

Design, Accuracy, and Calibration of Analog to Digital Converters on the MPC5500 Family

by: Mike Garrard, TSPG Powertrain Systems Engg.
Pauline Ryan, TSPG MCD Applications Engg.

1 Introduction

A new architecture of analog to digital converter (ADC) was introduced with the MPC5554 device, to meet requirements of increased speed and accuracy. This architecture consists of a converter engine with improved linearity, and correction hardware, to remove linear gain and offset errors. This correction hardware is the ADC calibration subsystem.

1.1 Abstract

This application note outlines the advantages of the new redundant signed digit (RSD) architecture of ADC introduced with the MPC5554, compared to the successive approximation architecture used in the MPC500 family. It covers the gain and offset errors introduced by the RSD ADC, and how the hardware multiply accumulate (MAC) unit is used to remove them. Design features of the ADC, used to recover full range, are explained, along with design features in the MAC required to avoid lost codes, and the half-count offset adder that increases accuracy. Levels of calibration, and

Table of Contents

1	Introduction	1
1.1	Abstract	1
1.2	Objective	2
2	ADC and MAC Design Features	2
2.1	ADC Design Requirements	2
2.2	ADC Design Solutions	2
2.3	ADC Accuracy	7
3	Applying Calibration to the ADC	11
3.1	Outline of Power Up Calibration	11
3.2	Outline of Run Time Calibration	13
3.3	Determining Noise Level	15
3.4	ADC Internal Registers	17
3.5	Calculating the Calibration Coefficients	18
4	Conclusion	20
5	References	20
6	Glossary	21
	Appendix A Software Example	22
	Appendix B Example Coefficient Calculations	24

their resulting accuracy, are described, and the parameters affecting total error are discussed.

Section 3 describes how to accomplish the calibration. A procedure for power up calibration is provided, as well as a procedure for run time calibration, including when to schedule it. A test is given to evaluate a system, in order to tune the calibration. The relevant ADC internal registers are described. Calculation of the coefficients is explained with both integer and floating point examples. Software to achieve the calibration function is described, and provided in Appendix A.

1.2 Objective

The objective of this application note is two-fold. Section 2 aims to educate on the structure and design features of the MPC5500 family ADC that make it suitable for embedded control. Section 3 is a tutorial on using the calibration feature to obtain best performance from the ADC.

2 ADC and MAC Design Features

2.1 ADC Design Requirements

When the MPC5500 family of devices was developed, the ADC was required to be faster, with improved linearity and monotonicity. A low capacitance input was required. It had to maintain full 0V to 5V, and d.c. capability, and it had to be an architecture capable of further development. A new RSD architecture of ADC was introduced to replace the successive approximation (SA) architecture used in the previous MPC500 family. Compared to the MPC500 family ADC, the ADC on the MPC5500 family typically offers:

- 4x the resolution
- 3x the linearity
- 4x the conversion speed
- At least the same absolute error
- 0.8 pF input capacitance

2.2 ADC Design Solutions

2.2.1 MPC500 Successive Approximation ADC

The MPC500 ADC is a successive approximation design. The core operation of such a design is quite straightforward, as shown in [Figure 1](#). The input voltage is converted one bit at a time, starting with the most significant bit. The DAC (digital to analog converter) is set to 50% of the voltage range, and the comparator reports whether the input voltage is higher or lower than this. The most significant bit (MSB) is set to 1 or 0 as appropriate. Let us say that the voltage is higher than midpoint, so $MSB = 1$. The second most significant bit is now determined: the DAC is set to 75% range, 50% for $MSB = 1$ plus 25% as the midpoint for second most significant bit. The comparator reports whether the input voltage is higher or lower than this, and bit 2 is set, as appropriate. The conversion continues, in this manner, until all bits are

determined. The input voltage remains connected during the conversion, and the MPC500 uses a buffer to charge this up initially. This architecture of ADC runs at moderate speeds and is economical to implement, but relies on the DAC and good amplifier offset for accuracy.

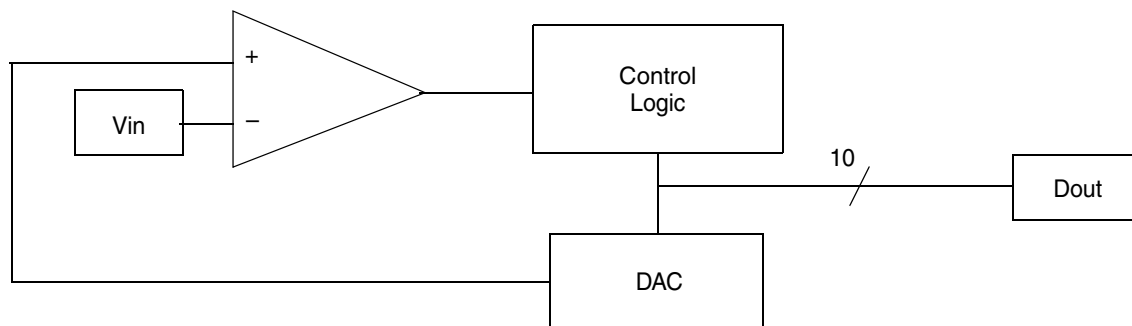


Figure 1. MPC500 Architecture Successive Approximation ADC

2.2.2 MPC5500 Redundant Signed Digit ADC

To meet the requirements for the MPC5500 family, the ADC was changed to the RSD architecture. This also uses successive approximation, but, instead of a DAC, an accurate x2 amplifier is used. The SA ADC makes comparisons at progressively smaller voltage differentials for each bit, whereas the RSD doubles the voltage each time and compares at the same level. The diagram below outlines a simplified structure of such a converter, which is sufficient to explain the significant features of this type of ADC. The actual design implemented on MPC5500 uses two reference points, 25% and 75%, with a three-option summation and a two-bit digital answer per phase, which is combined in the digital logic to provide the result. These enhancements provide greater noise immunity. For a description refer to the MPC5553/54 microcontroller reference manual.

In the RSD architecture (Figure 2), the input voltage is connected only for the initial sample via S1. This is compared to the 50% reference voltage to provide the MSB. If the current bit is 1, then the reference voltage is subtracted, using S2; otherwise, no subtraction occurs. For subsequent bits, the residue of the sum is multiplied by two, using S3, then compared to the reference voltage again using S4. Assume V_{in} is 65%. It is greater than 50%; therefore MSB = 1, and 50% is subtracted leaving a residue of 15%. This is multiplied by two, so the next compare is 30%. This is less than mid voltage (50%); so the next bit is zero. Nothing is subtracted either, so it is doubled to 60%. So the next bit is 1, and so on. The table below shows an example for a 6-bit conversion of 65%: $65\% \text{ of } 6 \text{ bits} = 0.65 * 64 = 41 = 0b101001$. The input voltage is connected only for sampling (MPC5500 default two clocks), and, as the capacitance on MPC5500 is kept small, no buffer is required.

