
AVR120: Characterization and Calibration of the ADC on an AVR

APPLICATION NOTE

Introduction

This application note explains various ADC (Analog to Digital Converter) characterization parameters given in the datasheets and how they effect ADC measurements. It also describes how to measure these parameters during application testing in production and how to perform run-time compensation for some of the measured deviations.

A great advantage with the Flash memory of the Atmel® AVR® is that calibration code can be replaced with application code after characterization. Therefore, no code space is consumed by calibration code in the final product.

Features

- Understanding Analog to Digital Converter (ADC) characteristics
- Measuring parameters describing ADC characteristics
- Temperature, frequency, and supply voltage dependencies
- Compensating for offset and gain error

Table of Contents

Introduction.....	1
Features.....	1
1. Theory.....	3
1.1. General ADC Concepts	3
1.2. Conversion Ranges.....	4
1.2.1. Single-ended Conversion Range.....	4
1.2.2. Differential Conversion Range.....	5
1.3. The Need for Calibration.....	5
1.4. Absolute Error.....	5
1.5. Offset Error.....	6
1.5.1. Offset Error - Single-ended Channels.....	6
1.5.2. Offset Error - Differential Channels.....	7
1.6. Gain Error.....	8
1.7. Non-linearity.....	9
1.7.1. Differential Non-linearity.....	9
1.7.2. Integral Non-linearity.....	10
1.7.3. Measurements and Compensation.....	10
1.8. Temperature, Frequency, and Supply Voltage Dependencies.....	10
1.9. Bandwidth and Input Impedance.....	10
2. Implementation.....	11
2.1. Fixed-point Arithmetic for Offset and Gain Error Compensation	11
2.1.1. Calibration.....	12
2.1.2. Compensation.....	13
3. Literature references.....	15
4. Revision History.....	16

1. Theory

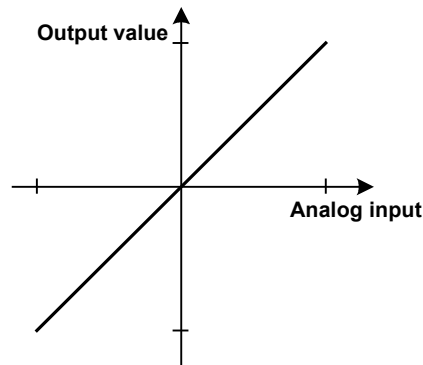
Before getting into the details, some central concepts need to be introduced. The following section ("General ADC Concepts") can be skipped if quantization, resolution, and ADC transfer functions are familiar to the reader

1.1. General ADC Concepts

The ADC translates an analog input signal to a digital output value representing the size of the input relative to a reference. To get a better basis for describing a general ADC, this document distinguishes between ideal, perfect, and actual ADCs.


An ideal ADC is just a theoretical concept, and cannot be implemented in real life. It has infinite resolution, where every possible input value gives a unique output from the ADC within the specified conversion range. An ideal ADC can be described mathematically by a straight-line transfer function, as shown in the figure below.

Figure 1-1. Transfer Function of an Ideal ADC



To define a perfect ADC, the concept of quantization must be introduced. Due to the digital nature of an ADC, continuous output values are not possible. The output range must be divided into a number of steps, one for each possible digital output value. This means that one output value does not correspond to a unique input value, but a small range of input values. This results in a staircase transfer function. The resolution of the ADC equals the number of unique output values. For instance, an ADC with eight output steps has a resolution of eight levels or, in other words, three bits. The transfer function of an example 3-bit perfect ADC is shown in the figure below together with the transfer function of an ideal ADC. As seen on the figure, the perfect ADC equals the ideal ADC on the exact midpoint of every step. This means that the perfect ADC essentially rounds input values to the nearest output step value.

