DALLAS SEMICONDUCTOR

03/07/00

Dear Weather Station experimenter:

The thermometer used in the Dallas Semiconductor WS-1 Weather Station Experimenters Kit is a Dallas DS1820, 1-Wire™ digital thermometer. Many issues about the DS1820 have been posted to the 1-Wire™ Weather Interest Group, and we feel it is important to try to explain the inaccuracies you may be encountering. Even more importantly, we would also like to introduce the next generation 1-Wire™ Digital Thermometer, the DS18S20, that addresses many of the issues with the DS1820 and offers smaller packaging options at a lower cost.

By design, the DS18S20 corrects for environmental stress-induced drift and a state machine glitch that has caused many of you confusion in the past months. DS18S20 accuracy and stability data will be presented to hopefully convince you the stability problems with the DS1820 have been eliminated. In addition to correcting stability problems, the DS18S20 design also allows for a wider power supply range and smaller packaging options. Incompatibilities between the DS18S20 and the DS1820 are minimal, but a few are discussed in detail. An incompatibility that will have a minor, yet noticeable effect on Weather Station users is the lower resolution of the DS18S20 data converter. We will very soon release a software patch to the WS-1 that performs waveform smoothing. Overall, we think you will be quite pleased with the performance of the newly redesigned and improved DS18S20.

Before getting into detail of the DS18S20, let us attempt to explain some of the peculiarities you may have experienced with the DS1820 in the past. The DS1820 was the first of its kind to offer an integrated solid state temperature sensor, an “analog-to-digital” converter, and a 1-Wire™ Network interface. Before describing the cause for some of the inaccuracies of the DS1820, a brief tutorial on the temperature-sensing algorithm used by the DS1820 would be beneficial. Refer to Fig.1 below:

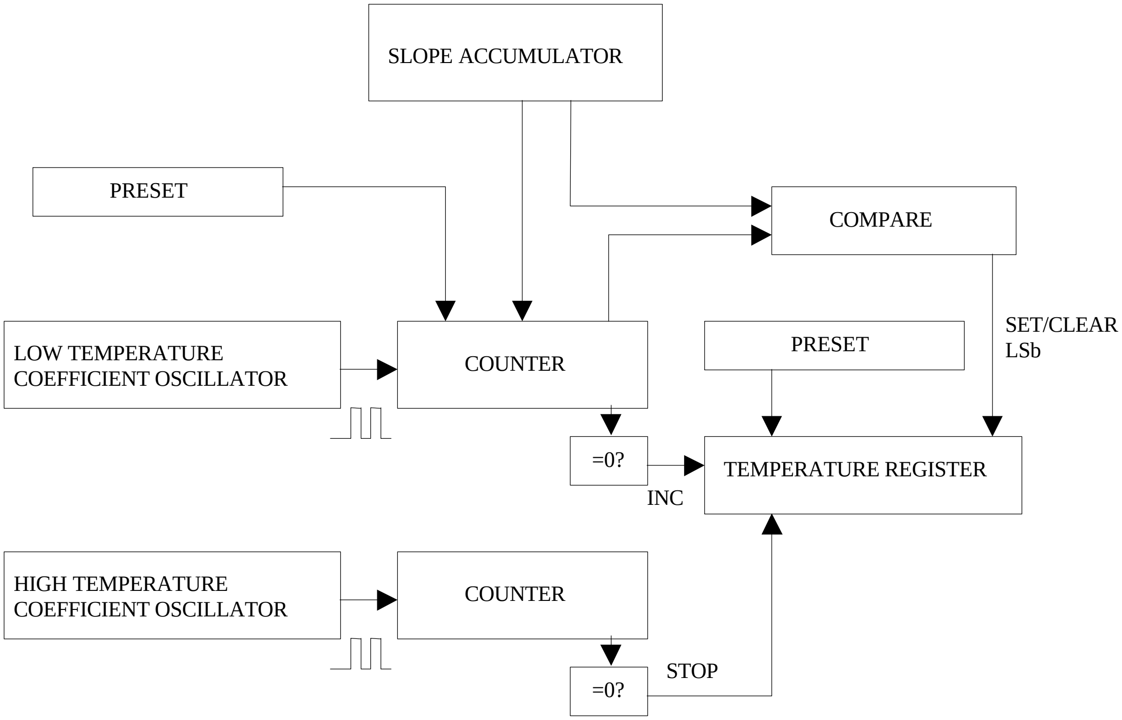

Fig.1. Temperature Sensing Block Diagram

The DS1820 measures temperature by counting the number of clock cycles that an oscillator with a low temperature coefficient goes through during a gate period determined by a high temperature coefficient oscillator. The counter is preset with a base count that corresponds to -55°C. If the counter reaches zero before the gate period is over, the temperature register, which is also preset to the -55°C value, is incremented, indicating that the temperature is higher than -55°C.

At the same time, the counter is then preset with a value determined by the slope accumulator circuitry. This circuitry is needed to compensate for the parabolic behavior of the oscillators over temperature. The counter is then clocked again until it reaches zero. If the gate period is still not finished, then this process repeats.

The slope accumulator is used to compensate for the nonlinear behavior of the oscillators over temperature, yielding a high-resolution temperature measurement. This is done by changing the number of counts necessary for the counter to go through for each incremental degree in temperature. To obtain the desired resolution, therefore, both the value of the counter (COUNT_REMAIN) and the value of the slope accumulator (COUNT_PER_C) at a given temperature must be known. Therefore, attainable resolution (1/COUNT_PER_C) is a function of temperature, and will vary from part to part. A typical COUNT_PER_C ranges from 80 –120 over the DS1820 temperature range, translating to a 14-15-bit resolution.

There are two issues with this algorithm. The first manifests itself as a drift in accuracy, and explains many complaints about the DS1820's temperature error. The solid state oscillators in Fig.1 contain polysilicon resistors, which are extremely difficult to control over temperature and mechanical stress. The temperature coefficient of the resistors is not a problem because we calibrate out their effect in the liquid bath. Mechanical stress, however, is something we cannot control or predict so to account for its effect in the slope accumulator. As the mechanical stress on the DS1820 die changes, so too does the value of the poly resistors in the oscillators. Resistance changes cause oscillator frequency changes, which will ultimately cause changes in the temperature register value. We assemble the DS1820 in a non-hermetic plastic package and calibrate the parts in a liquid bath. Liquid is free to enter the cavity above the die in the non-hermetic package, putting a nominal stress on the die (and resistors). The device is calibrated under this very controlled environment. As the device is removed from the liquid and allowed to dry out, the stress on the die changes and the temperature readout changes. This is why many of you note temperature readout variations in extremely humid environments. Another stress that will cause a drift is a high temperature process such as a vapor phase. Expansion/contraction of the plastic package/leadframe will cause the stress on the die to shift, and again a readout drift will result. A typical drift following a vapor phase is 2.0°C (approximately 4.0°F), and generally the part will read low following the stress. A typical drift following an Autoclave process (device is subjected to pressurized steam (2atm) at 121°C for 168 hours) is over 3.0°C (6.0°F); Autoclave has proven to be the worst case stress insofar as drift is concerned.