Residue	Compare	Adder
Sample 65%	>50% = 1	$65\% - 50\% = 15\%$
$2 \times 15\% = 30\%$	<50% = 0	$30\% - 0 = 30\%$
$2 \times 30\% = 60\%$	>50% = 1	$60\% - 50\% = 10\%$
$2 \times 10\% = 20\%$	<50% = 0	$20\% - 0 = 20\%$
$2 \times 20\% = 40\%$	<50% = 0	$40\% - 0 = 40\%$
$2 \times 40\% = 80\%$	>50% = 1	

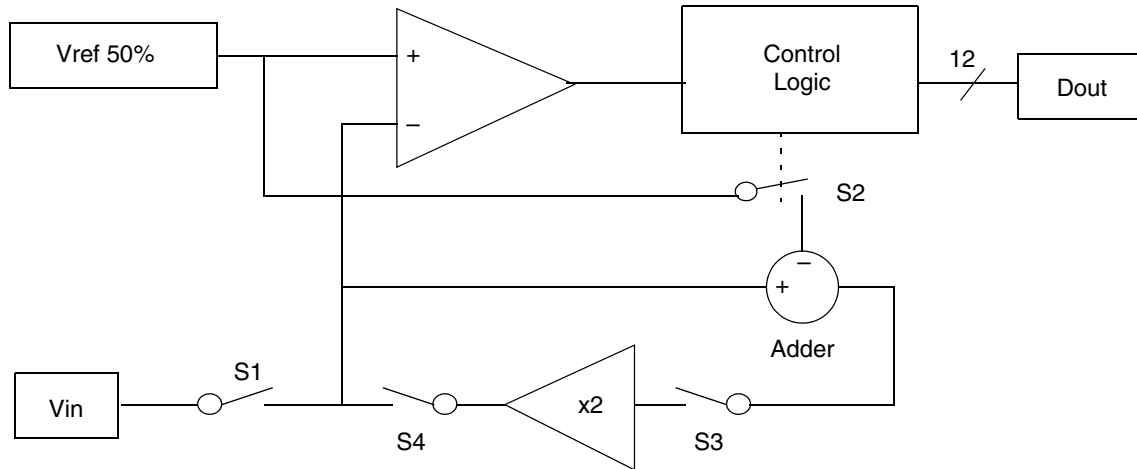


Figure 2. MPC5500 architecture RSD ADC engine

The precision elements of the RSD architecture are the x2 amplifier and adder. The advantage is that fewer accurately matched components are required, and these components are physically larger, so their matching is better. Since the same components are used at each bit, the design offers great linearity. Also, because the comparisons are made on scaled up voltages via the amplifier, the design is less sensitive to settling of the amplifiers and can be driven faster.

The first source of linear error for this design is capacitor matching, because the x2 amplifier is based on capacitors. Any error away from the ideal gain of two is doubled each time another bit is computed. For example, the sample voltage might go through 1.99 x 1.99 x 1.99 x ... This results in an overall gain error for the ADC transfer function. To obtain the performance achieved by the MPC5500, the capacitors must be matched to about 0.0003%. The second source of linear error for this design is amplifier input offset. Again, this is doubled each time another bit is computed. To obtain the performance specification of the MPC5500, the input offset must be less than 18 nV. Special techniques are used to achieve these impressive figures.

2.2.3 Multiply Accumulate Unit

The digital output of the RSD stage is passed to the second stage of the ADC subsystem, the Multiply Accumulate (MAC) unit. The MAC unit is invoked whenever a conversion command message with calibration enabled (CAL = 1) is received by an ADC. To calibrate a conversion, the MAC unit requires two prior pieces of information; the Offset Calibration Coefficient (OCC) and the Gain Calibration Coefficient (GCC). The offset of an ideal ADC is 0. The gain (slope) of an ideal ADC is 1. The purpose of this unit is to remove the gain and offset errors discussed at the end of [Section 2.2.2, “MPC5500 Redundant Signed Digit ADC”](#), using a simple $y = mx + c$ linear correction:

$$\text{Calibrated result} = \text{Gain} * (\text{Uncalibrated result}) + \text{Offset} \quad \text{Eqn. 1}$$

This correction is shown graphically in [Figure 3](#).

NOTE

The ADC is engineered by design to have a slope less than one and a positive offset (of codes). This means that the full 0V to 5V range is always available. Although this change was achieved through adjustment to the reference taps used in the RSD algorithm, the reference channels read for calibration reference remain at precisely 25% and 75%. However, the 50% channel is the analog ground and does show the 20 mV offset. The expected code after calibration is not 2048 but 2032

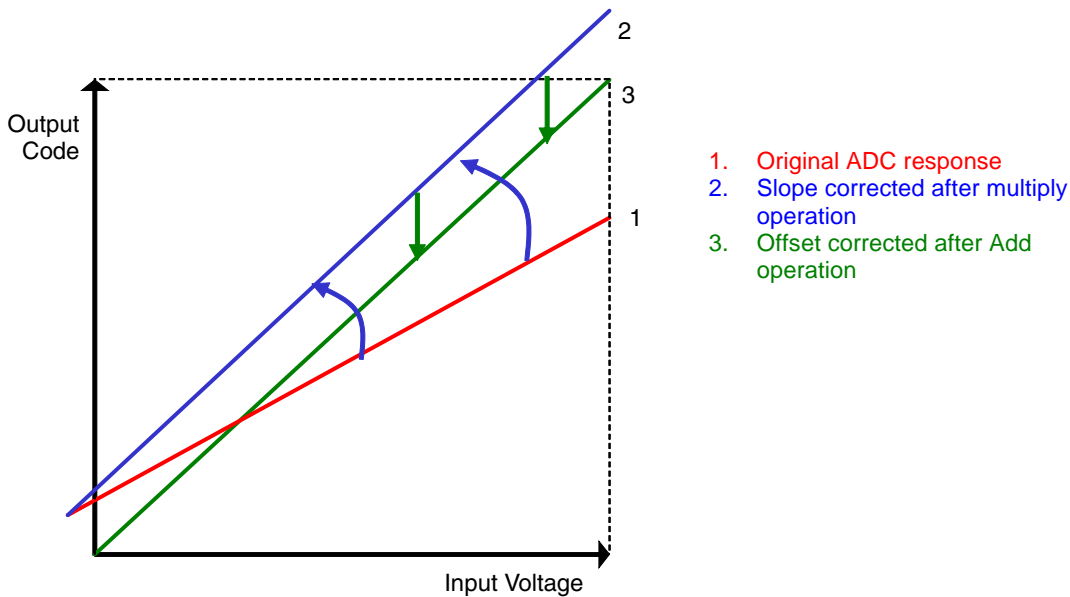


Figure 3. MAC Correction to ADC Response

A block diagram of the MAC is shown in [Figure 4](#).

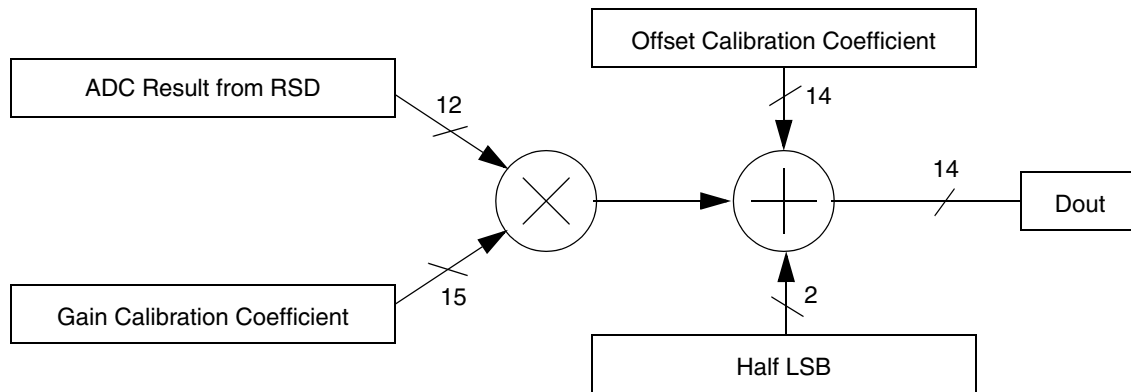


Figure 4. MPC5500 MAC Architecture

It can be seen that although the ADC provides a 12-bit result, the MAC itself uses fifteen bits for the gain coefficient and fourteen bits for the offset coefficient. This is to avoid lost codes or stuck codes. For example, the ADC in [Figure 3](#) has 4096 output codes. This staircase is then stretched by the gain correction, and pushed downwards by the offset correction. In doing this, a few codes (steps on the

staircase) will fall outside the input voltage range, either below 0V or above 5V. Therefore, the ADC between 0V and 5V will now deliver, say 4080 codes. However, the ideal response has 4096 available steps which must be covered. The only way for the MAC to do this is to take one extra step in sixteen places. This is undesirable because in these places, an input *voltage* change of one LSB (least significant bit) to the ADC will result in an output *code* change of two LSB; the ADC will have lost codes.

