Low Power Sensor Fusion for Wearables and Wireless Devices

Author(s) Christopher Piggin, Wessel Lubberhuizen

Revision Rev. A

Release Date 1/5/2018

Dialog Semiconductor

Syggrou Av. 143,
17121, N.Smirni, Athens
Greece

Phone: +30 210 93 10 580
Fax: +30 210 93 10 581

Neue Straße 95
73230 Kirchheim/Teck
Germany

Phone: +49 7021 805-0
Fax: +49 7021 805-100

Introduction

Over recent decades there have been many advances in the miniaturization, processing capability and power efficiency of microelectronic and wireless technologies. These have led to the advent of smart computing devices that operate independently and that are small, light and power efficient enough to be easily handled or worn about one's person on a daily basis. Many of these devices can sense their environment as well as various aspects of their wearer's state and movement using sensors which detect temperature, pressure, humidity, heart rate, proximity, touch, acceleration, rotation and magnetic field.

While each individual sensor has its own specific use, the real value comes when combining the information from multiple sensors using intelligent sensor fusion algorithms to derive more complex metrics that provide a better indication and overall understanding of what is happening. Sensor fusion algorithms are extremely useful for tracking peoples' activities and behavior as well as for allowing people to interact with their devices using gestures, touch and other real world stimuli.

The integration of smart devices into the real world is set to revolutionize how we live our lives and how we interact with the virtual world in a way that will eclipse the personal computer revolution of a few decades ago.

Wireless Sensor Devices and their Applications

Wearables and other types of wireless sensor devices are typically battery powered so minimizing the amount of power they consume is a critical requirement. They also often need to communicate with other devices across a wireless interface. Therefore the amount of data they need to transmit, and the power they consume while doing so, is also an important consideration.

Typically such devices contain multiple sensors like accelerometers, gyroscopes, magnetometers and barometers to detect motion, orientation and environmental conditions. The raw data from these sensors is generally of limited use by itself so must be processed, combined and analyzed using a process called sensor fusion in order to extract useful information and statistics.

A wide range of wearable and wireless devices and countless applications could potentially be enabled and enhanced by sensor fusion. They include:

- Wearables such as smartwatches and activity trackers
  - Activity/fitness tracking (e.g. step counting, distance travelled, floors climbed, calories burned, classification)
  - Heart rate monitoring
  - Pedestrian dead reckoning and indoor navigation
  - Gesture recognition for user interface control
- Wireless controllers such as TV/media center remote controls, game/VR controllers and air mice
  - 3 DoF (Degrees of Freedom) orientation tracking
  - 6 DoF controller positioning/dead reckoning
  - Gesture recognition

Figure 1: Examples of Wearable Devices

Wearable and Wireless Sensor Device Architectures

The architecture of a typical wearable sensor device includes a central microprocessor or microcontroller acting as a sensor hub which captures data from one or more sensors and processes it. It might also write the data to the device's local storage or display it to the user by some means.

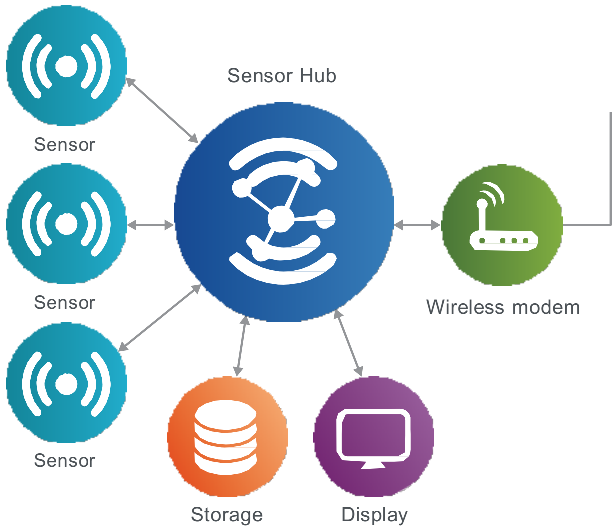

Figure 2: Wearable Device Architecture

Generally, wearables also have a wireless interface so that information can be transmitted to a host device such as a smartphone, tablet or personal computer for display or further processing and/or storage. The host device can then in turn upload the data to the cloud for more permanent storage and sharing.

