5. Theory of Servo Operation

5.1 Overview of Servo Concepts

The RCP8 provides two independent, and nearly identical, motion servos for the azimuth and elevation axes of the radar antenna. These servos are implemented digitally within the RCP8 microprocessor. This chapter describes the operational theory of the antenna servos and fail safe algorithms as well as how the theory relates to the various setup parameters, as described in Chapter 4.

In this chapter:

The servo software takes, as input, the digital antenna position and analog tachometer velocity and provides, as output, an analog drive signal for the motor power amplifiers. The interface between the processor and the tachometer and drive signals is made using 12-bit A/D and D/A converters. The servo software is periodically scheduled at 10 millisecond intervals and, in principal, has the capability of controlling antennas that have a significant response—such as 20Hz. In practice, most weather radar antennas are much more sluggish than this. Aside from the presence of limit switches in the elevation axis, the two servos are identical both in the configuration and in the operation.

There are three ways the RCP8 servos can operate:

1. Open loop
2. Velocity servo
3. Position servo

The Open Loop is not really a ”servo” at all. This simply applies fixed drive levels to the motor to measure the antenna performance. At installation is the only time when Open loop is run—when the antenna's characteristics are measured to set up the actual control parameters. This is a manual procedure that requires the local TTY and is described in Chapter 3.

The velocity and the position servos are interrelated—each mechanism uses parts of the other during normal operation. The velocity servo always runs once either servo is activated. To achieve a particular velocity, the servo is used directly. To achieve a particular position, a non-linear position error is fed into the velocity servo from the position servo.

The position servo is implemented in the following two stages:

1. To convert the position error into a requested velocity, and
2. to convert the requested velocity into a drive signal.

Theoretically, it can be shown that this two-stage position servo can always be made stable—the position will always be reached without overshoot or oscillation. The non-linear feedback function can also be tailored to achieve, not only stability, but a high performance as well. This means that a requested position is reached in the shortest possible time.

5.2 Velocity Servo Theory

The block diagram of the velocity servo is illustrated in the figure below.