Since the MAC uses 14-bit arithmetic, it is able to take 0.25 LSB steps (0.3 mV). Therefore, the 4080 codes can be distributed more evenly over the 0V to 5V range. Each time the ADC delivers a result, the MAC is able to select the most appropriate of the 16384 numbers available to it. The difference between a 12-bit result and a 14-bit result is shown in Figure 5, where blue is the input analog voltage, green is the response of a low-gain ADC and red is the output of the MAC. The lost codes can be seen as jumps in the red staircase.




Figure 5. Comparison of 12-bit MAC Result (left) vs. 14-bit MAC Result (right)

One final design feature is worthy of mention – the half LSB compensation. This reduces the error by centring the quantized ADC response over the linear analog input. The effect can be seen in Figure 6. Half LSB on the 12-bit ADC result is an adder of two in the 14-bit MAC calculations.

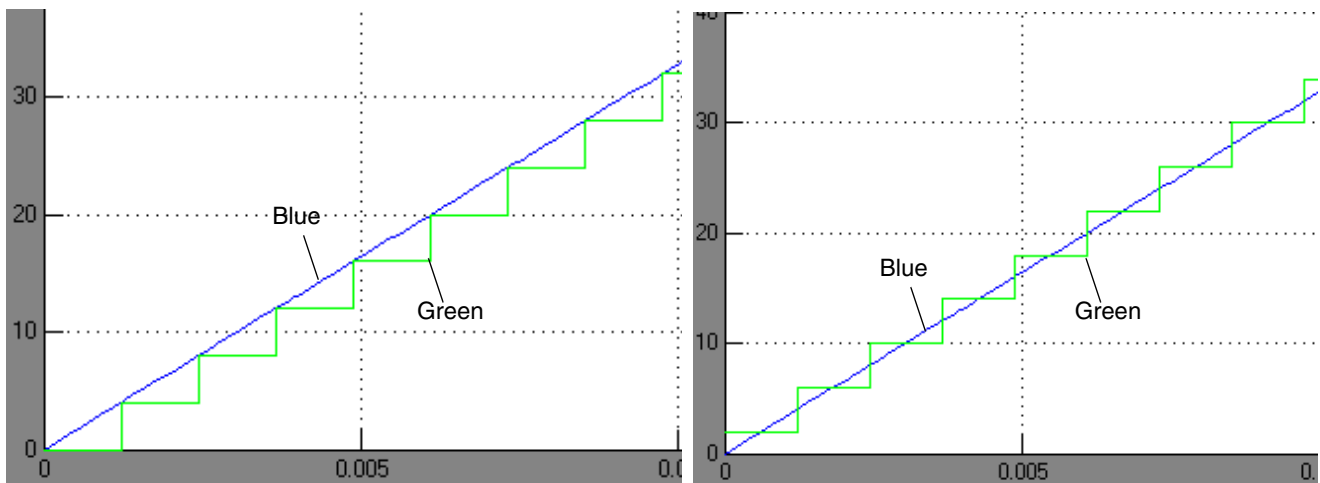


Figure 6. MAC Output Without Half LSB Adder (left) and With (right)

When a conversion command message is received that does *not* request calibration ($CAL = 0$), the MAC continues to add the half LSB correction of two, and continues to multiply by four. This provides consistency of 14-bit result format for both calibrated and uncalibrated conversions.

2.3 ADC Accuracy

2.3.1 Levels of Calibration

Calibration compensates for the errors below (listed in order of decreasing magnitude).

1. Reference offset to permit full 0V to 5V range
2. Silicon manufacturing variation
3. Temperature drift

Three levels of calibration are possible.

2.3.1.1 No Calibration

This will result in an absolute error (TUE — total unadjusted error) of around six to seven bits of accuracy. It will be dominated by the uncalibrated gain and uncalibrated offset terms in the data sheet. Although most devices will be well within the data sheet limits, it should be noted that they are not centred around zero, to ensure full 0V to 5V range. Therefore, an uncalibrated ADC will always show gain and offset error. It is expected that systems that require the performance of MPC5500 will always require some level of calibration.

2.3.1.2 Power Up Calibration

A single calibration on power up will remove the most significant errors — those engineered in for full voltage range capability. The ADC will measure about eight to nine bits accurately. This will suit systems requiring MPC500 levels of accuracy. One advantage of power up calibration is that it can occur before the control module begins to run large, noisy loads, such as ignition coils and transmission solenoids. The power supplies and references of the ADC will, thus, be less noisy, and a more precise calibration is likely to be obtained.

2.3.1.3 Run Time Calibration

The highest absolute accuracy is obtained by run time calibration since it will also remove temperature drift. The ADC will measure nine to ten bits TUE. It is also the most complex to design and implement. The procedure for power up and run time calibration is provided in [Section 3, “Applying Calibration to the ADC”](#).

2.3.2 What Level of Calibration is Required?

This depends upon the accuracy required by the system, the accuracy available to the ADC, and the types of sensors being measured.

2.3.2.1 Signals Dependent Upon ADC Absolute Accuracy TUE

The ADC is capable of nine to ten bits TUE. To measure this requires highly accurate and linear test equipment (preferably analog), stable low noise references, a quiet circuit board without peripheral activities, and a test environment without external interference. To make use of this level of accuracy in a production environment, further processing of the signal read by the control module will normally be necessary, due to the noise levels present on the signal itself and inside the engine control unit (ECU), particularly the ADC references. For example, if the ADC references have 3 mV of noise, and the signal being read has 10 mV of noise, then even a perfect ADC could only provide eight bits of accuracy. In this case, multiple reads and some form of averaging could remove the noise, but there are other error sources for which compensation is more difficult. Say, for example, the tracking supply to a ratiometric sensor has 20 mV regulation for line and load. To achieve better than eight bits, even with a perfect ADC, would require reading the supply directly, rather than relying on the tracking function. If the sensor plus filter is high impedance (above 100k Ω), a low offset buffer IS required to avoid a voltage drop that reduces accuracy to nine bits. If a sensor requires a voltage divider, then these two resistors must be better than 0.1% over temperature and age, for an ideal ADC to read nine bits of accuracy. Finally, there are error sources for which compensation is not possible. For example, the absolute reference voltage supplied to the ADC in the control module might vary by 15 mV over manufacture, temperature and aging. Note that, although the Freescale data sheet provides results from single ended conversions, the ADC is differential in nature and can be used in this mode. The ultimate ADC accuracy will be achieved using differential sensors read differentially.

