zilog®

Embedded in Life

An **IXYS** Company

White Paper

ZMOTION^{TM} – A New PIR Motion Detection Architecture

Dave Coulson,
WW Field Applications Manager, Zilog, Inc.

WP001703-1110

ZMOTION^{TM}

Revision History

Each instance in Revision History reflects a change to this document from its previous revision. For more details, refer to the corresponding pages and appropriate links in the table below.

Date	Revision Level	Description	Page Number
November 2010	03	Updated for new Zilog/IXYS logo	All
August 2010	02	Added ZMOTION product logo.	1
June 2010	01	Initial Release	All

Table of Contents:

Introduction

The Industry needs a better motion detector. The security industry is losing credibility with government and private enforcement agencies due to the significant number of false alarms caused by Passive Infra Red (PIR) motion detectors. End users are becoming frustrated with a trend of no response policies and heavy fines for false burglary alarms that are already in place within many jurisdictions. Some false alarms are user related, but the majority of false alarms originate from PIR motion detectors. Logic follows that if enforcement agencies penalize the end user for false alarms or cease responding altogether; there is less reason for the end user to pay for a monitoring service. This is threatening the revenue stream of the overall intrusion industry. These same issues and limitations are now being seen in the growing energy management industry where motion detection plays a large role in applications ranging from lighting control to vending and display systems.
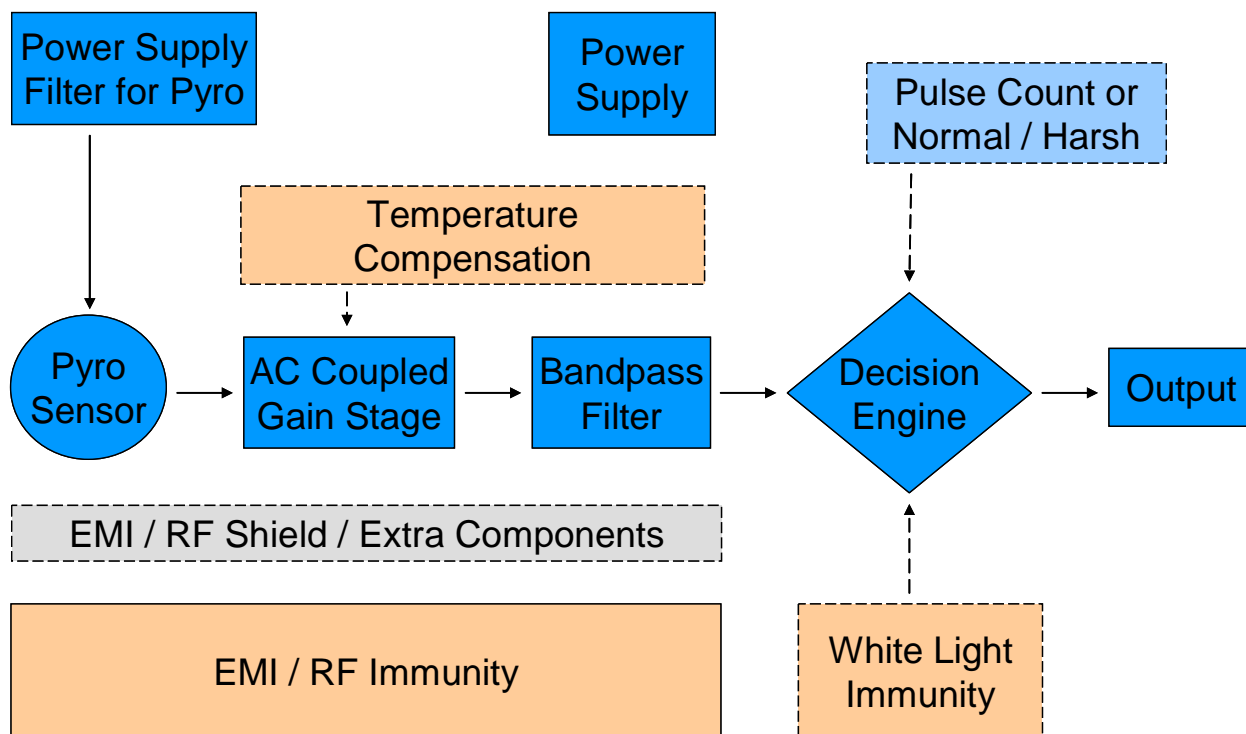
Traditional Approaches

Figure 1 - Traditional Motion Detector Architecture

The traditional PIR motion detector uses a pyroelectric sensing element that has a very low level analog signal output. The low but still usable AC signal is in the order of 1 to 2 mVp-p with a much larger ~10mVp-p high frequency noise component. This signal rides on a DC component of 400mV to 1800mV that will change with temperature, aging and even vary between devices. A Fresnel or similar type lens directs infrared energy onto this sensing element creating a usable frequency component ranging from 0.1Hz to 10Hz. The element's output is traditionally fed into a tight band pass filter stage to reduce high frequency noise and strip the DC element that the signal rides on. It is then fed into a high gain stage (~72db) so that the signal can be used by either discrete components or by a microcontroller to make decisions and act upon them (see figure 1).