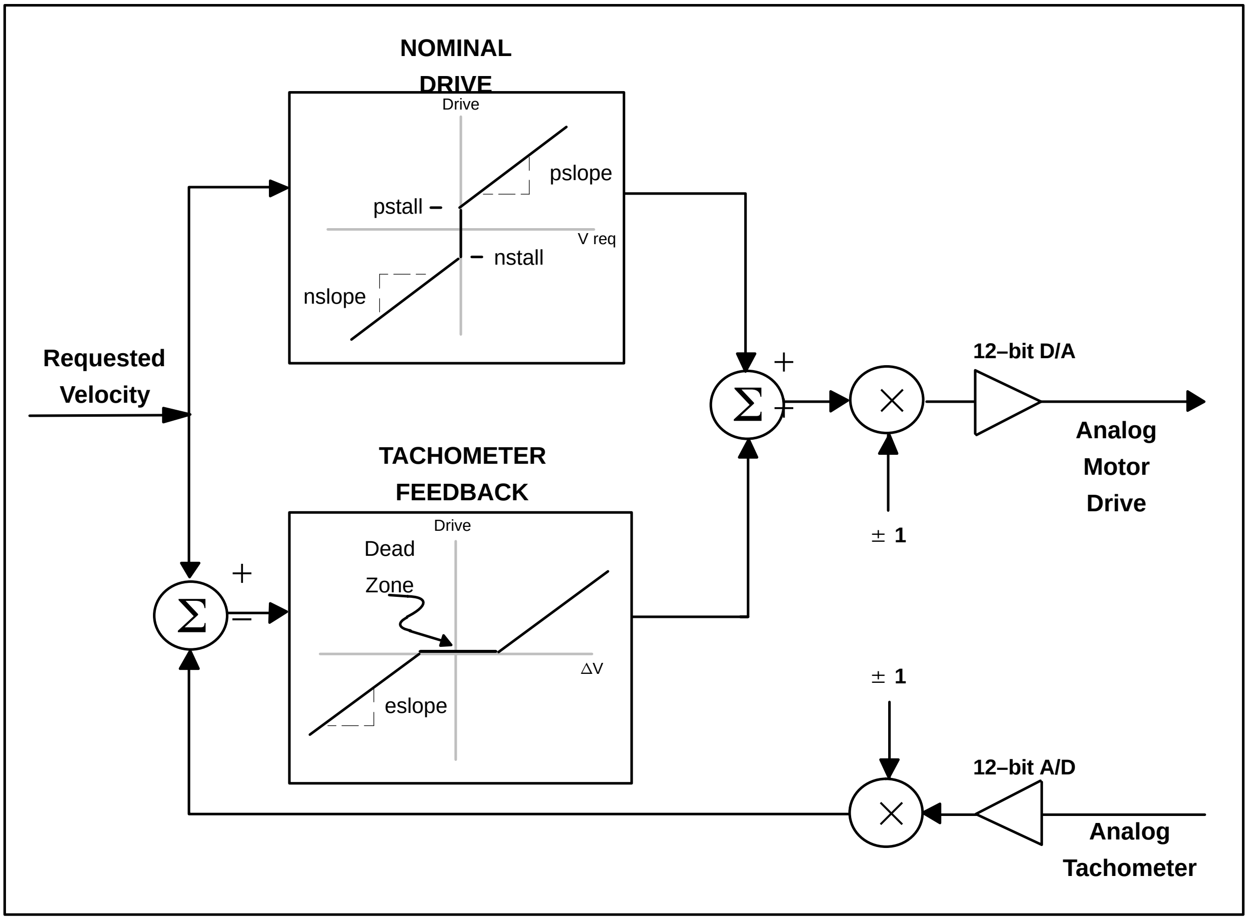

Figure 4–1: Digital Velocity Servo

5.2.1 Tachometer Input

The tachometer signal, from the motor gear box, is applied to a differential receiver and a 30Hz analog, low-pass filter. The signal is then digitized to 12-bits and added to the processor. The differential receiver ensures that any common-mode signal, on both of the tachometer leads (e.g. power-line noise) will not be falsely interpreted as antenna motion. In the case when there is no analog tachometer, a virtual tachometer based on the differentiated position can be selected as described in Section 4.5.

The drive levels, that are computed by the processor, are applied to a 12-bit D/A converter, scaled by the external analog amplifiers, then applied to the motor power amplifiers.

---

i **Note: The D/A and A/D convertors are signed. Thus they generate and accept both positive and negative voltage levels.**

In Figure 4–1, on the previous page, the drive signal is defined as the sum of two components:

1. A nominal level based solely on the requested velocity, and
2. a feedback term based on the difference between the requested and actual velocities.

The two graphed transfer functions indicate how the drive levels are derived for each of these components.

5.2.2 Nominal Drive Slope

The nominal component is an "initial guess" of the drive level that would sustain a given velocity in the steady state. For a requested velocity of zero, the upper-transfer graph indicates that no drive was applied. Without a drive, the motor will eventually come to rest. For non-zero velocities, most motors exhibit a dead zone in which the armature magnetization is insufficient to overcome the starting friction. Therefore, the nominal drive graph takes a discontinuous jump from zero. Due to the antenna's imbalances, this dead zone can also be asymmetric for both directions of motion.

These positive and negative starting drives are designated as "pstall" and "nstall" on the graph. Once the motor is started, a nominal slope is designated as "pslope" for the positive velocity and "nslope" for the negative velocity. Both are used to predict the required drive for large requested velocities.

5.2.3 Velocity Feedback Slope and Dead Zone

The feedback component of the motor drive is based on the difference between the requested and the actual (tachometer) velocities. The lower transfer graph demonstrates that the output is essentially linear, with a velocity error, except for the possible inclusion of a deadzone around zero. The slope is designated as "eslope" on the diagram. The deadzone, between $-V$ and V , is used to minimize motor "chatter" that can result from uncertainty in the LSB of the tachometer voltage samples. Typically, the "eslope" is fairly large in order to achieve a tight velocity servo however, this large value also magnifies the A/D errors. This small inactive region (dead zone) in the feedback loop, typically two 1 or 2 T-units, will eliminate the problem.