Figure 1-2. Transfer Function of a 3-bit Perfect ADC



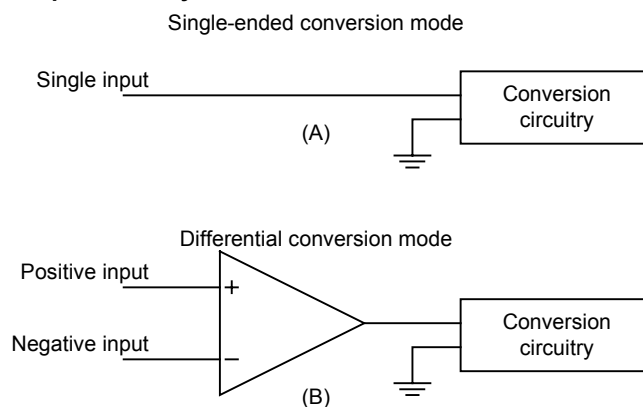
The maximum error for a perfect ADC is $\pm\frac{1}{2}$ step. In other words, the maximum quantization error is always $\pm\frac{1}{2}$ LSB, where LSB is the input voltage difference corresponding to the Least Significant Bit of the output value. Real ADCs have other sources of errors, described later in this document.

1.2. Conversion Ranges

The ADC in Atmel AVR devices can be configured for single-ended or differential conversion. Single-ended mode is used to measure input voltages on a single input channel, while differential mode is used to measure the difference between two channels. Regardless of conversion mode, input voltages on any channel must stay between GND and AV_{CC} .

When using single-ended mode, the voltage relative to GND is converted to a digital value. Using differential channels, the output from a differential amplifier (with an optional gain stage) is converted to a digital value (possibly negative). A simplified illustration of the input circuitry is shown in the figure below.

Figure 1-3. Simplified ADC Input Circuitry



To decide the conversion range, the conversion circuitry needs a voltage reference (V_{REF}) to indicate the voltage level corresponding to the maximum output value. According to the datasheets, V_{REF} should be at least 2.0V for standard devices, while devices operating down to 1.8V can use a reference voltage down to 1.0V. This applies both to single-ended and differential mode. Consult the datasheet for details.

1.2.1. Single-ended Conversion Range

Single-ended conversion feeds the input channel directly to the conversion circuitry, as shown in the figure [Figure 1-3 \(A\)](#) above. The 10-bit ADC of the AVR therefore converts continuous input voltages from GND to V_{REF} to discrete output values from 0 to 1023.

Any applied input voltage greater than the reference voltage V_{REF} will return the maximum value (1023 using 10-bit ADC), and any negative input voltage will return 0.

1.2.2. Differential Conversion Range

Differential conversion feeds the two input channels to a differential amplifier with an optional gain stage. The output from the amplifier is then fed to the conversion logic, as shown in the [Figure 1-3\(B\)](#). Voltage differences from $-V_{REF}$ to $+V_{REF}$ therefore results in discrete output values from -512 to $+511$. The digital output is represented in 2's complement form. Even when measuring negative voltage differences, the voltages applied to the input channels themselves must stay between GND and AV_{CC} .

Voltage differences smaller than $-V_{REF}$ will return the minimum value (-512 using 10-bit ADC), and voltage differences greater than $+V_{REF}$ will return the maximum value (512 using 10-bit ADC).

Note: Some devices cannot measure negative differences, e.g. Atmel ATtiny26.

1.3. The Need for Calibration

The total error of the actual ADC comes from more than just quantization error. This document describes offset and gain errors and how to compensate for them. It also describes two measures for non-linearity; differential and integral non-linearity.

For most applications, the ADC needs no calibration when using single ended conversion. The typical accuracy is 1-2 LSB, and it is often neither necessary nor practical to calibrate for better accuracies.

However, when using differential conversion the situation changes, especially with high gain settings. Minor process variations are scaled with the gain stage and give large parameter differences from part to part. The uncompensated error is typically above 20 LSB. These variations must be characterized for every device and compensated for in software.

At first sight 20 LSB seems to be a large value, but does not mean that differential measurements are impractical to use. Using simple calibration algorithms, accuracies of typically 1-2 LSB can be achieved.

1.4. Absolute Error

The absolute error is the maximum deviation between the ideal straight line and the actual transfer function, including the quantization steps. The minimum absolute error is therefore $\frac{1}{2}$ LSB, due to quantization.

Absolute error or absolute accuracy is the total uncompensated error and includes quantization error, offset error, gain error, and non-linearity. Offset, gain, and non-linearity are described later in this document.

