JPL-33-CR
169388
53P

Functional Description of Signal Processing in the Rogue GPS Receiver

J. B. Thomas

June 1, 1988

NASA

Functional Description of Signal Processing in the Rogue GPS Receiver

J. B. Thomas

June 1, 1988

ABSTRACT

Over the past year, two Rogue GPS prototype receivers have been assembled and successfully subjected to a variety of laboratory and field tests. This report presents a functional description of signal processing in the Rogue receiver, tracing the signal from RF input to the output values of group delay, phase, and data bits. The receiver can track up to eight satellites, without time multiplexing among satellites or channels, simultaneously measuring both group delay and phase for each of three channels (L1-C/A, L1-P, L2-P). The Rogue signal processing described in this report requires generation of the code for all three channels. Receiver functional design, which emphasized accuracy, reliability, flexibility, and dynamic capability, is summarized. A detailed functional description of signal processing is presented, including C/A-channel and P-channel processing, carrier-aided averaging of group delays, checks for cycle slips, acquisition, and distinctive features.

CONTENTS

ACKNOWLEDGMENTS

PRECEDING PAGE BLANK NOT FILMED

SECTION 1

INTRODUCTION

Over the past year, two Rogue GPS prototype receivers have been assembled and successfully subjected to a variety of laboratory and field tests. This report presents a functional description of signal processing in the Rogue receiver, tracing the signal from RF input to the output values of group delay, phase, and data bits. The receiver can track up to eight satellites, without time multiplexing among satellites or channels, simultaneously measuring both group delay and phase for each of three channels (L1-C/A, L1-P, L2-P). The Rogue signal processing described in this report requires generation of the code for all three channels.

As explained in Section 2, the signal-processing design emphasized accuracy, reliability, flexibility, and dynamic capability. Section 3 presents a high-level description of the signal processing steps that evolved from these design goals. Subsequent sections provide a more detailed functional description of signal processing, including C/A-channel processing, P-channel processing, carrier-aided averaging of group delays, checks for cycle slips, acquisition, and distinctive features. Among the distinctive features summarized in Section 7 are purely digital operations in the baseband processor, fast feedback loops driven by L1-C/A phase, and high-accuracy measurement of phase and delay.

SECTION 2

FUNCTIONAL DESIGN

2.1 DESIGN GOALS

Design of signal processing in the Rogue receiver was based on a number of goals. First, the receiver should be able to track up to eight satellites, simultaneously measuring both group delay and phase delay from each of the three channels (L1-C/A, L1-P, L2-P). For experimental purposes, the capability of simultaneously tracking L3-C/A along with these standard channels should also be an option. The output interval for these observables should be user-selectable over a wide range (e.g., 1 to 300 s), with time tags placed on integer seconds of Global Positioning System (GPS) time. Data bits should be extracted and interpreted for each satellite.

Point positioning solutions should be carried out by the receiver to measure receiver location and velocity and to synchronize the receiver clock in offset and rate with GPS time. At a minimum, such synchronization should be carried out once at the beginning of a session of data taking, and should set the receiver clock with an accuracy of 0.1 to 3 μs.

The phase lock loops should be very reliable, with very infrequent cycle slips, and they should be able to track phase dynamics up to 300 Hz/s (6 g). Tests should be made at frequent intervals to determine if a half-cycle slip has occurred, and, if so, the track should be corrected or terminated. At the start of the track for each satellite, the half-cycle ambiguity in carrier phase should be removed through inspection of the sign of the synchronization pattern in the data bits.

Accuracy of delay and phase should be a primary consideration in the design of each signal processing stage. System-noise errors should be 1 cm or less for the P group delays and less than 0.01 cm for phase delays, given nominal satellite-signal power, a 5-min integration time, a 125-K system temperature, and an antenna gain of 3 dB. Intersatellite systematic errors, the errors remaining after subtracting the delays simultaneously measured for two satellites, should be less than 1 cm for the P group delays and 0.1 cm for the phase delays, excluding multipath and antenna errors.

An important design goal was to make the receiver highly digital in order to gain the well-known advantages of a digital system: compactness, reliability, accuracy, and flexibility. In digital components, all operations should be purely digital in the sense that the same input data (i.e., sampled bits) will result in the same output observables, to the last decimal place, each time that input data are processed. Digital errors due to roundoff and discrete sampling should be made negligible.

2.2 DESIGN DECISIONS

Due to the limits of digital technology available during design, signal processing from RF to baseband (the "front end") is carried out by analog components, while all subsequent baseband processing is based on digital operations. Front-end operations were made as simple as possible and involve only a single down conversion from RF to baseband for each L-band channel. Simplicity in the front end led to simplicity in the frequency system, which is only required to generate three frequencies: a single L1 LO signal, a single L2 LO signal, and the P-code baseband sample rate (15.374 MHz).

Selection of a sample rate of 15.374 Msamples/s for the two P channels and 1.5374 Msamples/s in the C/A channel was based on two considerations. First, hardware restrictions at the time limited the P sample rate to less than 16 Msamples/s, even though a rate close to 20 Msamples/s would have been preferable. Within this constraint, commensurability considerations were used to set the P rate at 15.374 Msamples/s, as discussed in Section 5.6.

In the digital baseband processing, all high-speed operations (i.e., at 1.5374 MHz or at 15.374 MHz for carrier counterrotation and code correlation) are carried out by integrated circuits specially designed and implemented for the Rogue receiver, while all of the slower operations (e.g., tracking and averaging operations at 50 Hz or less) are performed in software. The tracking processor was implemented in software in order to provide ample flexibility in designing and testing the wide range of algorithms available for the slower tracking operations. Fast microprocessors with substantial memory were employed so that design would not be hampered by constraints of memory and computing speed.

Because of gate-count limits in chip technology at the time, the high-speed digital processing was based on minimum-bit algorithms that allow much more processing to be packed into a single chip. In the minimum-bit approach, the baseband signal is sampled with two levels (-1, +1) and the phasors used for carrier counterrotation are represented by three levels (-1, 0, +1). With this choice, the products produced by carrier counterrotation, as well as subsequent code correlation, are only three-level numbers (-1, 0, +1). Since only two bits are required to represent the processing at each stage up to the accumulators, complexity and gate count are greatly reduced. In this design, there is a voltage SNR loss of 4 percent due to the three-level quantization of the down-conversion sinusoids and a 20-percent loss due to two-level sampling of the baseband signal, leading to a total loss of 23 percent. In exchange for this modest loss in SNR compared to multilevel sampling, one can implement the minimum-bit design and arrive at digital logic with a minimum number of logical elements, thereby greatly reducing size, power consumption, and cost of the chip(s) fabricated to carry out these operations.

Another minimum-bit option that was considered, three-level sampling (-1, 0, +1) of the baseband signal, would have improved voltage SNR by 10 percent relative to two-level sampling. This slight improvement in SNR was not enough to justify the increased complexity in the sampling circuitry and in the front end hardware that was required by that option. In contrast to the

three-level option that requires a three-level A/D converter and accurate gain control in the front end, the two-level option requires sampling circuitry consisting of only a clipper and sign test, with no automatic gain control.

Code values for both the C/A and P channels are generated in a purely digital fashion through use of a chip advancer, which, in effect, relates each sample point with a chip count number. With the chip advancer, a sample rate can be selected that is essentially incommensurate with the chip rate.

In the high-speed digital processing, separate signal processors are provided for each satellite, with no time-sharing among satellites or among the C/A, P1, and P2 channels. This capability substantially improves SNR relative to multiplexed designs. The same three sampled signals are processed by each satellite signal processor. Since the baseband processor is purely digital and since the same sampled data are processed by all satellite processors, intersatellite errors introduced by the baseband processor can be reduced to a negligible level.

In order to obtain better tracking for low SNR signals, fast feedback for phase and delay for all channels is based on the phase of the strongest channel, the L1-C/A channel. Because it determines lock reliability, the C/A channel is sampled in quadrature to provide better lock in cases of low signal strength. Even though it would have been preferable, the P channels were not designed as quadrature channels due to gate-count limits on chip design at the time. Dropping quadrature for the P channels provides a significant simplification in instrumentation. The price paid for this simplification is an increase of $\sqrt{2}$ in system noise error in the phases and delays output by the P channels. The phase-error increase is insignificant, and the importance of the delay-error increase can be minimized in many applications by proper techniques in subsequent processing. Because of the design of the phase-lock algorithms, the extra noise on the P channels should not affect lock reliability on the P channels. That is, the P channels should not lose lock if the C/A channel maintains lock, given nominal signal strengths.

With the system parameters presented in the last subsection, the maximum voltage SNR for an 18-ms integration falls in the range of 25 to 40 for the quadrature C/A channel. Based on such relatively strong SNR values, the C/A phase-lock loop could be designed to track high dynamics, with high accuracy and extremely infrequent cycle slips or data-bit flips. The initial phase-lock loop (PLL) in the prototype receiver is based on arctangent extraction of tracking error and two-point linear extrapolation of measured phase. The use of an arctangent extractor provides accurate phase over a wider range than a sine extractor. This PLL has a voltage SNR cutoff of about 7 for an 18-ms integration, which provides a margin of 10 to 15 dB relative to maximum SNR. With this SNR margin and a typical omnidirectional antenna, an Earth-fixed Rogue receiver will not lose lock until a satellite falls well below 10° in elevation. The extrapolation loop can track with high accuracy and high reliability under high dynamics, given the high SNRs anticipated for Rogue applications.

In addition to the extrapolation PLL, a future version of the tracking processor software will provide the option of using a conventional digital PLL. This optional PLL, which is not described in this report, will track to lower SNR, but with lower dynamic capabilities.

Receiver time is kept by a system clock driven by the P sample clock. Based on the system clock, the receiver can specify, to the sample point, the start and stop time for the correlation sums, thereby exactly controlling the samples entering each sum. Consequently, time tags for the correlation sums (and for the output values for phase and delay) can be computed exactly in terms of receiver time. (The accuracy of receiver time relative to an external standard such as GPS time will depend on the accuracy of the technique used to synchronize receiver time with the external standard.)

In-lock data-bit synchronization is enforced by offsetting the start time of each correlation sum with the most recently measured C/A group delay. This approach maintains data-bit synchronization with an accuracy better than 10 μs, without the use of special synchronization circuitry or software.

The "averaging" operations to produce low-rate output values for phase are based on a least-squares fit of a polynomial to a selected number of the values generated by the tracking loops. The output "averaging interval" is set independently of tracking-loop parameters.

SECTION 3

TOP-LEVEL FUNCTIONAL DESCRIPTION

A top-level functional block diagram for signal processing in the Rogue receiver is shown in Figure 3-1. Signals from all visible GPS satellites are picked up by an omnidirectional antenna and then filtered and amplified at RF. Using fixed-frequency LOs, the front-end hardware down-converts the GPS signals from RF to baseband, where the signals are low-pass filtered and digitally sampled. Four digitally sampled signals are supplied to the digital baseband processor: the P1 and P2 signals (each nonquadrature at 15.374 Msamples/s) and the C/A signal (in quadrature at 1.5374 Ms/s).

As schematically illustrated in Figure 3-2 for a given channel (C/A, P1, or P2), the digital baseband processor contains two major processing components for each satellite: a high-speed digital processor and a tracking processor. Control information sent by the tracking processor to its associated digital hardware consists of the start time for the next correlation interval, as well as model values for phase rate, initial phase, chip rate, and initial chip for each channel (C/A, P1, P2). (The term "chip" will be used in this report to refer to a code sign generated by a code generator, with a duration of $1/1.023$ μs for C/A and $1/10.23$ μs for P. The terms "c-sample" and "p-sample" will refer to the interval between sample points, with the C/A channel at $1/1.5374$ μs and the P channels at $1/15.374$ μs, respectively.)

The system clock is driven by the P sample clock and consists of two registers: an integer-second register and a fractional-second register in units of p-samples. The fractional-second register is incremented by one with each cycle of the P sample clock. When the fractional-second register reaches 15,374,000, the integer-second register is incremented by one and the fractional-second register is zeroed. These two registers are accessed by the high-speed digital hardware to establish timing.

Based on this control information, the high-speed digital hardware accepts the four input data streams (C/A-I, C/A-Q, P1, P2), counterrotates the carriers, generates delayed pseudorandom codes, correlates the pseudorandom codes with the data streams, and accumulates the results over the specified sum interval (e.g., 18 ms). A processing channel with the above operations is provided for each of the two C/A channels and each P channel. Separation of signals between satellites is provided by orthogonality of the pseudorandom codes.

The output of the high-speed processing is a complex correlation sum for each lag for each channel. (The term "lag" in this report will refer to the correlation delays applied in steps of a sample interval.) These correlation sums are collected by the tracking processor and reduced to extract measured values for phase, delay, and data bits. These phase and delay values are extrapolated forward in time to obtain feedback values to drive the digital hardware over the next interval. The measured phase and delay values are also "averaged" over longer time intervals (e.g., three minutes) in order to reduce the output data rate.

