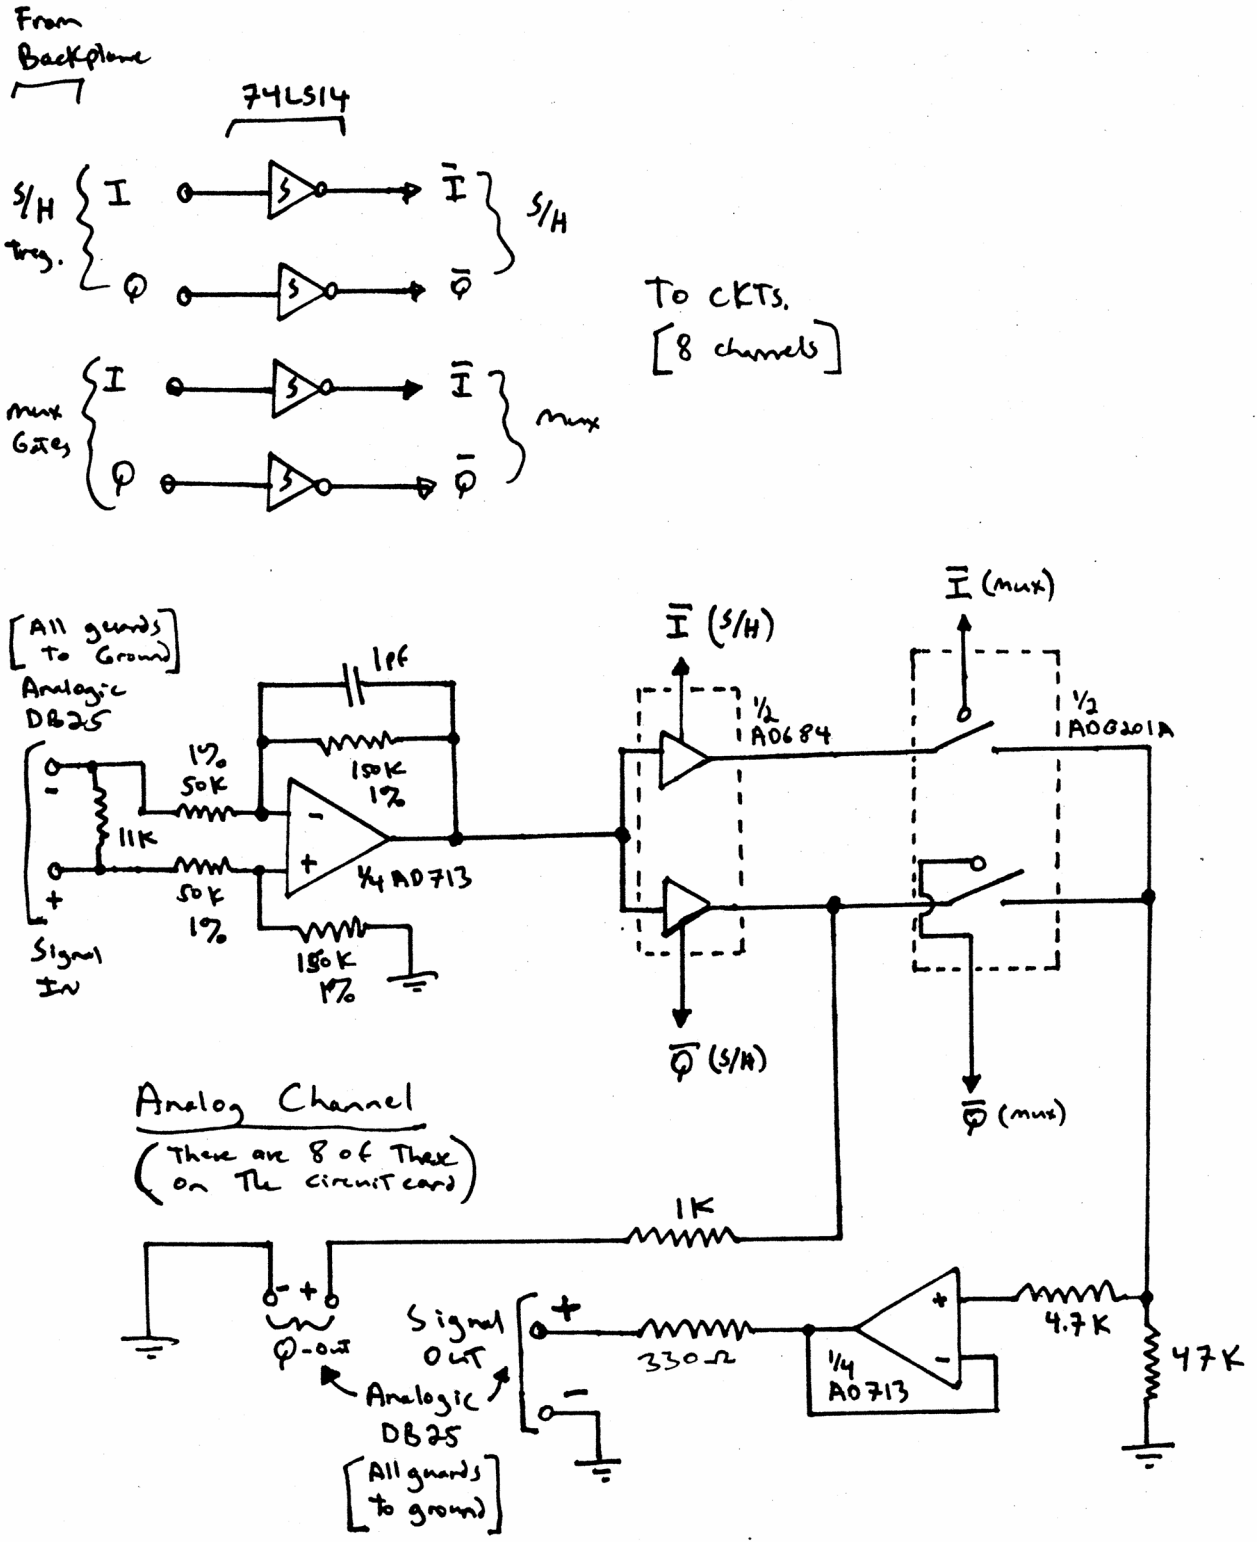

Figure 6: Sampler Module Schematic

The clocking source for the "8-Pulses" strobe can be selected to come from either a 400 Khz reference derived from the master clock chain (extracted from the 74LS93 in A3), or a user-adjusted VCO based on a CMOS 4046 PLL chip. This VCO is synchronized to the sampling clock by its "Inhibit" input, which is driven by an "OR" of the I and Q S/H gates. This locks the strobe oscillator to the sampling clock, and prevents strobe pulses from being issued until the I and Q signals have been sampled and transients have died away.

The input range gate can be derived directly from a Schmidt-conditioned TTL input, or the dual timer chain (two cascaded monostables in a 556 package), which implements the delayed gate. It's synchronized to the sampling clock in the flip-flop of C6, as discussed above.

The Clock Driver circuit that resides in the right-hand edge of the Sampler Crate is given in Fig. 5. This is a very simple circuit. I apologize for the untidy hand-drawing, but the circuit is so basic that I feel that my poor lab-note original will provide sufficient detail, and a re-draw is unnecessary. The I and Q S/H gates are buffered by Schmidt triggers in the 74LS244, and directly applied to the backplane. The I and Q multiplex gates are first Schmidt conditioned by the 74LS244, then applied to the 74LS153 selector, which prevents the "1,1" state from appearing (this will turn both I and Q simultaneously on in all the sampler multiplexers, shorting the sample-hold outputs together). This state is forbidden, and is converted to a "0,0" state (basically an exclusive-OR operation), which turns all