Absolute error can be measured using a ramp input voltage. In this case all output values are compared against the input voltage. The maximum deviation gives the absolute error.

Note: Absolute error cannot be compensated for directly, without using e.g. memory-expensive look-up tables or polynomial approximations. However, the most significant contributions to the absolute error - the offset and gain error - can be compensated for.


Be aware that the absolute error represent a reduction in the ADC range, and one should therefore consider the margins to the minimum and maximum input values to avoiding clipping keeping the absolute error in mind.

1.5. Offset Error

The offset error is defined as the deviation of the actual ADC's transfer function from the ideal straight line at zero input voltage.

When the transition from output value 0 to 1 does not occur at an input value of $\frac{1}{2}$ LSB, then we say that there is an offset error. With positive offset errors, the output value is larger than 0 when the input voltage approaches $\frac{1}{2}$ LSB from below. With negative offset errors, the input value is larger than $\frac{1}{2}$ LSB when the first output value transition occurs. In other words, if the actual transfer function lies below the ideal line, there is a negative offset and vice versa. Negative and positive offsets are shown in Figure below.

Figure 1-4. Examples of Positive (A) and Negative (B) Offset Errors



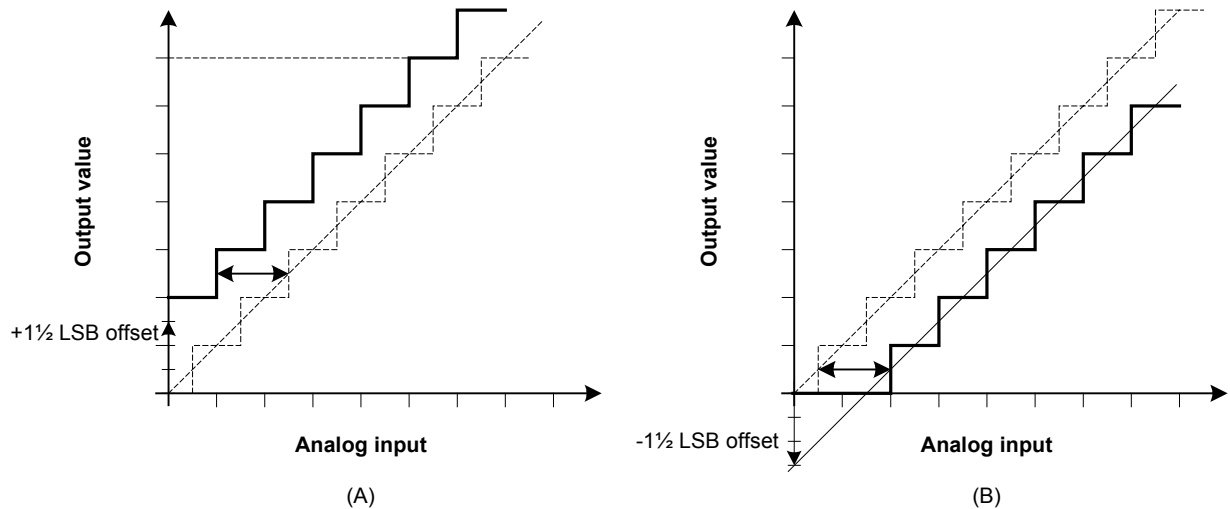
Since single-ended conversion gives positive results only, the offset measurement procedures are different when using single-ended and differential channels.

1.5.1. Offset Error - Single-ended Channels

To measure the offset error, increase the input voltage from GND until the first transition in the output value occurs. Calculate the difference between the input voltage for which the perfect ADC would have shown the same transition and the input voltage corresponding to the actual transition. This difference, converted to LSB, equals the offset error.