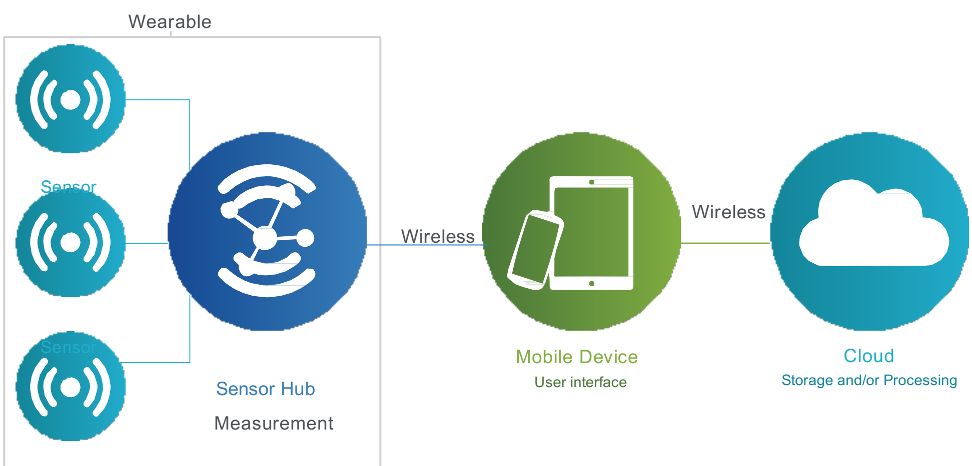

Figure 3: Typical architecture for wearable data processing and storage.

When deciding on the system architecture for such a device it is important to consider where the sensor fusion processing needs to be done and where the added value of the information that it provides is required. There is a range of options for this and determining which is the most suitable largely depends on the application's requirements. Processing can be done either within the wearable device, externally on a host device, in the cloud, or by some or all of these in combination.

Perhaps the simplest architecture is where the raw sensor data is either streamed to the host device via a wireless interface or to local storage for later transmission to the host. This requires only minimal processing capabilities within the wearable device, making it a fairly cost efficient solution. However, the sensors generate a lot of data requiring either a lot of memory to buffer it within the device, adding to the cost, or bandwidth to transmit it across the wireless interface, using a lot of power and therefore shortening battery life.

This architecture is adequate for devices where the information is ultimately utilized on the host or in the cloud, but it is not suitable for applications where processed information must be reported to the user directly on the device's display. It is also not suitable for applications involving multiple and/or high data rate sensors, as the ability to stream this data wirelessly or store it locally is limited. Suitable applications include low data rate activity and health tracking on low end devices that do not have display capabilities and are generally quite limited in their functionality.

Figure 4: Example of a low end and low data rate device which does not perform sensor fusion locally.

An alternative architecture is one in which some or all of the sensor fusion processing is performed within the wearable/wireless device. This approach has the advantage of being able to analyze, combine, numerically integrate and decimate the raw sensor data close to the source. This means only the essential processed information needs to be transmitted across the wireless interface, thereby minimizing the amount of bandwidth and power required to transmit it. The downside to this approach is that more resources and processing power is required on the device, potentially adding to the

This type of architecture is preferable for devices where the processed sensor information is utilized and/or displayed locally on the device, or where the amount of raw sensor data generated would be too excessive to store locally or transmit across the wireless interface. By performing sensor fusion within the wearable/wireless device, more sophisticated device applications can be enabled that would otherwise be impractical. This includes applications such as high accuracy motion controllers and higher end activity trackers.

Figure 5: A high end device which performs sensor fusion locally.

Sensor Fusion Algorithms

There are many applications and algorithms associated with wearable and wireless devices. Perhaps the most common are those that attempt to detect, interpret and classify the user's motion.

While it is possible to detect and track simple motion using data from just a single sensor - such as counting steps using data from an accelerometer - more accurate or complex motion tracking requires algorithms that look at and combine data from multiple sensors to get a more complete indication of what is occurring.

One such algorithm fuses data from a device's inertial and magnetic sensors to work out its real-world orientation. This is called an Attitude and Heading Reference System (AHRS).