MULTIPLE SATELLITES

RF L1, L2

FRONT END ←—— LOs ——— FREQUENCY SYSTEM

BASEBAND SIGNALS

ANALOG TO DIGITAL CONVERSION ←—— SAMPLING CLOCK

DIGITAL BASEBAND SIGNALS | C/A, P1, P2

DIGITAL BASEBAND PROCESSOR

τ, ϕ FOR C/A, P1, AND P2 FOR MULTIPLE SATELLITES
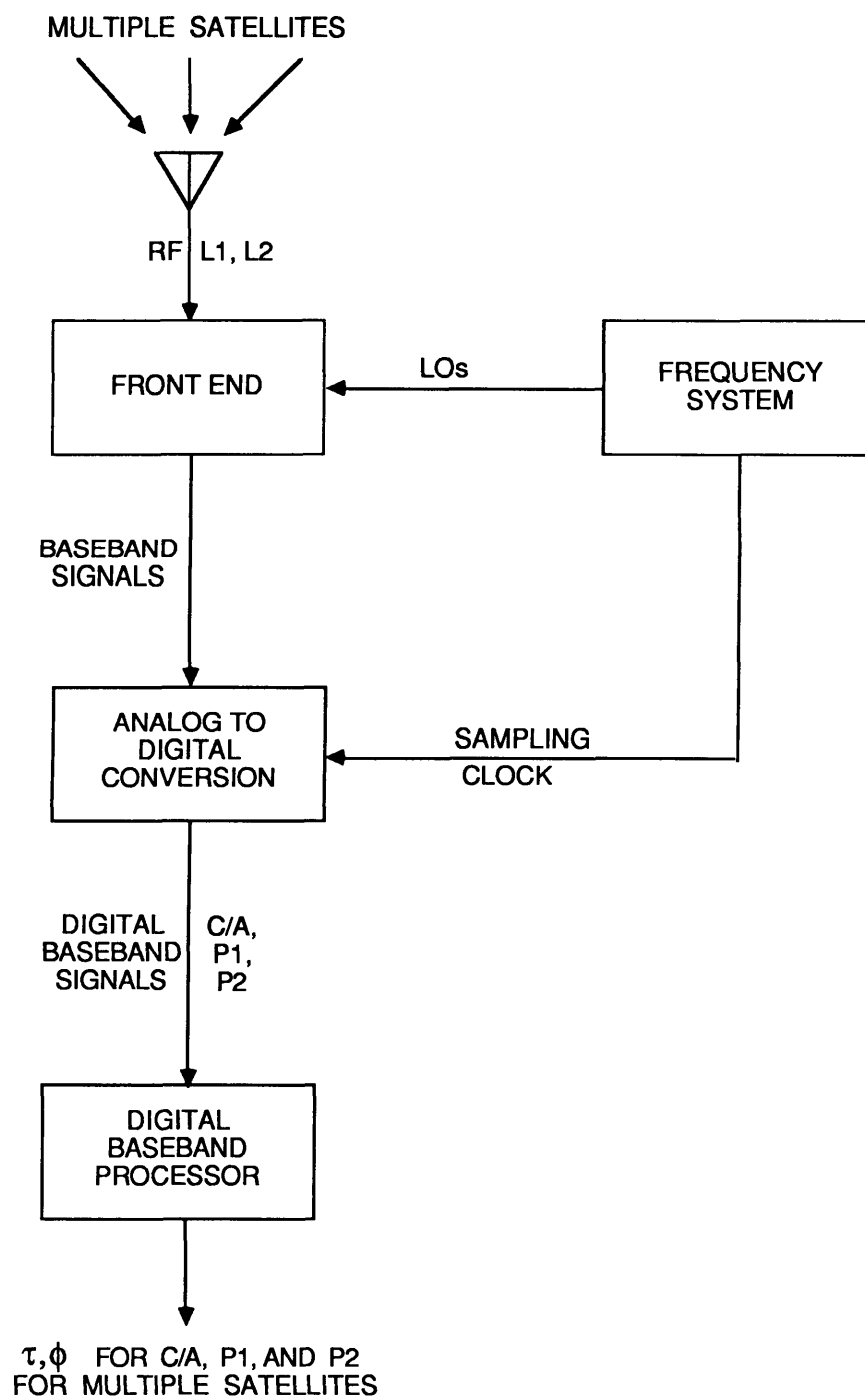
Figure 3-1. Top-Level Block Diagram for the Rogue Receiver
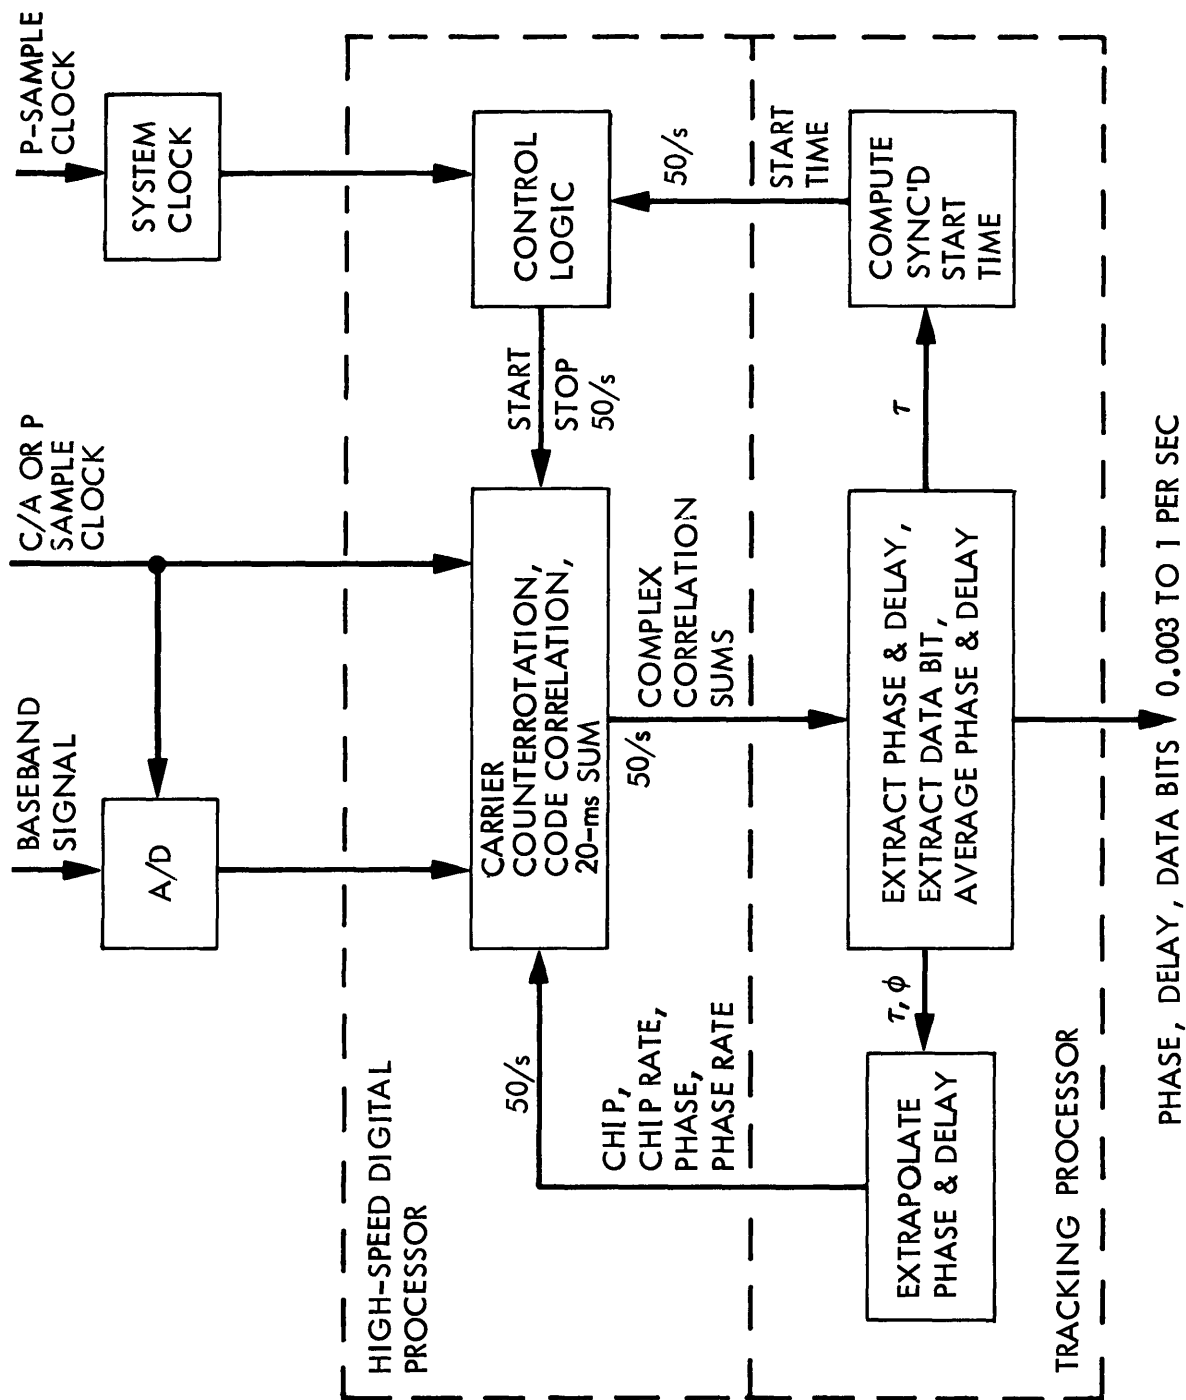
Figure 3-2. Schematic Block Diagram of the Digital Baseband Processor

SECTION 4

FRONT END

As shown in Figure 4-1, the output from the antenna is passed through a bandpass filter spanning 1140 to 1640 MHz in order to eliminate out-of-band RF interference. After amplification by a gallium arsenide field-effect transistor (GaAs FET) with a noise temperature of about 80 K, the signal is divided into L1 and L2 channels, and each channel is down-converted to baseband in the double sideband mode with fixed-frequency LOs. The L1 branch is down-converted in a quadrature mixer in order to generate quadrature signals for the C/A channel. In addition to supplying the C/A channel, one of the outputs of the quadrature mixer also supplies the P1 channel. The four resulting baseband signals are each passed through a low-pass filter whose width is 720 kHz for the C/A channels and 7.2 MHz for the P channels. The resulting signals are then each sampled in a two-level mode with a rate of 1.5374 MHz for each of the two C/A channels and 15.374 MHz for each of the two P channels.

To carry out the preceding processing, three frequencies are required: the sample frequency of 15.374 MHz and two LO frequencies equal to 1575.40768 MHz and 1227.5904 MHz. The C/A sample clock is obtained by dividing down the P sample clock. The three required frequencies are derived from a 5-MHz reference by the frequency subsystem of the Rogue receiver.

Care must be taken in selecting the sampling frequencies and the LO offsets, as discussed in Section 5.6. The sampling frequency must be highly incommensurate with the fundamental code chip rate (e.g., 15.374 MHz vs. 10.23 MHz) so that discrete-sampling errors will be negligible. The baseband carrier frequency should be about 5 kHz or greater so that carrier down-conversion in the baseband processor will work with high accuracy. To meet this requirement, the LO frequencies are offset from the corresponding L-band carriers by 12.32 kHz for L1 and 9.6 kHz for L2.

BP FILTER
BW = 500 MHz
f_c = 1390 MHz

f_{L1} = 1575.42 MHz

f_{L2} = 1227.6 MHz

AMPLIFIER
T = 80 K

L1

L2

QUADRATURE MIXER

f_{L1} − 12.32 kHz

f_{L2} − 9.6 kHz

Q

I

LPF
720 kHz

LPF
720 kHz

LPF
7.2 MHz

LPF
7.2 MHz

A/D

A/D

A/D

A/D

SAMPLING CLOCKS

±1

±1

±1

±1

L1 - C/A, Q
1.5374 MHz

L1 - C/A, I
1.5374 MHz

L1 - P
15.374 MHz

L2 -P
15.374 MHz
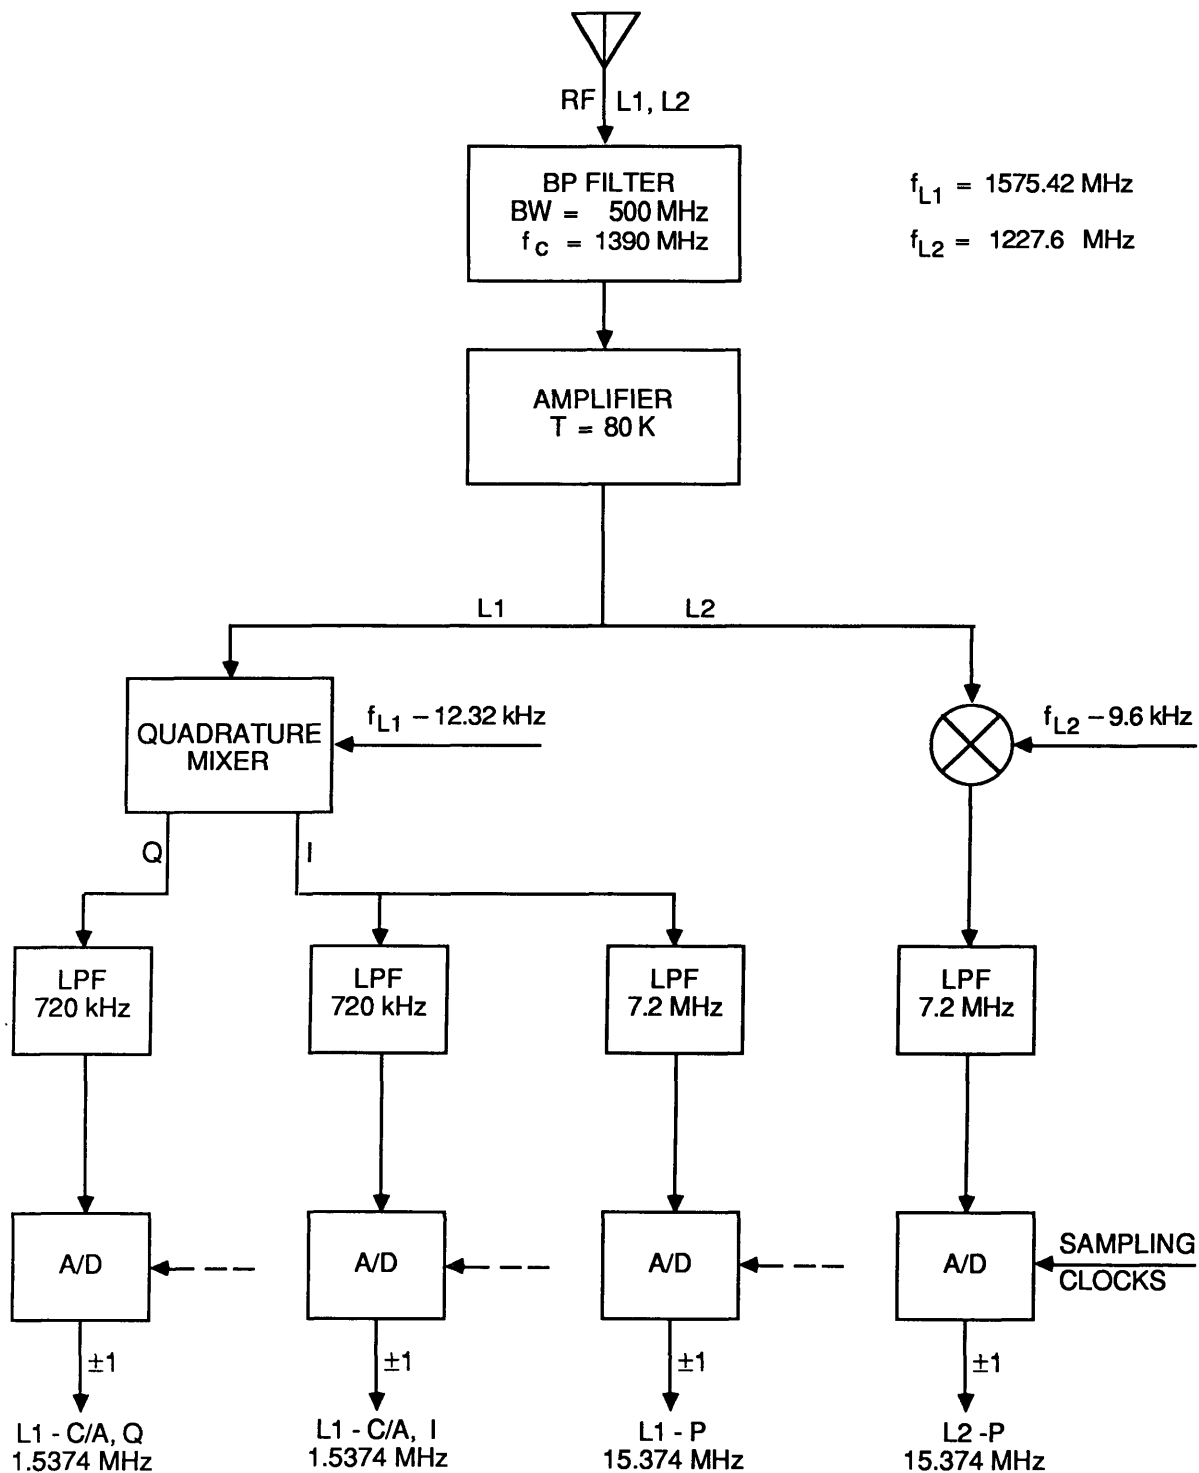
Figure 4-1. Block Diagram for the Front End of the Rogue Receiver

SECTION 5

HIGH-SPEED DIGITAL PROCESSOR

5.1   BLOCK DIAGRAM

A functional block diagram of the high-speed digital processor is shown in Figure 5-1. As indicated, the major steps in high-speed processing of the digitized baseband signal are carrier counterrotation, code cross-correlation, and correlation accumulation. This process produces a complex correlation sum for each of four lags. Figure 5-1 is valid for both the C/A channel and the P channels and is shown for nonquadrature sampling. In principle, quadrature processing could be accomplished by replicating for the quadrature channel the counterrotation mixer, the code mixers, and the accumulators, while sharing the phase and code generators. However, for the prototype receiver, quadrature processing is accomplished for the L1-C/A channel by replicating all of the circuitry in Figure 5-1. With this approach, one has the option of either processing the L1-C/A channel in quadrature or processing both the L1-C/A and the L3-C/A channels in nonquadrature. In the prototype receiver, both P channels are processed in the nonquadrature mode.