This is obviously an extremely frustrating problem. We have tried various layout implementations of the resistors (if the resistance shift in the high tempco oscillator matches that of the resistor in the low tempco oscillator, the effect would cancel), different plastic mold compounds, and the drift could never be completely eliminated.

Many of you have tried numerous experiments in an attempt to minimize the DS1820 thermometer error. Although some of the experiments posted may offer a minimal effect, the dominating factor in the error is the mechanical stress-induced drift described above, and very likely not related to anything you are doing wrong. The cause of the drift is a physical phenomenon that cannot easily be removed from the DS1820. It has been shown that the drift diminishes as the device is baked at 125°C for several hundred hours. After the bake, the user can compare the DS1820 reading against an accurate reference at any temperature (recalibrating the device), and add this offset to future readings. Generally speaking, the drift magnitude is consistent over the entire temperature range, thus allowing for a single-temperature post-bake offset correction.

This however requires access to an oven, a controlled temperature chamber, an accurate reference sensor, and the ability to modify the WS-1 software to add the resulting offset to the DS1820's readout. This solution is understandably a severe restriction for most users. The most feasible solution is for Dallas to redesign the DS1820 thermal algorithm so that mechanical stress will have a negligible effect on stability. The aforementioned DS18S20 does just that by measuring temperature with a completely different algorithm, one that virtually eliminates the effect mechanical stress has on stability. The DS18S20 will be discussed in detail later in this document.

The second problem many of you have noted with the DS1820 is also an artifact of the dual-oscillator algorithm and is also eliminated with the DS18S20. It manifests itself as a "glitch" in the readout of a maximum of 1.0°C. This issue is described in detail in an attached document. This document makes reference to the DS1620, but is directly applicable to the DS1820, as both Dallas sensors use the same thermal-sensing algorithm.

Solving these two issues was of primary concern in redesigning the DS1820 1-Wire™ Digital Thermometer. Technology has improved since the DS1820 was designed, allowing us fabricate the redesign with a 0.6µm minimum geometry process. So, not only are the drift and glitch problems eliminated, but also the new die is significantly smaller which allows for smaller packages and a lower cost sensor. The redesigned thermometer, the DS18S20, will be assembled in a TO-92 package (replaces the DS1820 PR-35) and a 150mil 8-lead SOIC (replaces the DS1820 SSOP).

The block diagram of the DS18S20 thermal sensing algorithm is illustrated below in Fig. 2.

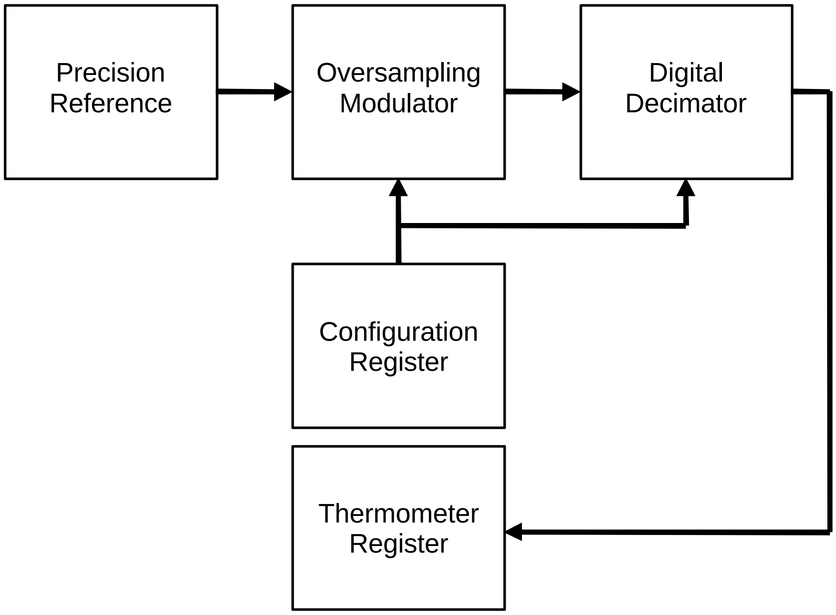

Fig. 2. Bandgap / Sigma-delta modulator temperature sensor

A very precise, on-chip bandgap reference is the first stage of the signal chain. Bandgap voltage references are inherently very stable and are commonly used in solid state temperature sensors, power monitoring products, etc. The sigma delta modulator serves as the data converter in addition to providing chopping which eliminates stress-induced errors in the differential Vbe bandgap. Dallas has this algorithm in production with several digital thermometers – DS75, DS1721, DS1722, DS1775, DS1780, DS1615, and the DS1921 Thermochron. Many of you have tested the DS1921 against the DS1820 and have noted the superior performance of the DS1921.

The DS18S20 will actually go one step further than all these sensors in that it can be bath calibrated (because of its on-chip EEPROM). The most accurate of the non-bath calibrated bandgap-based sensors above are the DS1721, DS1615, and DS1921 at ±1.0°C max error. Bath calibration drastically reduces the error at the calibration temperature; the maximum error over the -10°C to +85°C range for the DS18S20 is ±0.5°C, before, during and after environmental stress.

Fig 3 shows thermometer accuracy data for 4 different samples of DS18S20. Each sample of approximately 50 parts was subjected to one of 4 stresses: High Voltage Life (1000 hours at 6.0V at T=125°C), 3 consecutive passes through a vapor phase, Autoclave (unbiased at 2atm steam T=121°C for 168 hours), and Temperature Cycle (unbiased under 1000cycles from -55°C to +125°C). The data was taken before and after the respective stress and the plots show the mean and ±3σ error (relative to platinum RTD reference in liquid bath) after stress. Mean drift (average difference between post stress and pre stress error) is also plotted. This is a significant improvement over the DS1820. The average part drifts less than 0.1°C (roughly 1 LSb), with insignificant dependence on the type or duration of stress it was subjected to. Recall from earlier in the document, the average drift of the DS1820 following the Autoclave stress (the worst case stress) was over 3.0°C, representing better than a 99% reduction in stress-induced drift.