The sum of the nominal and feedback terms is clamped within the -100 to $+100$ drive unit range and is applied to the D/A converter to produce the motor drive voltage. It is important to realize that the nominal term does not need to be calculated with great accuracy. In traditional, hard-wired, analog velocity servos, this term is not even used.

The term is included in the digital servo for the following reasons:

- It provides a simple way to take motor stall currents into account, and,
- it helps reduce the mean error that appears in the feedback term necessary to maintain a given velocity.

Every feedback system requires a non-zero error component to maintain control of a non-equilibrium position. By predicting the equilibrium drive requirements, the nominal term helps to ensure that the mean steady-state value of the velocity error will be zero.

5.2.4 Drive and Tach Sign Correction

There are two optional sign inversions that can be introduced in the velocity servo loop: one for the tachometer input and one for the drive output.

These two inversions must be set in the manner to ensure that :

1. the overall servo is stable, and
2. the positive requested velocities results in the positive tachometer velocities.

If the first condition is not met, then flipping either sign will result in stability. If this leads to a violation of the second condition, then both signs must be reversed together. Therefore, both conditions can be met by a suitable choice of multipliers. The need for the stability condition is obvious but the need for the correct tachometer sign results from the requirements that the position servo imposes when it is running.

5.3 Position Servo Theory

Unlike the velocity servo, the position servo is implemented as a simple extension, as shown in Figure 4–2 on the following page. To reach a given position, the position error is used to calculate the velocity that is necessary to correct it. This velocity is then fed into the velocity servo, which continues to operate as described in the previous sections.

The elevation position servo will work properly over the complete 360-degree interval from –90 to +270 degrees. Servo motion will always be directed "over the top" when the antenna moves from one position to another. For example if the antenna is at +200 degrees, and a request is made to move to –30 degrees, then the antenna will traverse the 230-degree sector passing through 90-degrees. This is different from what would happen on the azimuth axis, where the shorter 130-degree path would be taken.