2.3.2.2 Signals Dependent Upon ADC Linearity INL

There are signals where the ADC will be able to deliver ten or even eleven bits of accuracy – these are from closed loop or a.c. sensors. The irony is that, for these, no calibration is required at all. The only data sheet specifications that are significant are INL (integral non linearity), which gives linearity, and DNL (differential non linearity), which indicates monotonicity. (Note that these specifications should not be added for total error.) Three examples are given below.

Knock sensor. This is an a.c. sensor, where the important information is the energy contained in a frequency band. On both counts the d.c. component is removed, so ADC offset is irrelevant. At first inspection, it might appear that the ADC gain is significant, because it affects the energy read. However, engines vary in noise initially and over time; knock sensors vary in output amplitude; and the amplitude with which a knock sensor ‘hears’ each cylinder varies due to distance and mounting effectiveness. The knock phenomenon itself, being combustion related, has considerable variation. Each of these variations is many times greater than the fraction of a percent variation in ADC gain, which, in itself, would reduce the need for calibration. However, as a result of this variation, most knock algorithms include a means to dynamically adjust the threshold at which knock is declared, thus negating any gain calibration. The ADC when used with a knock sensor operates at the accuracy of the INL specification, which is about eleven bits of accuracy.

Throttle potentiometer. This is a d.c. ratiometric sensor, which means it provides a voltage that is a ratio of its supply voltage. It is a critical sensor, particularly at low throttle openings, but there is a mechanical error as fitted. To increase precision, the interface algorithm will usually determine the closed voltage, since it occurs at idle and no lower voltage is normally seen. Likewise, the fully open position is determined as the maximum normal voltage, and when the engine manifold also reads maximum pressure.

Again, the significant accuracy for the ADC is only linearity, and for identification of changes in position and monotonicity.

Airbag accelerometer. A common sensor type for airbag systems is a micro-electromechanical system (MEMS) capacitor. There is a significant d.c. (zero g) tolerance, as well as significant gain (mV/g) tolerance, making gain and offset calibration of the ADC irrelevant. Linearity of the ADC is also much better than that of the sensor. However, in this safety critical application, it is essential to determine correctly a rising or falling g; hence, monotonicity (DNL) is the most significant factor.

Exhaust oxygen sensor (HEGO/Lambda sensor). This is an active sensor that provides a voltage as a function of the partial pressure of oxygen across it. Nominally, the ideal air-to-fuel ratio causes it to produce around 450 mV, but the absolute voltage depends upon sensor temperature, which varies significantly, even with a heated sensor. There is also a temporary sensor condition, CSD (characteristic shift downward), that causes the voltage to shift. Thus, often the algorithm searches for the midpoint between high and low peaks to determine the stoichiometric point. Again, the ADC delivers the information with INL levels of accuracy.

2.3.2.3 Summary Examples

This table makes some assumptions about system accuracies, noise elimination and impedances and so should be used as a guide only.

Table 1. Summary Examples

Sensor	Requirement	Calibration Level	Comments
Battery voltage	Absolute accuracy	Power up	0.1% divider resistors and 10mV reference
Battery voltage	Absolute accuracy	Run time	0.05% resistors over age and temperature, 3 mV reference, averaging
Temperature	Absolute accuracy	Power up	0.1% load resistor and 10mV reference
Temperature	Absolute accuracy	Run time	0.1% sensor, 3 mV tracking supply, 0.05% load resistor and 2 mV reference
Injector PWM current	Absolute accuracy	Power up	0.2% sense resistor
Injector current inflexion	Monotonicity (DNL)	None	Only waveform shape is important
Knock	Linearity (INL) Monotonicity (DNL)	None	a.c. sensor with wide amplitude variation
Throttle position sensor	Linearity (INL) Monotonicity (DNL)	None	Algorithm closes the loop to remove d.c. errors
Manifold pressure sensor	Absolute accuracy	Power up	0.2% sensor
HEGO for AFR	Absolute accuracy,	Power up	Fixed sense threshold voltage for stoichiometric
HEGO for AFR	Linearity (INL)	None	Algorithmic determination of stoichiometric
HEGO for individual cylinder	Monotonicity (DNL)	None	Only waveform shape is important

2.3.3 Total Unadjusted Error (TUE) as a Component in System Error

The ADC is measured, in accordance with IEEE Standard 1241-2000, for gain, offset, DNL and INL. TUE is not one of the specifications in this standard but is provided for consistency with the MPC500 family. For MPC5500, TUE applies to results using the full ADC subsystem including the MAC.

Unfortunately, it is not possible to calculate TUE from the individual errors provided in the data sheet. This is because the errors are not independent or simple (linear). A simplified model of the errors is shown in Figure 7 to explain this. INL is the divergence from the dark blue/diamond ideal line. The yellow/triangle line shows the output after calibration using the 25% and 75% reference points. The cyan/unmarked line shows the resulting linear response used to determine the calibrated gain (green/+, vertical top right) and offset (red/circle) lines.

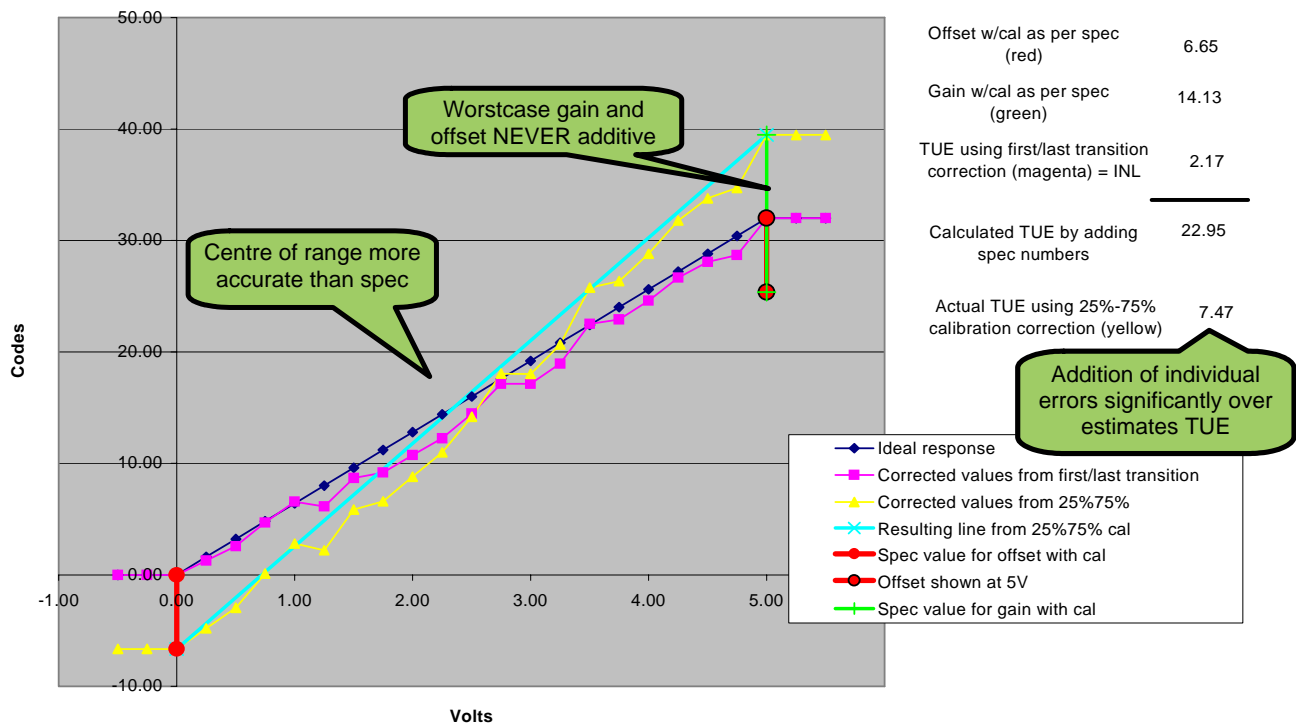


Figure 7. Model of ADC Errors

Some observations from the graph:

TUE is dominated by gain and offset. This is true because the model is built to reflect MPC5500 performance.