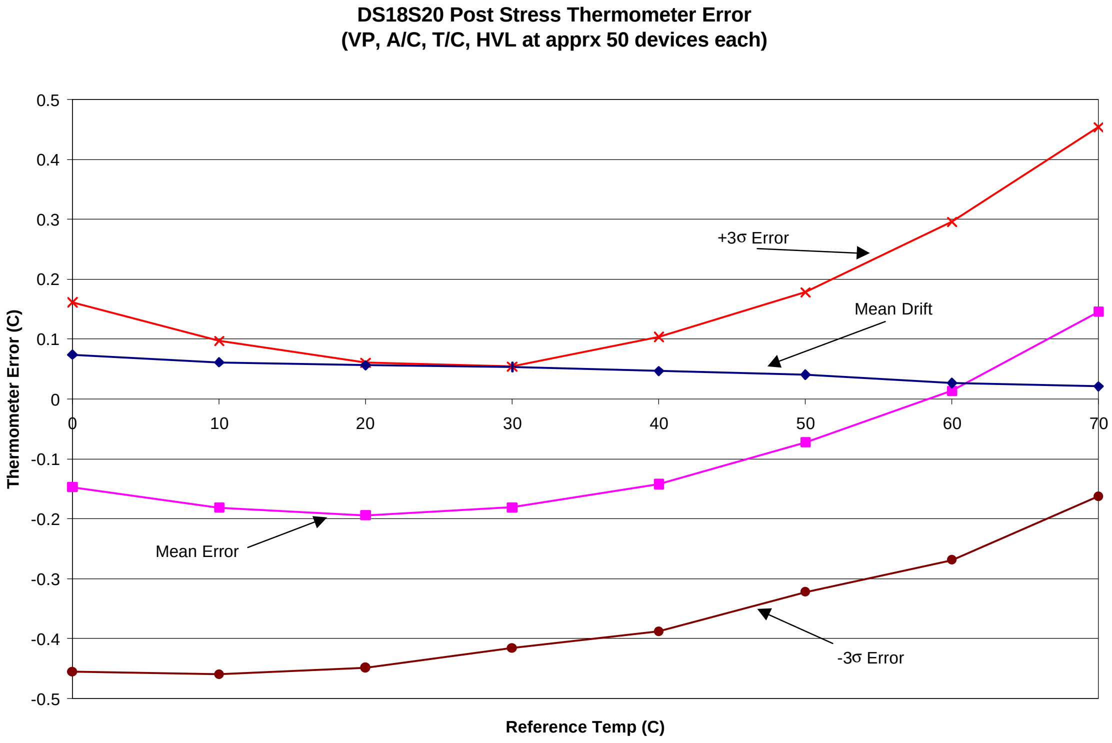

Fig. 3. DS18S20 Performance

Clearly, the DS18S20 has solved the inaccuracies many of you in the interest group have noted in the past. The other problem with the DS1820, the readout glitch that was described in the attached document, is totally an artifact of the dual-oscillator thermal sensing algorithm. The bandgap / Sigma-delta ADC sensing core of the DS18S20 by design will not be subject to such glitches.

The design goals of the DS18S20 were numerous:

1.) Eliminate stress-induced drift.
2.) Eliminate the state-machine glitch.
3.) Expand the operating voltage from 4.5V-5.5V of the DS1820 to 2.7V-5.5V.
4.) Significantly shrink the die so that it can fit in smaller, faster responding plastic and flipchip packages.
5.) Satisfy 1-4 while allowing the DS18S20 to drop into every DS1820 application, with no noticeable adverse differences.

All 5 of these goals are practically impossible to meet for such a drastic redesign of a complex mixed-signal IC. This document has described how we have met goals 1 through 4, but there are some noteworthy differences between the DS1820 and the DS18S20 that may not allow for a seamless transition into 100% of applications.

1.) Packaging. Smaller/lower cost packages was one of the design goals, and we met that. However, package differences may have an effect on some customers. The PR-35 of the DS1820 will no longer be offered with the DS18S20; it will be replaced by the TO-92 package. The pinout and lead spacing is identical between the two packages, so most customers will not have a problem. Customers who depend on the height of the PR-35 package in their application will be forced to make mechanical changes to their design. The surface mount option is more drastic. The huge DS1820 die was packaged in a 16-lead SSOP (with 13 no-connects); that was the smallest SMT package we could fit the DS1820 in. The SSOP is large and expensive. The SSOP will not be offered with the DS18S20, and the replacement SMT package is the 8-lead SOIC. Obviously, all SMT DS1820 customers will have a change. Dallas is offering the DS1820 in both the PR-35 and SSOP packages through 2000 on last time buy bases (but again, neither package will be offered with the redesigned DS18S20). Refer to the DS18S20 datasheet on the website for pinouts and mechanical specs of the packages.

2.) Thermometer Resolution. Recall the discussion about the temperature-to-digital conversion technique we use in the DS1820. Due to the nonlinear characteristic of the oscillators over temperature, we designed the slope accumulator with a high resolution (COUNT_PER_C generally over 100, yielding a sub-millidegree C resolution) so that the nonlinearity can be trimmed out. A conversion would always start with the counter preset to -55°C, and would count up to the temperature of the environment. Thus, conversion time depends significantly on the absolute temperature with the fastest at very low temperatures. This algorithm allowed for a conversion time of less than a half second, over the entire temperature range. The resolution attainable considering the conversion time and active current was outstanding with the DS1820, yet it was the dual oscillator technique that caused the drift. Solving the drift meant the dual-oscillator algorithm had to be designed out, while trying to minimize adversely affecting other specs. Sigma-delta modulators like that used in the DS18S20 are by design, slow data converters. There is a tradeoff between supply current, resolution, and conversion time. High resolution can be obtained in a relatively short amount of time if enough supply current is available. If supply current is crucial, high resolution requires long conversion times. You see the dependence.