multiplexers safely off. The outputs of the 74LS153 are buffered by the 74LS244, which then drives the crate's backplane. The I/Q multiplexer operation can also be defeated by the toggle switch on this unit, which will cause the I signals to be continuously asserted at the multiplexer outputs. The four input lines from the Clock Generator are terminated by 330 Ω loads. The backplane lines are also terminated to ground by 330 Ω loads at the edge of the crate opposite to the Clock Driver slot.

The Sampler schematics are given in Fig. 6. This simple diagram essentially represents all of the analog circuitry in this device. It's nearly a trivial circuit, with a differential-input front-end (set to a gain of 3) driving a pair of sample-holds. These are selected by a dual JFET switch, which chooses either the I or Q sampled signals. This result is buffered by a voltage follower, and brought to the "I/Mux" output port. The sampled Q signal is also available directly at its own output port. These signals are buffered by 330 Ω and 1K Ω resistors (respectively) before being output, in order to decouple the cable loads and avoid difficulty with oscillation. The input amplifier is compensated with the 1 pf feedback to minimize overshoot and flatten the high-frequency response (which is nearly flat out to roughly 350 Khz). The inputs are loosely terminated with an 11K Ω resistor, which doesn't afford too much common-mode rejection for high-impedance induced signals (but it's perhaps better than nothing....). This value couldn't be reduced further due to load limitations of the WQS1 output electronics, and there is no room on the Sampler circuit card for additional voltage followers.

