

Title: Accelerometer Technologies, Specifications, and Limitations

Author: Jeffery Miller

Institution: ActiGraph

Introduction: Technology advances and manufacturing efficiency improvements have drastically increased the options and flexibility available to users of accelerometer-based products. These advancements have led to confusion among activity monitor users, and many have been led to believe that device output normalization is an obvious and easy step. In truth, there are many engineering obstacles that make normalization extremely complicated.

Purpose: The purpose of this presentation is to educate users on the differences between integrated circuit accelerometer technologies and the tradeoffs and burdens that may be incurred by selecting one type over another, regardless of manufacturer. This paper will work to dispel the myth that accelerometers and/or activity monitors can easily be used interchangeably while furthering the understanding of accelerometers and their capabilities within the arena of human activity monitoring.

Summary: Until recent years, researchers interested in monitoring and quantifying human activity only had one viable option. The technology of Microelectromechanical systems (MEMS) was not yet cost effective for everyday use, and as a result, accelerometer-based activity monitors relied primarily on piezoelectric bimorph beams. These beams, while functional, are typically expensive to manufacture, require periodic calibration, and are limited to measurement of time varying acceleration, which precludes positional information such as subject posture. Advances in silicon wafer and manufacturing processing have enabled MEMS based accelerometers to become prevalent in many applications, including nearly every smartphone manufactured today. This technology is extremely stable, exhibits negligible measurement drift due to temperature, and requires only a single calibration. Furthermore, because they are capable of measuring static acceleration, positional information can be harvested for various applications.

These advances in the field of MEMS technology, coupled with the decreasing cost per bit for non-volatile memory, have led to a fundamental shift in the way activity data are collected. Researchers are no longer limited by the long standing filtered/epoch level data collection that dominated the arena for so many years. They are now free to collect raw acceleration data. This approach maximizes flexibility, allowing researchers to post process and reprocess data as new algorithms become available. With this additional flexibility comes a growing interest by the research community in normalizing outputs across devices, thereby removing the unique value provided by individual manufacturers. While this idea holds great promise among those tasked with harvesting useful information from the collected data, there are numerous hurdles that prevent it from being easily achieved. Proper education can establish appropriate data collection expectations, reduce confusion about the data collected, and allow for realistic comparisons between different accelerometer based products.

Introduction

The measurement of human activity continues to be a growing area of interest within research, clinical, and personal health arenas. Information gathered and derived from devices that capture activity related measurements can provide valuable insight into many health related outcomes. Within the research and clinical markets, this information is often captured subjectively by way of patient reported outcomes. Not only is this subjective method inherently unreliable, but the information collected lacks the intrinsic long term value associated with objective data sets that may be shared, processed, and studied.

As with any tool, the intended use and desired output/results play a very large role in the selection process. Over the course of the last several years, a wide array of miniature, low cost activity monitoring tools targeted at varying markets have become available. Primarily driven by improvements in manufacturing efficiencies and overall technology advances, these devices have grown in complexity and feature set to better accommodate a rapidly evolving field of study. As complexity increases, use cases begin to diverge, measurement populations become more focused, and the need for improved user education on the varying parameters and nuances that differentiate one device from another becomes more imperative.

Such a rapid evolution in product capability and offerings is often coupled with confusion by those charged with harvesting useful information from the devices on a day to day basis. This white paper will work to alleviate this confusion and educate users on the

differences between integrated circuit accelerometer technologies and the tradeoffs and burdens that may be incurred by selecting one type over another. Proper education can establish appropriate data collection expectations, reduce data confusion, and allow for realistic comparisons between different accelerometer based products. Furthermore, this paper will discuss how the increasing desire among the research community to normalize outputs is difficult at best and puts device manufacturers in the undesirable position of having to share intricate design nuances and potentially proprietary trade secrets.

Accelerometer Technologies

There are many technology types that yield transducers capable of converting acceleration into a quantifiable and measurable signal, with each varying in maturity, capability, cost, and manufacturability. In the broadest sense and most relevant to those interested in human activity, accelerometers can be classified into one of two primary categories: AC coupled, which are only capable of measuring time varying accelerations, and DC coupled, which measure both static and dynamic accelerations.