    If the DS18S20 met the resolution (14-15bit) and active current (1.5mA) spec of the DS1820, the conversion time required would be between 4 and 8 seconds! That was obviously not an option. A tradeoff was made in the DS18S20 design. We kept the active current spec constant at 1.5mA, and allowed for a max conversion time increase from 500ms (DS1820) to 750ms (DS18S20). This will yield 12-bits of resolution (LSb weight of 0.0625°C). In the time/temperature display of the Weather Station, you will certainly notice a difference in the resolution. In a relatively stable temperature environment, the DS1820 curve would look smooth compared to that of the DS18S20. The conversion time increase will be transparent to the Weather Station because the WS-1 software polls the sensor for a “conversion complete” flag. If any of you are using the DS1820 in other designs, be aware of the resolution and conversion time changes with the DS18S20.

Although your display may appear more “jagged” with the lower resolution DS18S20, do not confuse that for inaccuracy. Data has been presented on the superior accuracy and stability of the DS18S20. The DS18S20 is accurate to ±0.5°C. That means the 9 most significant bits of the thermometer have zero error. Anything less significant than the ninth bit does not contribute to the accuracy of the sensor, but they are quite useful for monitoring minute temperature changes on a relative basis (relative to the last measurement, for example). For example, if bit 12 of the readout increases by 1 from one readout to the next, that does not necessarily mean the absolute temperature increased by the LSb weight of 0.0625°C. The part is not accurate to the 12th bit. It does imply however, that there was an increase of some absolute magnitude less than 0.5°C (the rated accuracy) from the last reading. Also, do not confuse the lower resolution graph of the DS18S20 with the state machine glitching of the DS1820.

Some of you in the interest group may choose to modify the Weather Station code to add algorithms that essentially smooth the temperature vs. time graph for the DS18S20. Smoothing algorithms will have a negligible effect on the accuracy or time constant of the thermal sensing system. The WS-1 has a large amount of thermal mass (and thus a very long thermal time constant) and the sensor is not exposed to moving medium (i.e. wind). Dallas will also release a patch to the WS-1 software that will perform curve smoothing.

3.) Error notification. The DS18S20 relies on strict protocol for predictable operation. Unfortunately, in the real world, less than ideal activity on the 1-Wire™ interface due to noise, lightening, etc may cause the DS18S20 to enter an ambiguous state. The absence of an error flag in memory left us to decide to notify the user of an error condition by reserving 1 state of the 12-bit (4,096 possible) thermometer readout for that purpose. The error state for the DS18S20 results in a temperature of +85.0000°C, which corresponds to a value of AAh in byte0 (TEMPERATURE LSB), 00h in byte1 (TEMPERATURE MSB), 0Ch in byte6 (COUNT_REMAIN), and 10h in byte7 (COUNT_PER_C). This reading indicates an error state and the reading should be discarded. The patch to the WS-1 software will ignore such readings and display that last completed conversion.

To summarize, Dallas is continually trying to make improvements to its thermal management product line. The DS1820 provided a solution to users who wanted a digital I/O thermometer with the minimum number of interface signals. Although it did fill a niche quite well, its algorithm contained some inherent inaccuracies and instabilities. Removal of those inaccuracies, in addition to widening the voltage range, and offering new/smaller/lower cost packages was of primary concern in the design of the DS1820’s replacement, the DS18S20.

The B5 revision of the DS18S20 will be sampled to customers soon, and we would like for you to be among the first to “try them in the field”. We will replace your DS1820 inventory with DS18S20 on a “2-for-1” basis. The DS18S20 design has been submitted for reliability testing, which we expect to be completed in May2000; at that time, the DS18S20 will be released to mass production. We do expect that not all applications will be able to transition to the new device without some system modification, but we expect these cases will be the exception and not the rule. The only noticeable adverse difference you should see with the DS18S20 in a Weather Station with unmodified code is lower thermometer resolution and a possible error notification reading of 85.0000°C. Dallas will very soon release a patch to the WS-1 software that will make both issues transparent to the Weather Station display. Thank you for your involvement in the Weather Station interest group and in the Dallas temperature sensor product line. If you have any questions or comments, please contact Dan Awtrey at dan.awtrey@dalsemi.com.

DALLAS
SEMICONDUCTOR

Dear Weather Station experimenter:

An anomaly has been detected in the temperature conversion algorithm of Dallas Semiconductor's thermal and battery management products that may affect some applications. The anomaly manifests itself as a random erroneous temperature measurement that can be as high as 1.0°C. The subsequent explanation of the problem and the suggested work-around assumes knowledge of the high resolution algorithm explained in Application Note 105, which is available on the Dallas website or faxback.

EXPLANATION OF THE PROBLEM

A digital state machine is used to control the temperature conversion algorithm. In normal operation, the state machine cycles between two states, watching for a control signal that signals the end of the conversion. Once this signal is detected, the next clock of the state machine sends it to the final state, suspending the conversion. As the state machine cycles between these two intermediate states, it increments the temperature register and reloads the slope counter, with the updated value in the slope accumulator.

Occasionally, through a mechanism still under investigation, the temperature conversion algorithm causes the state machine to cycle between the two intermediate states one more time after the end of conversion control signal is asserted. And, as is expected, this additional cycle increments the temperature register and reloads the counter from the slope accumulator for that new temperature. On the additional clock that sends the state machine to its final state, the slope counter is decremented. The result from the sensor is a temperature conversion result with a positive error in the range of:

$$+\frac{1}{COUNT\_PER\_C} \le Error(\deg C) \le +1.0$$

The magnitude of the error is dependent on the absolute temperature. Because the slope counter is decremented only once by the additional clock that sends the state machine to its final state, COUNT_PER_C - COUNT-REMAIN will be 1 when this error occurs. Therefore, the high resolution measurement will yield a result of:

$$Temp_{HIGHRES} = xxx.75 + \frac{1}{COUNT\_PER\_C}\ .$$

The error will be smallest for absolute temperatures slightly less than xxx.75°C (where COUNT_PER_C and COUNT_REMAIN approach their maximum difference), and approach 1°C for absolute temperatures slightly above xxx.75°C (when COUNT_PER_C and COUNT_REMAIN are nearly equal).

This behavior is illustrated in the following example. This example illustrates the result of a "good" measurement and that for a "bad" measurement in a simulated temperature sweep from 21°C to 24°C. The first figure, Fig. 1 gives the 9-bit direct reading for both a "good" and "bad" measurement. Note that the error, the difference between the two curves, ranges from 0.5°C to 1.0°C, depending upon the absolute temperature. This will be illustrated in Fig. 6.