5.2   PROCESSING CONTROL

As indicated in the lower right-hand corner of Figure 5-1, the high-speed hardware is started and stopped by control logic that must be updated by the tracking processor every 20 ms with a start time in the form of an integer-second part (in seconds) and a fractional-second part (in units of p-samples). When receiver time, as kept by the system clock, has advanced to equal the start time, the start line is raised to activate the high-speed digital hardware. Deactivation occurs when the number of summed sample points reaches the integer value in a sum-interval register supplied at the beginning of the tracking. With this method of sum timing, the tracking processor can control exactly which samples are included in each "18-ms" correlation sum and can therefore compute the sum time tag exactly in terms of receiver time. During in-lock tracking, the tracking processor accurately synchronizes sum interval with data-bit edge by offsetting the start time with the most recently measured C/A group delay, as explained below.

5.3   PHASE ADVANCER AND CARRIER COUNTERROTATION

As mentioned above, the signal input to the high-speed hardware signal is digitized to two levels (+,-) with a Nyquist sample rate of 1.5374 MHz for C/A and 15.374 MHz for P, for the current hardware implementation. Other sample rates could be also be processed by the high-speed hardware and tracking processor.

In the high-speed processor, each signal carrier is counterrotated by means of complex multiplication with a three-level (+,0,-) phasor generated in the high-speed hardware on the basis of feedback. This counterrotation removes both the LO offset mentioned above and Doppler rate. Over each sum interval, carrier phase is approximated by a linear function whose initial

Figure 5-1. High-Speed Digital Baseband Processing in the Rogue Receiver

5-2

value and rate are reset every sum interval with feedback from the tracking processor. Baseband carrier rates range between 6 and 17 kHz. The initial phase value is supplied in the form of an integer which is obtained by multiplying the starting fractional phase by 2^{24} and then rounding. Similarly, the phase rate value is obtained by multiplying phase rate by 2^{24} and rounding to the nearest integer.

For each new sample point, the phase advancer increments phase by adding the phase-rate register to the fractional-phase register. As phase is updated, integer cycles that appear as overflow are discarded since they have no effect on computation of sinusoids. The three-level phasors used for counterrotation at each sample point are obtained by a table lookup based on the four most significant bits of the (fractional) phase register, as illustrated in Figure 5-2 for a particular phase value. The form of the three-level amplitude quantization is shown in Figure 5-3, together with tick marks along the phase axis to establish the 16 table-lookup bins.

Multiplication of the sampled input signal with this quantized phasor results in a three-level product (+,0,-). After counterrotation, the signal is sent along four paths and each path is multiplied by a delayed code value from the code generator.

5.4   CHIP ADVANCER AND CODE GENERATOR

By supplying the proper initial integer chip (i.e., initializing the code generator), initial fractional chip, and chip rate, the tracking processor can make the chip advancer and code generator in the high-speed hardware produce a code sequence that starts at a fraction of a given chip and advances at the selected chip rate. The initial chip value (integer and fractional parts) are based on start time and feedback delay, while the chip rate is equal to the sum of the delay rate (in chip per sample) and the fundamental chip rate divided by the sample rate (i.e., 10.23/15.374). The output code sequence is generated at the data sample rate (e.g., 1.5374 MHz for C/A). As implied, the model delay consists of a linear time function over an 18-ms correlation interval.

The chip advancer for the C/A channel is illustrated in Figure 5-4. A similar figure would describe the chip advancers for the P channels. Every 20 ms, the tracking processor initializes the fractional-chip register and the adjacent chip-rate register. These registers are 24 bits long for the C/A channel and 28 bits long for the P channel.

The tracking processor also initializes registers in the code generator so that the correct code sign in the pseudorandom sequence will be produced for the first sample point. Initial values for these registers are derived from the integer part of the model (feedback) chip. Every 20 ms, two register values must be derived and supplied to the C/A-code generator and seven to the P-code generator.
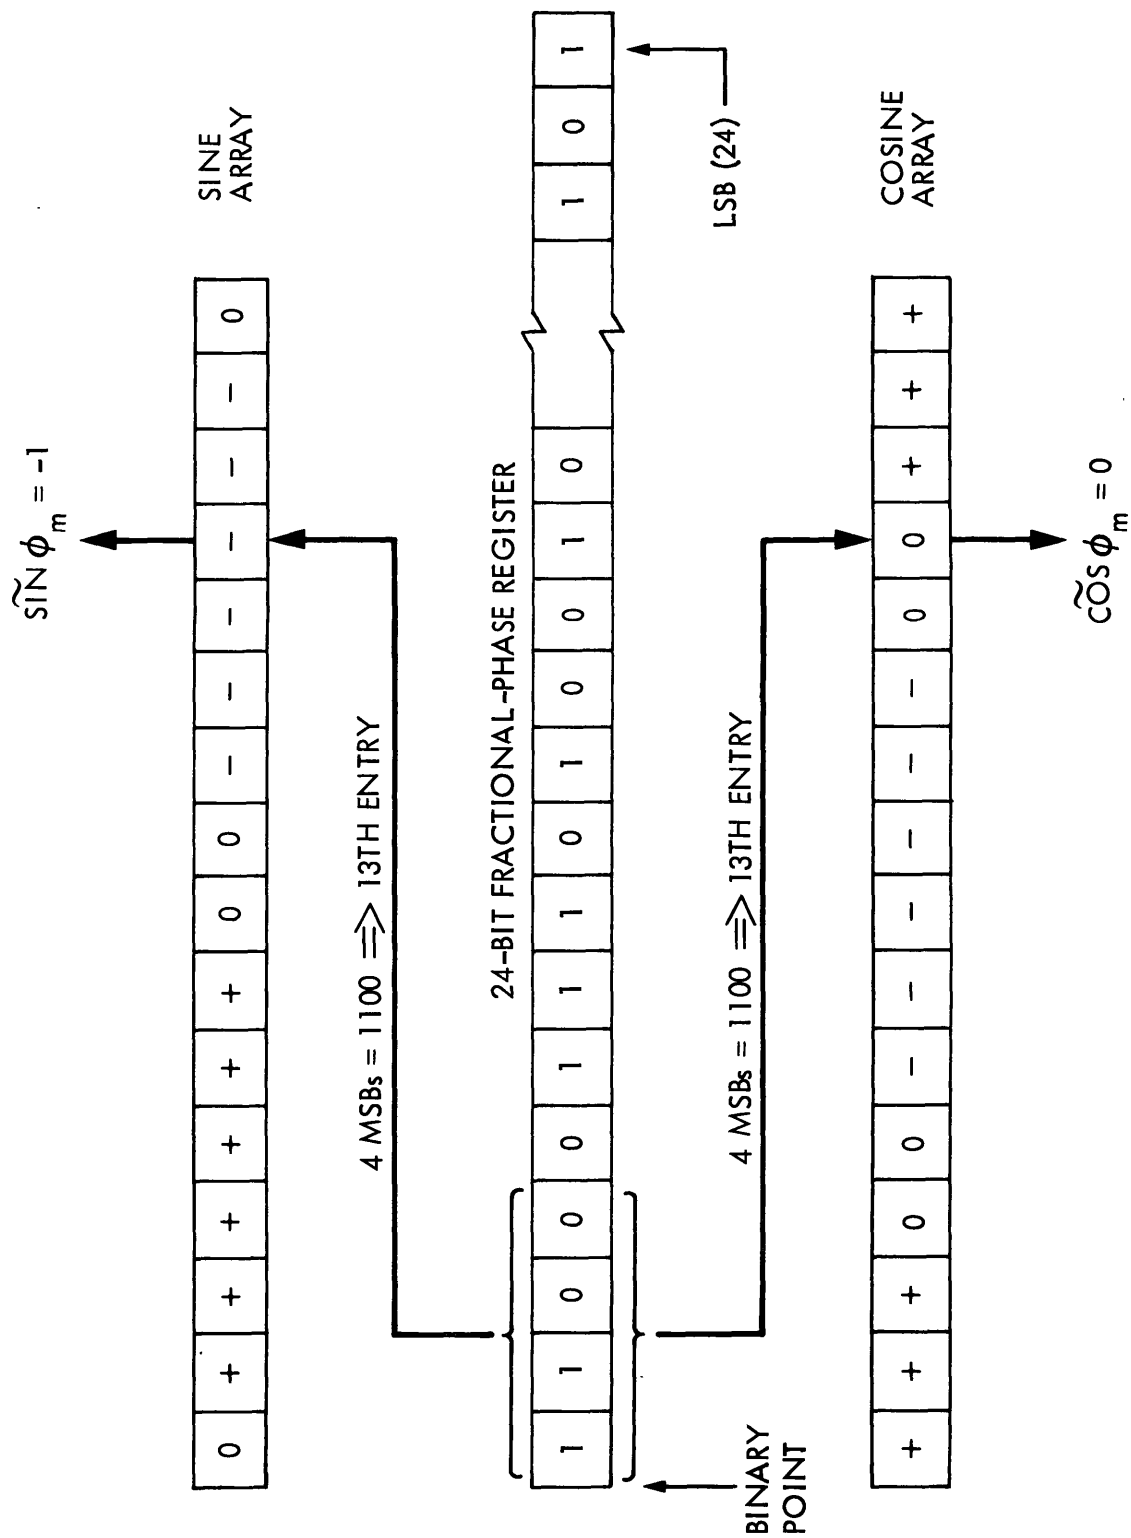
Figure 5-2. Schematic Illustration of the Generation of Three-Level Counterrotation Sinusoids
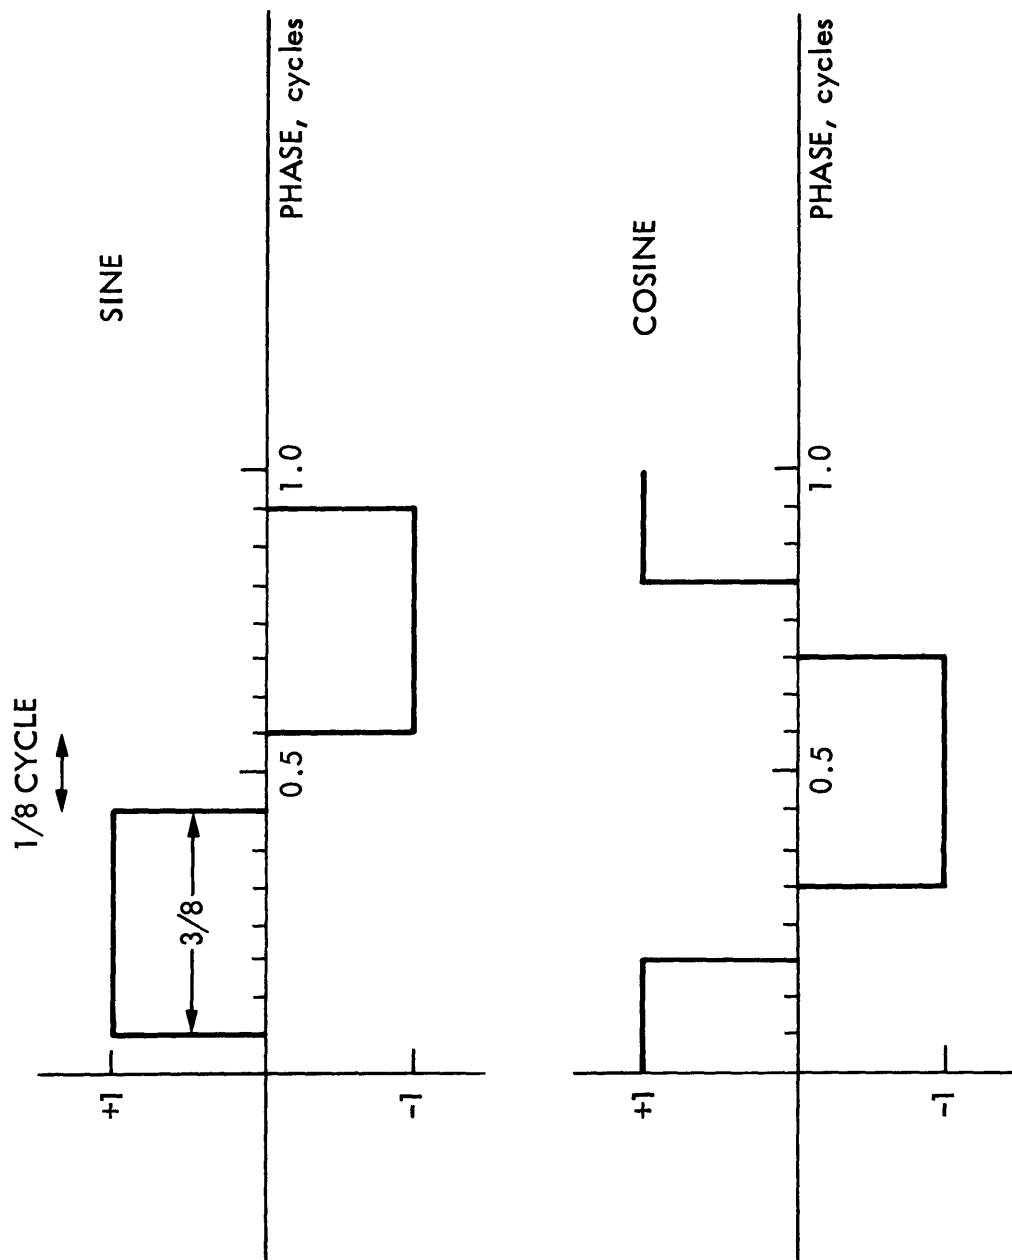
Figure 5-3. The Three-Level Counterrotation Sinusoids
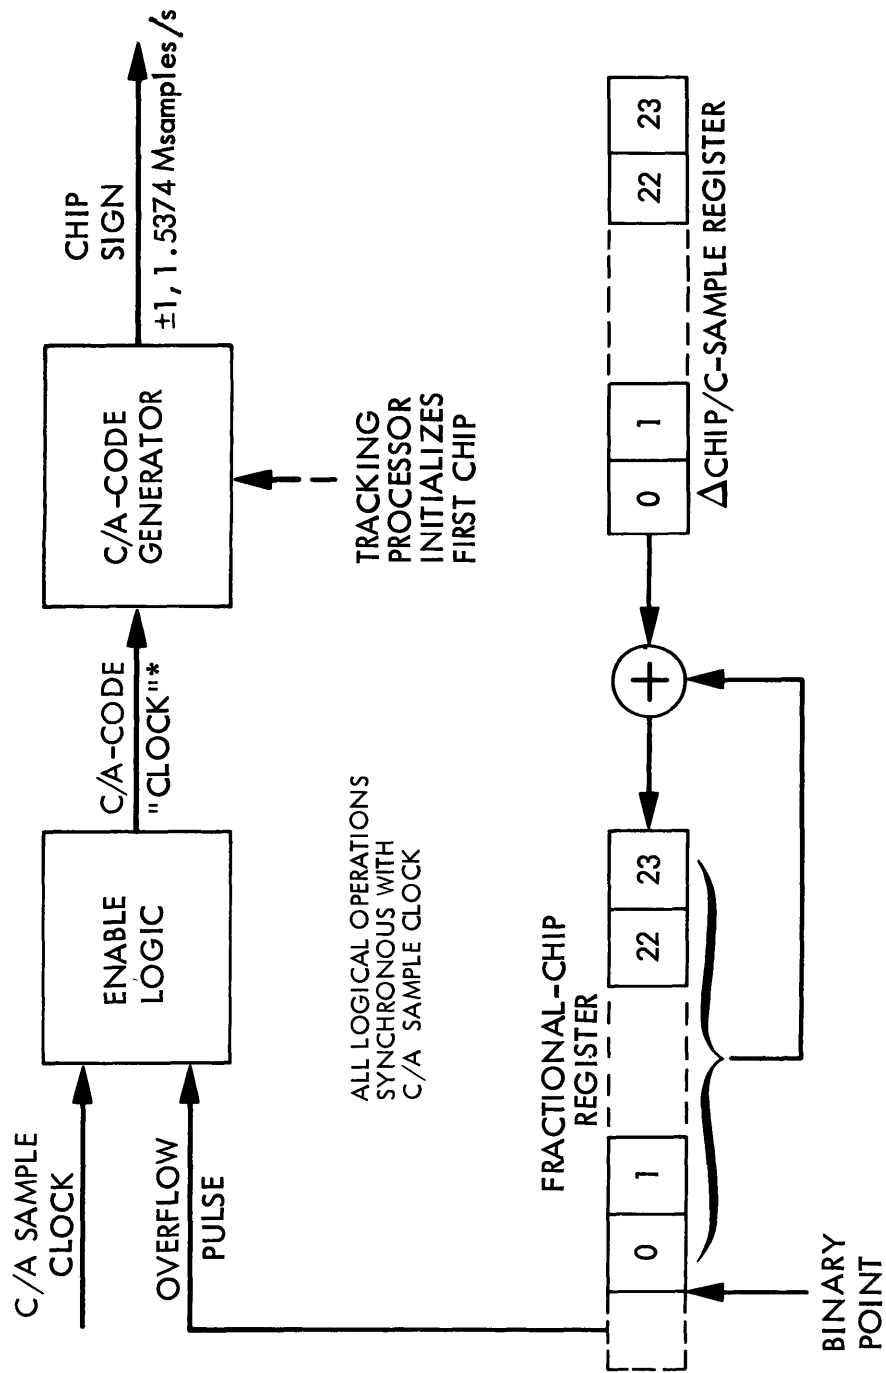
Figure 5-4. Schematic Illustration of the Chip Advancer for the C/A Channel