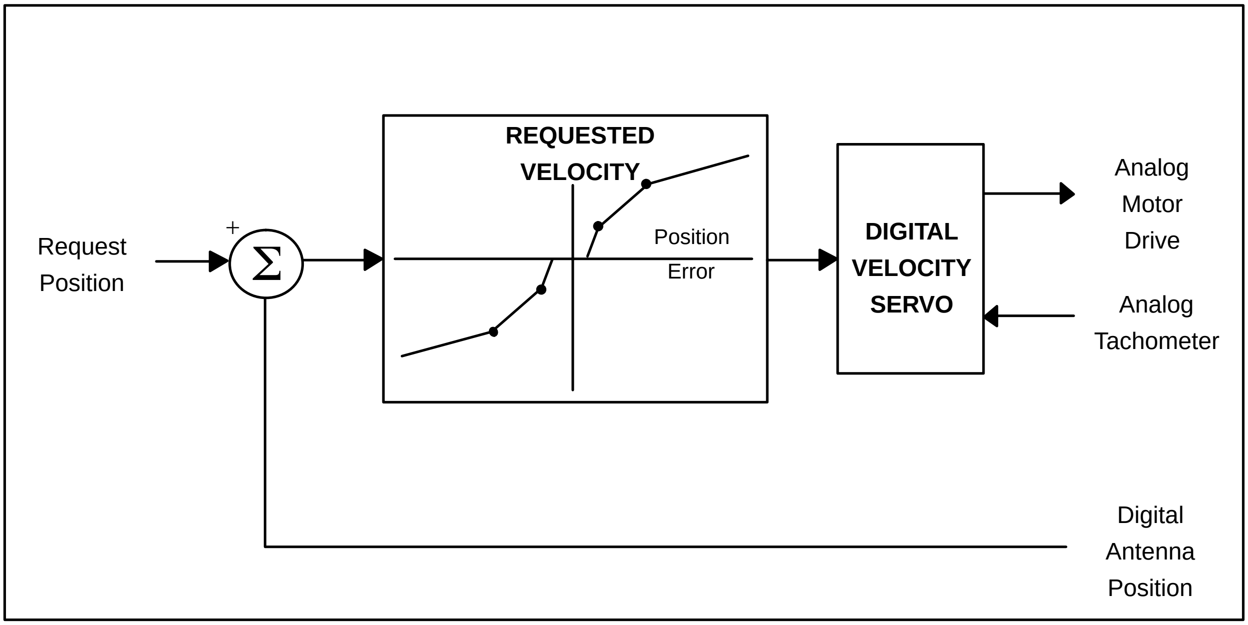

Figure 4–2: Digital Position Servo

5.3.1 The Position Servo Response Curve

As illustrated in the figure above, the mapping between the position error and the requested velocity, known as the Position Servo Response Curve, is quite non-linear and takes account of:

- Stored kinetic energy in the antenna mass
- Non-ideal, power driver characteristics
- Inductive and regenerative motor effects
- Friction

The details of determining actual values for the curve are discussed in Section 4.7, but the overall, concave downward shape can be understood as follows: The angular velocity of the antenna cannot be changed instantaneously, but rather is limited to a rate that may depend on the velocity itself. The time integral, of such velocities, produces roughly "quadratically shaped" positions. When approaching from a far distance at high speed, the distance covered—in the time required to reduce the velocity by half—will be roughly three-quarters of the initial position error.

5.4 Fail-safe Antenna Features

Radar antenna systems consist of the reflector, pedestal, gears, motors and drive amplifiers. These are expensive components that must be protected in the event of failures. The RCP8 has several features that are designed to protect the antenna system from various types of failures.

In the event that a critical failure is detected, the RCP8 diverges into a shutdown state in the following ways:

1. The drive output voltage is zeroed.
2. The drive control relay signal is set to low. If a drive control relay is used, then the RCP8 drive outputs will be physically disconnected from the servo amplifiers. Depending on the installation, this may switch in an alternative drive system such as handwheels.
3. An error message is sent to the front panel display.
4. An error bit is set in the output to the host computer.

⚠ **CAUTION: In the case where the alternative drive system may attempt to move the antenna, it may be undesirable to automatically switch to the alternative drive when a shutdown occurs. In this case, the RCP8 drive relay signal should not be implemented and a manual switching approach should be used instead.**

When a shutdown occurs, the operator should investigate the reason for the shutdown, either by viewing the front panel for a detailed message or by viewing the Control/Monitoring menu. After the fault has been corrected, a reset command can be issued, either from the RCP8 menus or over the host computer serial line.

One of the most potentially damaging situations is when the antenna operates out of the specified elevation range. There are several limits that are typically imposed to protect against this. These limits are illustrated in Figure 4–3. The example shows the lower elevation limits for a typical system. Upper elevation limits are analogous.

Elevation Limit Switch Shutdown Algorithm

Elevation limit switches can be set to force an antenna shutdown, as described in Section 4.4.5. The algorithm checks 40 times per second for limit-switch contact.

Elevation "Shutdown Limit" Shutdown Algorithm

The Elevation Axis setup, as described in Section 4.5, allows the user to specify upper and lower elevation limits that, if exceeded, will cause the antenna to shutdown. The limits are checked 100 times per second. There is no tolerance for this test.

Elevation "Soft Limit" Watchdog Algorithm

⚠ **CAUTION: This algorithm should only be activated after the elevation position servo has been configured and tested. Refer to Section 4.5.**

In order to enforce the soft position limits on the elevation axis, the velocity servo calls a few of the position servo subroutines on each iteration. This is done in order to determine whether the currently requested velocity, which may not have come from the position servo, is such that the antenna could still be stopped before encountering the limit. If the requested velocity is too great, then it is replaced by the velocity that the position servo would have used in order to just reach the limit. This safeguard ensures that the antenna speed is reduced in plenty of time to to reach a controlled stop before encountering the specified soft limit.

When the soft limit algorithm is activated, it will ensure that the antenna is brought to a safe stop at the soft limits, regardless of the servo mode (i.e., open loop, velocity, or position.)

Note: Analogous soft limits can be set for azimuth as well, but these are rarely used since most antennas can rotate freely in azimuth.