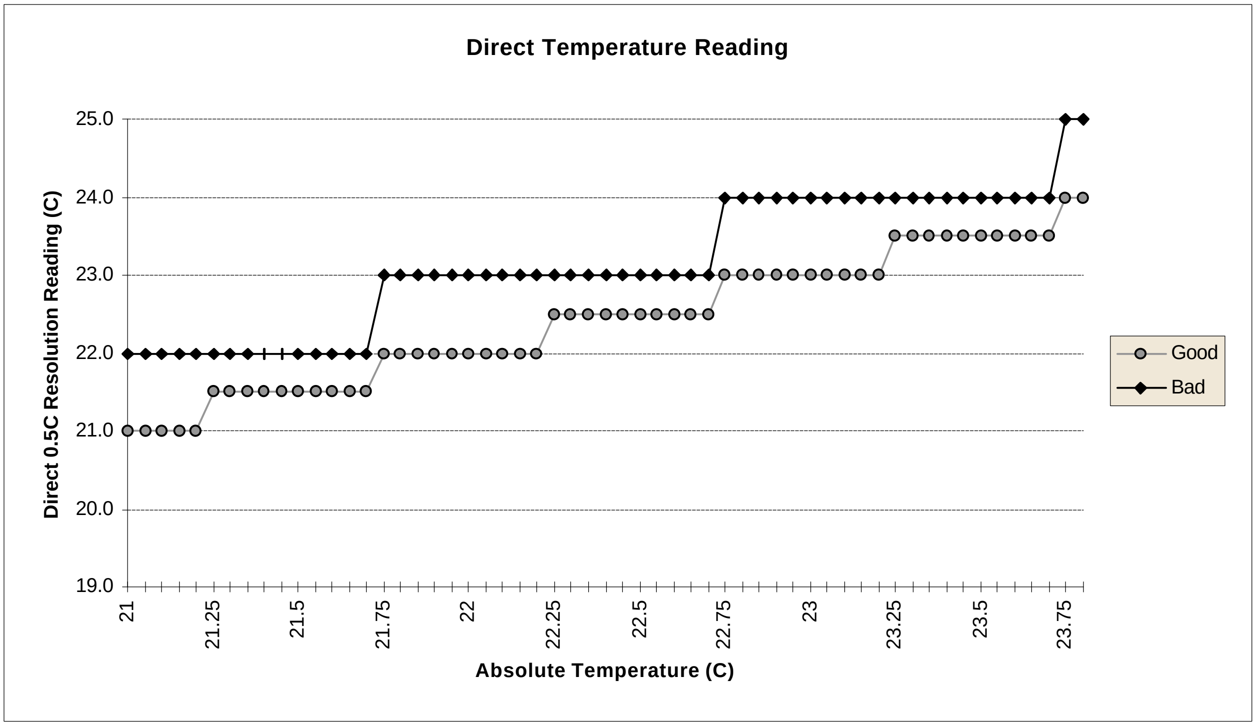

Fig. 1. Direct 9-bit Temperature Reading

The high resolution algorithm requires the truncation of the LSb from the direct 9-bit reading. Fig. 2 illustrates "good" and "bad" measurements after this truncation.

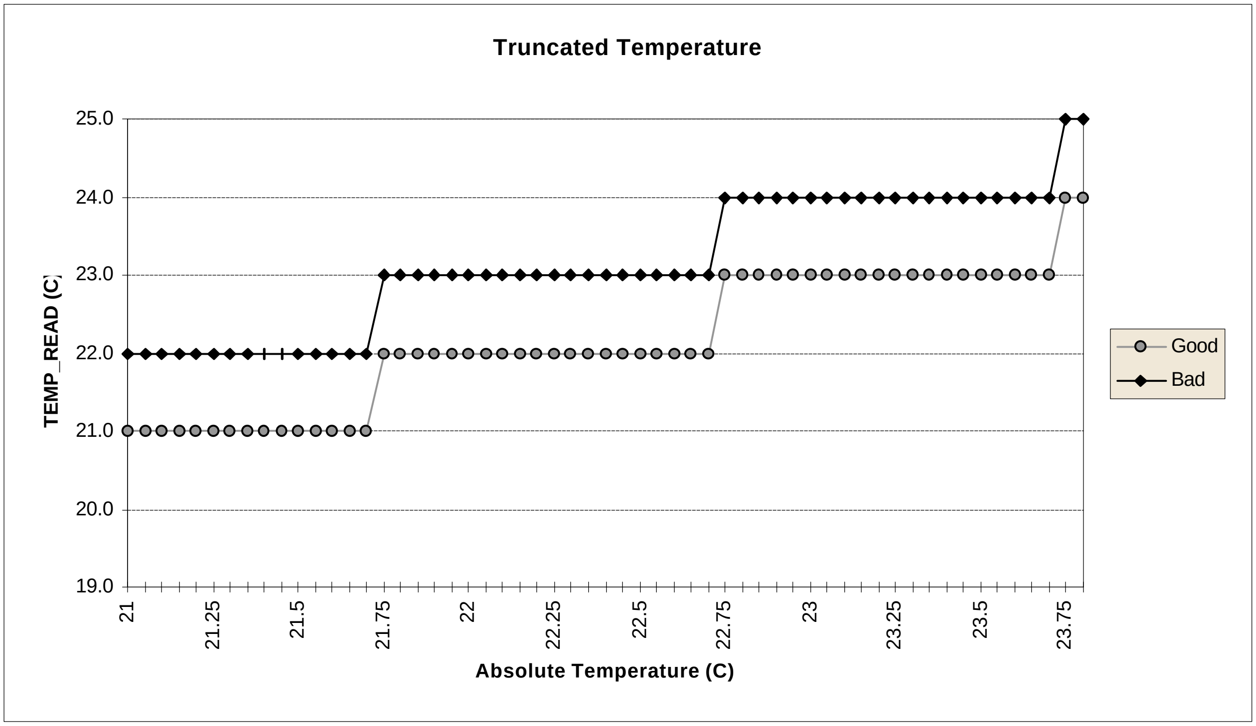

Fig. 2. Temperature Reading after Truncation (TEMP_READ)

The COUNT_PER_C register contains data that is programmed into the device at calibration. It sets the value of the slope accumulator at a given temperature. For a given temperature, this value is not the same for all devices. The absolute value depends upon the thermal characteristics of the oscillators. The values for this particular example were extracted from measured data on a DS1620, but may not be representative of the values in other DS1620 parts or other temperature sensors, in general. Nevertheless, it does illustrate the comment made earlier in that when the error occurs, the temperature register is incremented and the appropriate slope accumulator value is loaded. In this case, but not necessarily in all, the COUNT_PER_C value is 1 higher when the error occurs. This would not be true in a part which is calibrated such that two adjacent temperature values point to the same value for the slope accumulator.