When processing is started, the chip advancer is clocked by the sampling clock (e.g., 1.5374 MHz for the C/A channel). For each new sample point, the chip advancer adds the chip rate register to the fractional chip register. When the fractional chip register overflows, indicating a transition to a new chip, a pulse is sent to advance the code generator to a new chip. That pulse triggers enable logic that lets the concurrent sampling clock edge advance the code generator to the next chip. If the code generator does not receive a new chip advance for a given sample point, the code sign persists for that sample. In this manner, a stream of code signs is generated at the sample rate and passed on for cross-correlation with the counterrotated data stream.

The stream of code signs produced by the code generator can be viewed as discrete samples of a piecewise-constant function of time consisting of a sequence of chips with essentially instantaneous transitions at sign changes. The rate of this chip sequence is equal to the fundamental chip rate with a Doppler shift (e.g., $1.023 (1 - \dot{\tau})$ MHz).

5.5 CORRELATION AND ACCUMULATION

The output of the code generator is passed into a tapped four-bit delay line. The four outputs of the delay line are multiplied by the complex output of carrier counterrotator at the sample rate (1.5374 MHz for C/A or 15.374 MHz for P). The complex three-level output of each of these multiplications is then accumulated over an 18-ms interval. At the end of the sum interval, the contents of the 32 accumulators are passed to the tracking processor and then set to zero for the next 20-ms interval.

5.6 ERRORS INTRODUCED BY THE HIGH-SPEED DIGITAL PROCESSOR

Because all operations, including carrier counterrotation, are purely digital, phase and delay can be tracked and measured with extremely small errors. Errors due to roundoff, level quantization, and commensurability are reduced to negligible levels. Given 24-bit chip and phase advancers, for example, roundoff errors are less than 100 nanochips and 100 nanocycles. The design of the chip advancer makes the generation of the code sequence purely digital and extremely accurate. The code sequence produced by the chip advancer has no amplitude distortion and is delayed with a cumulative 20-ms accuracy better than 100 nanochips with respect to the chip model generated by the tracking processor. Similarly, the three-level counterrotation phasor is generated exactly as presented in Figure 5-3 with a cumulative 20-ms phase accuracy better than 100 nanocycles with respect to the phase model of the tracking processor.

A subtle effect not explained to this point is the error buildup due to roundoff in the phase-rate or chip-rate registers. When the tracking processor calculates the integer value to initialize a rate register for a 20-ms interval, a roundoff error is incurred. The tracking processor calculations are designed to anticipate the error buildup resulting from this roundoff and to cancel its cumulative effect on phase or delay across a 20-ms interval. The nature of this adjustment can be outlined for phase as described below.

By carrying out the calculation of start phase described in Section 6.2.2.1, one reduces to a negligible level the error buildup due to roundoff in phase rate across a correlation interval. The phase shift used to calculate start phase from center phase is based on the rounded rate passed to the rate register. Since the subsequent use of this rate register by the phase advancer between the start time and center interval will advance phase back to the center value by exactly this phase shift, the model phase value applied at interval center by the phase advancer differs from the model phase of the tracking processor by less than 100 nanocycles, in spite of rate roundoff error.

With this approach, the only effect of rate roundoff is to reduce amplitude. This amplitude reduction is due to increased mismatch between actual and model phase rate across the interval. Since the maximum frequency mismatch due to rate rounding is about 30 nanocycles/sample for a 24-bit rate register, the maximum amplitude loss will be about 0.0001 percent for C/A and 0.01 percent for P, given an 18-ms integration interval. These calculations show that the effect of roundoff errors in the phase rate registers is negligible for the Rogue receiver. Analogous arguments demonstrate that the error resulting from rate rounding in a chip-rate register is also negligible.

Commensurability errors depend on the specific value of sample rate relative to baseband carrier frequency and to chip rate (with both offset by Doppler effects). Due the large difference in sample rate and baseband carrier frequency (15.374 MHz vs. about 10 kHz), the commensurability error in phase will be less than a millicycle for the P channels, given the quantization of the counterrotation phasors. As Doppler changes, this error will average down to an even smaller value.

On the other hand, to reduce the error caused by commensurability of sample rate and chip rate, one must be more careful. If the sample rate is selected so that it is effectively incommensurate with the fundamental chip rate, the integrated error due to discrete sampling of the model code sequence can be reduced to negligible levels after integration over a fairly short time interval. For the selected P sample rate of 15.374 MHz, the commensurability ratio becomes

$$15374000/10230000 = 7687/5115$$

after common factors are canceled. One can readily show that, after 7687 samples (0.5 ms), the placement sequence of sample points relative to nearest chip edges begins to repeat so that the integrated error due to discrete sampling becomes 1/7687 after 0.5 ms. In other words, if a chip sequence were shifted in time (i.e., by changing the initial fractional chip register sent to the chip advancer), the entire sequence of 7687 code signs would not change until the magnitude of the shift reached 1/7687 chip. In this case, delay could not be measured any more accurately than 1/7687 chip (0.4 cm for P). This error is actually the worst-case error incurred in the case of zero Doppler and zero Doppler rate. In practice, even this small error will be made much smaller by the smearing caused by changing Doppler. For example, given an RF Doppler rate of 0.7 Hz/s for a ground-based receiver, the quantization error of 1/7687 chip will be traversed in only 0.24 s due to

continuous chip shifting caused by changing Doppler. In selecting the sample rate, a value has been chosen that is sufficiently removed from highly commensurate ratios (e.g., 2/1, 3/2, etc.) so that Doppler cannot change the chip rate to a value highly commensurate with the sample rate.

Another error found in this system is the sum-note error associated with nonquadrature down-conversion of carrier. This sum-note error will average down to less than a millicycle over an 18-ms integration if the carrier frequency is more than 5 kHz at baseband. A large positive baseband carrier frequency of 6 to 17 kHz has been obtained by setting the LO frequency used in down-conversion equal to a value that is less than the minimum Doppler-shifted carrier frequency at RF, given Earth-fixed antennas. With such a large baseband frequency, effects of unwanted harmonics in the three-level counterrotation sinusoids also average down to negligible levels over an 18-ms interval.

SECTION 6

TRACKING PROCESSOR

6.1   TOP-LEVEL BLOCK DIAGRAM

A top-level block diagram of the tracking processor is shown in Figure 6-1, including the interface with the high-speed hardware. The figure illustrates the C/A channel and one P channel, where the P channel represents either P1 or P2.

The following description will refer to the fast feedback interval as the "20-ms" interval even though, in the tracking software, the interval can be set to 1, 2, 4, 5, 10, or 20 ms (in practice, current hardware and microprocessor speeds limit the choices to 10 or 20 ms). For a setting less than 20 ms, logic is included to account for the fact that the same data bit is sampled more than once and to reduce those multiple measurements into one per 20 ms. Between "20-ms" intervals, dead-time operations are required to complete feedback calculations and to initialize hardware. For the current Rogue implementation, these "dead-time" operations take 2 ms, which leaves 18 ms for the correlation sum when 20-ms feedback is selected (or 8 ms for 10-ms feedback).

Processing in the C/A channel will be outlined first. The input to the C/A processing, the C/A correlation sums, are analyzed every 20 ms to extract L1 carrier phase. This phase is used to extrapolate ahead in time to produce fast (20 ms) feedback values for all delays and for the L1-P and L2-P phases. To obtain the data bit for each 20-ms interval, only the correlation output for the prompt delay of the C/A channel is tested. Delay feedback based on C/A phase is so accurate over short time intervals that the C/A correlation-sum amplitudes can be accumulated over one second. Lag dependence of these 1-s sums is analyzed to extract residual delay, which is then used in two ways.

First, residual delay is used as a slow (1 s) feedback correction to the fast (50/s) delay feedback based on C/A phase. This correction accounts for slow drifts of group delay relative to phase delay, drifts mainly due to ionospheric effects. Second, residual delay is combined with averaged model (feedback) delay to produce a total measured delay each second. Once per second, this total delay is used to readjust the data bit synchronization of the "20-ms" start time that activates the digital hardware.

The 20-ms phase values for the C/A channel are fit with a quadratic time function over each 1-s interval to produce one phase value per second. Both the 1-s phase values and the 1-s delay values are then subjected to a polynomial fit over an N_s-s interval to produce one phase and one delay value every N_s seconds. These fit values, along with the data bits, are the primary output of the C/A channel.

In parallel with the C/A processing, the two P channels are also reduced, as symbolized by the one P channel in Figure 6-1. As indicated above, the fast feedback for both delay and phase for the P channels is based on L1-C/A phase. Once each 20 ms, phase and delay extrapolated in this manner

HIGH-SPEED DIGITAL PROCESSOR

START TIME, SUM LENGTH $t_s, \Delta t$ 50/s

C/A FEEDBACK: PHASE, PHASE RATE, CHIP, CHIP RATE 50/s

C/A CORRELATION SUMS I_ℓ^c, Q_ℓ^c 50/s

P FEEDBACK: PHASE, PHASE RATE, CHIP, CHIP RATE 50/s

P1 OR P2 CORRELATION SUMS I_ℓ^p, Q_ℓ^p 50/s

COMPUTE DATA-BIT-SYNCHED START TIME

t_s 50/s

EXTRAPOLATE FEEDBACK PHASE AND CHIP, C/A

ϕ_{L1}^c 50/s

EXTRACT L1-C/A PHASE

ϕ_{L1}^c 50/s

EXTRAPOLATE FEEDBACK PHASE AND CHIP, P

t_s 50/s

ϕ_{L1}^c

I_ℓ^c, Q_ℓ^c 50/s

$\Delta\tau_p, \Delta\phi_p$ 1/s

$\tau_{c/a}$ 1/s

(A)

$\Delta\tau_{c/a}$ 1/s

EXTRACT L1-C/A DATA BIT

t_s 50/s

(A)