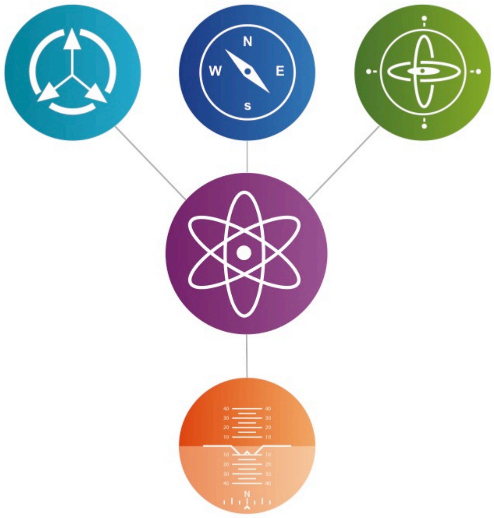

Complementary Filtering

By combining data from multiple sensors, the usefulness, accuracy and stability of the derived information can be enhanced by taking advantage of their individual strengths while compensating for any weaknesses.

For example, inertial sensors (accelerometers and magnetometers) work well at low frequencies; together they can track absolute orientation but they are fairly noisy and do not track rapid rotations very well. Gyroscopes on the other hand work well at high frequencies; they track rapid rotations very accurately and with little noise but they have no awareness of their absolute orientation.

By combining low-pass filtered information from the inertial sensors with high-pass filtered information from a gyroscope, the sensor fusion algorithm has a much better overall indication of what is happening over the entire frequency band. This means it can now to accurately track both absolute orientation and rapid rotations at the same time.

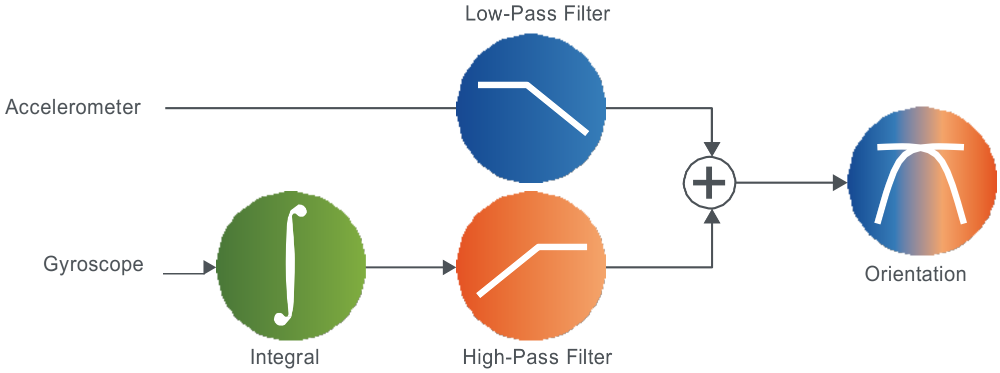

Figure 8: Complementary Filtering

Sensor fusion algorithms which filter and combine data from multiple sensors in this way are known as complementary filters.

Kalman Filtering

Complementary filtering works well in situations where the data from various sensors is indeed 'complementary' in the sense that the movement that each detects is consistent between them. It doesn't work so well when the data is inconsistent, such as when the sensors have large distortions or the real-world absolute references, such as direction of gravity or magnetic field, are unreliable. These inconsistencies can result in distortion and drift artefacts in the fused output data.

To resolve this, more sophisticated sensor fusion algorithms have been developed that rely on Kalman filtering to compensate for these inconsistencies. Kalman filters work by modelling the expected behavior of the sensors in response to movement, taking into account noise and distortions, and comparing this with what actually happens in order to iteratively refine the model until it accurately reflects the true behavior.

The advantage of Kalman filters is that they are much more robust and less prone to generating artefacts than complementary filters. Not only are they able to compensate for the noise and distortions exhibited by the sensors but they can also detect scenarios where data from one of the sensors is unreliable and should be ignored.

The disadvantage of Kalman filtering is that it is much more complex and computationally demanding than complementary filtering. The mathematical calculations that are needed to implement a Kalman filter generally require high precision and dynamic range, and are typically implemented in floating point. They also require storage of large amounts of data. Systems that implement Kalman filter based sensor fusion therefore require computing platforms with more complex architectures and much greater processing and memory capabilities. As a consequence they are much more expensive and less power efficient.