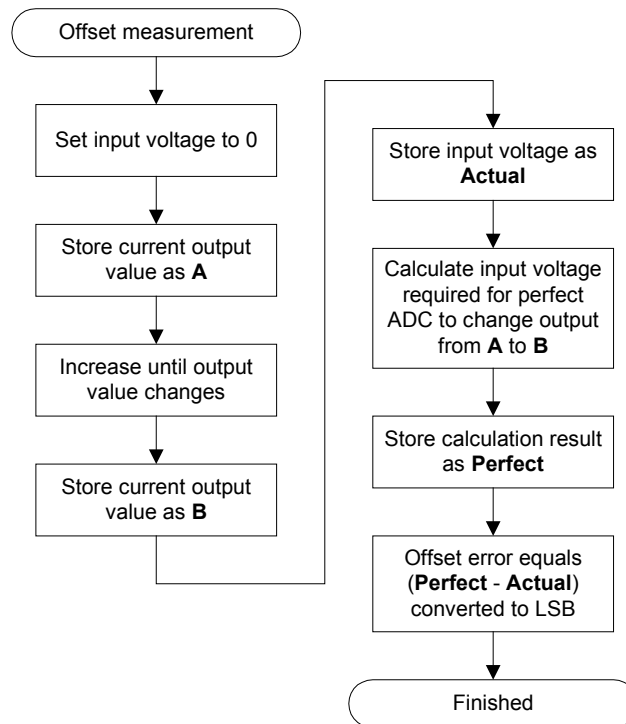
In the [Figure 1-4 \(A\)](#), the first transition occurs at 1 LSB. The transition is from 2 to 3, which equals an input voltage of $2\frac{1}{2}$ LSB for the perfect ADC. The difference is $+1\frac{1}{2}$ LSB, which equals the offset error. The double-headed arrows show the differences. The same procedure applies to the [Figure 1-4 \(B\)](#). The first transition occurs at 2 LSB. The transition is from 0 to 1, which equals an input voltage of $\frac{1}{2}$ LSB for the perfect ADC. The difference is $-1\frac{1}{2}$ LSB, which equals the offset error.

Figure 1-5. Positive (A) and Negative (B) Offset Errors in Single-ended Mode



The measurement procedure can be formalized as described in the figure below.

Figure 1-6. Flowchart for Measuring Single-ended Offset Errors



To compensate for offset errors when using single ended channels, subtract the offset error from every measured value. Be aware that offset errors limit the available range for the ADC. A large positive offset error causes the output value to saturate at maximum before the input voltage reaches maximum. A large negative offset error gives output value 0 for the smallest input voltages.

1.5.2. Offset Error - Differential Channels

With differential channels, the offset measurement can be performed much easier since no external input voltage is required. The two differential inputs can be connected to the same voltage internally and the resulting output value is then the offset error. Since this method gives no exact information on where the first transition occurs, it gives an error of $\frac{1}{2}$ to 1 LSB (worst case).

To compensate for offset errors when using differential channels, subtract the offset error from every measured value.

1.6. Gain Error

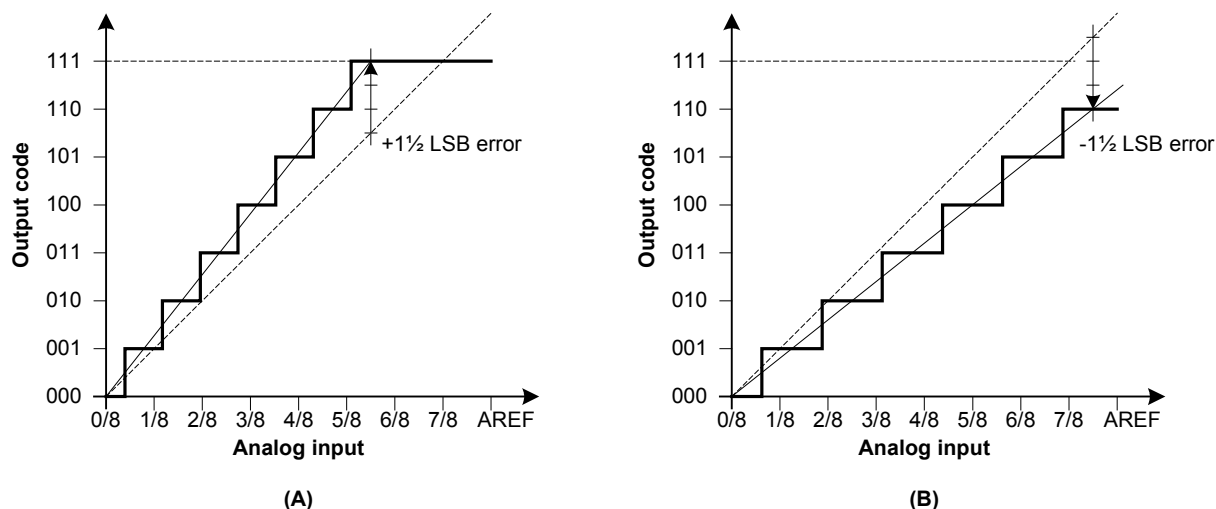
The gain error is defined as the deviation of the last output step's midpoint from the ideal straight line, after compensating for offset error.