SIGN 50/s

CORRECT P CORRELATION SUMS FOR DATA-BIT SIGN

SIGN-CORRECTED I_ℓ^p, Q_ℓ^p

I_ℓ^c, Q_ℓ^c 50/s

SUM C/A CORRELATION SUMS OVER ONE SECOND

ϕ_{L1}^c 50/s

QUADRATIC FIT TO L1-C/A PHASE OVER ONE SECOND

ϕ_{L1}^c 1/s

SUM P CORRELATION SUMS OVER ONE SECOND

SUMS 1/s

SUMS 1/s

COMPUTE RESIDUAL C/A DELAY

$\Delta\tau_{c/a}$ 1/s

$\Delta\tau_p, \Delta\phi_p$ 1/s

COMPUTE RESIDUAL P DELAY AND PHASE

$\tau_{c/a}$ 1/s

COMPUTE TOTAL C/A DELAY

$\Delta\tau_p, \Delta\phi_p$ 1/s

COMPUTE TOTAL P DELAY AND PHASE

$\tau_{c/a}$ 1/s

ϕ_{L1}^c

POLYNOMIAL FITS TO C/A DELAY AND PHASE OVER N SECONDS

$\tau_{c/a}, \phi_{L1}^c$ DATA BITS ONCE PER N SECONDS

τ_p, ϕ_L^p 1/s

POLYNOMIAL FITS TO P DELAY AND PHASE OVER N SECONDS

τ_p, ϕ_L^p ONCE PER N SECONDS

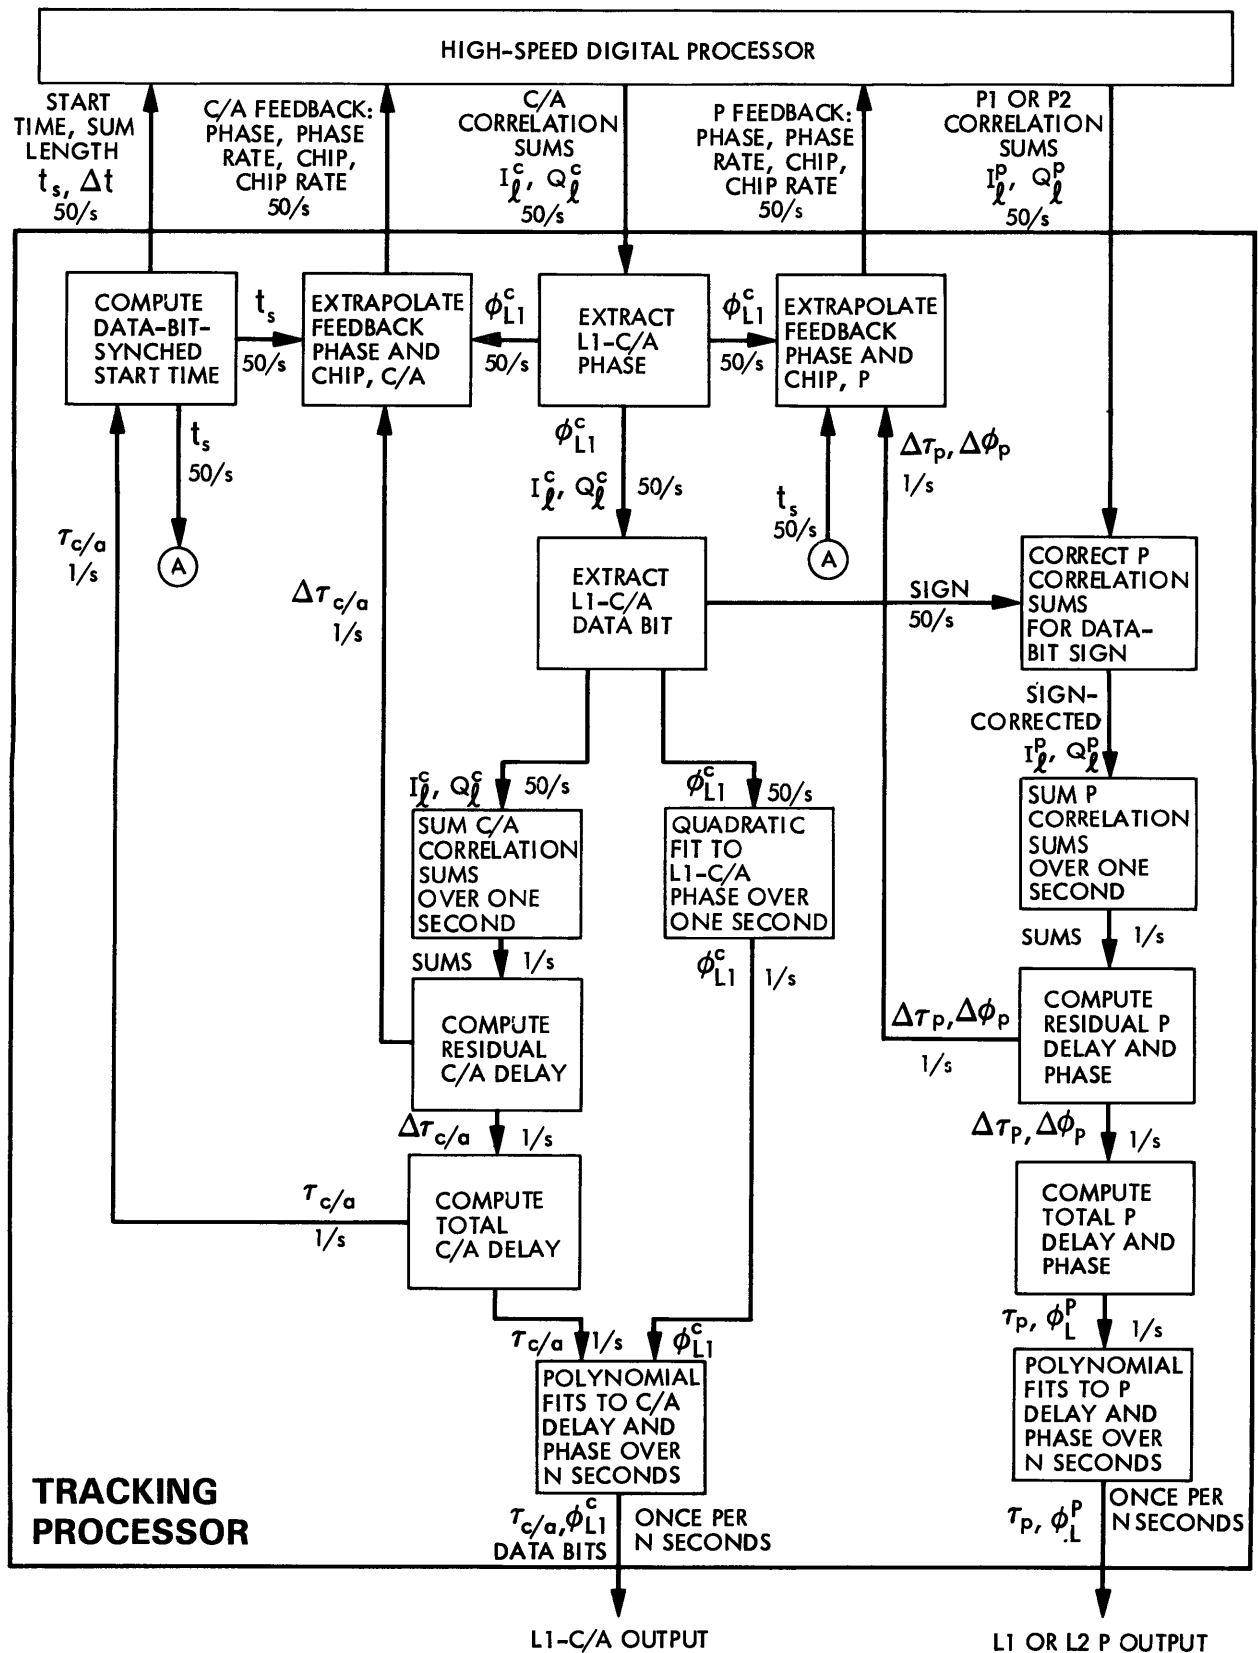TRACKING PROCESSOR

L1-C/A OUTPUT

L1 OR L2 P OUTPUT

Figure 6-1. Top-Level Block Diagram of Signal Processing in the Tracking Processor of the Rogue Receiver

6-2

are used to calculate quantities needed to initialize the digital hardware: initial chip (integer and fractional parts), chip rate, phase (fractional part), and phase rate. The P-channel correlation sums produced by the digital hardware are summed over one second. A sign correction based on the data-bit sign from the C/A channel is applied to the P 20-ms sums as needed. This sign correction expands the ambiguity range of P1 and P2 phase from 180 to 360 degrees, provided the C/A channel does not slip half cycles. After a 1-s sum is completed, an algorithm based on early/prompt/late amplitudes extracts residual delay while an arctangent operation based on the in-phase and quadrature components for the prompt lag extracts residual phase. These residual delay and residual phase values are both fed back once per second to correct the fast feedback values predicted by L1-C/A phase. This slow feedback correction accounts for slow drifts of P delay and P phase relative to the values based on L1-C/A phase. Residual delay and residual phase are also combined with model delay and model phase, respectively, to obtain total measured P delay and P phase. As in the C/A channel, the resulting 1-s values are then fit with a polynomial time function to produce one phase and one delay value every N_s seconds.

In real-time systems, the time at execution becomes an important consideration. In the feedback process, some operations must be performed during the dead time between 18-ms sum intervals if information from the most recent interval, n, is to be incorporated in the feedback for the next interval, n+1. Since the dead-time interval should be made as small as possible in order to minimize lost data, the number of dead-time operations should be minimized. In order of sensitivity to feedback promptness, the feedback quantities are phase, phase rate, delay, delay rate. (A large reduction in sensitivity to phase rate is obtained by center-interval modeling as explained in Section 6.2.2.1.) Of these four, only feedback phase is based on the most recent feedback interval. The other three can be set up on the basis of earlier intervals (e.g., n-1, n-2). Thus, as indicated in the figures that follow, only phase is updated during the dead-time interval. Consequently, the other feedback values can be computed during a 18-ms sum interval in parallel with high-speed digital operations. This compromise, which is of negligible consequence in performance and accuracy in anticipated applications, greatly relaxes timing constraints in the design of the feedback operations.

Voltage SNR can be calculated for a given channel and a given integration interval by computing the root-sum-square of the real and imaginary components of the complex correlation sum and dividing it by expected root-mean-square noise. For two-level sampling, noise variance can be calculated very accurately from the expression $3N_p/4$, in which N_p is the number of sample points over the integration interval. The factor of 3/4 accounts for blanking in the three-level counterrotation sinusoids.

6.2 C/A-CHANNEL PROCESSING

A block diagram of signal processing operations for the C/A channel is presented in Figure 6-2, while the associated feedback operations are shown in Figure 6-3. Note that each operation is labeled to allow identification of its multitasking priority, with dead-time operations having the highest priority, followed in order of priority by 20-ms operations and 1-s operations. The signal processing operations will be explained first.

6.2.1   In-Lock C/A Signal Processing

Every 20 ms during the dead-time between 18-ms sum intervals, the tracking processor collects the correlation sums produced by the high-speed digital hardware and supplies new phase and chip values, together with a new start time, for the next 20-ms interval. The start time has been precalculated during the preceding 20-ms interval by incrementing the start time variable by 20 ms. In addition, once per second, start time is adjusted on the basis of measured C/A group delay to maintain data-bit synchronization. For ground-based receivers, data bit synchronization can be maintained in this fashion with an accuracy better than 10 μs.

As mentioned above, the C/A channel is processed in quadrature, producing two sets of correlation sums, one complex sum for each quadrature channel. Because the phase, phase rate, chip, and chip rate for both of these quadrature channels are initialized with the same values, the real component of one quadrature channel corresponds with the imaginary component of the other quadrature channel, and vice versa. (Hardware design and tests verify that the phase of the two channels is separated by 90 degrees to within a degree, and that the receiver delays to baseband are the same for the two channels to better than 0.01 C/A chip.) Thus, the first step in processing the C/A channel is to combine with appropriate sign corresponding components to obtain one complex correlation sum for the C/A channel. This combining, which is not shown in Figure 6-2, is carried out for each lag offset.

Processing follows three major routes: carrier-phase processing, data-bit extraction, and group-delay processing. Phase processing will be discussed first.

6.2.1.1   C/A Carrier-Phase Processing. Along the phase route, the prompt (in-phase and quadrature) components are combined during the dead time between 20-ms intervals to extract residual phase. If feedback (model) phase is accurate to better than 0.25 cycle with respect to all errors (e.g., system noise, acceleration, oscillator instability), then a two-quadrant arctangent will yield a value for residual phase with the correct 1/2-cycle ambiguity. Also estimated during the dead time is the total measured phase, which is equal to the residual phase plus the model phase previously computed for the middle of that 20-ms interval. The time tag for this phase value is the 20-ms start time plus one-half of the "18-ms" sum interval.

The resulting phase value is used to drive the feedback loop (described below). In addition, these 20-ms phase values are subjected to a quadratic fit over a 1-s interval with the fit reference time placed at the integer-second point at interval center. This fit provides values for phase, phase rate, and phase-rate rate once per second. The phase rate and phase-rate rate are used to make corrections to the other observables as explained below.