A significant problem with this traditional design approach comes from the filter and gain stage. By filtering the signal, information that is sometimes critical to making a reliable decision is removed. Any signal discontinuity due to external electrical factors and/or forces between the element and the filter stage will look no different from a low level infrared energy signature at the output of the gain stage. This impacts the detectors maximum range, pet immunity, sensitivity and stability. There are many open and patented approaches to determining the signature of motion from this highly modified signal including rate of rise, amplitude and RMS (root mean square) energy under the curve measurements. Although combining some or all of these techniques can provide a more stable motion detector than using a single method, most low cost Microcontrollers simply do not have the processing power to implement multiple processing techniques. For example, the RMS calculation is complicated by a frequency component that will vary from 0.1 to 10Hz and will change with movement. There is often not even a single full cycle of any given frequency to use.

With the limitations created by the signal pre-conditioning, almost all detectors include a "pulse count" feature. This "feature" basically acknowledges the fact that the detector can and will false trigger under normal operating conditions. It also reduces the speed of detection which makes it an unacceptable solution for lighting control where fast response to real motion is a requirement. Higher end (and higher priced) detectors can include a micro wave sensor where one technology is used to confirm the other in the decision making process.

Other Approaches

One way to overcome these design limitations is to feed the signal directly into a Digital Signal Processor (DSP) with a high resolution Analog to Digital Converter (ADC). A DSP is capable of working very well with low signal levels, and high frequency components. Aside from significant cost increases with this approach, technical drawbacks such as power consumption still remain.

A DSP is designed to work on signals in the frequency domain. It is uniquely tailored to support Fourier math analysis of signals at high frequencies. The problem here is that PIR signals exist predominantly in the time domain. There is no consistent signal frequency to analyze.

Digital filtering of the significant high frequency noise component would be required to detect the discontinuities in the signal that constitute motion.

Even though the signals to be analyzed are low frequency in nature (0.1Hz to 10Hz) anti-aliasing issues also need to be addressed. An external filter can be used, but this adds cost and component count to an already highly cost sensitive product. Super-sampling the input signal (sampling the data at significantly higher frequencies than those of interest) can be used, but these samples need to be taken for time durations up to one second in order to detect the changes in the low frequency signals. These samples would be stored in the DSP RAM which increases the memory footprint requirements, increasing overall cost.

Goals of a New Detector Architecture

Any new PIR detector design must meet certain goals in order to be successful.

1. **Reduce or at least maintain product and installation cost**
   The industry is very cost competitive and any increase in cost would not be acceptable. In addition to the manufactured cost of the product, installation cost plays a large role as this is certainly part of the overall deployment cost.

2. **Improve performance and reliability**
   In spite of the commoditization of the PIR detector market, most responsible manufacturers truly attempt to develop products that meet or surpass industry standards like EN50131. These standards set the minimum performance requirements, but there is certainly room to greatly exceed them.

3. **Reduce false alarms**
   This might seem obvious and most of the industry standards use false alarms as failure criteria, but there are so many influences that can trigger false alarms that complete coverage by any test specification becomes nearly impossible. Manufacturers rely on field experience to determine which products perform better than others.