After compensating for offset errors, applying an input voltage of 0 always give an output value of 0. However, gain errors cause the actual transfer function slope to deviate from the ideal slope. This gain error can be measured and compensated for by scaling the output values.

Run-time compensation often uses integer arithmetic, since floating point calculation takes too long to perform. Therefore, to get the best possible precision, the slope deviation should be measured as far from 0 as possible. The larger the values, the better precision you get. This is described in detail later in this document.

The example of a 3-bit ADC transfer functions with gain errors is shown in the [Figure 1-7](#). The following description holds for both single-ended and differential modes.

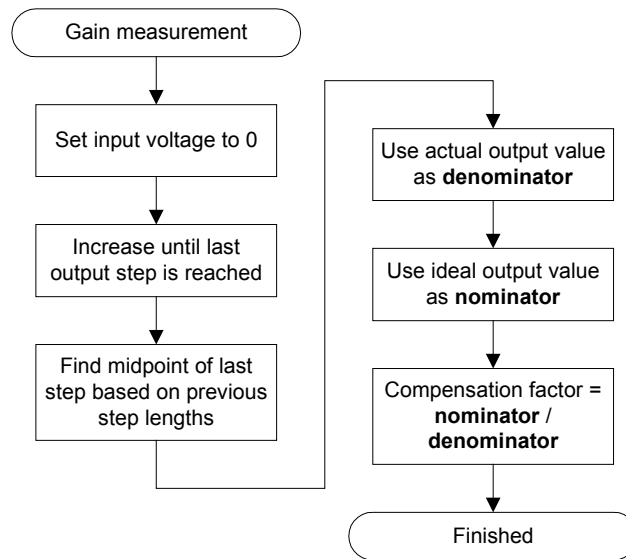
Figure 1-7. Examples of Positive (A) and Negative (B) Gain Errors



To measure the gain error, the input value is increased from 0 until the last output step is reached. The scaling factor for gain compensation equals the ideal output value for the midpoint of the last step divided by the actual value of the step.

In the [Figure 1-7 \(A\)](#), the output value saturates before the input voltage reaches its maximum. The vertical arrow shows the midpoint of the last output step. The ideal output value at this input voltage should be 5.5, and the scaling factor equals 5.5 divided by 7. In the [Figure 1-7 \(B\)](#), the output value has only reached 6 when the input voltage is at its maximum. This results in a negative deviation for the actual transfer function. The ideal output value for the midpoint of the last step is 7.5 in this case. The scaling factor now equals 7.5 divided by 6. The measurement procedure is illustrated in the flowchart below.


Figure 1-8. Flowchart for Measuring Gain Errors



1.7. Non-linearity

When offset and gain errors are compensated for, the actual transfer function should be equal to the transfer function of perfect ADC. However, non-linearity in the ADC may cause the actual curve to deviate slightly from the perfect curve, even if the two curves are equal around 0 and at the point where the gain error was measured. There are two methods for measuring non-linearity, both described below. The figure below shows examples of both measures.

Figure 1-9. Example of a Non-linear ADC Conversion Curve



1.7.1. Differential Non-linearity

Differential Non-Linearity (DNL) is defined as the maximum and minimum difference between the step width and the perfect width (1 LSB) of any output step.

Non-linearity produces quantization steps with varying widths. All steps should be 1 LSB wide, but some are narrower or wider.

To measure DNL, a ramp input voltage is applied and all output value transitions are recorded. The step lengths are found from the distance between the transitions, and the most positive and negative deviations from 1 LSB are used to report the maximum and minimum DNL.