Worst case TUE occurs when one reference goes high and the other low. In this case, offset *always* cancels out some of the gain error. This is one reason why addition of the gain and offset numbers does not provide TUE. The sum of actual errors is shown top right at 22.95 LSB, but the real TUE is only 7.47.

Worst case errors always occur at the ends of the response. This is good for control systems because the active operating range is usually centred, with the ends being used for failure mode and effect analysis

(FMEA). This means that the ADC will, in practice, be operating at a level more accurate than the data sheet suggests.

It might be thought that, if the absolute accuracy of the references were known, then the TUE could be calculated. This is not the case. The references are not just points for calibration; they are used by the RSD algorithm, so the relationship is complex. The references are also subject to INL when read, which is not a calculable error. To determine their accuracy directly, they would have to be tested, but the action of bringing the references outside the device for measurement would significantly degrade the ADC performance. Instead, the method used to test TUE includes any error in the references.

It might also be thought that external references could be used to increase the TUE accuracy of the ADC. In theory, this is possible, but, in practice, there are several hurdles. It would be preferable to do this with the ECU in a production-ready state, but this means that there must be access to the internal ADC reference supplies. The sensor tracking supply would be available but, even unloaded, would be expected to have a few mV difference. The ADC inputs would also have protection resistors, so the input leakage of the ADC will cause a few mV difference (+/-) between the pin and the ADC. If the ECU is open, then there are still problems. Software must drive the calibration, disturbing the power supplies and references. Connection must be made directly to the ADC references and test channels. The test voltages must be able to respond to the measured ADC reference voltage (both VRH and VRL) to provide the calibration points. To make the exercise worthwhile, the calibration points must be delivered and read with greater accuracy than the internal points, which, based on the measured TUE, is less than 4 mV. This is not the end of the problem, because internal calibration software still must be written, to adjust the values with temperature; and the exercise is negated if the module is re-flashed at the dealer.

3 Applying Calibration to the ADC

This section describes the method for ADC calibration. Two levels of calibration are considered: on power up only, and during run-time. The details of calibration depend upon the characteristics of the system in which the ADC is used, so experiments to optimize calibration are explained. Calibration software is provided in Appendix A.

3.1 Outline of Power Up Calibration

Figure 8 is a flow chart of the procedure for power up calibration. The steps are described in more detail below. Although it is written for a single ADC, both ADCs will require calibration. This can be performed in parallel, if desired.

1. Power is applied to the module. The device will come out of reset at default frequency with ADCs disabled. Run initialization code, which would include ramping to final frequency and configuring the DMA channels to run the ADC queues.
2. Enable the ADC. This involves writing to registers associated with the ADC that are not memory mapped. A queue can be used for this. The ADC should also be set to final frequency.
3. Allow a delay to permit the power supply to settle after the current change caused by changing the device frequency. The delay will depend upon the power supply, typically no more than 100 μ s. The ADC also takes about 10 μ s to charge internal voltages.

Applying Calibration to the ADC

4. Take multiple readings of the 25% and 75% reference channels. Long sample time should be used, but the calibration bit should not be set. Read [Section 3.3, “Determining Noise Level”](#) to determine numbers. Note that it is possible to write to both ADCs with one queue. It is also possible to have the results sorted by the DMA, by using the tag field in the Conversion Command Word. Reads from each reference could be sent to a separate results queue.
5. Wait for the conversions to finish. This will be a calculable duration based on ADC setup. It is not necessary to be idle, but it is sensible to avoid activities that create circuit board noise near the ADC.
6. Average the results of each reference using a mean or median filter, see [Section 3.3, “Determining Noise Level”](#).
7. Calculate the calibration coefficients (see [Section 3.4, “ADC Internal Registers”](#)). These should be checked against limits to ensure correct system operation. The maximum value for uncalibrated offset is given in the data sheet at 12 LSB resolution, so 4x this value should contain the offset, as calculated by software (14-bit), plus a margin for the noise as determined on the target system. The maximum value for uncalibrated gain is also given in the data sheet again at 12 LSB resolution. This must be converted to a ratio to compare against the calculated GCC. The ADC has a range of 4096 counts, so the ratio limit becomes $(4096 \pm (\text{offset error})) / 4096$.
8. The calculated values of OCC and GCC can now be written to the ADC.
9. It is advisable to re-read the reference channels using the new calibration. They should match the ideal 25% and 75% results, plus a margin that depends upon the noise on the system.

This completes power up calibration. The process must be performed at least once per device. However, at that point, one option would be to store the calibration using EEPROM emulation. This might be used to shave a small amount of time from boot up, or, if run time calibration is also employed, to begin that without having to repeat the power up calibration.

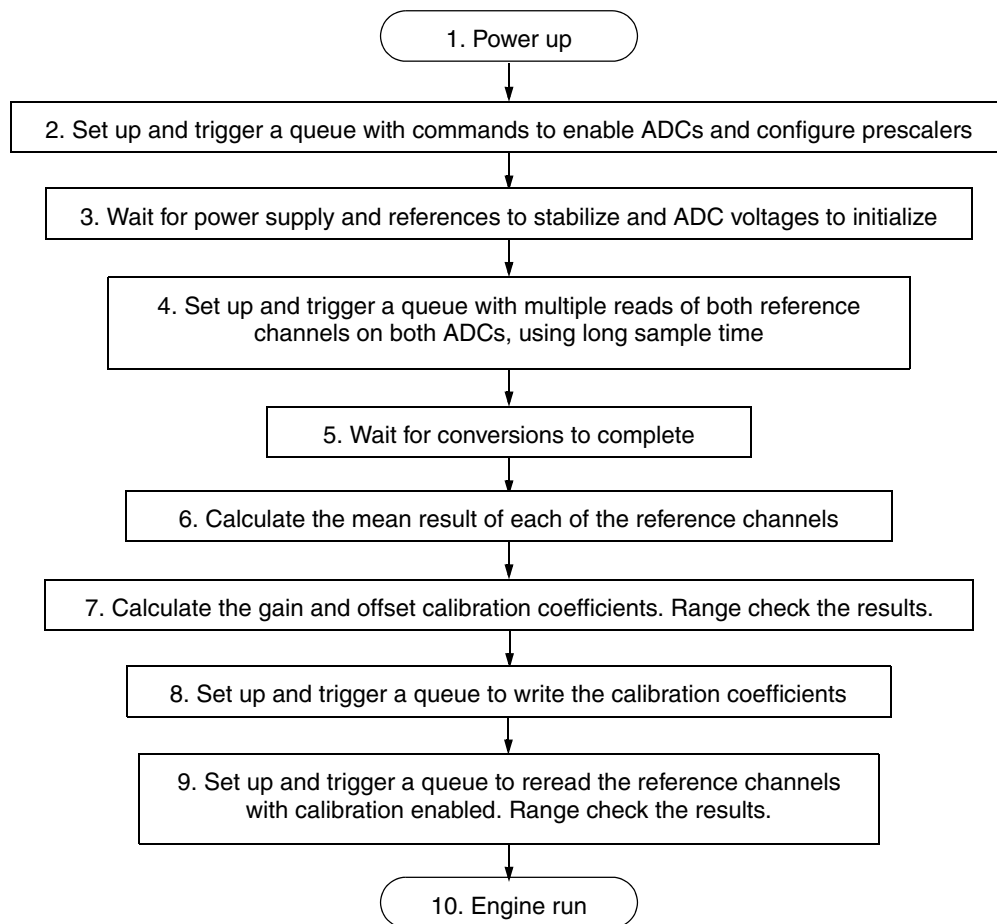


Figure 8. Power Up Calibration Procedure

3.2 Outline of Run Time Calibration

Figure 9 is a flow chart of the procedure for run time calibration. The first two steps describe power up calibration, and the subsequent steps describe new software that must be written.