A New PIR Motion Detector Architecture

With the introduction of Zilog's Z8F082A – Z8 Encore! XP® Microcontroller family, the path was cleared to create a completely new motion detection methodology. The device is based on the high performance eZ8 Core which is a modified Harvard architecture and provides a linear address space, instruction pipelining, and very fast interrupt response. Other features like a Sigma/Delta ADC and 1KB of RAM create the ability to completely redesign the PIR motion detector.
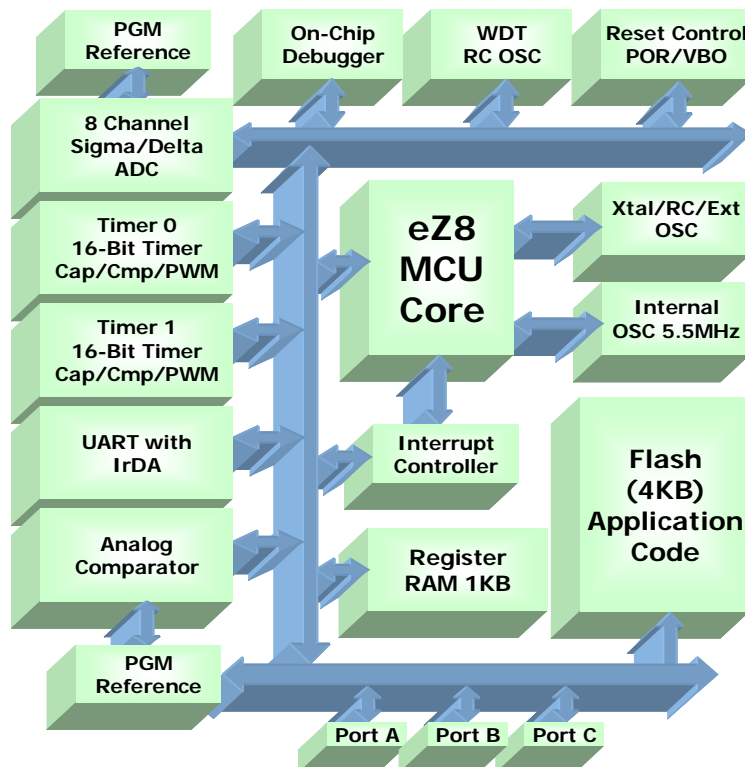
Figure 2 – Z8F082A Block Diagram

If Fourier math is the method of choice for frequency domain signal analysis, then statistical math would be the logical choice for time domain data analysis. In fact this exists in every day life today, but is used more as a means to work with real world data such as stock trends, financial forecasting, process control, gambling, sports, etc. If statistical techniques work well for data analysis in these applications, then it follows that similar techniques can be used for real time analysis of other time domain signals like those used in a PIR motion detector. The math involved is much simpler which would not require a high end processor to accomplish.
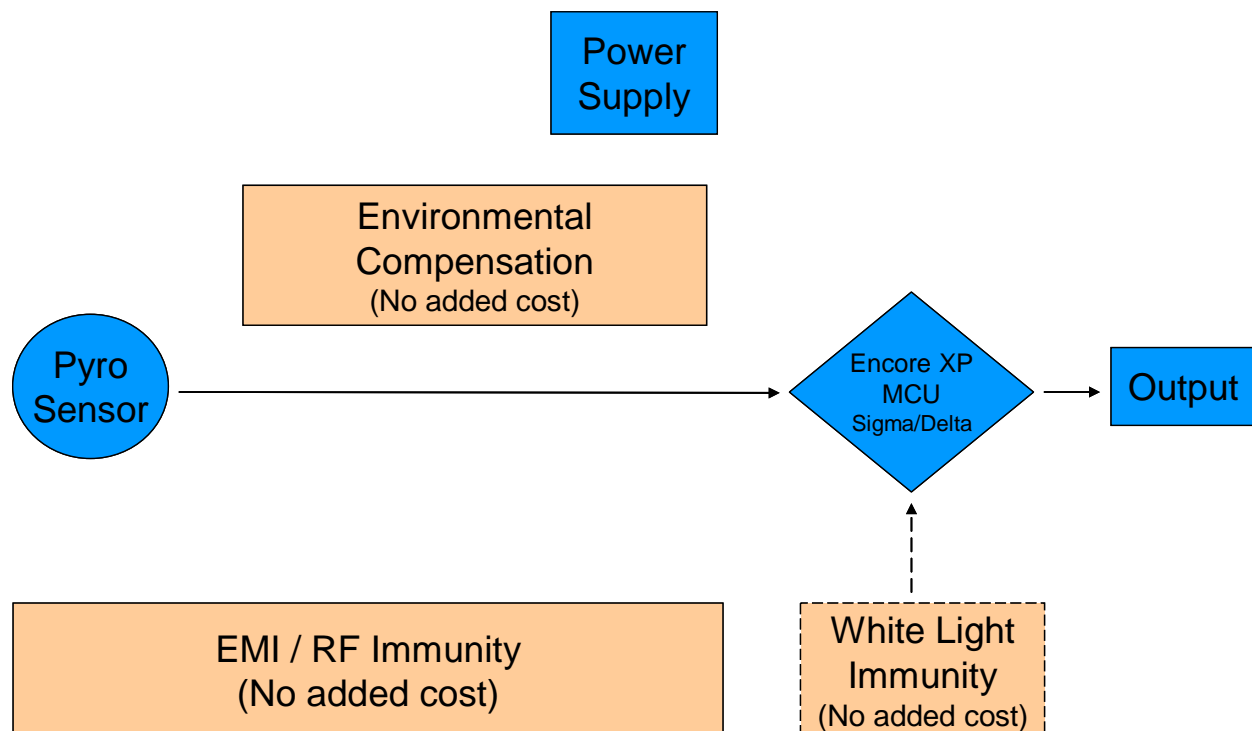