SmartFusion™ Library

Dialog Semiconductor's SmartFusion library provides functionality for calibrating, processing and fusing 3D inertial and magnetic data from MEMS sensors.

Sensor Calibration

Optimized routines are provided for applying calibration efficiently to captured raw sensor data in order to compensate for the distortions which are characteristic of these types of MEMS sensors. The library can not only be used to apply externally derived static calibration coefficients but also to determine sensor distortions 'on-the-fly' and calculate and update the calibration coefficients automatically.

The SmartFusion Auto-Calibration routine is a sophisticated algorithm that quickly and efficiently determines and compensates both for offset and scaling distortions seen in all sensors, and for hard and soft-iron spherical distortion seen in magnetometer data.

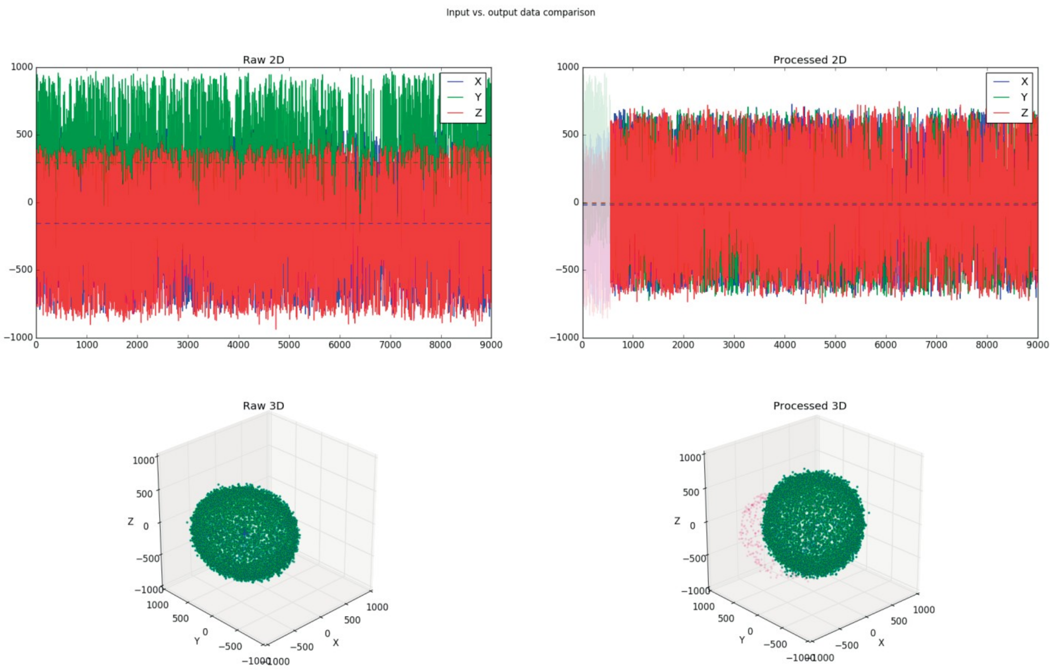

Figure 9: Comparison of Raw vs. Calibrated Magnetometer Data

Sensor Fusion

The library also includes the SmartFusion AHRS algorithm. This fuses data from a MEMS accelerometer, gyroscope and/or magnetometer in order to compute orientation relative to the Earth frame of reference.

The SmartFusion AHRS algorithm takes a different approach to sensor fusion from other solutions available on the market. It has been specifically designed to perform fusion of sensor data in the most computationally efficient way possible, removing any unnecessary complexity while still maintaining good performance and accuracy.

Traditionally sensor fusion algorithms represent orientations and rotations using matrices or Euler angles. By using quaternion geometry, the SmartFusion AHRS algorithm can perform calculations far more efficiently and accurately than is possible using the matrix based calculations used by some other solutions. It also avoids numerical issues such as gimbal lock.

The output orientation generated by the algorithm is represented in 16 bit fixed point unit quaternion form. This allows orientation information to be transmitted in a data efficient form using the minimum amount of available bandwidth for subsequent wireless transmission to an external device.

The SmartFusion™ Advantage

Dialog Semiconductor has designed the SmartFusion library from scratch to ensure all of its functionality has been implemented as efficiently as possible. A particular focus was placed on keeping the algorithms' footprint, in terms of processing power, memory usage and power consumption, to an absolute minimum.