In order to maintain brevity, only the technologies that are currently used or have value within the human activity markets will be discussed. In the most basic of terms we will discuss piezoelectric and MEMS based capacitive accelerometers.

Piezoelectric

Piezoelectric based accelerometers have been commercially available since the first half of the twentieth century, and they have been used for a broad range of applications with great success. They are based on the phenomenon of

crystalline structures, which yield an electrical signal proportional to the amount of acceleration they experience.

These sensors are typically manufactured with a mass carefully situated at the end of a cantilevered beam. This weighted beam deflects in relation to the acceleration experienced, yielding an electrical signal proportional to the acceleration experienced. This is a very effective solution for measuring human activity, but there are a number of factors that limit its functionality within the current activity economy.

Until recent years, most devices intended for end user consumption utilized a piezoelectric bi-morph beam based accelerometer. These devices have a significant advantage over their counterparts in that they require zero power to operate, providing for an end user device with a very long battery life. Unfortunately, bi-morph beam based devices are difficult to manufacture, which typically increases the cost to the end user. Additionally, because they are susceptible to drift due to environmental conditions and mechanical shocks, regular field calibration is generally required. Another drawback to the bi-morph beam type accelerometer is the fact it provides an AC coupled output only. In practice this means only transient, or time varying, accelerations can be measured, thus precluding the ability to harvest positional or posture information.

MEMS Capacitive

Recent years have witnessed the proliferation of Microelectromechanical Systems (MEMS) based accelerometers. Most, if not all, MEMS devices utilized in human activity monitors are

based on a capacitive response. With these devices, an internal capacitance changes proportionately to the acceleration experienced, yielding a very stable output signal. Unlike bi-morph beams, which are considered passive components, these accelerometers are active semiconductors. As a result, they require power in order to operate, yielding a net negative effect on activity monitor battery life. The MEMS accelerometer manufacturing process produces a very stable transducer that does not suffer from the same environmental drift as its piezoelectric counterpart. This translates into a single calibration requirement that is typically achieved during the activity monitor manufacturing process, resulting in a reduced operational burden to the end user.

MEMS accelerometers have very broad usage and adoption across many markets, including air bag technology, home appliance, mobile phones, and game controllers, as well as the developing human activity market. Because they are packaged in standard semiconductor packaging, they are easily machine placed during the manufacturing process, lowering overall product cost. MEMS accelerometers are also typically DC coupled outputs, allowing for the measurement and capture of both static and transient accelerations. From this, it is possible to pull positional or posture related information. Furthermore, these devices can be purchased in a myriad of configurations and options that include analog or digital output and programmable dynamic ranges.

Data Collection

Unfortunately, as the number of accelerometry device options and choices increase, so does confusion. Historically, researchers only needed to concern themselves with a handful of parameters easily achieved by most devices. However the increased availability of MEMS devices has fueled a shift in paradigm related to measuring and collecting data in recent years. Rather than the previously limiting method of storing the activity data in filtered format in predetermined epochs, researchers are taking advantage of significant increases in memory densities and battery life improvements to store data in a raw format. This method offers far more flexibility because the data can be post processed repeatedly as new algorithms become available. Researchers are just now beginning to reap the benefits of this shift in measurement style, through the emergence of new areas of interest such as activity pattern recognition.

Regardless of the technology type employed or the source of the acceleration to be measured, all accelerometer devices **must** perform the two basic functions of pre-filtering, or anti-aliasing filter, and sampling to achieve activity measurements.




Figure 1 Filtering and Sampling Process

The manners in which these functions are performed are determined by the designer and

are typically driven by a multitude of variables including the intended measurement requirements, cost, battery life requirements, device size, environmental demands, and interoperability with previously manufactured devices. There are many tradeoffs made by the designer during the process in order to achieve the end result while simultaneously meeting as many of the design requirements as possible.

The remainder of this paper will focus on providing education on the key specifications of accelerometers, filtering methods, and how they impact various measurements. After presentation of these parameters, it will be obvious how device output normalization is not a trivial endeavor.