Figure 3 – ePIR™ Motion Detector Architecture

The Z8F082A architecture provides a greatly simplified design compared to the traditional approach. All external filters and gain circuits are removed (no electrolytic capacitors), the power supply design can be simplified and no temperature compensation is required. This results in a lower component count design.

The pyroelectric sensor is interfaced directly to the Z8F082A MCU without any AC coupling, gain or filtering. This allows the MCU to work with the true, unmodified signal and see the real time effects caused by shifts in DC offset, transients and other non-motion based signal changes. For example, a transient or white light event can generate a fast rise in output voltage from the pyroelectric sensor. The ePIR™ architecture can sense this event and respond to it accordingly. In traditional design, the signal would be modified by the low pass filter which would create a slow rising edge that looks similar to a 'real' motion event; potentially causing a false alarm.

Challenge #1: Gathering the Data

A method is needed to recover the low signal level data from the pyroelectric sensor. There are two issues: resolution and DC offset. Simply sampling the signal with a standard 8 to10 bit SAR ADC will not provide the required resolution or noise immunity. The large DC offset present with the signal needs to be considered (but not eliminated) and brought within the limits of the ADC's operating range. The DC offset can actually provide valuable information pertaining to health of the pyroelectric sensor.

Using Zilog's Z8F082A MCU device the signal from the pyroelectric sensor can be connected directly to the positive input of the differential Sigma/Delta ADC. An internally generated 1V reference is delivered to the negative input of the ADC. This creates a +/-1V range for the pyroelectric sensor input. Although specified for 10 bit accuracy, the Sigma/Delta ADC has a 16 bit result register of which one bit is used for overflow indication and another bit for sign. By over-sampling and averaging, it provides up to15 bits of resolution. In this application, accuracy is not as important as resolution. Based on the conversion rate used, constructed sample values have a +/-16384 count range which provides a usable resolution of 61uV/count. Since the ADC is a sigma delta type, it provides a high degree of linearity and resolution for a trade off in sample/conversion speed. By over-sampling and averaging, a very high resolution raw signal input that does not require any hardware pre-conditioning is provided.

The ADC's effective resolution is +/-16384 steps over a +/-1 volt range. As data is sampled and buffered, the maximum and minimum values are tracked. The reason for this is to limit the number of floating point operations required. By keeping the minimum and maximum readings, data samples can be normalized back into 8 bit integer data without loosing resolution information. This also reduces the amount of data buffering required – it's a form of compression.

Challenge #2: What to do with the Data

Figures 4 and 5 show the sampled data from a pyroelectric sensor (signal + noise + DC component). Analysis of this data shows that it is normally distributed which allows a shortcut to be used when calculating the standard deviation since 68.26% of the data will fall within 1 standard deviation of the mean, 95.46% of the data will be within 2 standard deviations, and 99.73% will fall within 3 standard deviations.
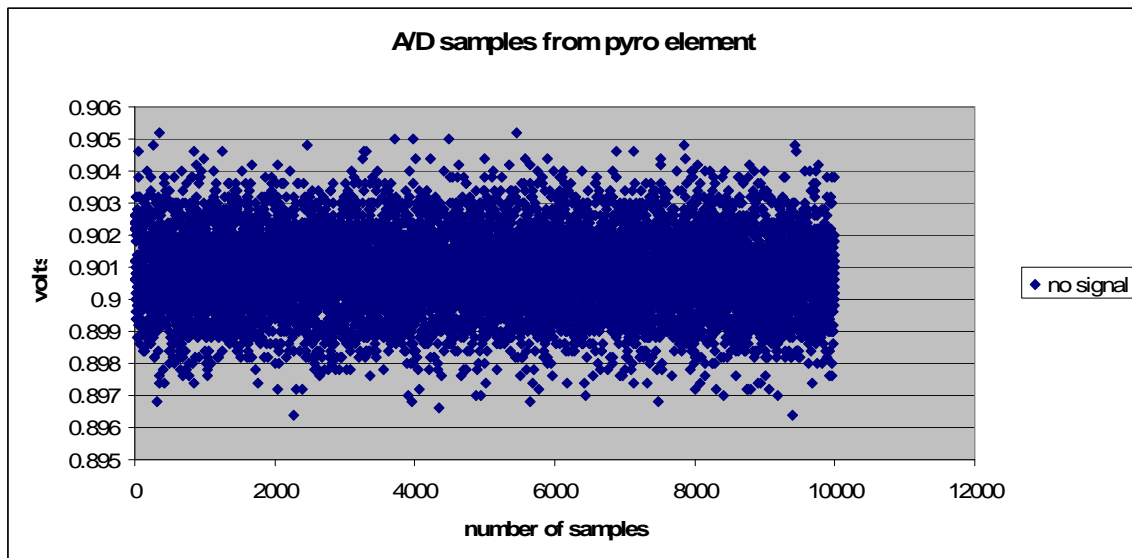