Maximum Velocity Watchdog Algorithm

The RCP8 performs the following two types of checks on the velocity to ensure that the antenna is operating within the safe limits:

1. A velocity request limiter that clamps any of the out-of-bounds velocity requests from the host computer or indirectly from the position servo at the maximum value.
2. A continuous check on the antenna velocity to determine that it is operating within safe limits.

The following setup parameters, as described in Section 4.6, are defined as:

- Maximum Absolute Velocity — 80 Tach units
- Velocity Shutdown Safe Margin — 5 Tach units
- Velocity Shutdown Time Check — one second

The watchdog will force an antenna shutdown if the velocity exceeds the Maximum Absolute velocity, by more than the Safe margin, during a time period longer than that of the Time check.

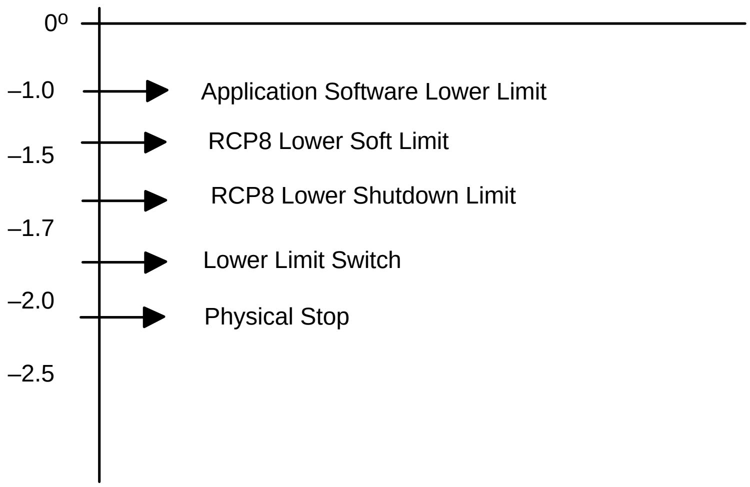

Figure 4–3: Example of the Lower EL LIMITS.

Tach/Position Consistency Shutdown Algorithm

Note: This algorithm requires that the tach be calibrated in degrees per second. If the tach gain potentiometer is adjusted, then the tachometer calibration should be redone.

Note: If virtual tach is used rather than an actual tachometer, this algorithm is disabled.

When the Tach is properly calibrated, the observed change in the antenna position should match the integrated velocity. If these are inconsistent, this could indicate failure of either the Tach or the position sensing and continued operation could lead to antenna damage.

The following algorithm parameters, as described in Section 4.5, are defined as:

- Permissible fixed error — 1.5°
- Permissible relative error — 10.00%

The watchdog algorithm computes the expected difference in position by integrating the velocity and comparing this to the observed position difference over the prior one second. The algorithm will force a shutdown of the antenna if the difference between the observed and the computed antenna displacements exceeds the larger of the fixed and the relative error.

For the antenna displacements greater than 15 degrees—in one second—the relative error of 10%, for an example, would be used as the standard for the test, while for the displacements of less than 15 degrees, the fixed error of 1.5 degrees would be used.

i **Note: The algorithm integrates over the prior one second interval and is updated 16 times per second.**

Unresponsive Antenna Watchdog Algorithms

When drive is applied to the antenna, then the antenna will generally accelerate. Failure of the antenna to accelerate could be the result of one, or more, of the following reasons.

- Servo amp. failure or servo amp. turned off.
- RCP8 drive output failure.
- Drive cable failure.
- Catastrophic gear failure of the antenna drive.
- Obstacle impeding the antenna motion, such as a person, a ladder, or a stow pin inadvertently left in the antenna.

With the exception of the servo amplifiers simply being turned off, any of of these events warrants an antenna shutdown. However, if the antenna is scanning at its equilibrium velocity, the output drive will not cause the antenna to accelerate since it is just balancing against frictional losses. This must be taken into account to avoid false alarms.

The unresponsive antenna algorithm is based on a linear model of the antenna velocity, with a constant moment of inertia, and frictional losses that are proportional to velocity. Under this model, the expected change in velocity can be calculated by numerical integration. The expected change is then compared to the actual change in velocity.