The last step in processing phase (see Figure 6-2) is to compress data volume further by fitting the 1-s values with a polynomial time function over N_s seconds, which leads to one output phase value each N_s seconds. The reference time (i.e., time tag) in each fit is placed on an integer-second
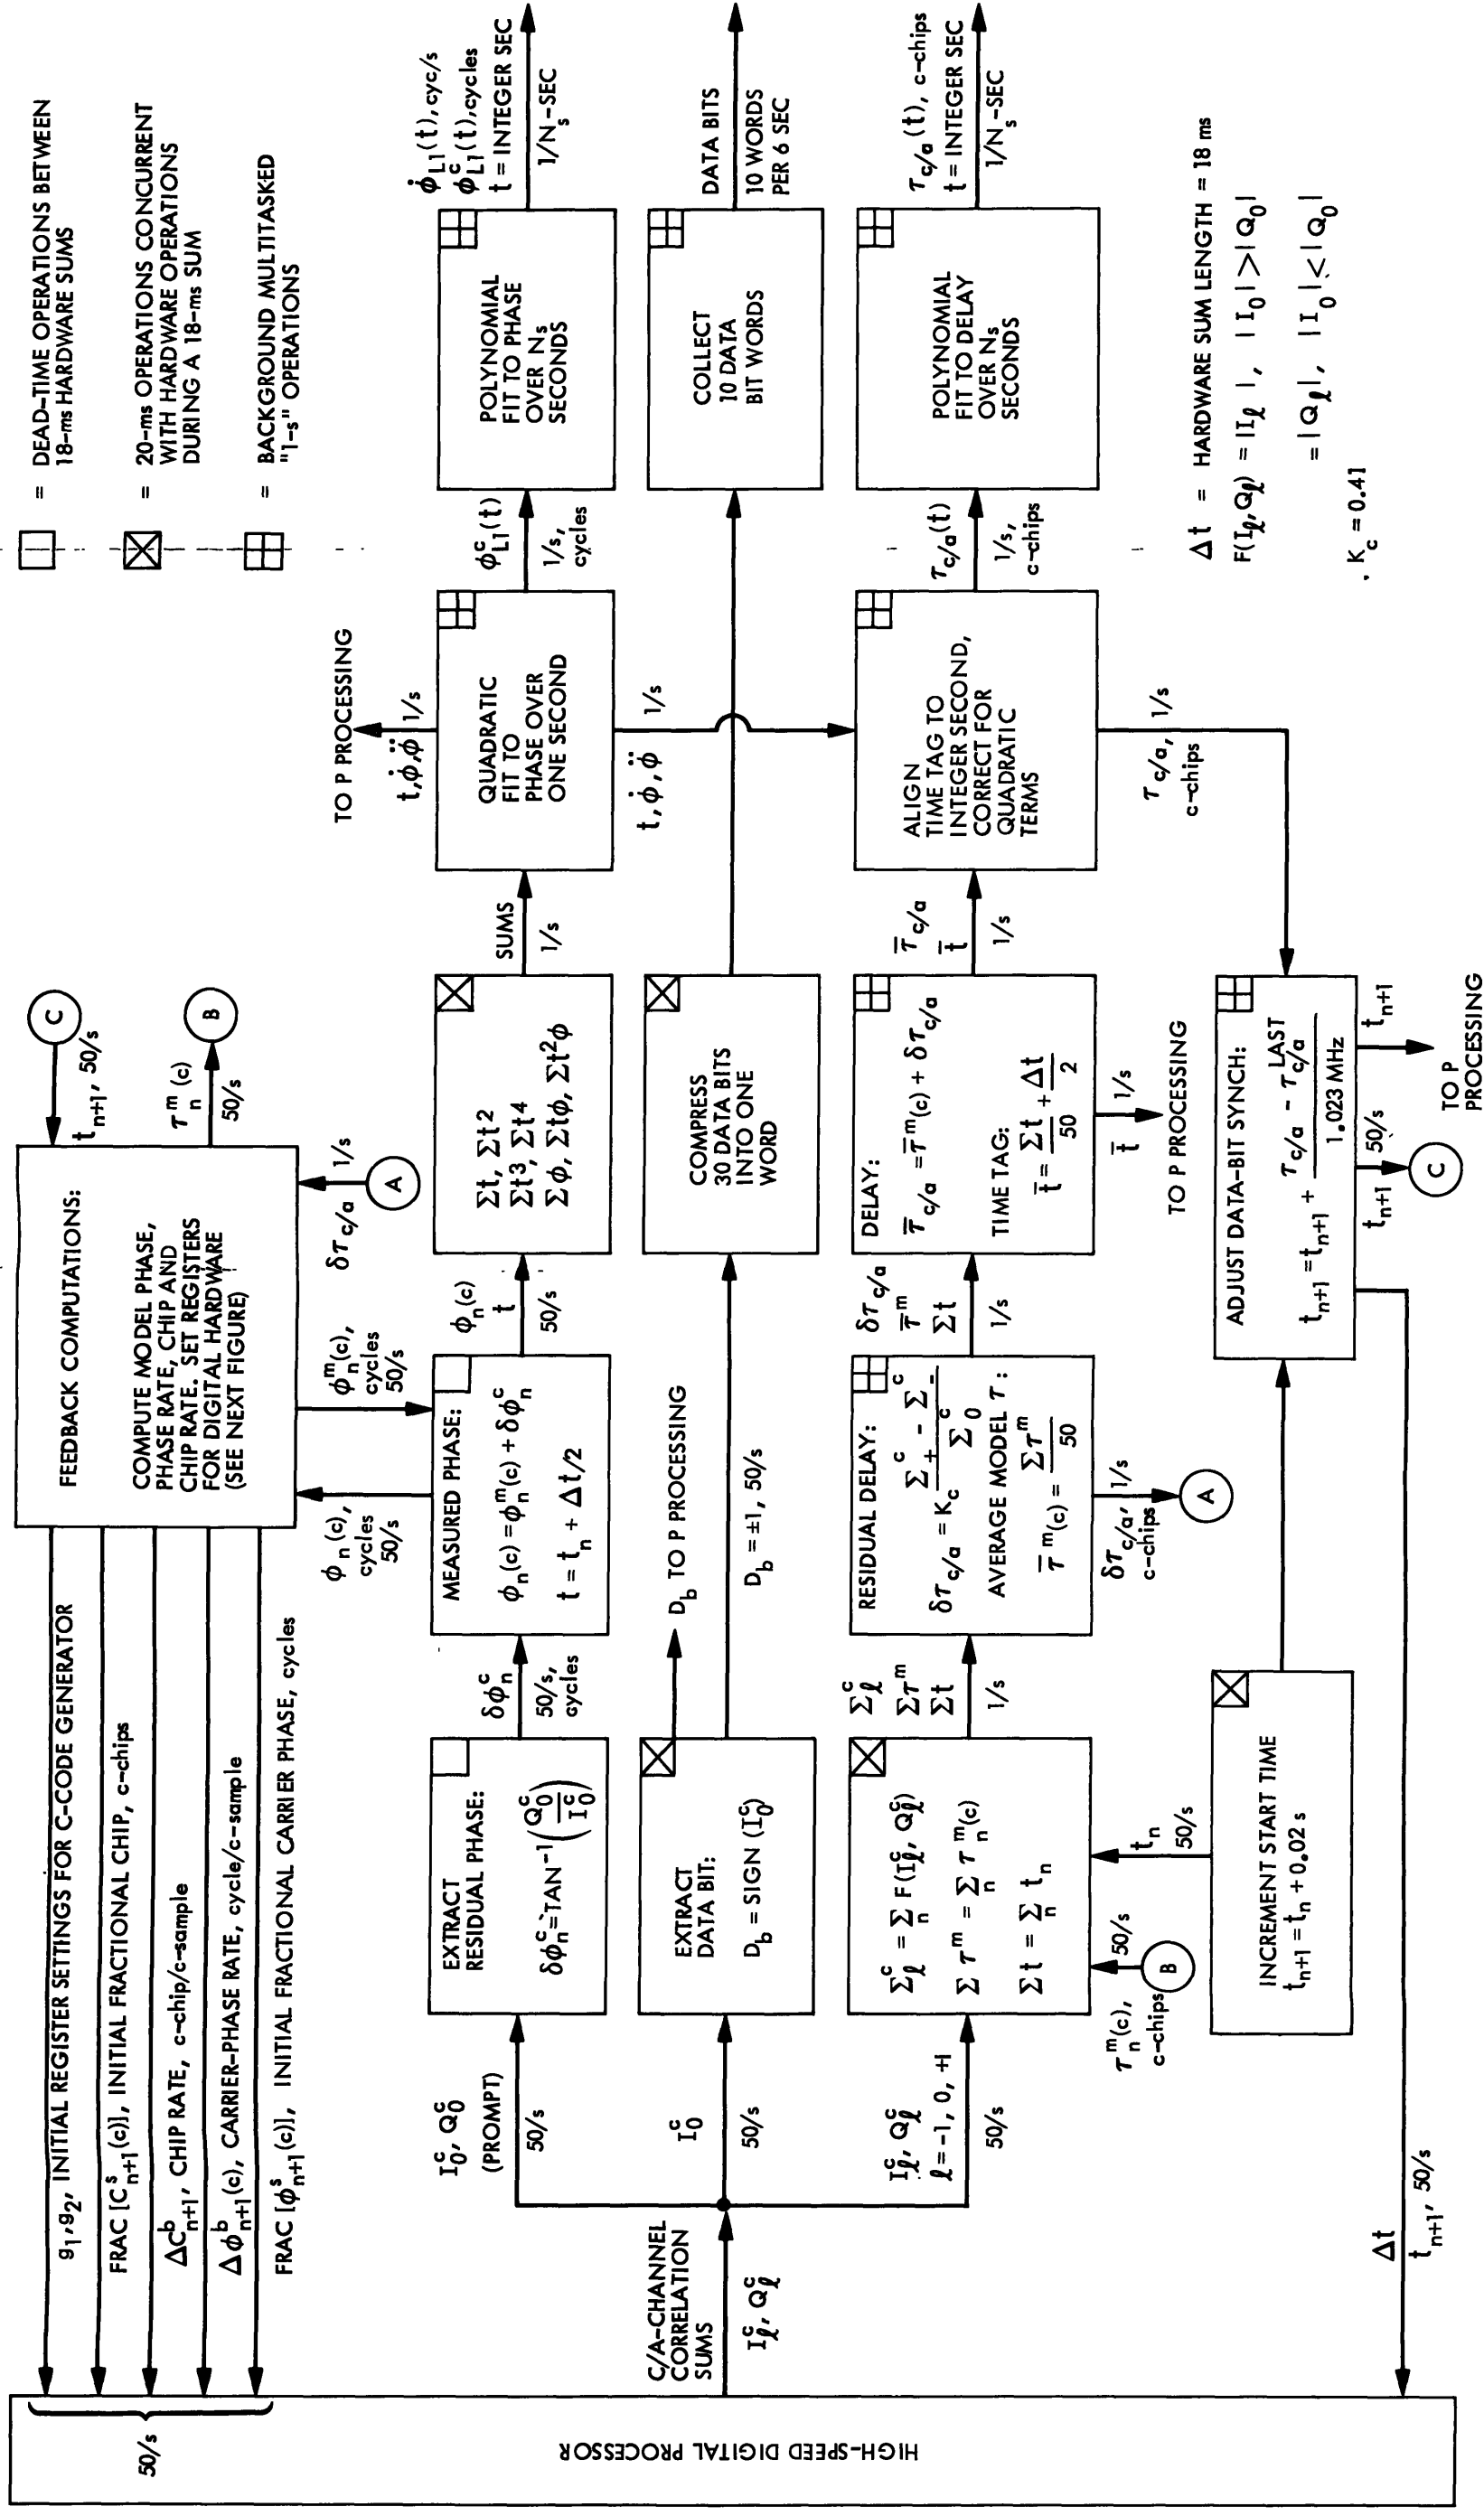
Figure 6-2. In-Lock Signal Processing in the Tracking Processor for the C/A Channel

value that is at interval center and is an integer multiple of N_s seconds. If point positioning has been carried out to find the offset between receiver clock and GPS time, a time shift equal to that offset is applied in the fit so that the output time tag will refer to GPS time. The user can select both the fit interval (N_s = 1, 2, ... or 300 s) and the order of the polynomial (order $\leq$ 5). Estimated parameters in this fit are phase and phase rate at the reference time.

Not shown in Figure 6-2 is a final correction needed to remove from measured phase and phase rate the effect of the known frequency offset introduced during down conversion (i.e., 12,320 Hz for L1). To remove the effect from the rate observable, the sign of measured phase rate is reversed and the positive LO offset is subtracted. For the phase observable, one reverses the sign and then subtracts a linear phase function with a positive slope equal to the LO offset. To set the arbitrary initial phase of this linear LO phase function, its first point in the track is set to zero. After subtracting this LO function, the integer-cycle part of the first corrected phase value of the track is extracted and subtracted from all corrected points. The resulting phase values are equal to negative integrated Doppler, with the first point set to the measured fractional cycle and with no cycle ambiguities between time points. The resulting phase rate and phase values become the output for the C/A channel. Phase delay and phase-delay rate, each with the same sign convention as group delay, can be obtained from the output phase and phase rate by dividing by the L1 carrier frequency (1575.42 MHz).

6.2.1.2   Data-Bit Processing. As shown in the second path of signal processing, the data bit is extracted by testing the sign of the prompt in-phase component of the C/A channel. As data bits are collected, they are compressed into words with 30 data bits per 32-bit word. Every 6 s, 10 of these data-bit words are produced.

6.2.1.3   C/A Group-Delay Processing. The third path of C/A processing extracts group delay. In this path, three types of information are subjected to a 1-s sum-and-dump operation: start times, model delays, and correlation-sum amplitudes. The algorithms employed along this path are designed to reduce computation time and, simultaneously, to increase SNR on the correlation output. The sums for time and model delay are used later for computing average values. As indicated in Figure 6-2, early, late, and prompt correlation sums are each accumulated using a compromise algorithm that is extremely simple computationally. For each of these three lags, if the magnitude of the in-phase component is larger than the magnitude of the quadrature component (which is typical), then the in-phase magnitude is summed. Otherwise the quadrature magnitude is summed.

In high-dynamic cases (e.g., > 2 g), this compromise algorithm can cost these amplitude sums as much as $\sqrt{2}$ in voltage SNR relative to optimal processing. The SNR loss in high-dynamic cases does not affect tracking reliability since lock depends on C/A phase, but precision of measured delay can be reduced. Simplicity of the algorithm, its good performance under low dynamics, and its respectable performance under high dynamics make the algorithm an attractive choice for the 1-s averages. Another obvious option is to sum the 20-ms amplitudes, computed as the root-sum-square (RSS) of the in-phase

and quadrature components. This option is considerably more complicated computationally, but is better for high-dynamic cases. Other algorithms can be used, but they will not be discussed here.

Once per second, the sums are analyzed to calculate residual delay, average model delay, and average start time by means of the formulas shown in Figure 6-2. The resulting residual delay is sent to the feedback portion of the code as a slow delay feedback. Residual delay is also combined with average model delay to obtain total measured C/A group delay. The time tag for this delay is the average start time plus 1/2 of the "18-ms" sum interval.

The resulting total delay observable, which is an average delay for the 1-s interval, requires two corrections to produce the instantaneous delay at midinterval. Based on the phase rate from the 1-s fit to C/A phase, the first correction shifts the time tag from the average time tag to the integer-second time tag, t. Based on the phase-rate rate from the same fit, the second correction removes the bias in delay caused by averaging the quadratic term (delay-rate rate term) over one second.

The last step in processing delay is to compress data volume further by fitting the 1-s values with a polynomial time function over N_s seconds, which leads to one output delay value each N_s seconds. The reference time (i.e., time tag), fit interval, and polynomial order in each fit conform with the C/A phase fit described above.

6.2.2   C/A Feedback Loops

Separate but interconnected feedback loops are implemented for C/A phase and for C/A group delay. Since C/A phase controls the fast feedback for all other loops, that loop will be explained first.

6.2.2.1   C/A Phase Feedback. The phase-lock loop in the current version of the Rogue software is based on an extrapolation technique that employs a simple algorithm to form a linear model based on total measured phase of the two most recent intervals. As illustrated in Figure 6-3, the feedback loop for the C/A phase consists of several steps. First, during the dead time between 20-ms intervals, total carrier phase measured for the center of the (n-1)th 20-ms interval (the previous interval completed) is subtracted from the total measured phase for the center of the interval just completed, the nth interval, to obtain an estimate for the most recent phase change per 20 ms, the nth. To predict model phase for the center of the next 20-ms interval, the (n+1)th interval, this nth phase change is added to the total measured phase at the center of the nth interval. The resulting projected center-interval value is saved to be combined later with the residual phase of the (n+1)th interval.

As shown in Figure 6-3, an adjustment is applied to the (n+1)th center-interval value to obtain the (n+1)th start phase. This adjustment is computed as the phase change per sample times the number of samples in half of an 18-ms interval. The phase rate used in this calculation must be the rounded rate used to set the H/W rate register for the (n+1)th interval. As explained below, the (n-1)th 20-ms phase change is used in the calculation of the rate for the (n+1)th interval.

Figure 6-3. Feedback Loop for the C/A Channel in the Tracking Processor

6-9/6-10

FOLDOUT FRAME

The fractional-cycle part of the extrapolated start phase obtained in this manner is extracted, multiplied by 2^{24}, and then rounded to the nearest integer. The resulting integer is sent to the high-speed hardware to initialize starting phase register for both of the quadrature channels for C/A.

The 20-ms phase change estimated above is also used to generate the phase-rate register for the digital hardware, except that the previous phase change,the (n-1)th, is used rather than the nth. This approximation allows the rate operations to be performed concurrently with digital hardware operations rather than during the dead time between 20-ms intervals. Phase change per sample is calculated by dividing the 20-ms quantity by the number of samples per 20 ms. To obtain the feedback for the high-speed hardware, phase change per sample (in cycles per sample) is multiplied by 2^{24} and rounded to the nearest integer. The resulting integer is sent to the high-speed hardware to initialize the phase-rate registers for both of the quadrature channels for C/A.

Every second, the C/A SNR is checked to verify signal strength is adequate for reliable tracking. When the 20-ms voltage SNR falls below 7 for any satellite, the track for that satellite is terminated.

In a future version of the tracking processor software, the C/A phase-lock loop will be made more versatile by providing an optional conventional loop that will track to lower SNR values, but with lower dynamic capabilities.

6.2.2.2 <u>C/A Chip Feedback</u>. Fast feedback for C/A delay and chip is based on L1-C/A phase, as illustrated in Figure 6-3. First, the (n-1)th 20-ms phase change obtained from C/A phase is converted to phase delay change by removing the effect of the LO offset (12320 Hz for L1 in the prototype receiver) and then dividing by the number (1540) of L1 cycles per C/A chip. This approximate 20-ms delay change is then used to determine for the next 20-ms interval the model delay and model chip as well as the change per sample in these quantities.