Figure 4 – ADC Samples from Pyroelectric Sensor without Motion
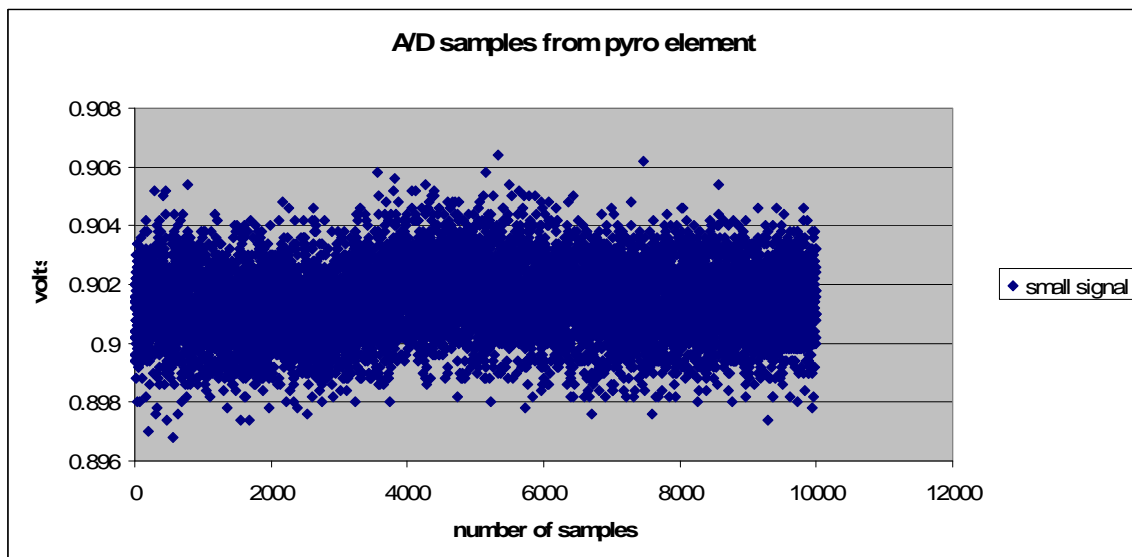

Figure 5 – ADC Samples from Pyroelectric Sensor with Motion

From control chart theory, control limits (figure 6) can be generated for the data using simple math. This is done by taking a sample size and then generating sub groups within the sample.

- Calculate the averages of each sub sample.
- Calculate the range for each sub sample.
- Calculate the sample average as the average of all the sub sample averages.
- Calculate the range average as the average of all the sub sample ranges.