All the routines within the library have been implemented using fixed-point calculations exclusively, thereby eliminating the need for higher footprint floating point hardware or external software libraries for floating-point emulation. They can therefore be executed on the smallest, least complex and most power efficient microprocessors/microcontroller architectures that exist in the market.

The library's low processing requirements allow a much lower processing clock speed than is necessary for most competing solutions. Running on an ARM Cortex-M0 clocked at 16 MHz, it can easily process sensor data from an inertial measurement unit (IMU) at output data rates of over 1 kHz. Other solutions typically require a Cortex-M4 or M7 clocked at over 50 MHz as a minimum to process data in a similar way, even at lower rates.

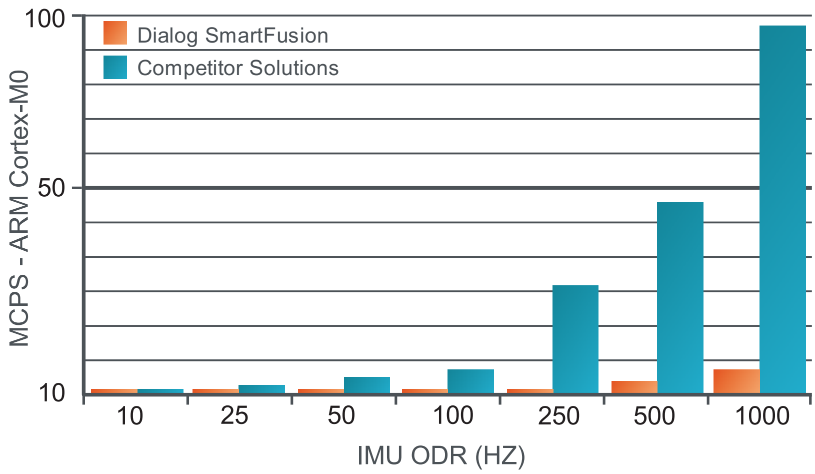

Figure 10: Comparison of Dialog and competitor sensor fusion solutions with regards to processing efficiency

Devices using Dialog Semiconductor's SmartFusion library can therefore be designed with much more power and cost efficient hardware than would otherwise be possible. The minimal processing requirements of the library allow sensor calibration and fusion to be performed much more quickly than is possible with other solutions. This makes it very well suited to applications that require low latency.

The library's memory requirements are also tiny. The entire library, including all sensor calibration and fusion functionality, has a memory footprint of less than 5 kB. Other solutions typically require over 100 kB of memory.

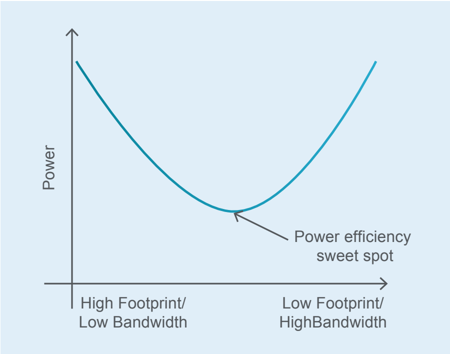

Figure 11: Power vs. Footprint/Bandwidth

This novel approach combines the best characteristics of high footprint/low bandwidth solutions with low footprint/high bandwidth solutions to provide a solution which is both low footprint and low bandwidth. It enables sensor fusion to be performed within a wearable device in order to minimize the latency and bandwidth of the data transmitted across the Bluetooth Low Energy interface but using the minimum amount of computing resources necessary, thereby hitting the power efficiency sweet spot.

When combined with Dialog Semiconductor's range of Bluetooth Low Energy SoC solutions, it provides the most power efficient and low cost wireless sensor fusion solution on the market today.

To illustrate, a wireless sensor fusion design using the SmartFusion library running on Dialog's DA14583 SoC in combination with suitably low power sensors can draw as little current as 1.5 mA while in operation. Designs based on other solutions typically operate with more than ten times this amount of current.

The SmartFusion library has been designed with flexibility and scalability in mind. It is capable of supporting a wide range of sensor types, combinations, sample rates and use cases. It is not dependent on the use of sensors from any particular manufacturer and can be tuned to achieve the best performance according to the individual behavior and characteristics of specific sensors.