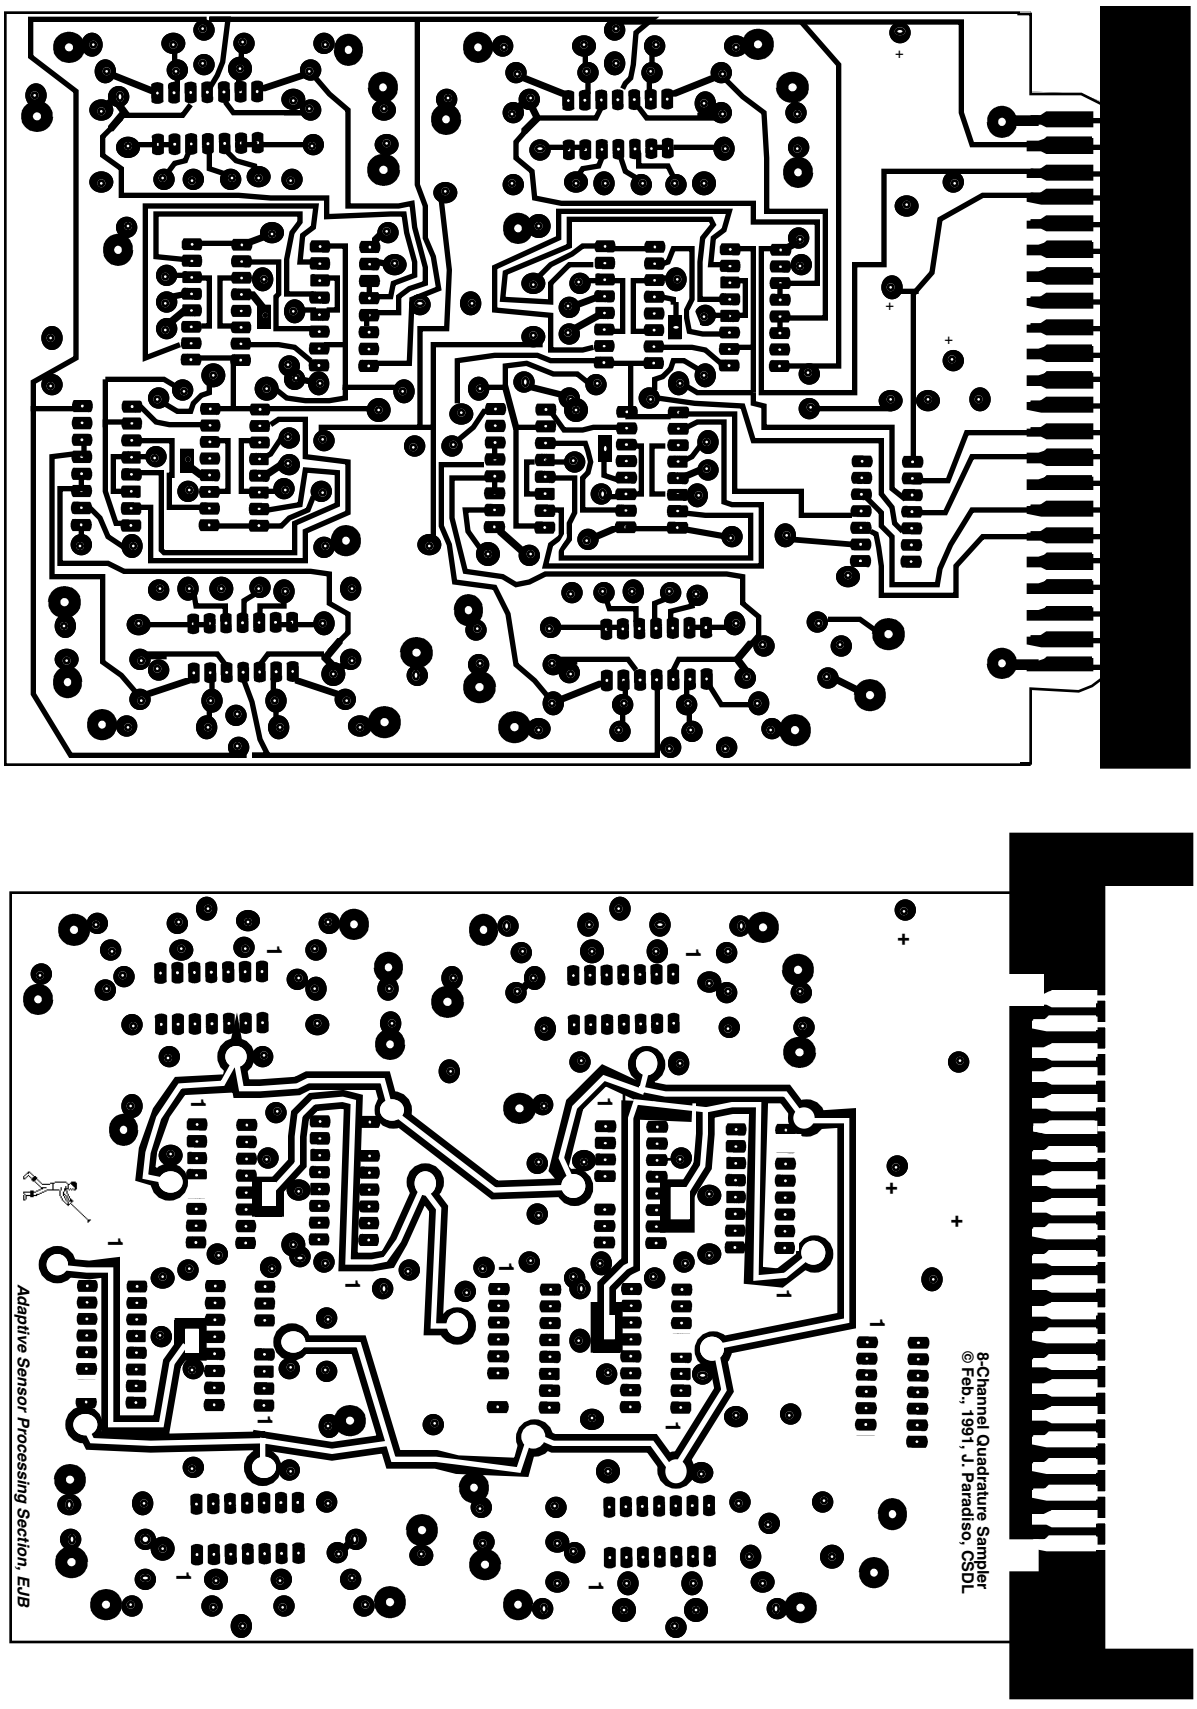

Figure 7: Circuit Layout for Sampler Module

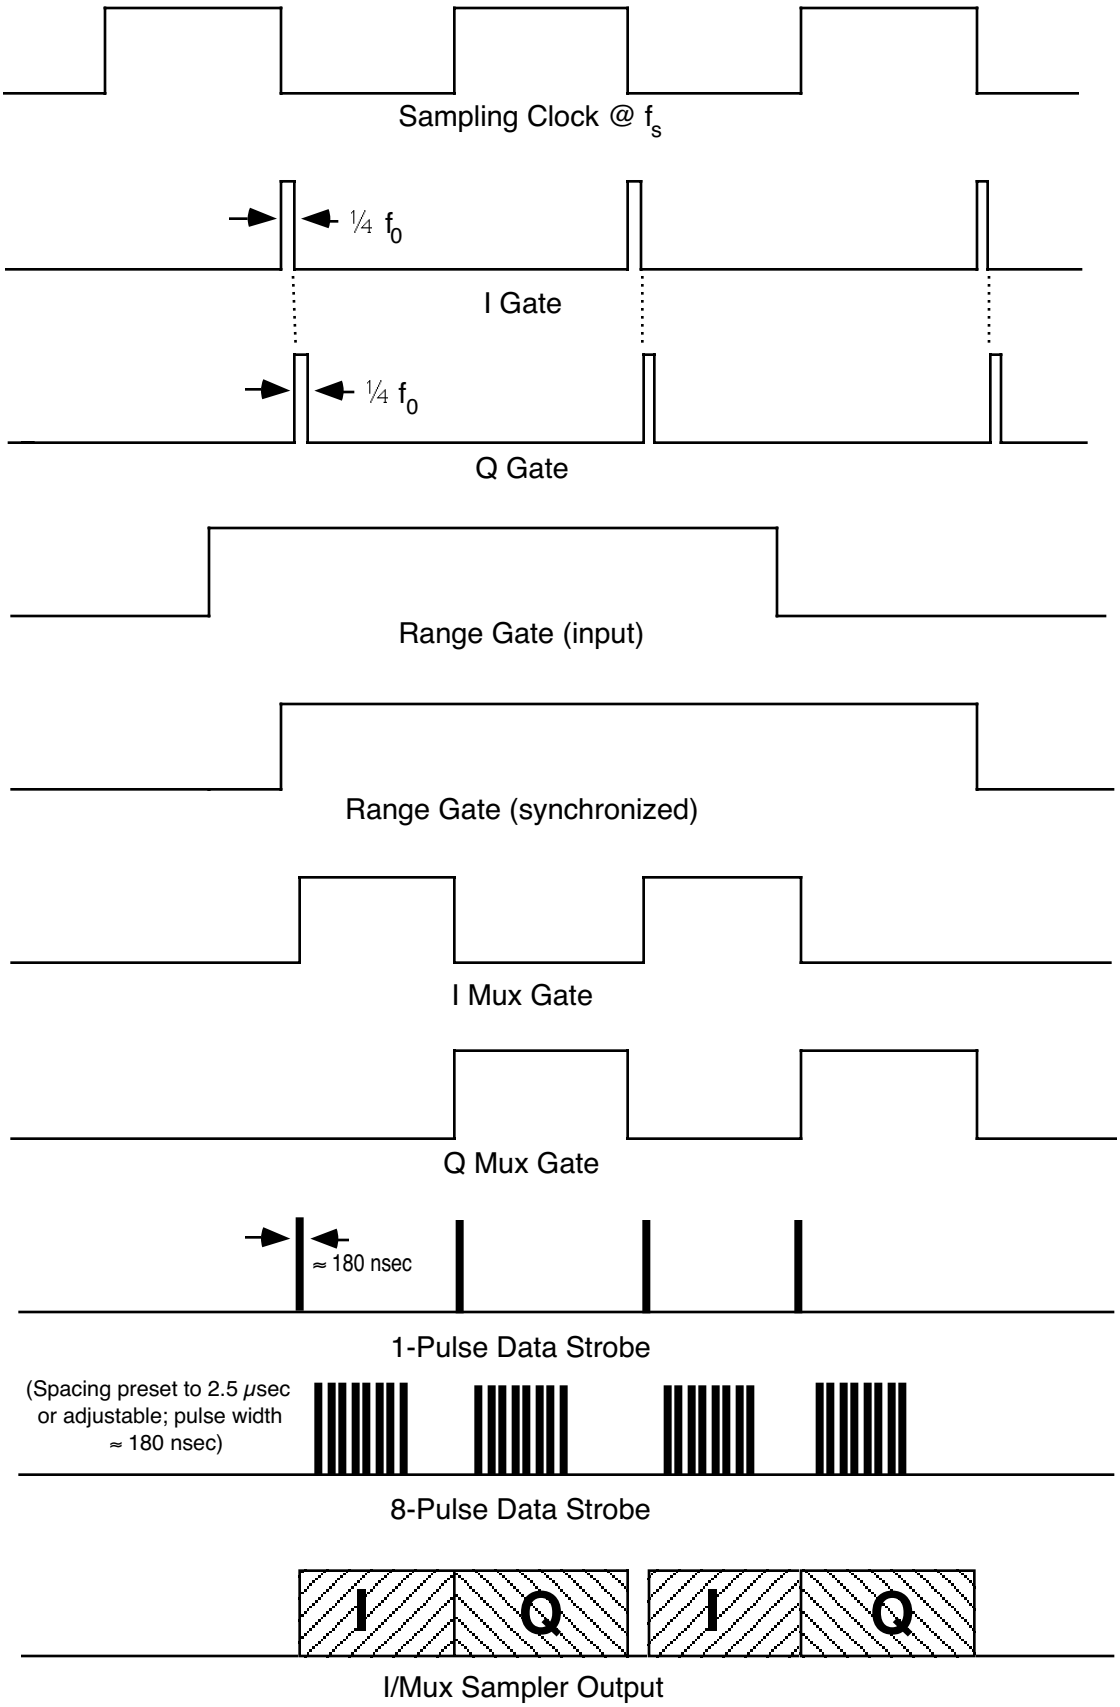

Figure 8: Timing Diagram for Quadrature Sampler System

The analog electronics shown at the bottom of Fig. 6 are replicated 8 times to form the 8 channels resident on a sampler card. A pair of channels are designed around a single chip set, using the AD713 quad OP-AMP, the AD684 quad sample/hold, and the ADG201A quad analog switch. The 2-layer circuit card layout is shown in Fig. 7, along with the channel designations. The card was quickly done up in MacDraw II; one chip-set was laid out for a single channel pair, and duplicated four times to fill a single vectorboard-sized card. Interconnections