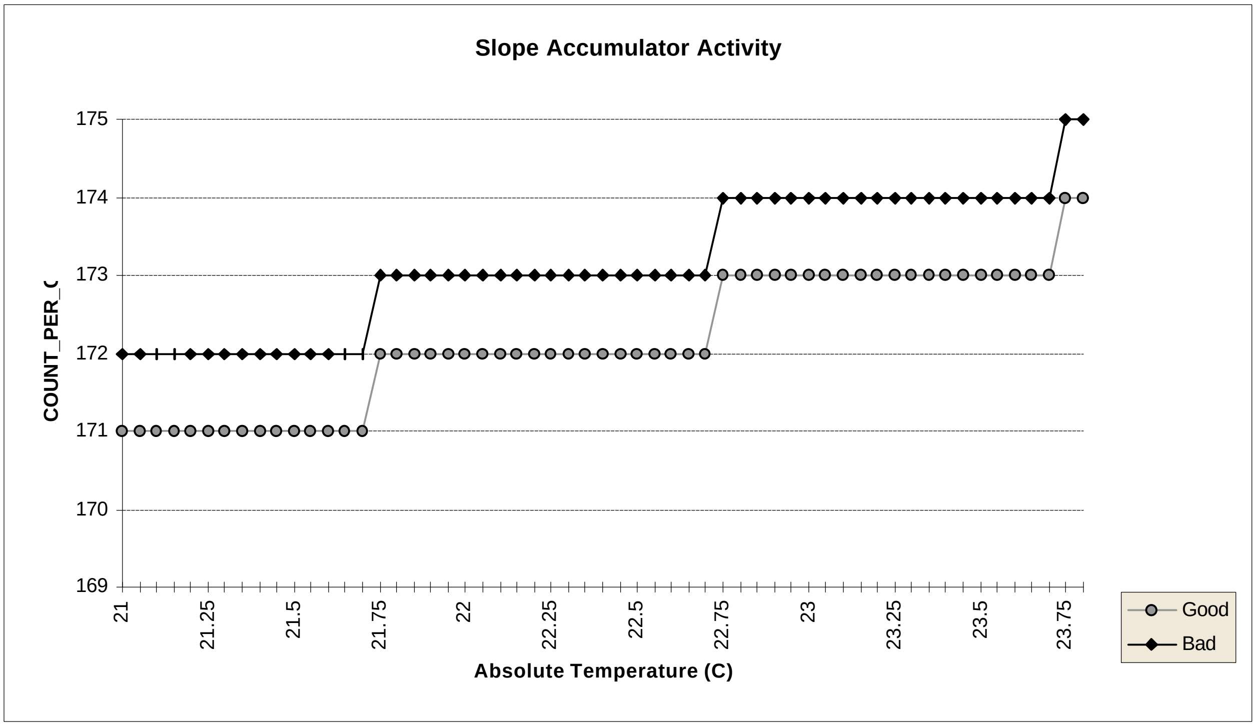

Fig. 3. Slope Accumulator Contents

Fig. 4 shows the value of the COUNT_REMAIN register in this temperature sweep example. It is the value left in the counter that was initialized to COUNT_PER_C for that given temperature. Recall that one additional clock cycle occurs in the event of an erroneous measurement; therefore, COUNT_REMAIN will be 1 less than COUNT_PER_C at a given temperature.

Using the calculation highlighted in Application Note 105 to obtain the high-resolution measurement, the measurement of Fig. 5 results. Note that the temperature always reads approximately xxx.75°C for an erroneous measurement. The final plot, Fig. 6, shows the error that one can expect when this anomaly occurs. As stated in the inequalities, the error varies from 1/COUNT_PER_C to 1°C in the high resolution mode. If only the 9-bit temperature is used, i.e., no high-resolution calculations are performed, the error will either be 0.5°C or 1.0°C, depending on the fractional value of the absolute temperature.