1.7.2. Integral Non-linearity

Integral Non-Linearity (INL) is defined as the maximum vertical difference between the actual and the perfect curve.

INL can be interpreted as a sum of DNLs. E.g. several consecutive negative DNLs raise the actual curve above the perfect curve as shown in the [Figure 1-9 \(A\)](#). Negative INLs indicate that the actual curve is below the perfect curve.

The maximum and minimum INL are measured using the same ramp input voltage as in DNL measurement. Record the deviation at each conversion step midpoint and report the most positive and negative deviations as maximum and minimum INL.

1.7.3. Measurements and Compensation

It is important that DNL and INL values are measured after offset and gain error compensation. If not, the results will be infected by the offset and gain error and thus not reveal the true DNL and INL.

Non-linearity cannot be compensated for with simple calculations. Polynomial approximations or table lookups can be used for that purpose. However, the typical DNL and INL values are $\frac{1}{2}$ LSB for the 10-bit ADC of the AVR, and are rarely of any concern in real life applications.

1.8. Temperature, Frequency, and Supply Voltage Dependencies

When using the internal voltage reference with the ADC, the accuracy of this reference must be considered. The internal voltage reference is proportional to the bandgap voltage, which is characterized in the devices' datasheets. The characteristics show that the bandgap voltage is slightly dependent on supply voltage and operating temperature.

The ADC accuracy also depends on the ADC clock. The recommended maximum ADC clock frequency is limited by the internal DAC in the conversion circuitry. For optimum performance, the ADC clock should not exceed 200kHz. However, frequencies up to 1MHz do not reduce the ADC resolution significantly.

Operating the ADC with frequencies greater than 1MHz is not characterized.

1.9. Bandwidth and Input Impedance

When using single-ended mode, the ADC bandwidth is limited by the ADC clock speed. Since one conversion takes 13 ADC clock cycles, a maximum ADC clock of 1MHz means approximately 77ksps (kilo samples per second). This limits the bandwidth in single-ended mode to 38.5kHz, according to the Nyquist sampling theorem.

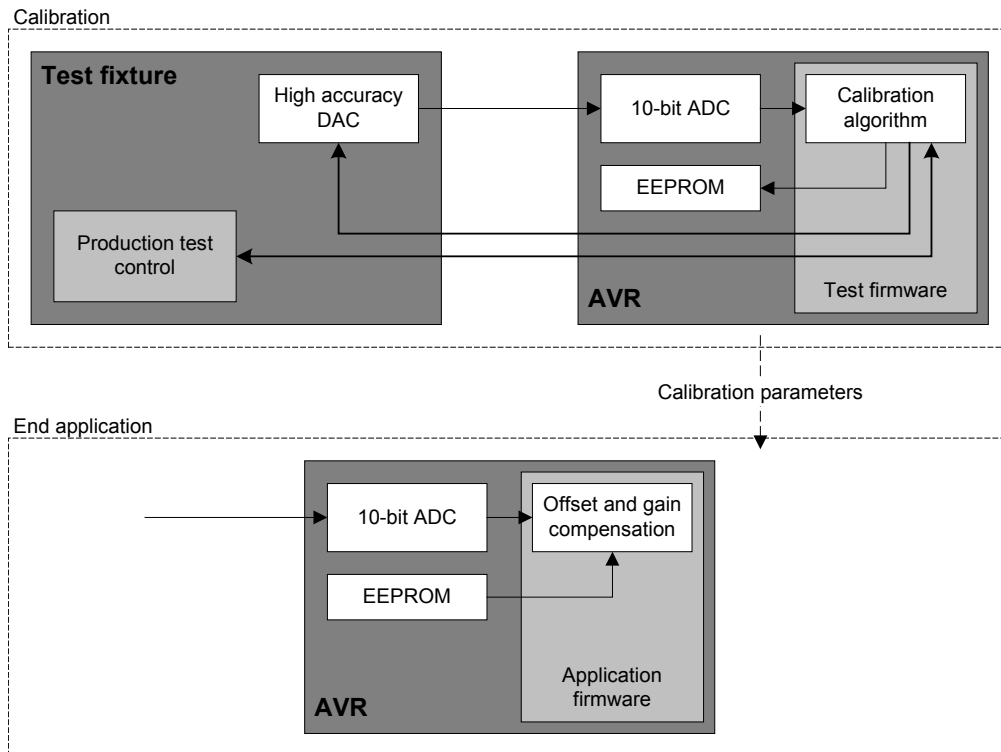
When using differential mode, the bandwidth is limited to 4kHz by the differential amplifier. Input frequency components above 4kHz should be removed by an external analog filter, to avoid non-linearity's.