for clocking and supply lines were then added, plus a 74LS14 Schmidt trigger buffer to pull the 4 I/Q signals off the backplane and fan them out to all of the 8 channels (as depicted at the top of Fig. 6).

The AD684 Sample/Holds are low-cost quad monolithic devices that boast an acquisition time of well under 1 μsec and are seen to adequately take their samples within the 1.25 μsec S/H gate width (assuming f_0 = 200 Khz).

A timing diagram is drawn in Fig. 8, showing the interrelation between the various clocks, gates, and signals generated by the circuitry. The shaded blocks at the bottom of Fig. 8 represent the analog I/Mux output of the sampler modules

Actual scope photographs of the generated waveforms are given in Fig. 9. All traces are triggered on the falling edge of the sampling clock. The upper photograph shows an entire sampling period. The lower two traces show the I and Q S/H gates; the signal is acquired while the gates are high, at the left edge of the figure. The second trace down shows the $1\overline{\text{Pulse}}$ data strobe. It is asserted as soon as the multiplexer outputs change. The top trace shows the $8\overline{\text{Pulses}}$ strobe. Two groups of 8 pulses are seen to be produced, each after the multiplexer switches I and Q outputs. If the multiplexing were defeated by turning the "I/Q Mux" switch off, only the first $1\overline{\text{Pulse}}$ strobe and first group of $8\overline{\text{Pulses}}$ would be present.