The control limits are then calculated by taking the sample average and adding/subtracting the sample range averages multiplied by the A2 factor. The A2 factor is a constant that is based on the sub sample size. This A2 factor eliminates the need to do actual standard deviation calculations. It can only be used with the distribution of data that is normal (i.e. Gaussian / Bell Curve).
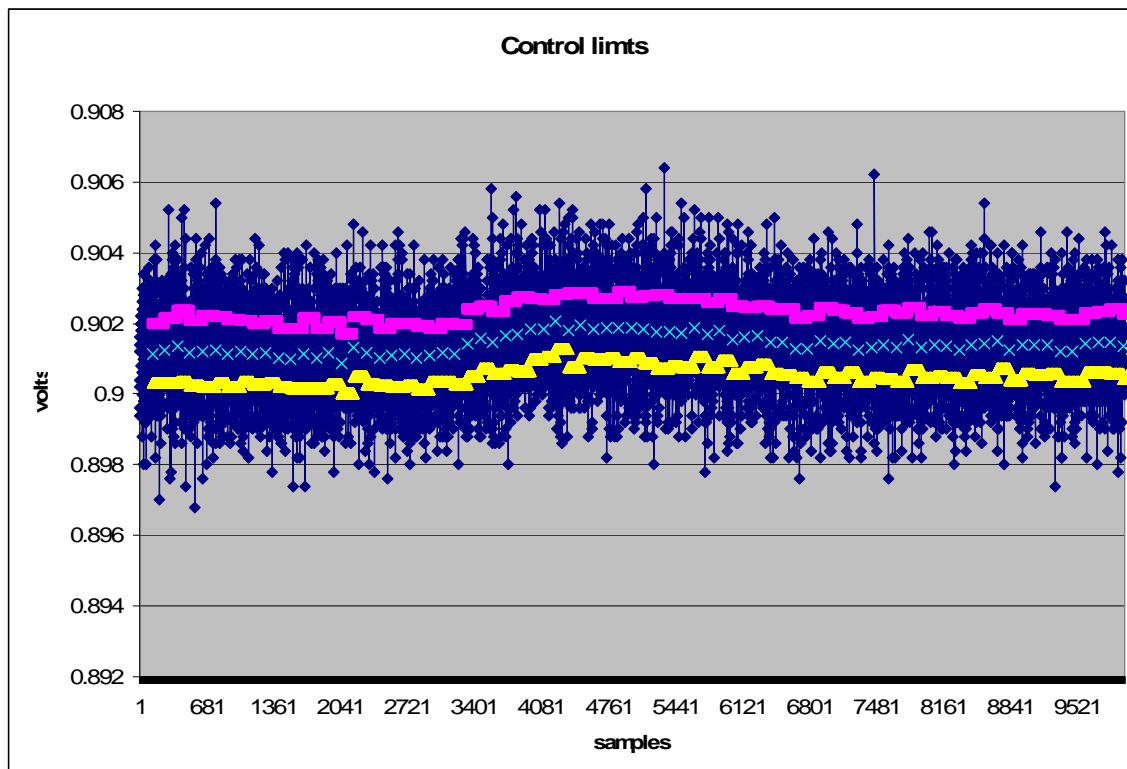

Figure 6 – Control Limit Calculations

If there are data points (sub group averages) that tend to gravitate to the area between 2 and 3 standard deviations from the center line, then an abnormality exists. From probability it can likewise be said an abnormality exist when 2 of 3 or 3 of 7 or 5 of 10 consecutive points lie in this region. It's a relatively simple task to analyze the data after the fact, but how does the system make use of it in making real time determinations?

Challenge #3: How to use the control limits in real time

The approach taken is to have the system create a minimum of 4 different sets of sample control limits at differing time intervals. The control limits can be compared between time indexes to find small movement trends in the data. To find a trend, 3 or more points are required. With 4 control limits it can be determined if 2 out of 3 of the control limits have trended up or down; which would signal a change in the system. This requires the system to buffer a lot of data which is the reason the raw samples are normalized from 16 bit floating point to 8 bit data values.

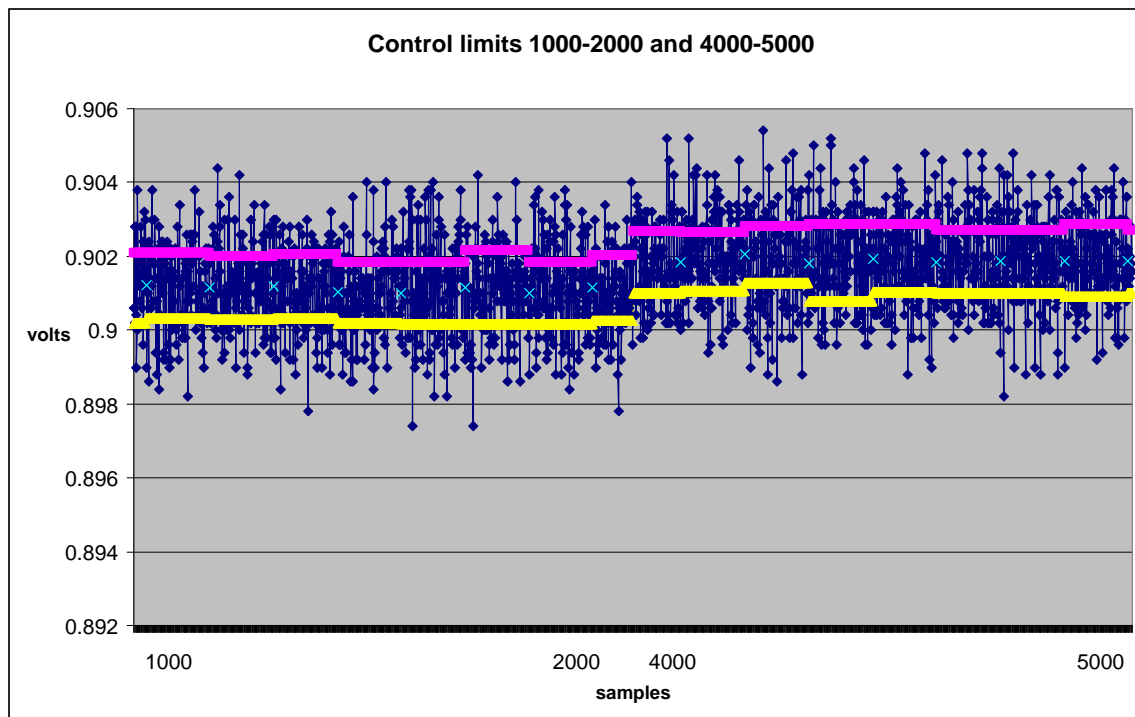The Z8F082A device has 1K Bytes of internal ram which provides the large buffering requirements.