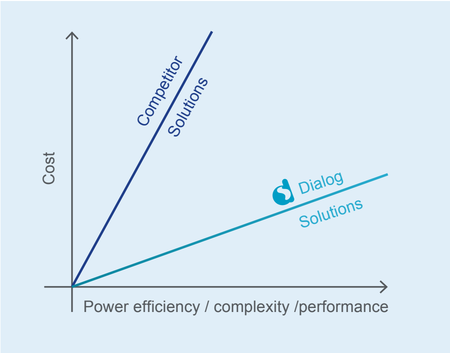

Figure 12: Comparison of Dialog and competitor wireless sensor fusion solutions
with regards to power, cost, complexity and performance

The library can be configured to support different sensor combinations in various 3, 6 & 9 degrees of freedom scenarios.

Mode	Degrees of Freedom	Gyroscope	Accelerometer	Magnetometer	Supported
Gyroscope Only (G)	3	✓	✗	✗	✓
Gyroscope + Accelerometer (GA)	6	✓	✓	✗	✓
Gyroscope + Magnetometer (GM)	6	✓	✗	✓	
Accelerometer + Magnetometer (AM)	6	✗	✓	✓	✓
Gyroscope + Accelerometer + Magnetometer (GAM)	9	✓	✓	✓	✓

Figure 13: Supported Sensor Combinations

It also supports a wide range of sensor sample rates and even allows use cases where each sensor is sampled at a different rate. Its tunable API allows it to be configured in a variety of ways so that it can meet the performance requirements of a broad array of potential applications and use cases.

SmartFusion™ Hardware Platforms and Development Kits

The SmartFusion library is provided free of charge for customer designs which use the Dialog Semiconductor SmartBond range of Bluetooth Low Energy SoCs. These include:

- DA1458x - World's smallest, lowest power and most integrated Bluetooth Low Energy SoC
  - Beacons & Proximity, Health & Fitness, HID, Smart Home
- DA1468x - Highest integration, flexibility and security
  - Wearables, Virtual Reality, Smart Home, Apple HomeKit, HID, Other rechargeable device

The SmartFusion library is available on and can be evaluated with the following development kits:

- **DA14583 IoT Sensor Development Kit**

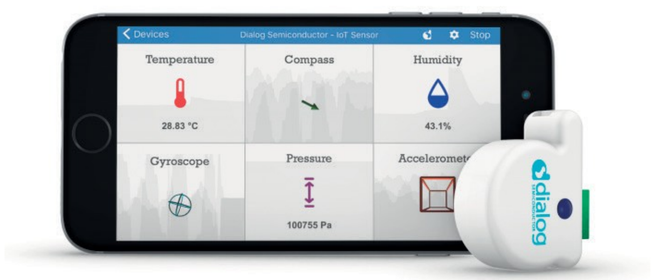

Figure 14: IoT Sensor Dongle with Associated Mobile App

Figure 15: IoT Dongle with Associated Interface Board

► DA14681 Wearable Development Kit

Figure 16: Wearable Device with Associated Mobile App

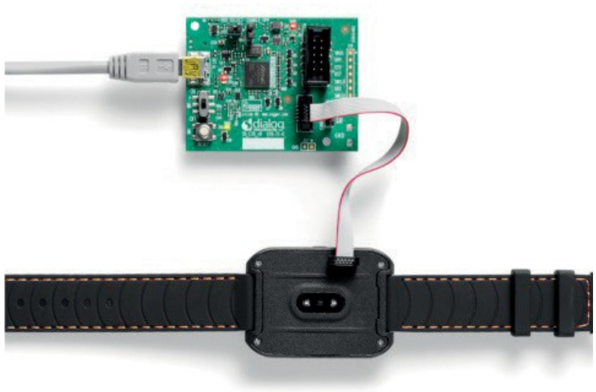

Figure 17: Wearable Device with Associated Interface Board

► DA14585 IoT Multi-Sensor Development Kit

Figure 18: IoT Multi-Sensor Development Kit with Associated Mobile App

Further details about supported devices and developments kit can be found on the Dialog Semiconductor website:
https://www.dialog-semiconductor.com/bluetooth-low-energy

For access to datasheets, documentation, reference designs and software tools, please create an account at the SmartBond™ support website:
https://support.dialog-semiconductor.com/connectivity