The input impedance to V_{CC} and GND is typically 100M Ω . Together with the output impedance of the signal source, this creates a voltage divider. The signal source should therefore have sufficiently low output impedance to get correct conversion results.

2. Implementation

An example setup for calibrating the ADC is shown in the figure below.

Figure 2-1. Production Calibration Setup



During the production test phase, each device's ADC must be characterized using a test setup similar to this. When the test fixture is ready for calibrating the AVR, the tester signals the AVR to start calibrating itself. The AVR uses the test fixture's high accuracy DAC (e.g. 16-bit resolution) to generate input voltages to the calibration algorithm. When calibration is finished, the parameters for offset and gain error compensation are programmed into the EEPROM for later use, and the AVR signals that it is ready for other test phases.

Note: This requires the EESAVE fuse to be programmed, so that the Flash memory can be reprogrammed without erasing the EEPROM. Otherwise, the ADC parameters must be temporarily stored by the programmer while erasing the device.

2.1. Fixed-point Arithmetic for Offset and Gain Error Compensation

Floating-point arithmetic is not an efficient way of scaling the ADC values. However, the scaling factor for gain error compensation will be somewhere close to 1, and needs a certain precision to yield good compensated ADC values. Therefore, fixed-point numbers represented by integers can be used.

Since the gain compensation factor certainly never exceeds 2, it can be scaled by a factor of 214 to fit exactly in a signed 16-bit word. In other words, the scaling factor can be represented by two bytes as a 1:14 signed fixed-point number.

The equation for both offset and gain error compensation is shown in Equation 1.

Equation 1:

$$\text{realvalue} = (\text{adcvalue} - \text{offset}) \cdot \text{gainfactor}$$

When the calculation result is truncated to an ordinary integer later, it is always truncated to the largest integer less than or equal to the result. In order to achieve correct rounding to the nearest integer, 0.5 must be added before truncating. Adding 0.5, scaling the equation by 2^{14} and moving the offset correction outside the parentheses gives Equation 2.

Equation 2:

$$2^{14} \cdot \text{realvalue} = 2^{14} \cdot \text{adcvalue} \cdot \text{gainfactor} + 2^{14} \cdot 0.5 - 2^{14} \cdot \text{offset} \cdot \text{gainfactor}$$

Since the gain factor and offset correction value are constants, further optimization can be achieved. In addition, if the result is scaled by 2^2 , giving a total scale of 2^{16} , the upper two bytes of the result equals the truncated integer without the need for a 16-step right-shift.

We introduce some constants and summarize it all in Equation 3.

Equation 3:

$$\text{factor} = 2^{14} \cdot \text{gainfactor},$$

$$\text{correction} = 2^{14} \cdot (0.5 - \text{offset} \cdot \text{gainfactor}),$$

$$2^{16} \cdot \text{realvalue} = 2^2 \cdot (\text{adcvalue} \cdot \text{factor} + \text{correction})$$

Using this method, the calibration software calculates the constants and stores them in EEPROM memory. At execution time the compensation firmware only needs one integer multiplication, one addition and two left shift operations to correct the ADC values. Using the IAR™ C compiler with highest optimization for speed, this takes 42 CPU cycles.