Figure 7 – Control Limit Usage

The data in figure 7 represents a subset of the data shown in figure 6. In the chart, points from 1000 to 2000 in the first half are plotted (window 1), and in the second half of the graph are the points from 4000 to 5000 (window 2). You will also notice that more than 2 of 3 and 3 of 7 of window 2's sub window averages (blue x's) are all around the same level as the upper control limits of window 1 (pink lines). For larger signal changes, the lower control limit of window 2 may even be higher than the window 1 upper control limit or vice-versa depending on the direction of the energy swing. In either case, this constitutes a change in the stable steady state data system. This is what we are looking for. A change in the steady state system means there is a change in the infrared energy being seen. The smaller the range of the 2 sample windows is, the less sensitive the algorithm will be to small changes in energy detection. The larger the range, the better it can detect small changes. Likewise the larger the sample window size is then the less sensitive to change the algorithm will be.

This processing method benefits primarily from eliminating the need to perform significant processing/decision making on every data point while maintaining good frequency discrimination.

This method of data analysis also allows the system to deal with significant noise without filtering. Noise is Gaussian in nature so the amount of noise will not affect the control limits for the system. The data remains normally distributed. This also means that the "pulse count" feature **of traditional PIR approaches** is no longer required.

Another benefit of this method is the ability to detect a discontinuity in the data. This ability is lost when the signal is filtered as in traditional designs. A discontinuity could be attributed to an external influence such as RF energy, white light, unstable power, electrical disturbances, etc.
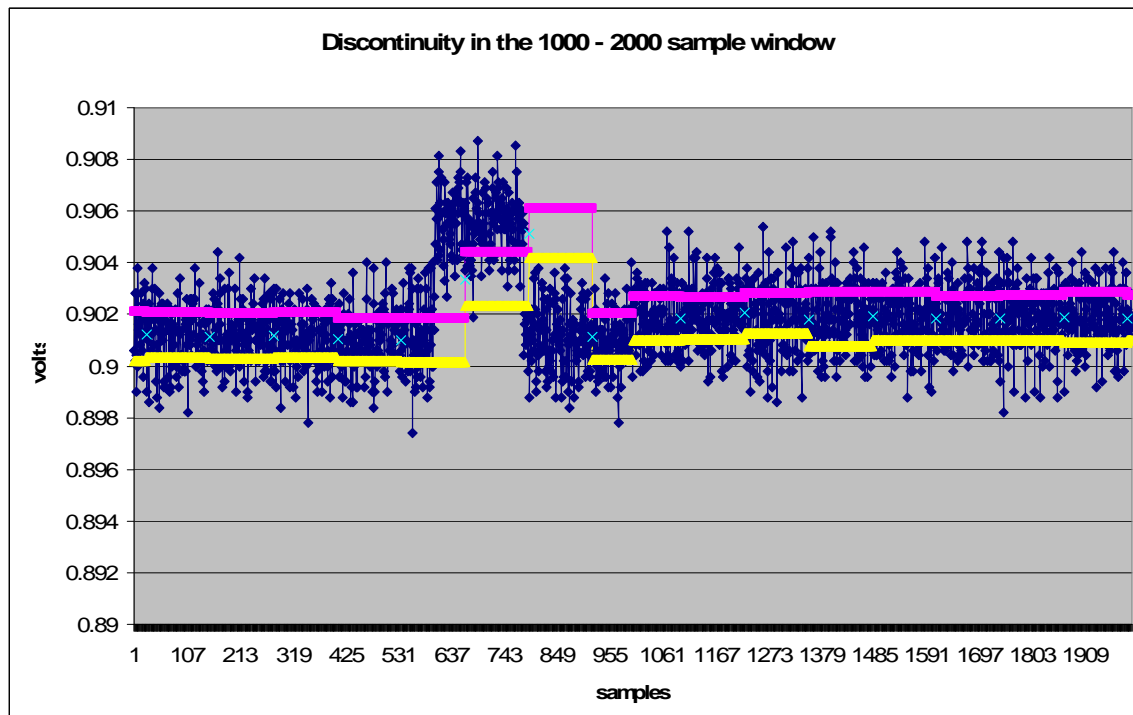