The delay change per sample is estimated by simply dividing the 20-ms delay change by the number of samples per 20 ms. To obtain the total chip change per sample, this delay change per sample is subtracted from the chip change due to the passage of time (i.e., 1.023 MHz/1.5374 MHz). The result is then multiplied by 2^{24} and rounded to the nearest integer to produce the integer value that the digital hardware expects for the chip rate. The rate register for both of the quadrature channels for C/A receives this same integer.

To estimate phase delay at the center of the next 20-ms interval (n+1), the (n-1)th 20-ms delay change calculated above is added to the phase delay for the present 20-ms interval (n). The model C/A group delay for the center of the the (n+1)th interval is then obtained by adding the slowly varying delay offset for C/A to this (n+1)th model phase delay. The resulting model group delay is saved to be summed over the current 1-s interval in preparation for model restoration.

Once each second (once per 50 feedbacks), the delay offset is subjected to a slow feedback adjustment in the form of the residual delay obtained from C/A group delay processing, as described above. This adjustment corrects for the slow drift of group delay relative to phase delay (typically <2 cm/s).

The starting value for the model chip for the (n+1)th interval is obtained by multiplying 1.023 MHz times the interval start time and subtracting the model group delay (in units of c-chips) computed above. As shown in Figure 6-3, an adjustment is needed in this calculation to shift the model delay from the center to the start of the 20-ms interval. This shift is equal to the delay change per sample, estimated above, times 1/2 the "18-ms" sum interval. An additional correction is applied to the starting chip value to cancel the anticipated error buildup caused by roundoff in chip rate. This last correction is not shown in Figure 6-3.

The resulting starting value for the model chip is then separated into integer and fractional parts. The fractional part is multiplied by 2^{24} and rounded to obtain the integer value sent to the digital hardware as the initial value in the chip advancer. The integer part is used to obtain (by table lookup) the g_1 and g_2 registers needed to initialize C/A code generator to the first (integer) chip. Accurate data-bit synchronization in the start times keeps the table lookup range small and near the beginning of the full 0-1022 range of possible starting chips (calculated modulo 1023). The registers of high-speed hardware for both quadrature channels for C/A are initialized with the same values for integer and fractional chip.

6.3   P-CHANNEL PROCESSING

A block diagram summarizing post-correlation processing for a P channel (either P1 or P2) is shown in Figure 6-4. Signal processing will be described first, followed by feedback operations.

6.3.1   In-Lock P Signal Processing

After collecting the P-channel correlation sums, the signal processing follows two paths: one path extracts phase, while the other extracts delay. In both cases, the correlation sums are accumulated for 1 s to collect statistics and reduce the computational load. The long integration time is possible due to the accurate short-term feedback derived from L1-CA phase.

6.3.1.1   P Phase Processing. Along the phase route, the in-phase and quadrature components of the prompt lag are summed after a sign correction based on the data-bit sign extracted for the same 20-ms interval from the prompt lag of the C/A channel. Model phase is summed along with the correlation sums. Each second, the sums are collected and the sum registers are set to zero in preparation for the next 1-s sum.

Figure 6-4. In-Lock Signal Processing in the Tracking Processor for a P Channel

6-13/6-14 FOLDOUT FRAME

Residual phase is extracted from the 1-s sums by applying an arctangent operation to the in-phase and quadrature components. Due to the data-bit sign correction, a four-quadrant arctangent can be applied with a range of -1/2 to +1/2 cycle. Average model phase is computed as the straight average of the 50 model phases in the 1-s interval.

As shown in Figure 6-4, measured residual phase is added each second to a phase-offset variable that keeps track of the offset and slow drift of P phase relative to C/A phase. Separate offsets are tracked for the P1 and P2 channels. The use of offset phase in feedback will be discussed below. Total measured phase, which is effectively an average across the 1-s interval, is computed as the sum of residual phase and averaged model phase. The time tag for the phase observable is the average time tag computed in the C/A channel. As with C/A group delay, two corrections are applied to obtain the instantaneous phase at interval center. To align the time tag with the integer-second time tag at the center of the 1-s interval, a correction based on the "1-s" phase rate from the L1-C/A channel is applied. In addition, a correction based on C/A phase-rate rate is applied to remove any bias caused by averaging quadratic phase variation. For L2, both of these corrections are scaled by RF frequency.

After these corrections, the phase values are fit with a polynomial over N_s seconds to produce one phase value and one phase-rate value for each N_s seconds. The fit interval length, time tag, and polynomial order conform with the fit to L1-C/A phase. As with C/A-phase, the effect of the known LO offset is removed to produce the output values for phase and phase rate.

6.3.1.2   P Group-Delay Processing. Along the delay processing path, the correlation-sum amplitudes for the early, prompt, and late lags are accumulated over 1 s, as are the associated model delays. The algorithm for summing P amplitudes is the same as the one used for C/A. Each second, the 1-s sums are collected to compute the residual delay according to the formula in Figure 6-4. The resulting residual delay is added to a delay-offset variable that keeps track of the offset and slow drift of P group delay relative to C/A phase delay. Separate offsets are tracked for the P1 and P2 delays. The use of delay offset in the feedback loop is discussed below.

Total measured group delay, which in effect has been averaged over 1 s, is computed as the sum of the residual delay and the average model delay. Two corrections are then applied to this delay in a manner analogous to the C/A-delay procedure to align the time tag and to remove the effect of averaging over 1 s. After these corrections, the P delays are fit over N_s seconds in the same manner as the other observables to produce the output values.

6.3.2   P Feedback Loops

In analogy with the C/A channel, the P feedback consists of two components, phase feedback and delay (chip) feedback.

6.3.2.1   P Phase Feedback. In the P feedback loop, the first step is to estimate fast feedback phase at the center of the next interval, the (n+1)th, by first scaling the C/A feedback phase by RF frequency and then adding the offset phase that accounts for offset and slow drift of P phase relative to L1-CA phase. The resulting center-interval model phase is saved to be averaged over 1 s. To obtain starting phase, an adjustment is applied to shift the center phase to the start time, as carried out for C/A phase. The fractional-cycle part of the resulting phase is then multiplied by 2^{24} and rounded to obtain the integer value needed to initialize the phase advancer of the digital hardware.

Phase-rate feedback for P is obtained by scaling the (n-1)th C/A phase change per sample according to RF frequency and then scaling by 10 to account for the higher P sample rate. The resulting phase rate is then multiplied by 2^{24} and rounded to obtain the integer value needed to initialize the phase-rate register in the phase advancer in the digital hardware.

6.3.2.2   P Chip Feedback. Estimation of P-chip-rate feedback is very simple since the C/A chip rate is the same as the P chip rate (i.e., p-chip/p-sample = c-chip/c-sample). To obtain the integer value needed to set the rate portion of the P-chip advancer in the digital hardware, however, the chip rate is multiplied by 2^{28} rather than 2^{24}. A higher multiplier is required for the P channels due to the increased effect of rate roundoff for the higher P sample rate.

Model group delay for the P channel at the center of the next 20-ms interval, the (n+1)th, is computed by multiplying the C/A phase delay for the (n+1)th interval by 10 to account for chip size and then adding the aforementioned delay offset (P1 or P2) that accounts for the offset of P group delay from C/A phase delay. This model delay is saved to be averaged over 1 s in preparation for model restoration. Model chip feedback, the projected chip value at the start of the (n+1)th interval, is computed by multiplying 10.23 MHz by the start time and subtracting the model group delay at the start of the interval. In analogy with the C/A chip feedback, corrections are needed to shift the model delay from center interval to the start time and to account for chip rate roundoff.

The resulting chip value is separated into integer and fractional parts. From the integer part are derived seven register values that initialize the P-code generator so that it will produce the correct initial chip sign. The fractional-chip part is multiplied by 2^{28} and rounded to obtain the integer value that sets the fractional-chip portion of the chip advancer in the digital hardware.

6.4   CARRIER-AIDED AVERAGING OF GROUP DELAYS

Receiver design provides another option for averaging the 1-s group delay values over N_s seconds, a method that avoids the increase in the error in delay caused by a polynomial fit, but requires concurrent measurement of both L1 and L2 carrier phases. (For example, the least-squares delay error resulting from a quadratic fit is greater than the error resulting from a

straight average by a factor of 1.5.) The new method would replace the last step in Figure 6-4, the polynomial fit step.

As shown in Figure 6-5, the 1-s phase values (ϕ_L^P) from the P1 and P2 channels (see the input values in Figure 6-4 provided to the polynomial fit step) are first adjusted to remove the LO rate offset (see Section 6.2.1.1) and are then converted to phase delay (L1 phase from the C/A channel could replace L1 phase from the P1 channel in this process). These phase values are then combined, as indicated in Figure 6-5, to produce, to within an unknown additive constant, very precise models for the time variations of the group delays from the P1 and P2 channels. The effect of this combination is to reverse the sign of the ionospheric component found in phase delay so that the resultant delay will have the right group-delay sign for all delay components, i.e., the ionosphere as well as the geometric and tropospheric components. The resultant model delays, which follow the time variation of the group delays with "phase-delay" accuracy, will be referred to as synthetic group delays. The presence of the unknown additive constant in the synthetic delay causes no problem since that constant component is implicitly subtracted and then implicitly added back in the processing.
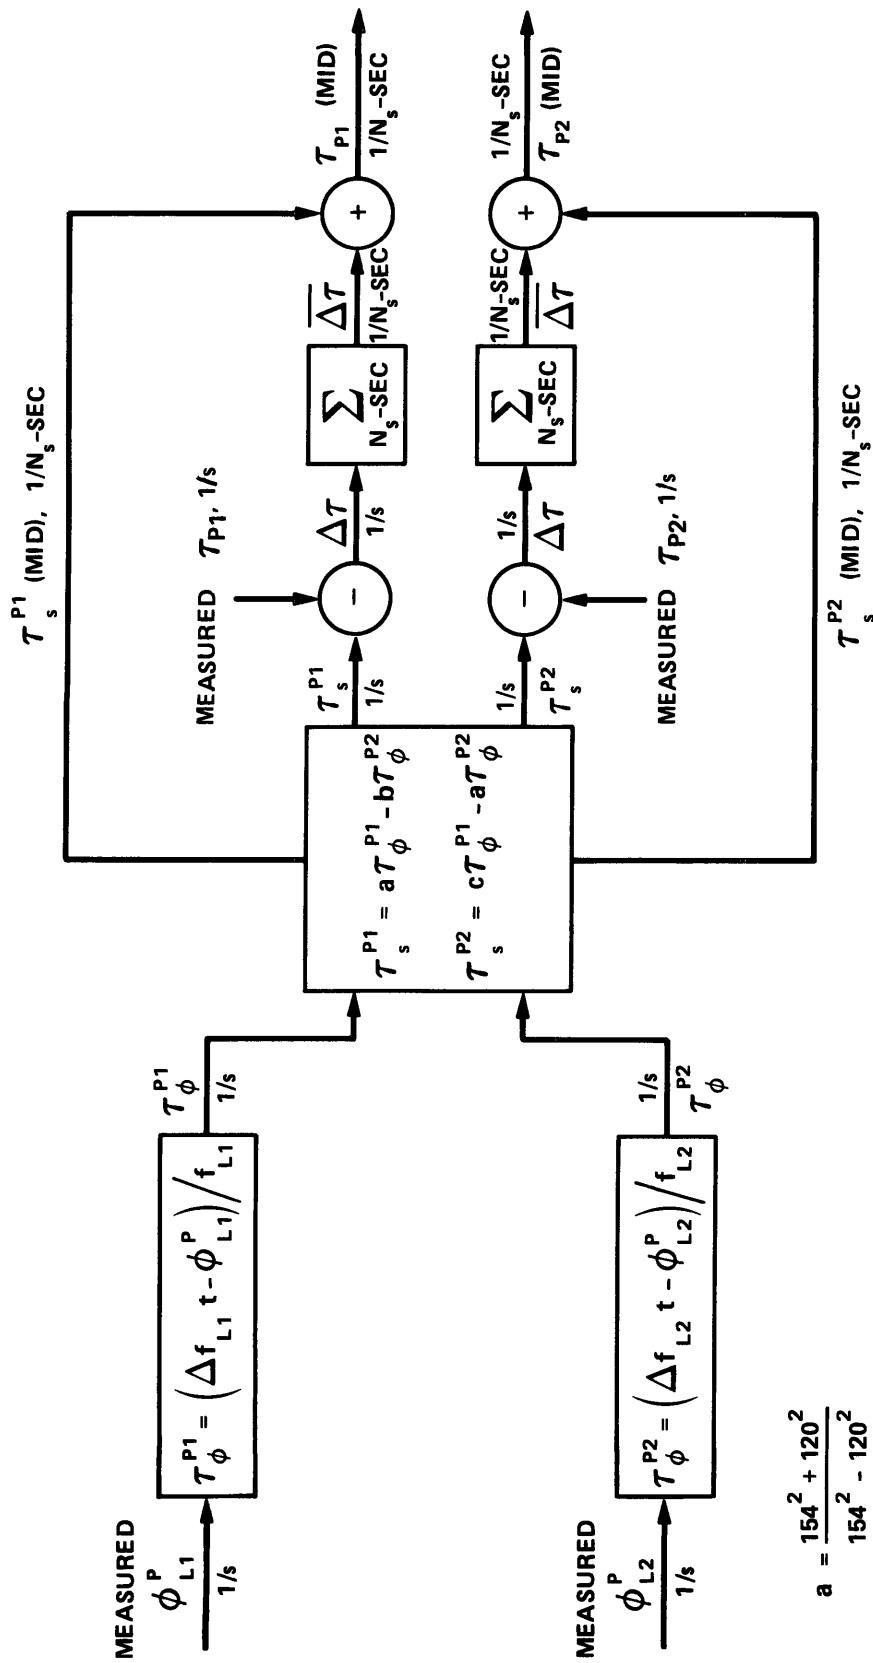
Each second, the synthetic group delays are subtracted from the measured P group delays (the 1-s values, τ_p, from Figure 6-4) to form residuals that are constant with respect to geometric, ionospheric, and tropospheric variations. These residual delays are then summed over an N_s interval to form an average residual delay for the interval. To obtain the final total measured group delays for the interval, the average residual delays are added to the synthetic group delays obtained at the middle of the interval. Alternately, instead of the midinterval value, one can fit the synthetic delays across the interval with a parameterized function such as a polynomial and obtain a least-squares estimate for the synthetic delay at the middle of the interval. Even though it is not illustrated, this carrier-aided averaging method could be applied in the same manner to the C/A group delay.

6.5   OUTPUT DATA

For each output interval for each satellite, the receiver output consists of data bits, a time tag and the observables of group delay, phase and phase rate for each channel (L1-C/A, L1-P, L2-P). The time tags fall on integer seconds of GPS time and are separated by an interval of 1, 2, 3,....., or 300 s. The delay and phase observables can be regarded as "instantaneous" values measured at the time indicated by the time tag, which is at the center of the output interval. For each channel, the data also include the 1-s voltage SNR, averaged over the output interval, and system noise error for delay and phase. As a measure of tracking performance, maximum and RMS tracking errors for each output interval are passed for both phase and delay for all channels.