Step 3. This represents the start of the operating system and repetitive application code.

Step 4. As with power up calibration, an extended sample time of eight clocks is recommended. Only a single read is required because averaging is carried out later. Neither timing nor coherency of the reference reads is critical, so they can be performed using low priority commands. The aim of run time calibration is to compensate for temperature drift. Although the range of temperature drift is small, self heating of the device can cause a relatively fast initial temperature rise (seconds), if the ambient temperature is low. The rate of reads must be several times faster than the rate of drift of the ADC. A time based queue at 100 ms would be a suitable for this, but see Step 5.

Step 5. The recommended filter for run time calibration is a standard digital (software) low pass. This filter uses few resources and is easy to tune with a single gain constant. The filtering must be sufficiently severe that the result is free from noise, leaving just temperature drift. To establish this, the gain constant must be reduced until it is sufficiently small that noise does not affect the calibration. This can be established by

monitoring the filter output when the device is at constant temperature. If this means that the filter is now too slow to respond to temperature changes, then the rate of reads must be increased; the level of noise, rather than the rate of temperature drift, is now the determining factor for the filter.

NOTE

The time constant of the filter depends upon both the gain constant and the rate of reads; therefore, the gain constant should be recalculated if the rate of reads is changed.

Step 6. Drift of the ADC can be seen from drift in the reference channel results. It is not necessary to calculate the calibration coefficients to identify drift. The main factor in the decision to declare drift is the accuracy required by the application. The fact that run time calibration is in use implies that accuracy to the last few LSB is the aim, and that other inaccuracies including noise on the system and references have shown themselves capable of this. In which case, declaring drift at only one or two LSB is advantageous – see step 7.

Step 7. There is a coherency to be avoided when re-writing calibration coefficients. This occurs when the coefficients are written using any queue other than the highest priority, Q0. What happens is a conversion is scheduled, due to its higher priority, after the update of one calibration coefficient (OCC, for example) but before the update of the other coefficient (GCC, for example). The issue is circumnavigated when a change in coefficients is declared with very little drift (Step 6). In this case, there will be only one or two LSB change in offset or gain. Therefore, even if a high priority conversion takes place in between the update of coefficients, the effect of the incoherency will be within the error of the system. Using this method, calibration coefficients can be updated using a low priority queue.

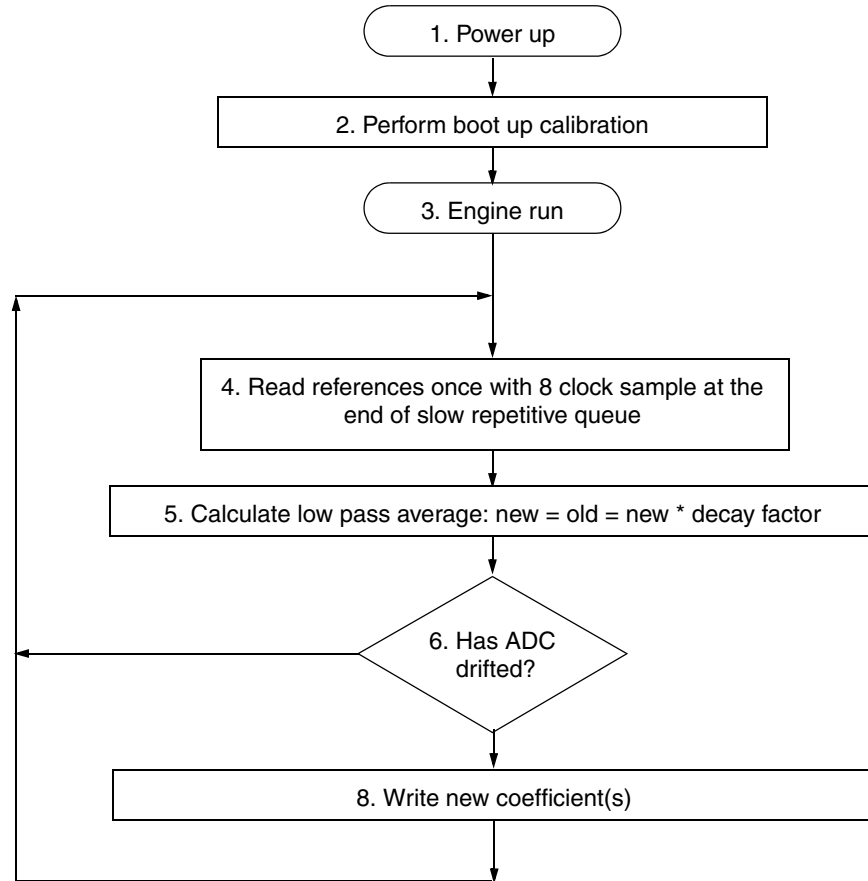


Figure 9. Flow Chart of Calibration Procedure

3.3 Determining Noise Level

Since the calibration coefficients will affect every subsequent reading by the ADC, it is desirable to make them as precise as possible. Precision can be improved by averaging out noise on the reference channels. This is achieved using multiple reads, then calculating an average in software. There are two items to be decided: the number of samples and the average to use.

To determine the noise on the system, all that is required is to take multiple readings on the reference channels and look at the data. [Figure 10](#) shows an example set of results obtain by 60,000 reads on the 75% reference channel. It can be seen that a single read might lead to poor calibration: the span of results is 29 codes; however, nine out of ten results are either 0xBF7 or 0xBF8, so it should not take many reads to see the trend. In this case, five or six reads should be sufficient.




Figure 10. Results of a Jitter Test

The results in Figure 10 suggest that this board has a generally low level of noise, because the majority of reads span only a few codes. They also suggest occasional large noise spikes, because the spread is large, but with low incidence. With a distribution like this, it can be seen that a mean average is not a good choice. A single extreme read could lead to a poor result. Looking for the modal (most common) result, or median (centre value), would be better signal processing techniques.

Significant variation has been found in different applications. If the distribution were closer to a Gaussian (bell shaped) curve over, say 10 codes, this would suggest more persistent noise, but at a lower amplitude. The greater number of common codes would mean that more reads would be required and, with a Gaussian response, a mean average might be most suitable.

Noise appearing on the references can come from a wide variety of sources, through capacitive or inductive coupling, or conduction via the reference pins VRH and VRL. Running an experiment to determine the level of jitter on the target system will generate data to allow an averaging scheme to be determined.