Figure 8 – Control Limits; Determining Non-Motion Events

If at any time in either sample window there is a shift in the data trend outside the control limits of its own sample window, it should not be considered valid data from the pyroelectric sensor. The pyroelectric sensor can not behave this way (based on the sample rate, the PIR element can not change the voltage this fast). These invalid shifts can be attributed to events such as EMI or white light. The traditional method of filtering and amplifying would see this external influence no differently than it would valid infrared energy moving quickly from one beam to another. With this additional information, the system can refrain from making a decision until the input data returns to a stable state. If the data does not return to stability within a certain time, then the unit can initiate a trouble condition or an alarm if no trouble state is supported.

Challenge #4: Temperature compensation

As the ambient temperature approaches 35 degrees C it becomes increasingly difficult to detect the difference in radiated infrared energy between a person and the surrounding background. Traditional detectors will increase the system gain as the temperature approaches 35 degrees C. Above this temperature, the gain should be reduced as the difference between body temperature and room temperature starts once again to increase. Many detectors continue increasing the gain in one direction only making these detectors extremely unstable and unreliable in hot environments. The ePIR™ algorithms' ability to work with an un-amplified low level signal negates the need to increase the signal

gain altogether. The algorithm always performs at maximum sensitivity that is reliable for any given environment that the unit is placed within. This eliminates the need for temperature compensation components.

Putting it All Together

With the introduction of the Z8F082A - Z8Encore! XP$^®$ MCU providing more computational horsepower, memory resources and unique on board peripherals, it is now possible to use new methods of signal processing to replace the traditional root mean square approach of measuring the energy under the curve.

By integrating the Z8F082A with this new motion detection methodology, Zilog has created a unique IC product series called **Z8FS040** and **Z8FS021** as part of its **ZMOTION™ Detection and Control Family**. The motion detection algorithms comprise the ePIR™ Engine and run in the background of the MCU while control and status of the Engine is accessed through a software API (Application Programmer Interface). This allows custom application software to be created while taking advantage of the ePIR™ Motion Detection Technology. A series of lenses and pyroelectric sensors are provided with optimized API settings for various applications.

With the ePIR™ statistical approach, there is greater control in making valid decisions on changes in analog signals that are predominately in the time domain. With the combination of the reduced hardware signal conditioning and the more powerful and flexible processing algorithm, it is now possible to overcome the short comings of the traditional PIR design while reducing the cost at the same time.

- All external gain and filter circuits are removed
- Power supply design is simplified
- No temperature compensation
- No dedicated white light circuitry
- No electrolytic capacitors (only ceramic chip decoupling capacitors)

The ZMOTION™ S/W API allows ePIR™ technology to be easily adapted to other applications too. Lighting Control, HVAC, Access Control, Vending, Display, Proximity, Power/Energy Management, Occupancy Sensing and any other application requiring both sensitivity and stability can take advantage of this new approach.

With this new class of motion detector, the false alarm issues plaguing the security industry can now be significantly improved.

For more information, please visit Zilog's website at: http://www.zilog.com and refer to the Z8FS040 ZMOTION™ product family.

Warning: Do Not Use in Life Support

LIFE SUPPORT POLICY

ZILOG'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS PRIOR WRITTEN APPROVAL OF THE PRESIDENT AND GENERAL COUNSEL OF ZILOG CORPORATION.

As used herein

Life support devices or systems are devices which (a) are intended for surgical implant into the body, or (b) support or sustain life and whose failure to perform when properly used in accordance with instructions for use provided in the labeling can be reasonably expected to result in a significant injury to the user. A critical component is any component in a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system or to affect its safety or effectiveness.

Document Disclaimer