The lower photograph in Fig. 9 shows a 10 X expanded view of the initial sampling cycle that will follow the receipt of a valid range gate. The top trace is the complemented synchronized range gate signal $\overline{\text{Synched}}$, which was used to trigger all traces here. Immediately upon its falling edge, the I gate is seen to pulse, followed by the Q gate 1.25 μsec later. The first of the $8\overline{\text{Pulses}}$ strobes is visible at right in the bottom trace, ocurring 4 μsec after the falling edge of $\overline{\text{Synched}}$.

The upper photograph in Fig. 10 shows the relation of the data strobe pulses to the I/Mux output signal. Again, we trigger on the falling edge of the sampling clock. The top trace shows the I/Mux output (a sine wave is input to the sampler at roughly 201 Khz; this photograph is exposed over many traces). The initial transient at left is due to multiplexer and Sample/Hold switching. Within 2→3 μs or so, the signal settles out to a steady-state value. A small glitch is present midway down the trace, where the I/Mux output switches to the Q signal (the Q multiplexer gate is also shown in the bottom trace).

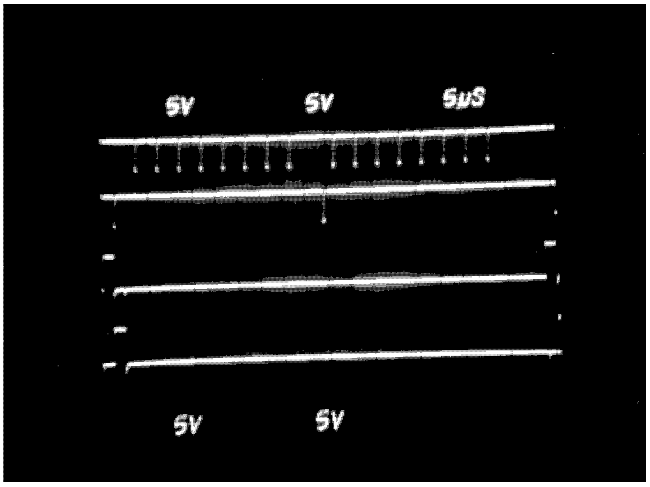

a) Clock Generator Signals; full Sampling Period

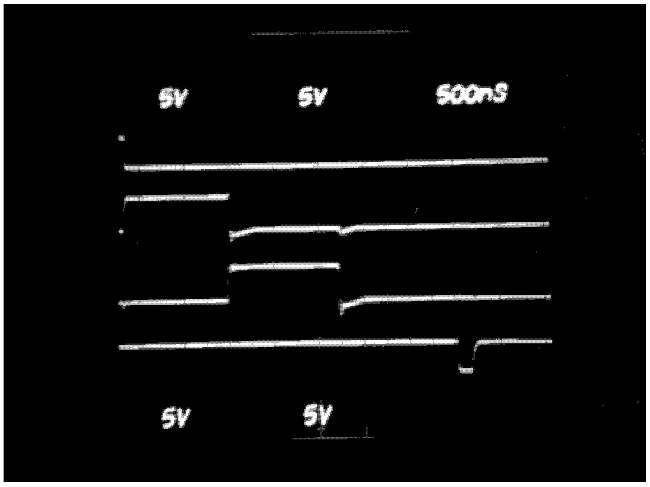

b) Clock Generator Signals; 10 X Expanded View

Figure 9: Clock Generator Signals

The 8 $\overline{\text{Pulses}}$ data strobe is shown in the second trace down. The I/Mux signal is seen to have well settled by the time the 8 $\overline{\text{Pulses}}$ strobes clock in the data, for both the I (first) and Q (second) components. The third trace down shows the 1 $\overline{\text{Pulse}}$ data strobe. As discussed previously, it does arrive somewhat early. The I/Mux signals take an additional 2→3 μsec to settle, particularly after the I signal is presented at the start of the trace. Of course, this is presently of little consequence, since the Analogic DAS uses the 8 $\overline{\text{Pulses}}$ strobes, but it should be kept in mind if the 1 $\overline{\text{Pulse}}$ strobe is used in future applications.