Sampling

One of the many new accelerometer features implemented in recent years is the introduction of varying sample rates. Rather than being limited to the historical sample rates of 10 or 30 Hz, mainstream devices now support rates up to and beyond 100 Hz. The introduction of this feature alone has created the misconception among end users that 'more information' is naturally contained within a signal that is sampled at a higher rate. Depending on the intended use and how it is implemented, this may not be the case and could result in unnecessarily large files and related data storage issues.

More activity monitors are beginning to utilize MEMS accelerometers with a digital output because of the flexibility and rapid design cycles they provide. One thing that must be accounted for is the fact that digital accelerometers generally have finite sample

rates to select from, and therefore numerical techniques must be performed in order to achieve the desired end user sample rate. For instance, if a desired end user sample rate of 30 Hz is required and the accelerometer only samples at octaves of 200 Hz (very common), sample rate shifting must be implemented to achieve the desired 30 Hz result. While further discussion of the varying shifting techniques is outside the scope of this white paper, it is important to realize that each method impacts the overall measurement output in its own way.

An additional consideration in regards to sample rate is the accuracy and stability of the timing source for the sampling. Many digital accelerometers have internal timing sources, yielding little to no flexibility to the designer. If not taken into consideration, these crude oscillators can vary from device to device by as much as +/-15%. This means that a desired 30 Hz sampling rate can yield values ranging between 34.5 Hz and 25.5 Hz, and two identical activity monitors can be configured identically and yet have sampling rates that vary by as much as 9 Hz. Having a repeatable and predictable sample rate is paramount for accurate post processing and further digital filtering.

Anti-Aliasing Filtering

An additional consideration is that of the anti-aliasing filter, or pre-filter. This filter sets the initial bandwidth of the signal to be sampled. It is important to note that **all** activity data collected is pre-filtered with an anti-aliasing filter. All data, even raw, is filtered at least once. This is driven by the requirement of the Nyquist-Shannon sampling theorem which

states that any frequency band-limited signal can be perfectly reconstructed from a sequence of samples if the bandwidth, B , is no greater than $\frac{1}{2}$ the sampling rate, f_s . Simply stated, the sampling rate must be twice the highest frequency component of the signal to be sampled.

$$f_s \geq 2B$$

If this criterion is not met, a phenomenon called 'aliasing' occurs, which corrupts the sampled data so that it no longer accurately represents the original signal. The ability to harvest useful activity information from an aliased signal is significantly compromised. In addition, these filters vary from one component manufacturer to another; digital accelerometers typically have a preset filter bandwidth much wider than is needed for human activity measurement, requiring further filtering within the digital domain to achieve the desired bandwidth. While this dual filtering approach is a perfectly valid solution in this circumstance, there are numerous ways in which it can be implemented. Each solution presents itself with a number of tradeoffs that must be made, and in the scenario of digital filtering the issue of the repeating spectrum is introduced. Unlike its analog filter counterpart, digital filters exhibit a repeating frequency spectrum that repeats at the sample rate. This must be taken into account during the system design to ensure these repeating spectrums do not inadvertently allow undesired acceleration noise to impact the final output.




Figure 2 Digital Filter Frequency Response

Oversampling

When both the bandwidth, which is set by the anti-aliasing filter, and the sample rate, which may or may not include sample rate shifting techniques, are considered, it becomes obvious that two sample rates may actually have the same bandwidth. If the Nyquist criterion is met in both cases, then both sample sequences can fully recreate the original signal.

There can be benefit, however, to oversampling (sampling more than twice the highest signal frequency) data, and this has been utilized in the audio industry for years. The primary benefit is an improvement in the Signal to Noise Ratio (SNR). In regards to day-to-day activities for the average person, this likely does not hold value. But for those subjects that do not generate typical acceleration values in their ambulation, such as elderly who may shuffle their feet, oversampling could improve results as more signal is discernible in the very low acceleration range.

It is clear from this brief discussion that these two parameters, bandwidth and sample rate, have the ability to significantly impact device output. Without knowing precisely how other devices solve these core problems, the ability to

consistently normalize outputs becomes very difficult. Further discussions will show that key specifications of the accelerometer Integrated Circuit (IC) itself can vary, also introducing output variations between devices.