Two final notes:

Firstly, a long sample time of eight counts is recommended, due to the impedance and settling time of the internal references.

Secondly, if the system proves to be very noisy, then the answer is to eliminate the coupling of the noise, not to run excessive sample numbers. If it appears that more than 30 reads are required, then it is likely that noise is getting onto the references. Even though averaging could be used to provide the ideal calibration, each read during operation would be subject to a similar level of noise. Such interference would negate the effort in determining the ideal calibration.

3.4 ADC Internal Registers

There are six 16-bit registers in each MPC5500 ADC engine: one for conversion command messages, and five control registers for configuring the ADC. These registers are non memory-mapped and are only accessible via read and write command messages. To read a value from an ADC register, a read command message with the address of the register is sent to the ADC. The contents are returned as a 16-bit result. To write to an ADC register, the 16-bit value is embedded in a write command message, along with the address of the ADC register. The header field of the read or write command message directs the register value to the appropriate ADC engine.

Table 2. ADC Engine Internal Register Map

Register Address	Register Name	Description
0x00	—	Used for Conversion Commands
0x01	ADCn_CR	Control Register
0x02	ADCn_TSCR	Time Stamp Control Register
0x03	ADCn_TBCR	Time Base Control Register
0x04	ADCn_GCCR	Gain Calibration Control Register
0x05	ADCn_OCCR	Offset Calibration Control Register
0x06 – 0xFF	Reserved	—

Two of these non memory-mapped ADC internal registers are involved in ADC calibration; the gain calibration constant and the offset calibration constant registers. These must be written by an application software calibration routine before the MAC calibration unit can be invoked by conversion command messages.

3.4.1 Gain Calibration Constant Register (ADCn_GCCR)

The gain calibration constant register is a 15-bit fixed point register. It has fifteen significant bits in two bit fields. Bit 1, GCC_INT, contains the integer portion of the gain value. Bits 2–15, GCC_FRAC, contain the fractional portion of the gain value. There is an independent GCC register for each ADC0 and ADC1.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	GCC_INT	GCC_FRAC													

Bit(s)	Bit Field Name	Bit Field Description
0	—	Reserved
1	GCC_INT	Integer part of the gain calibration constant for ADCn.
2–15	GCC_FRAC	Fractional part of the gain calibration constant for ADCn. GCC_FRAC can express decimal values ranging from 0 to 0.999938964: $2^{-1} + 2^{-2} + 2^{-3} + 2^{-4} + 2^{-5} + 2^{-6} + 2^{-7} + 2^{-8} + 2^{-9} + 2^{-10} + 2^{-11} + 2^{-12} + 2^{-13} + 2^{-14}$

3.4.2 Offset Calibration Register (ADCn_OCCR)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	OCC													

Bit(s)	Bit Field Name	Bit Field Description
0–1	—	Reserved
2–15	OCC	ADCn offset calibration constant. Contains the offset calibration constant used to fine-tune ADCn conversion results. Negative values should be expressed using the twos complement representation.

3.5 Calculating the Calibration Coefficients

3.5.1 Obtaining the Equation Terms

The following steps can be used to calculate GCC and OCC. The ideal 75% VREF and 25% VREF conversion values of a 12-bit ADC are known:

$$2^{12} * 75\% = 3072 = 0xC00; 0xC00 \ll 2 = 0x3000$$

$$2^{12} * 25\% = 1024 = 0x400; 0x400 \ll 2 = 0x1000$$

$$\text{IDEAL_RES75} = 0x3000$$

$$\text{IDEAL_RES25} = 0x1000$$

It is necessary to left-shift the ideal result twice to match the left shifts applied by the MAC on the actual reference reads. All terms are thus 14 bits. Note that, in the reference manual, the ideal results are called CAL_RES 75% VREF and CAL_RES 25% VREF.

Now we must determine the uncalibrated (raw) performance of the ADC. We do this by reading the two reference points, 25% (VRH-VRL) and 75% (VRH-VRL). These conversions must be done using conversion command words with CAL = 0.

```
0x00042B00; // BN=0 ADC0
             // CAL=0
             // Message Tag=0 Result in RFIFO 0
             // LST=0b01 8 ADC clks
             // Convert 75% VREF channel 43(0x2B)

0x00042C00; // BN=0 ADC0
             // CAL=0
             // Message Tag=0 Result in RFIFO 0
             // LST=0b01 8 ADC clks
             // Convert 25% VREF channel 44(0x2C)
```

These conversion command words will return the uncalibrated conversion results for the ADC: `raw_res75` and `raw_res25`. They will be compared against the ideal results for these reference voltages.

3.5.2 Gain Calibration Coefficient

Calculate the gain calibration coefficient (GCC) using the familiar formula for the slope of a line:

$$m = (y_2 - y_1) / (x_2 - x_1) \quad \text{Eqn. 2}$$

Using the terms of the ADC, the equivalent formula is:

$$\text{GCC} = (\text{IDEAL_RES75} - \text{IDEAL_RES25}) / (\text{raw_res75} - \text{raw_res25}) \quad \text{Eqn. 3}$$

GCC is an unsigned 15-bit fixed point value, with a single-bit integer portion and 14-bit fractional portion: `GCC_int.GCC_frac`. GCC can have the range 0 to 1.999938964 (see [Section 3.4.1, “Gain Calibration Constant Register \(ADCn_GCCR\)”](#)). Numerically, this is a 15-bit integer divided by 2^{14} .

The result of [Equation 3](#) must be converted from the native floating point or integer format into the 15-bit format of the GCC register. There are several ways in which this can be done.

The fastest way is using long (32-bit) integer arithmetic. We can use the formula:

$$(a * 2^{14})/b \quad \text{Eqn. 4}$$

The equivalent of this in C code is:

```
gcc = ((IDEAL_RES75 - IDEAL_RES25)<<14) / (raw_res75 - raw_res25)
```

It will be observed that the bracketed ideal half of this equation is a constant.

```
(IDEAL_RES75 - IDEAL_RES25) = 0d8192 = 0x2000
```

making the equation:

```
gcc = 134217728/(raw_res75 - raw_res25);
```

If we are using floating-point arithmetic, we can use the formula $(a/b) * 2^{14}$ and round the result to the nearest integer. This 15-bit integer can then be written directly into the GCC register:

```
gcc = (IDEAL_RES75 - IDEAL_RES25) / (raw_res75 - raw_res25) * 16384
```

Or, using constants:

```
gcc = 8192/(raw_res75 - raw_res25) * 16384;
```

which can be simplified further into the same code as used with integers.

3.5.3 Offset Calibration Coefficient

Solving the MAC equation for OCC, gives us the formulae below:

$$\text{OCC} = \text{IDEAL_RES75} - \text{GCC} * \text{raw_res75} - 2 \quad \text{Eqn. 5}$$

or

$$\text{OCC} = \text{IDEAL_RES25} - \text{GCC} * \text{raw_res25} - 2 \quad \text{Eqn. 6}$$

The -2 is there to take account of the half LSB correction always added by the MAC. Since it is present on the raw results, we must remove it to calculate the appropriate OCC.

In C code, using long integer arithmetic, this can be done simply as:

```
occ = IDEAL_RES75 - ((gcc * raw_res75)>>14) - 2;
```

The OCC can be a positive or negative number, centred around zero (ideal offset). Negative numbers are represented as two's complement. The ADC register for OCC is a 14-bit value. Writing the 16-bit or 32-bit result to the 14-bit register will simply ignore the leading bits.