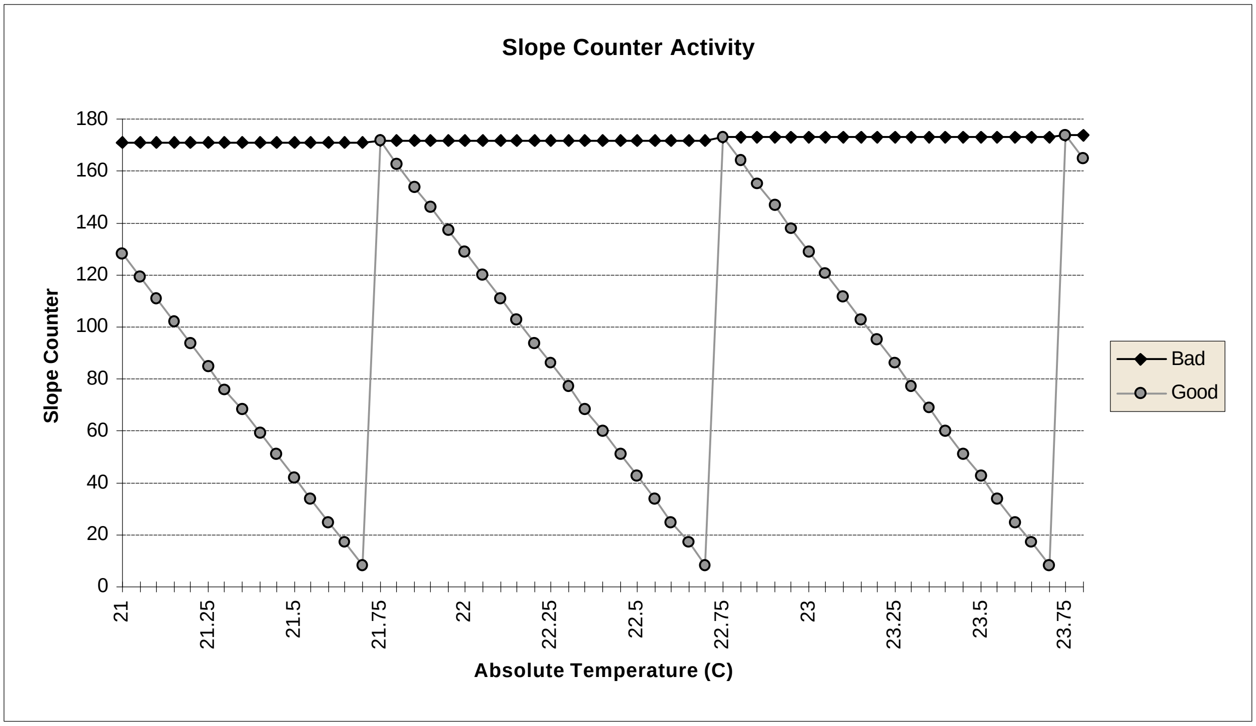

Fig. 4. Slope Counter Activity

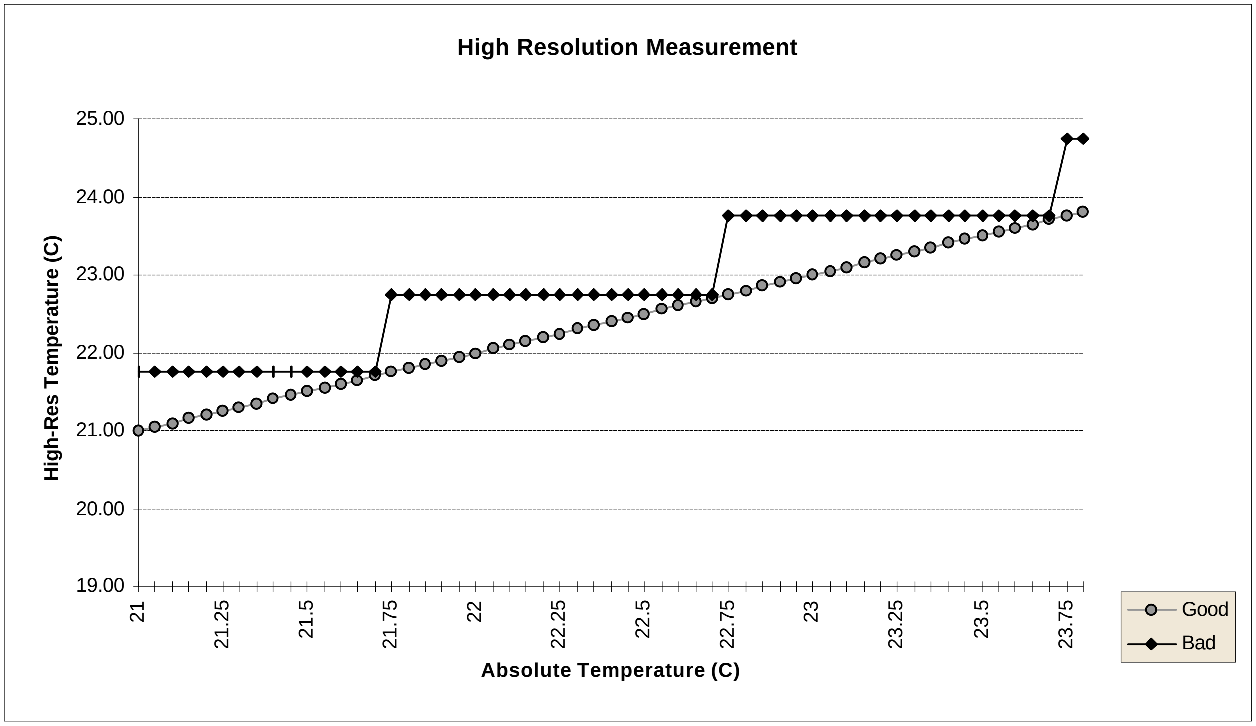

Fig. 5. Temperature Measurement after High-Resolution Calculation

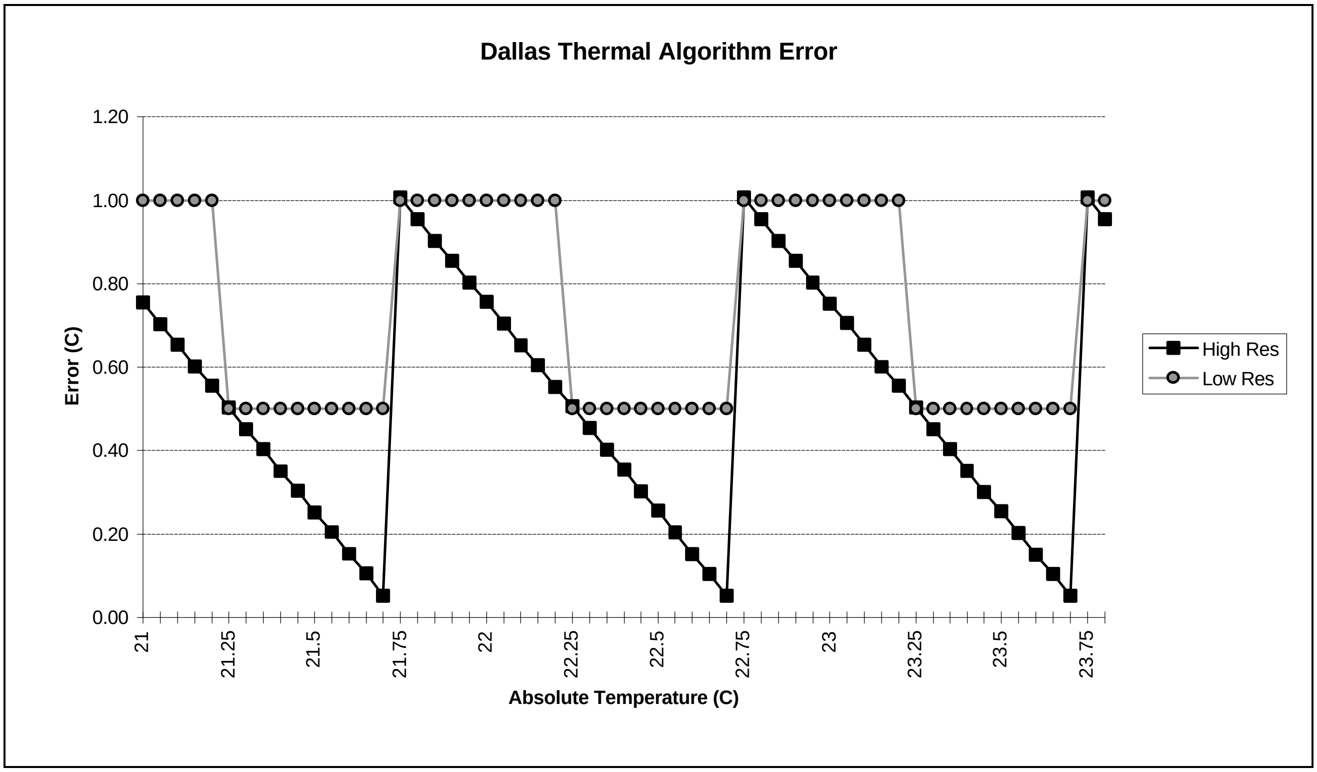

Fig. 6. Temperature Error Encountered when Anomaly Occurs

The occurrence of erroneous readings appear to be random in measurements we have performed thus far. Close examination of the example given, however, gives the user a means of understanding erroneous measurements based on the values of the registers read. Recall that only one additional clock cycle occurs to halt the state machine cycle; thus, COUNT_REMAIN will **ALWAYS** be one less than COUNT_PER_C for an erroneous measurement.