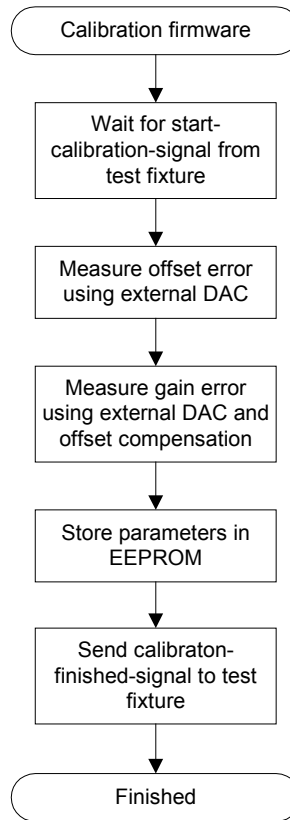
2.1.1. Calibration

The test fixture design is outside the scope of this application note. However, a flowchart of a calibration firmware in the AVR is given. It uses the external DAC through the test fixture and runs its own calibration algorithm.

There is no need for using several ADC channels, only switching between single ended and differential conversion. The ADC parameters are the same regardless of which channel is used. The multiplexer does not introduce any errors.

The firmware could be implemented as shown in the figure below.

Figure 2-2. Flowchart for Calibration Firmware

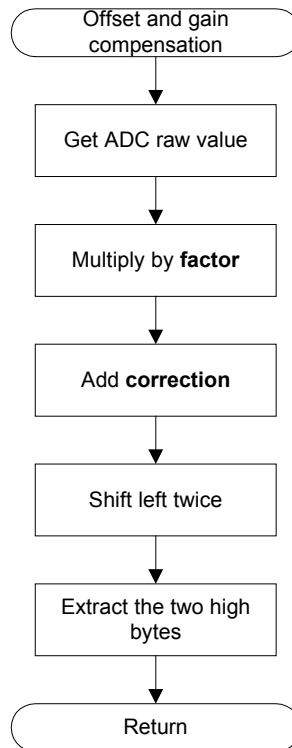


This piece of firmware is programmed into the AVR prior to calibration, and is replaced by the actual application firmware afterward. Once again, the EESAVE fuse must be programmed to preserve the calibration parameters in EEPROM during Flash reprogramming.

2.1.2. Compensation

The run-time compensation code can be implemented as a small function. Every ADC measurement is run through this function, which uses the constants 'factor' and 'correction'.

Figure 2-3. Flowchart for Offset and Gain Compensation



The calculation in the [Figure 2-3](#) can be implemented using the following C function, or alternatively as a macro:

```
signed int adc_compensate (signed int adcvalue, signed int factor, signed long correction)
{
return (((((signed long)adcvalue*factor)+correction)<<2)>>16);
}
```

The parameter 'constants' stored in EEPROM could be copied to SRAM variables during startup for quicker access later.

3. Literature references

- Application Note AVR201 – ["Using the AVR Hardware Multiplier"](#)

4. Revision History

Doc. Rev.	Date	Comments
2559E	08/2016	Converted to new layout
2559D	02/2006	

Atmel®, Atmel logo and combinations thereof, Enabling Unlimited Possibilities®, AVR® and others are registered trademarks or trademarks of Atmel Corporation in U.S. and other countries. Other terms and product names may be trademarks of others.

DISCLAIMER: The information in this document is provided in connection with Atmel products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Atmel products. EXCEPT AS SET FORTH IN THE ATMEL TERMS AND CONDITIONS OF SALES LOCATED ON THE ATMEL WEBSITE, ATMEL ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL ATMEL BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS AND PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF ATMEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Atmel makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and products descriptions at any time without notice. Atmel does not make any commitment to update the information contained herein. Unless specifically provided otherwise, Atmel products are not suitable for, and shall not be used in, automotive applications. Atmel products are not intended, authorized, or warranted for use as components in applications intended to support or sustain life.

SAFETY-CRITICAL, MILITARY, AND AUTOMOTIVE APPLICATIONS DISCLAIMER: Atmel products are not designed for and will not be used in connection with any applications where the failure of such products would reasonably be expected to result in significant personal injury or death ("Safety-Critical Applications") without an Atmel officer's specific written consent. Safety-Critical Applications include, without limitation, life support devices and systems, equipment or systems for the operation of nuclear facilities and weapons systems. Atmel products are not designed nor intended for use in military or aerospace applications or environments unless specifically designated by Atmel as military-grade. Atmel products are not designed nor intended for use in automotive applications unless specifically designated by Atmel as automotive-grade.