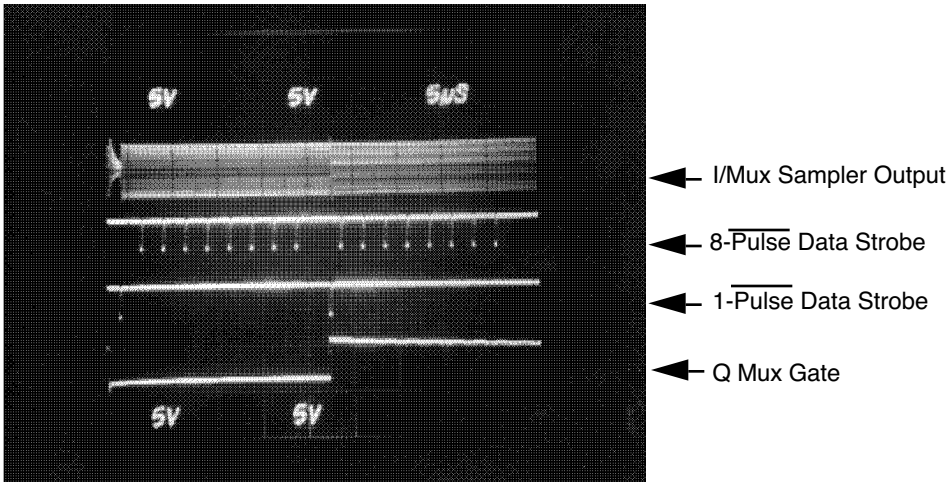

a) Data Strobes vs. I/Mux Output; full Sampling Period

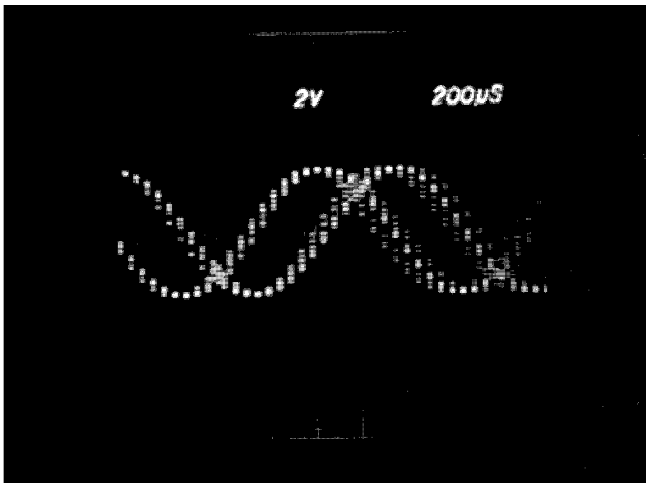

b) Interleaved I/Mux Output for input sinusoid near 200 Khz

Figure 10: Data Strobe Timing and Multiplexed Sampler Output Signals

The lower photograph of Fig. 10 shows the I/Mux signal produced in response to an input sinusoid quite close to the 200 Khz f_0. The trigger is derived from the displayed waveform (thus its a bit jittery), and the sweep time is quite expanded relative to the sampling period (which can be seen by the flat regions, which are of length $\Delta/2$). This shows a quadrature "beat", where the phasor determined by the complex envelope of the input signal (relative to the carrier f_0) undergoes a slow rotation in the complex plane. The expected sine/cosine nature of the interleaved I/Q samples (alternating through the toggling multiplexer) can be readily seen in this photograph.