Specifications

While there are numerous specifications that determine a MEMS accelerometer or activity monitor's performance, this white paper will only discuss dynamic range, sensitivity, noise density, and resolution. These four parameters and how they are used have the ability to impact device output in a significant manner.

Dynamic Range

The dynamic range of an accelerometer is defined as the range of accelerations a device can successfully measure while still maintaining a linear response. Typical accelerometers on the market today can measure +/- 2g, +/- 4 g, +/- 6g, and +/- 8g. In some instances, these are programmable on the device to maximize flexibility.

It should be noted that although an accelerometer may be specified to a range of +/-6 g, it is possible that it will report values outside of this range if the stimuli exceeds these range




Figure 3 Implications of Dynamic Range Non-Linearities

However, there is no guarantee of accuracy. Depending on the type of subject activity and wear location (wrist and foot worn devices generally experience higher accelerations than waist worn devices), it wouldn't be considered abnormal for the device to be subjected to accelerations outside of its printed specification range. While not damaging to the device, this could lead to erroneous data if not handled properly. Some manufacturers choose to limit the reported output to the printed specification range of the accelerometer while others report whatever value is collected. This difference in approach from manufacturer to manufacturer can further complicate matters.

Sensitivity

The sensitivity specification of an accelerometer is typically given in milli-volts/g for analog output devices and milli-g/LSB (Least Significant Bit) for digital output devices. It provides a measure of the device's sensitivity to the input it experiences. A more sensitive device will provide a 'larger' output for a given stimulus. Typically, the more sensitive the device, the better it is at measuring and discerning low amplitude signals as the output is 'amplified' to overcome the inherent noise.

To the end user, this implies that two devices with very similar overall specifications may yield different results in low acceleration environments.

Noise

Noise is an often overlooked specification within the activity monitoring world, but the importance of this parameter will continue to grow as end user requirements and subject

populations become more focused in order to study very specific areas.

Typically provided in $\mu\text{g}/\sqrt{\text{Hz}}$, noise density's impact on the overall output is dependent upon the system bandwidth, which is set by the anti-aliasing filter. If not properly taken into consideration, this parameter can impact the activity monitor's ability to reliably discern between internally generated noise and low acceleration stimuli.

Resolution

The final parameter of interest is the resolution in which the data are presented. This is a function of the dynamic range of the accelerometer and the number of bits of the Analog to Digital Converter (ADC) that are employed. It is a measure of the size, in g's, of the discrete levels used to represent the original signal and provides a measure for the lowest discernible acceleration that can be measured, in the absence of noise. This is important because two devices with differing resolutions will report different values while subjected to the same stimuli.

Conclusions

In recent years, many human activity monitoring solutions have become available, and their steadily increasing variety of features and options has the potential to result in end user confusion. It has been discussed that most activity monitors manufactured today utilize one of two types of transducers: piezoelectric or MEMS capacitive. Each type is a suitable solution with varying pros and cons, but the general consensus has been to move towards the use of MEMS based devices due to the

overwhelming availability of components on the market, ease of manufacturing, and lack of field calibration requirement.

The best representation of any measured signal is the original signal, and any function performed on the signal beyond its original analog state impacts the output in its own unique manner. It has been shown that there are numerous factors and parameters that impact the overall device output including bandwidth, filter strategy, sampling rate, sample rate accuracy, dynamic range, sensitivity, noise density, and resolution. These parameters work together in varying ways, often times working against other functional requirements such as battery life, size, or water resistance. Such interaction naturally requires tradeoffs to be made during device design.

Due to the number of parameters that have the potential to negatively impact the measurement result, the practicality of normalizing outputs from one manufacturer's device to another is very low. While the belief in doing so is that devices could be intermixed and data exchanged regardless of activity monitor manufacturer, the reality is a high level of collaboration would be required across device manufacturers. Furthermore, doing so would remove the unique and intrinsic value offered by different companies, reducing years of design and engineering knowledge to common and public knowledge.