The following Axis Setup parameters for this algorithm, as described in Section 4.5, are defined as:

- Permissible Tach Prediction Error — 15 Tach units
- Maximum duration of such error — two seconds
- Moment of Inertia — 4.00 Drive/Tach units

The moment of inertia is computed whenever the antenna is accelerating and is exhibited in one of the "alt" displays in the Control and Monitoring menu. A representative value is then entered in the setup.

Note: This algorithm does not require the tach to be calibrated in degrees per second. However, if the tach or drive potentiometers are adjusted, this algorithm will need to be reconfigured.

The algorithm performs a numerical integration, over the previous 2.5 seconds, to obtain the expected change in velocity (in Tach units.) If the difference between the expected velocity and the current velocity exceeds the "Permissible Tach error" for a period greater than the "Maximum duration," the Watchdog will force an antenna shutdown.

Note: The algorithm integrates over the previous 2.5 second interval and is updated 8 times per second.

5.5 Modification of Servos For Use on a Moving Platform

The RCP8 is most commonly used to control land-based weather radar antennas. However, with the addition of base motion inputs, the RCP8 can also carry out electronic stabilization of an antenna that is mounted on a moving platform. The position and velocity servos are modified so the antenna motion is referenced to the inertial (Earth) frame of reference. The positions and the velocities are requested by the user and reported back to the user, relative to the local horizon and local north, just as they would for a stationary pedestal. The RCP8 manages all of the coordinate transformations needed to convert between the Earth system of units and the pedestal system of units.

To stabilize an antenna on a moving pedestal, it is necessary to know the instantaneous roll, pitch, and heading of the pedestal base as well as the time derivatives of those three quantities. Pedestal orientation is necessary to convert between the two systems of coordinates. Less obvious, however, is the need to know the rate of change of pedestal orientation.

The following describe the components that contribute to the net Earth velocity of a scanning antenna:

1. The component results from rotation of the pedestal Azimuth and Elevation axes themselves.
2. The component is the result of projected motion of the entire pedestal assembly.
3. The component is computable from the rate of change in the base orientation angles.

The basic velocity servo, as described in Section 5.4, can be modified to work in Earth coordinates by adding a single coordinate conversion module to the block diagram. The modified servo is illustrated in Figure 4–4. Rather than accepting requests for particular pedestal speeds, the new servo responds to commands to move at real Earth-relative velocities. The conversion module receives pedestal attitude information from an inertial navigation unit mounted on or near the pedestal. Based on these data, and on the current pedestal azimuth and elevation, Earth velocity requests are converted into equivalent pedestal velocity requests. The latter are fed into the two "old-style" velocity servos that use tachometer feedback to compute appropriate motor drives.

The velocity conversion module requires the pedestal azimuth and the elevation angles as input in order to project the Earth velocity into the pedestal frame. As a secondary effect of these calculations, the Earth azimuth and the elevation angles are also computed. When a position servo is running on one or both axes, these computed Earth angles are used by the position servo in the same way that pedestal angles were used in the land-based case. Therefore, the position servo from Figure 4–2 is unchanged with the exception of the Earth angles, which are substituted wherever a pedestal angle previously appeared. Also, the angles can also be wired to a nearby signal processor and simultaneously sampled with the data from the radar.