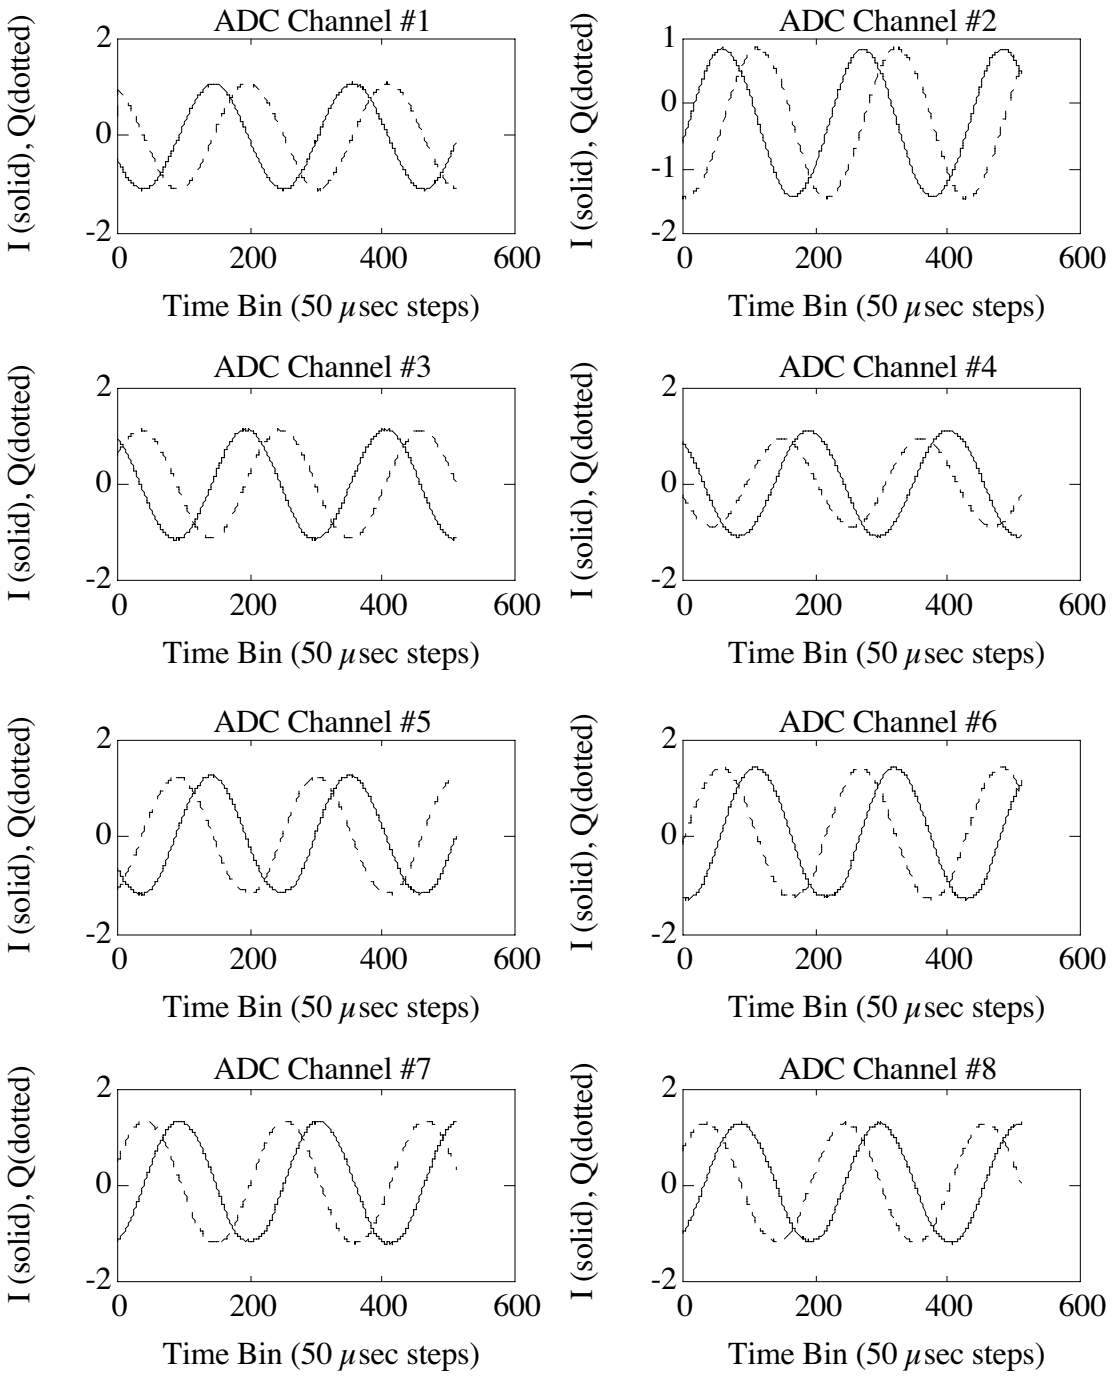

Figure 11: DAS Readout of 8 Sampler Channels for Input Sinusoid Close to 200 Khz

As a final test, all 8 channels of a Sampler module were driven by such a sine wave (near 200 Khz), and clocked into an Analogic DAS, as diagrammed in Fig. 2. The resulting I/Q waveforms were de-interleaved and plotted in Fig. 11 for all 8 channels. One can clearly see the sine/cosine quadrature pair undergoing complex rotation as the input signal beats with the carrier. The gain shift and offset change of the two plots at upper right were either caused by an out-of-spec Sample/Hold chip or DAS readout error (this plot resulted from the initial test of the system with the DAS, which

is still being debugged). Any effects of this sort due to the sample/holds should be quite static, and can be readily compensated by taking a calibration of this sort before using this device for data collection (of course, any IC's appreciably out of spec can easily be isolated and replaced).

The Clock Generator circuitry presented above was essentially wired with discrete (TTL) logic. At the time, this was the deemed fastest way to build, prototype, and modify the evolving design. The chip count could potentially be reduced by going to programmable logic devices (PLD's, PAL's, etc.), and avoiding the TTL. In any future (potentially deployed) implementation, a simplified design could be adopted, eliminating most of the user-friendly options and signal pick-offs (which lend considerable complexity to this design), and much of the discrete TTL could probably be replaced by programmable logic devices. Most of the circuitry would then be dedicated the analog front-end, where duplication of a very simple design is essentially the name of the game. In a deployed system, the sampling front-end would certainly hand-shake with an entirely different digitization and data acqusition system, which will probably entail alternate timing and interface requirements.

4) References

1) Paradiso, J., "Sonar Application of a Quadrature-Sampled Interpolation Beamformer," EJB 90-191, Dec. 31, 1990.

2) Grace, O.D. and Pitt, S.P., "Sampling and Interpolation of Bandlimited Signals by Quadrature Methods, " Journ. of the Acoustical Society of America, Vol. 48, No. 6 (Part 1), Nov. 3, 1969.