If OCC and GCC are defined as a floats, then the integer portion of occ

```
occ = IDEAL_RES75 - (gcc * raw_res75) - 2;
```

is the correct value to be programmed to the OCC register. Ideally it would be rounded rather than truncated.

Once the GCC and OCC values are known and formatted, they should be written to the appropriate ADC internal registers for use by the MAC when calibrated conversion results are required. Write GCC value to ADCn gain calibration registers and OCC value to ADCn offset calibration constant registers using write configuration commands.

NOTE

There will be different GCC and OCC values for each ADC; therefore, the correct values must be identified independently and written to the appropriate ADC.

4 Conclusion

The RSD architecture of ADC combined with the hardware MAC delivers a significant advance in ADC performance for the MPC5500 family. Accuracy is beyond previously achievable levels — in fact, beyond the capability of typical measuring equipment. Linearity is particularly good, providing closed loop sensors with precision of 2 or 3 mV. This available precision brings with it the need to understand the application. In many cases, the ADC accuracy will be better than the accuracy of the control system to which it is fitted; hence, the discussion on sensor types, total accuracy, linearity, and experiment to determine noise floor.

Once the task is understood, and the required accuracy determined, the code required to achieve calibration is fairly straightforward. Some simple arithmetic is required, and some reformatting of the results to suit the hardware MAC. The processes described achieve calibration either just on power up, or during run time.

5 References

1. MPC5553/MPC5554 Microcontroller Reference Manual (MPC5553/54RM)
2. ADC Definitions and Specifications, J. Feddeler and Bill Lucas (AN2438)

NOTE

Reference 1 uses integral non linearity (INL) for linear error, differential non linearity for monotonicity, and total unadjusted error (TUE) after calibration for absolute error. To understand these terms in detail, refer to the document reference 2.

6 Glossary

ADC	analog to digital converter
CSD	characteristic shift downward
DAC	digital to analog converter
DMA	direct memory access
DNL	differential non linearity
ECU	engine control unit
FMEA	failure mode and effect analysis
GCC	offset gain coefficient
INL	integral non linearity
LSB	less/least significant bit
MAC	multiply accumulate
MEMS	micro-electromechanical system
MSB	more/most significant bit
OCC	offset calibration coefficient
RSD	redundant signed digit
SA	successive approximation
TUE	total unadjusted error

Appendix A Software Example

```

#include "mpc5554.h"
#include "common.h"

#define REGSIZE 15

int cal_eqadc(uint16_t dac)
{
    const IDEAL_RES75 = 0x3000;
    const IDEAL_RES25 = 0x1000;

    unsigned long raw_res75;
    unsigned long raw_res25;

    unsigned long gcc, occ;
        float gccf, occf;

    unsigned short gain;
        short offset;

    unsigned long i, numerator, denominator;
    unsigned long bit, mask;

    // configure Queue0 and send conversion command messages
    EQADC.CFCR[0].B.MODE = SW_TRIG_SS;
    EQADC.CFCR[0].B.SSE = 1; /* software trigger */
    EQADC.CFPR[0].R = 0x00042b00 + (dac<<25); /* Convert channel 43 (75%VRH)with CAL=0 */
    EQADC.CFPR[0].R = 0x00142c00 + (dac<<25); /* Convert channel 44 (25%VRH)with CAL=0 */

    // wait for results
    while (EQADC.FISR[0].B.RFDF==0) {};
    raw_res75 = EQADC.RFPR[0].R; /* 75%VRH */
    while (EQADC.FISR[1].B.RFDF==0) {};
    raw_res25 = EQADC.RFPR[1].R; /* 25%VRH */

/* ----- Method 1: (Fastest) ----- */

    gcc = occ = 0;

    gcc = ((IDEAL_RES75-IDEAL_RES25)<<14)/(raw_res75-raw_res25);
    occ = IDEAL_RES75 - ((gcc * raw_res75)>>14) - 2;

    gain = gcc;
    offset = occ;

/*----- End of Method 1 ----- */

```

Software Example

```
/* ----- Method 2: Floating Point (Slowest) ----- */

gccf = (8192.0)/(raw_res75 - raw_res25);
gain = gccf * 16384;
occf = IDEAL_RES75 - (gccf * raw_res75) - 2.0;
offset = occf;

/*----- End of Method 2 ----- */

/* ----- Method 3 (Alternative) ----- */
bit = 1 << (REGSIZE - 1);
mask = (1 << REGSIZE) - 1;

gcc = occ = 0;

denominator = raw_res75 - raw_res25;
numerator = IDEAL_RES75 - IDEAL_RES25;

for(i = 0; i < REGSIZE; i++)
{
    if(numerator >= denominator)
        {
            numerator -= denominator;
            if(numerator > denominator)
                {
                    //GCC is greater than 1.999 and we have an overflow
                    return 0;
                }
            gcc |= bit;
            occ += raw_res75;
        }
    bit >>= 1;
    numerator <<= 1;
    occ <<= 1;
    numerator &= mask;
}
occ = occ >> REGSIZE;
occ = IDEAL_RES75 - occ - 2;

gain = gcc;
offset = occ;

/* ----- End of Method 3 ----- */

// now store Gain in ADCn_GCCR
EQADC.CFPR[0].R = (gain<<8) | 0x00000004 | (dac<<25);

// now store offset in ADCn_OCCR
EQADC.CFPR[0].R = ((unsigned short)offset << 8) | 0x00000005 | (dac<<25);

return 1;
}
```


Appendix B Example Coefficient Calculations

In this example, the results of multiple reads and averaging is:

$$\text{raw_res75} = 0x2FE7 = 12263, \text{raw_res25} = 0x100D = 4109$$

and we know:

$$\text{IDEAL_RES75} = 0x3000 = 12288, \text{IDEAL_RES25} = 0x1000 = 4096$$

and to recap on the two equations:

$$\text{GCC} = (\text{IDEAL_RES75} - \text{IDEAL_RES25}) / (\text{raw_res75} - \text{raw_res25}) \quad \text{Eqn. 7}$$

$$\text{OCC} = \text{IDEAL_RES75} - \text{GCC} * \text{raw_res75} - 2 \quad \text{Eqn. 8}$$

Floating point version:

Define variables gccf and occf as floats, with gcc and occ the values to be programmed into the registers. The integer ADC values in decimal are shown for visibility.

$$\text{gccf} = 8192 / (12263 - 4109) = 1.00466$$

$$\text{occf} = 12288 - (1.00466 * 12263) - 2 = -34.149$$

To get the register format:

$$\text{gcc} = \text{ROUND}(1.00466 * 16384) = 16460$$

$$\text{occ} = \text{ROUND}(-34.149) = -34$$

Integer version:

Define variables gcc and occ as ints. The integer ADC values in decimal are shown for visibility. Note that $(\text{IDEAL_RES75} - \text{IDEAL_RES25}) \ll 14$ is a constant 134217728.

$$\text{gcc} = 134217728 / (12263 - 4109) = 16460$$

$$\text{occ} = 12288 - (16460 * 12263) / 16384 - 2 = -33$$

Note that the integer gcc result was automatically truncated from 16460.354, giving an occ that is again truncated from -33.88, hence the slight difference in occ.

How to Reach Us:

Home Page:

www.freescale.com

E-mail:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
+1-800-521-6274 or +1-480-768-2130
support@freescale.com

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064
Japan
0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447 or 303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters that may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals", must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners.

© Freescale Semiconductor, Inc. 2005. All rights reserved.