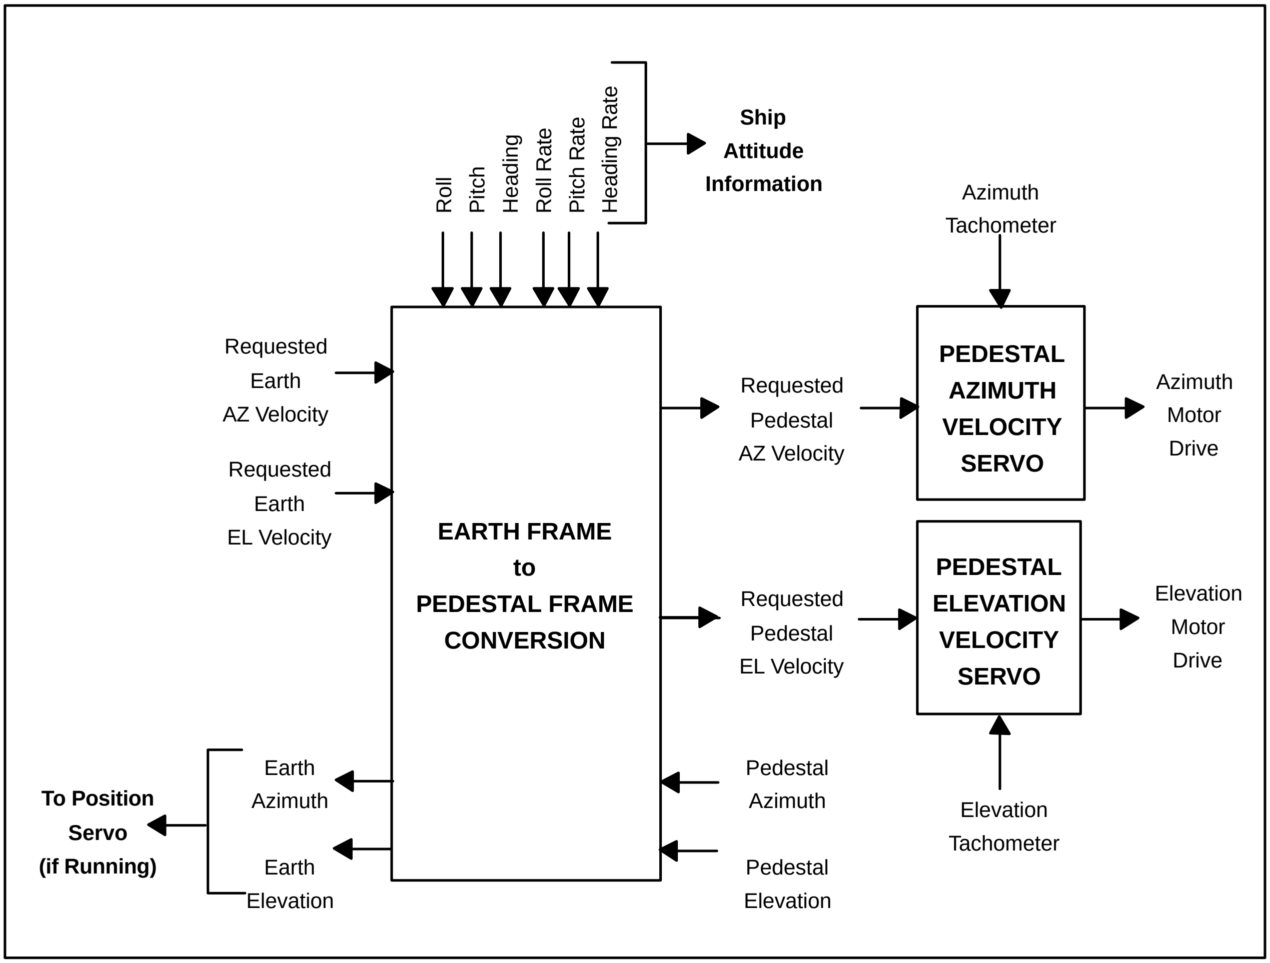

Figure 4–4: Modification of Velocity Servos

An interesting, complimentary addition to the RCP8 moving platform servos is its ability to scan co-planes. By introducing an artificial bias to the pedestal attitude information, a zero-degree elevation scan can be transformed into a planar scan in any orientation, not simply along the horizontal plane. This technique works equally well on either land or at sea.

The INU data stream may include status bits which convey the validity of the attitude angles. The RCP8 will "coast" for up to one second when it receives an invalid INU Roll/Pitch/Heading bit, or until the invalid bit is cleared, which ever occurs first. The last valid report of INU parameters will be used for stabilization during this time (including computation of the earth-relative output angles). Since it is unlikely that the antenna azimuth and/or the ship attitude will move more than 30 degrees in one second, the IRIS message "DSP AZ angles exceed 30 degrees" will not be triggered by very short bursts of invalid INU data.

Note that the option of continuing to use the new INU parameters for the one second interval (rather than coasting with the last valid ones) was specifically rejected for safety reasons, since there is a possibility that the new angles are truly bad.