6.6   ERRORS INTRODUCED BY THE TRACKING PROCESSOR

All operations in the tracking processor are carried out with sufficient precision to reduce roundoff errors to a negligible level. For

Figure 6-5. Carrier-Aided Averaging of Group Delays

$$\tau_\phi^{P1} = \left(\Delta f_{L1} t - \phi_{L1}^P\right)/f_{L1}$$

$$\tau_\phi^{P2} = \left(\Delta f_{L2} t - \phi_{L2}^P\right)/f_{L2}$$

$$\tau_s^{P1} = a\tau_\phi^{P1} - b\tau_\phi^{P2}$$

$$\tau_s^{P2} = c\tau_\phi^{P1} - a\tau_\phi^{P2}$$

$$a = \frac{154^2 + 120^2}{154^2 - 120^2}$$

$$b = \frac{2 \times 120^2}{154^2 - 120^2}$$

$$c = \frac{2 \times 154^2}{154^2 - 120^2}$$

Earth-fixed receivers, the models behind the tracking processor algorithms contribute very small errors to phase and delay (< 1 mcycle for phase and < 0.3 cm for P delay).

6.7   CHECKS FOR CYCLE SLIPS

If the C/A channel loses a half cycle at a given point, that channel can sometimes show a rapid divergence of phase delay and group delay, with phase delay linearly deviating from the group delay at a rate equivalent of 25 Hz at L band. Thus, by comparing the change in measured phase delay with the change in measured group delay relative to the start of the track, one can detect such a half-cycle slip. Since the deviation must exceed the maximum ionospheric difference expected for the two delay types, one must wait about 20 s for the difference to become significant. This check for a half-cycle slip is included in the prototype receiver.

Another check for cycle slips can be made on the basis of data bits. Every 6 s, the data bits include a known synchronization pattern (8B in hexadecimal notation) that is eight bits in length. Once data-bit synchronization and clock synchronization have been carried out, the location of each synchronization pattern can be exactly predicted. When a half-cycle slip occurs, the "sign" of this synchronization pattern will reverse. Thus, by checking the sign of the synchronization pattern every six seconds, one can determine if a half-cycle slip (or odd number of same) has occurred in the last six seconds. This check for half-cycle slips has been included in the prototype receiver.

6.8   ACQUISITION

The fast-acquisition scheme for the Rogue receiver is a simple version that requires relatively accurate values for delay and delay rate. It is assumed that the delay rate value is sufficiently accurate (error < 3 ns/s) to eliminate the need for a search over delay rate. With the delay rate fixed at the input value, the software searches for the C/A signal over a range (100 μs) of delay values around the input delay value. The maximum time required to search this range is about 1.5 s for the present implementation. After the peak in C/A amplitude has been detected, L1 phase is extracted and C/A group delay is derived from the dependence of amplitude on lag. This improved value for delay is then used for the next 20-ms interval. The output of that next interval is analyzed to further improve the delay value and to obtain a second value for L1 phase. The two values for L1 phase are subtracted and divided by the time separation to produce a very accurate estimate of phase rate and delay rate. Based on these new accurate estimates for delay, delay rate, phase and phase rate, lock can be initiated and maintained for subsequent 20-ms intervals.

With the C/A channel now running in lock in the background, the P channels are set up on the basis of the new C/A information. Because this information is accurate, the estimated delays for the P channels are sufficiently accurate to limit the search in delay to a few P lags. When the signals for the P channels are found, the delay and phase values are extracted and lock is initiated for those channels.

Special software for resolution of the 1-ms ambiguity and for data-bit synchronization is not necessary for this scheme because of the presumed accuracy of the a priori delay. Data-bit synchronization is initiated and maintained by offsetting sum start time, as discussed in Sections 6.2.1.

A future version of the Rogue receiver will allow a "sky" search for satellites, involving a search over all possible delays and delay rates. This acquisition scheme will take longer but will require no a priori values for delay and delay rate.

SECTION 7

DISTINCTIVE FEATURES OF ROGUE SIGNAL PROCESSING

The signal processing design described in this report possesses a number of distinctive features.

7.1 PURELY DIGITAL OPERATIONS

The purely digital design of the baseband processor, as discussed in Section 2.1, eliminates unpredictable errors arising from semi-analog steps found in some "digital" designs. Consequently, analysis of roundoff and quantization errors can provide exact assessment and virtual elimination of the errors contributed by the baseband processor. Because the same input bits should always produce the same output bits, this feature also provides an exact method for testing the health of a baseband processor through repeatability tests and comparisons of different satellite channels.

7.2 CHIP ADVANCER

The chip advancer described in Section 5.4 provides a purely digital method for generating a pseudorandom code that has no distortion in amplitude and is delayed with a cumulative 20-ms accuracy better than 100 nanochips. The output of the code generator can be viewed as exact equi-spaced discrete samples of a piecewise-constant function of time consisting of a sequence of chips, with virtually instantaneous transitions at chip sign changes. With the chip advancer, errors in delay caused by commensurability of sample rate and chip rate can be reduced to a negligible level through use of a sample rate that is effectively incommensurate with the chip rate.

7.3 FAST FEEDBACK LOOPS DRIVEN BY L1-C/A PHASE

Another distinctive feature of the tracking processor is the feedback approach: C/A phase drives all fast feedback loops. Several important advantages are gained by this approach. Since the C/A signal is the strongest, the P channels will maintain phase lock to a lower SNR value than separate P feedback loops would. Because of the 1-s integration time for the P channels, those channels are strengthened by a factor of 50 in power SNR relative to a 20-ms integration time. This improvement exceeds the difference in 20-ms SNRs between C/A and P2, even if C/A is quadrature and P2 is not. Thus, as signal strength decreases, loss of lock for the P2 channel will be determined by C/A lock. Based on the nominal power levels of Block I satellites and equal antenna gains for L1 and L2, this provides a 6-dB improvement in loss-of-lock SNR relative to an independent phase-lock loop for the P2 channel.

Another benefit of this approach is that system-noise errors in feedback for C/A and P group delays are reduced. For all feedback delays, the system-noise error is reduced by a factor of 7 relative to 20-ms delay-locked loops since the slow delay feedback is averaged over one second and since the C/A phase projection introduces virtually no system-noise error in delay.

Besides providing much more stable loops, this greatly reduces delay feedback error which, in turn, greatly reduces the error due to nonlinearity in loop response. With this approach, loop-response errors can be reduced to considerably less than 0.1 millichip (0.3 cm for P).

Another benefit of C/A-based feedback is that, for some applications, quadrature implementation would not be needed for the P channels. With fast feedback based on C/A phase, signal lock for the P channels would not be improved by adding the P-quadrature component even though the final output P-phase and P-delay precision would be improved by $\sqrt{2}$. Applications that demand good lock but do not require ultimate precision could implement quadrature only on C/A. Avoidance of quadrature implementation on the P channels substantially reduces the complexity and cost of both front end and baseband hardware.

7.4   C/A PHASE-LOCK LOOP

Design of the phase-lock loop for the C/A channel was based on the goals of high accuracy and very reliable lock, even under relatively high dynamics. The simple extrapolation technique initially adopted to meet these goals makes use of the relatively high SNR that should be obtained by the Rogue receiver in expected applications. As described in Section 6.2.2.1, feedback phase is obtained by means of a two-point linear extrapolation of the total phase measured for the two most recent 20-ms intervals. Extensive tests with Earth-fixed Rogue receivers using this PLL under normal conditions indicate that even one cycle slip is unlikely over many days of multi-satellite tracking.

Tests and theory indicate that high dynamics can be reliably and accurately tracked by this feedback technique. At the upper end of the expected SNR range (see Section 2.2), the C/A channel can track phase dynamics up to about 300 Hz/s (6 g) and 1000 Hz/s (20 g), with feedback intervals of 20 and 10 ms, respectively. (With a noiseless system, these limits become 625 Hz/s and 2500 Hz/s, respectively.)

In a future version of the tracking processor software, the C/A phase-lock loop will be made more versatile by providing an optional conventional PLL, with selectable gain, which will track to lower SNR values, but with lower dynamic capabilities.

7.5   P-CHANNEL DATA-BIT REMOVAL

The data-bit sign extracted from the C/A channel is used to remove the data-bit sign from both of the P channels (see Section 6.3.1.1). Provided the C/A sign is correct, this procedure expands the phase ambiguity range for the P channels from 180 to 360 degrees, thereby greatly reducing the possibility of cycle slips for those channels.

7.6   ANALYTIC DATA-BIT SYNCHRONIZATION

A simple and reliable analytic method is used to align the "18-ms" sum with the data-bit edges with an accuracy better than 10 µs. (See

Sections 2.2 and 6.2.1.) Special circuitry and software for data-bit synchronization are not needed during in-lock operation.

7.7   CYCLE COUNTING

Cycle counting and removal of phase ambiguities are carried out implicitly by the tracking processor. The simple technique used for cycle counting has proven to be very reliable in the prototype system. Carrier-cycle counting is implicitly carried out when total phase is accurately projected ahead to the next interval. That is, projected (model) phase analytically "keeps track" of the integer cycles in the integer part of the phase variable. This approach for cycle counting avoids the special cycle-count circuitry found in some receiver designs. For this process to work, projected phase for the C/A channel must be in error by no more than 90 degrees, which is a condition readily met for wide ranges of SNR and dynamics. For the P channels, projected phase must be in error by no more than 180 degrees. As in other respects, cycle counting for the P channels depends on the solid performance of the C/A channel, since the C/A cycle count is scaled to the other channels. Deviations between L1 phase and L2 phase caused by the ionosphere are typically a very small fraction of a cycle over 1 s and are therefore easily tracked by the slow feedback loop without causing cycle slips.

7.8   FLEXIBLE FEEDBACK CONTROL

Substantial flexibility in control is provided by the fact that, for each correlation interval, the high-speed hardware requires values for start time, integration length, phase, phase rate, chip, and chip rate. Since both phase (chip) and phase rate (chip rate) are updated every 20 ms, the buildup of rate roundoff error can be anticipated and nullified.

A multiplexing design that switches a correlator between satellites could be accommodated by this design since all parameters can be reset for each integration interval for each satellite. Fundamental chip rates other than the 10.23 MHz setting for GPS can be easily accommodated by properly setting the chip-rate register. A similar statement can be made for baseband carrier frequency, which can modeled over a wide range.

The correlation interval is variable and can be set to 1, 2, 4, 5, 10, or 20 ms in the software logic. For the prototype Rogue receiver, however, hardware and microprocessor speeds limit the options to 10 and 20 ms. This option allows a tradeoff between SNR and dynamics. For example, maximum allowed acceleration could be increased by as much as a factor of 4 by decreasing the feedback interval from 20 ms to 10 ms.

7.9   WIDE CHOICE OF OUTPUT DATA RATES

By providing a wide choice (1, 2, 3,... or 300 s) of time separations for output data points, the receiver can satisfy a wide range of users. The polynomial fit to obtain these output observables is set up so that the time tag of the estimated observable falls on an integer second (see Section

6.2.1.1). By analyzing measured group delays, the receiver can determine accurately the offset of receiver time from GPS time. When supplied with this offset, the tracking processor can account for receiver clock error and force the integer time tags to correspond to GPS time with an accuracy of 0.1 to 3 μs.

7.10   CARRIER-AIDED AVERAGING OF GROUP DELAYS

A carrier-aided technique for averaging group delays (see Section 6.4) reduces the system-noise error in those averaged delays relative to the error that would be produced by a polynomial fit to the group delays. For example, the system-noise error generated by the carrier-aided technique is 1.5 times smaller than for a quadratic fit to group delays.

7.11   USE OF RESIDUAL PHASE

Measured phase is computed as residual phase (tracking error) plus model phase (feedback phase) at interval center (see Sections 6.2.1.1 and 6.3.1.1). This method for computing measured phase is unlike conventional phase-lock loops, which take model (NCO) phase as measured phase. By doing so, such loops ignore the most recent error signal (residual phase), which represents the offset of the most recent feedback phase from true phase. Therefore, in applications where tracking errors are substantial and where SNR is adequate, conventional loops will produce less accurate phase measurements. (Given adequate SNR, the use of an arctangent phase extractor further improves the accuracy of this approach in dynamic applications, since an arctangent extractor is more accurate over the full unambiguous range of residual phase than a sine extractor.) The same approach is used to obtain measured group delays.

7.12   TRACKING CAPABILITY

The Rogue receiver can track up to eight satellites, without time multiplexing between satellites or channels, simultaneously measuring both group delay and phase delay from each of three channels (L1-C/A, L1-P, L2-P).

7.13   ACCURACY

Theory and tests to date indicate that phase and delay can be accurately measured by the Rogue receiver. Excluding system-noise errors and errors introduced outside of the receiver (e.g., multipath and antenna errors), intersatellite errors, as defined in Section 2.1, are about a millicycle (0.02 cm) or less for phase and about a centimeter or less for P group delays. Based on the system parameters specified in Section 2.1, system-noise errors are less than 0.01 cm in phase delay and a centimeter or less for P group delays, given a 5-min integration time.

SECTION 8

SUMMARY

This report has presented a functional description of signal processing in the Rogue GPS receiver and outlined some distinctive features relative to other designs. Processing steps from RF to the final output values for delay, phase, and data bits have been described. Subsequent reports are planned to present signal processing theory, system design and integration, error analysis, design and implementation of software (real-time multitasking, signal processing, control, interfacing) and of hardware (high-speed digital processor, tracking processor, front end, frequency system), and results of laboratory and field tests.

1. Report No. 88-15	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Functional Description of Signal Processing in the Rogue GPS Receiver		5. Report Date June 1, 1988
		6. Performing Organization Code
7. Author(s) J.B. Thomas		8. Performing Organization Report No.
9. Performing Organization Name and Address JET PROPULSION LABORATORY California Institute of Technology 4800 Oak Grove Drive Pasadena, California 91109		10. Work Unit No.
		11. Contract or Grant No. NAS7-918
		13. Type of Report and Period Covered JPL Publication External Report
12. Sponsoring Agency Name and Address NATIONAL AERONAUTICS AND SPACE ADMINISTRATION Washington, D.C. 20546		
		14. Sponsoring Agency Code RE 211 BG-314-40-50-55-09

15. Supplementary Notes

16. Abstract

Over the past year, two Rogue GPS prototype receivers have been assembled and successfully subjected to a variety of laboratory and field tests. This report presents a functional description of signal processing in the Rogue receiver, tracing the signal from RF input to the output values of group delay, phase, and data bits. The receiver can track up to eight satellites, without time multiplexing among satellites or channels, simultaneously measuring both group delay and phase for each of three channels (L1-C/A, L1-P, L2-P). The Rogue signal processing described in this report requires generation of the code for all three channels. Receiver functional design, which emphasized accuracy, reliability, flexibility, and dynamic capability, is summarized. A detailed functional description of signal processing is presented, including C/A-channel and P-channel processing, carrier-aided averaging of group delays, checks for cycle slips, acquisition, and distinctive features.

17. Key Words (Selected by Author(s)) electronics and electrical engineering	18. Distribution Statement unlimited distribution	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 41	22. Price

JPL 0184 R 9/83