SUGGESTED ERROR DETECTION

The occurrence of the error is always associated with the condition COUNT_PER_C - COUNT_REMAIN = 1. Therefore, one could simply check for this condition before performing any further high-resolution calculations, and discard any measurements associated with this condition. However, COUNT_PER_C - COUNT_REMAIN = 1 does not necessarily constitute an error condition, and you could be discarding a valid measurement if you use the corrective action described above. This could be an acute problem in applications where the temperature is extremely stable, and the sensor returns a valid measurement such that COUNT_PER_C - COUNT_REMAIN = 1. If the method described above were implemented in this environment, several valid measurements in a row could be discarded. In measurements we have performed thus far, the error occurs very infrequently, about once per 5000 conversions. Never have we observed the erroneous condition occurring in two adjacent measurements.

Therefore, we suggest the algorithm depicted in the flowchart in Fig. 7 to detect an erroneous measurement and subsequently discard it. Basically, if the condition COUNT_PER_C - COUNT_REMAIN = 1 occurs on two measurements an a row, there is an extremely high likelihood that the measurement is valid. Similarly, if a conversion returns the above condition and the next returns something different, there is an extremely high likelihood that the measurement is invalid, and should be discarded.

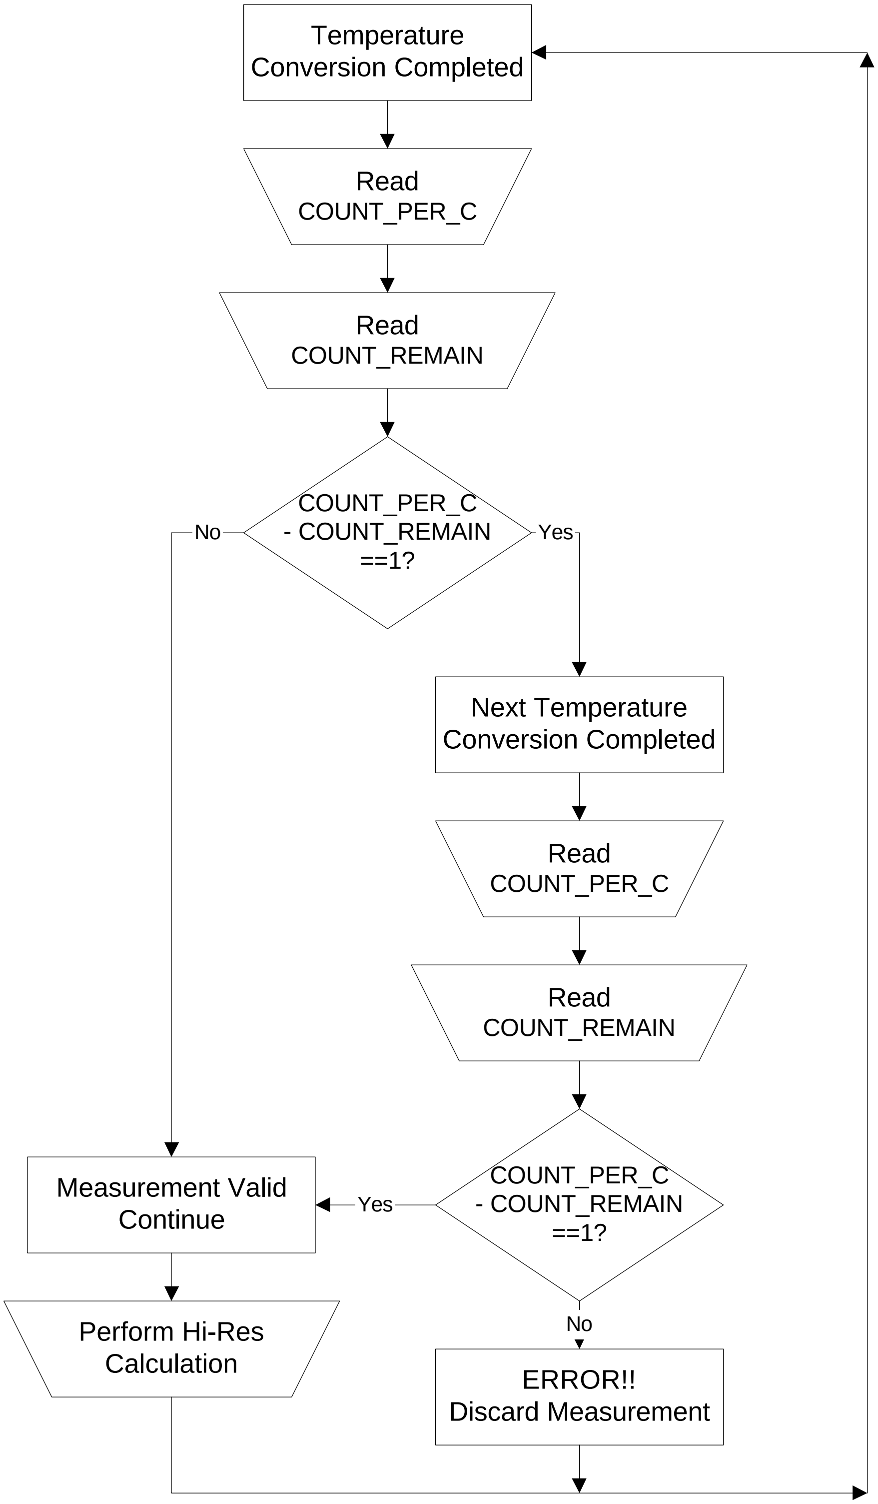

Fig. 7. Suggested Method of Error Detection