AZD004

Azoteq Sensing Technology Introduction

Contents

1 Introduction

Azoteq ProxSense® devices provide capacitive sensing solutions, offering highly accurate analogue measurements that can detect minute changes in capacitance ranging from picofarads to femtofarads. These devices leverage advanced digital signal processing to detect minute changes in capacitance. Focus is placed on detecting very fine differences in capacitance rather than absolute capacitance measurements. This is particularly advantageous in user interface applications such as capacitive touch buttons and trackpads, where absolute capacitance measurements are often unnecessary and pose challenges due to variations in PCB tolerances.

ProxFusion® expands upon the capabilities of ProxSense® devices by incorporating additional sensing techniques such as Hall-effect sensing and inductive sensing (for metal detection) into the analogue measurement process. This broadens the range of applications where these devices can be utilised.

The document introduces charge-transfer-based capacitive sensing and covers various fundamentals such as current consumption, component choices, long-term average (LTA) techniques, automatic calibration (ATI), and ProxFusion® technologies. For more detailed design guidelines related to capacitive sensing, refer to application note AZD125.

2 The Charge Transfer Method

Azoteq uses the charge transfer method to measure external charge with its ProxFusion® and ProxSense® ICs. Using a simple analogy, charge and capacitance can be represented by a liquid and a container, represented in Figure 2.1. The smaller container is the variable capacitance of the sensing electrode, while the larger container is a fixed reference capacitance (C_s). The C_s capacitor is internal to Azoteq ICs.

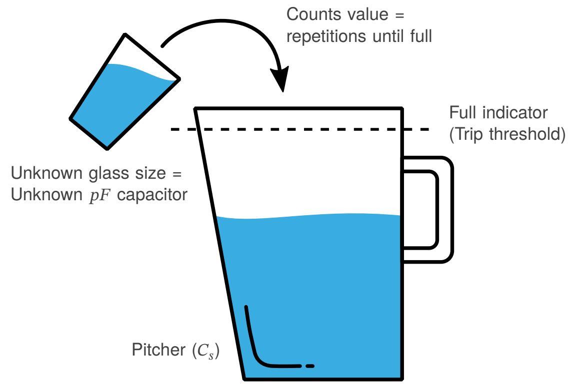

Figure 2.1: Charge Transfer Model

The smaller container is repeatedly filled (charged) and then emptied (transferred) into the larger container, until the larger container is filled. By counting the number of transfers it takes to fill the larger container, the size of the smaller container can be deduced. For example, if it takes 20 transfers to fill the large container, the smaller container is 1/20th the size of the larger container.

Capacitive sensing works similarly, relying on the formula that relates capacitance (C) to voltage (V) and charge (Q):

$$C = \frac{Q}{V}$$

The external sensing electrode, with an unknown variable capacitance, is charged up to a known voltage, storing some amount of charge. The amount of charge stored is dependent on the electrode's inherent capacitance to ground as well as the capacitance between the user and the electrode. As a user or object approaches the electrode, the capacitance, and therefore the amount of charge, will increase.

That charge is then transferred to the internal reference capacitor, which has a known fixed capacitance (typically in the range of 40 pF to 80 pF). This charge and transfer process is repeated until the voltage over the reference capacitor reaches a threshold level, and thus a known amount of charge. The Azoteq device counts the number of charge transfer cycles required to fill the reference capacitor; this number is simply referred to as *counts*. This counts value provides an indication of the capacitance of the sensing electrode – larger counts values implies that the electrode capacitance was small, and could only store a small amount of charge. Therefore, more charge transfer cycles were required to fill the reference capacitor. This process is shown in Figure 2.2.

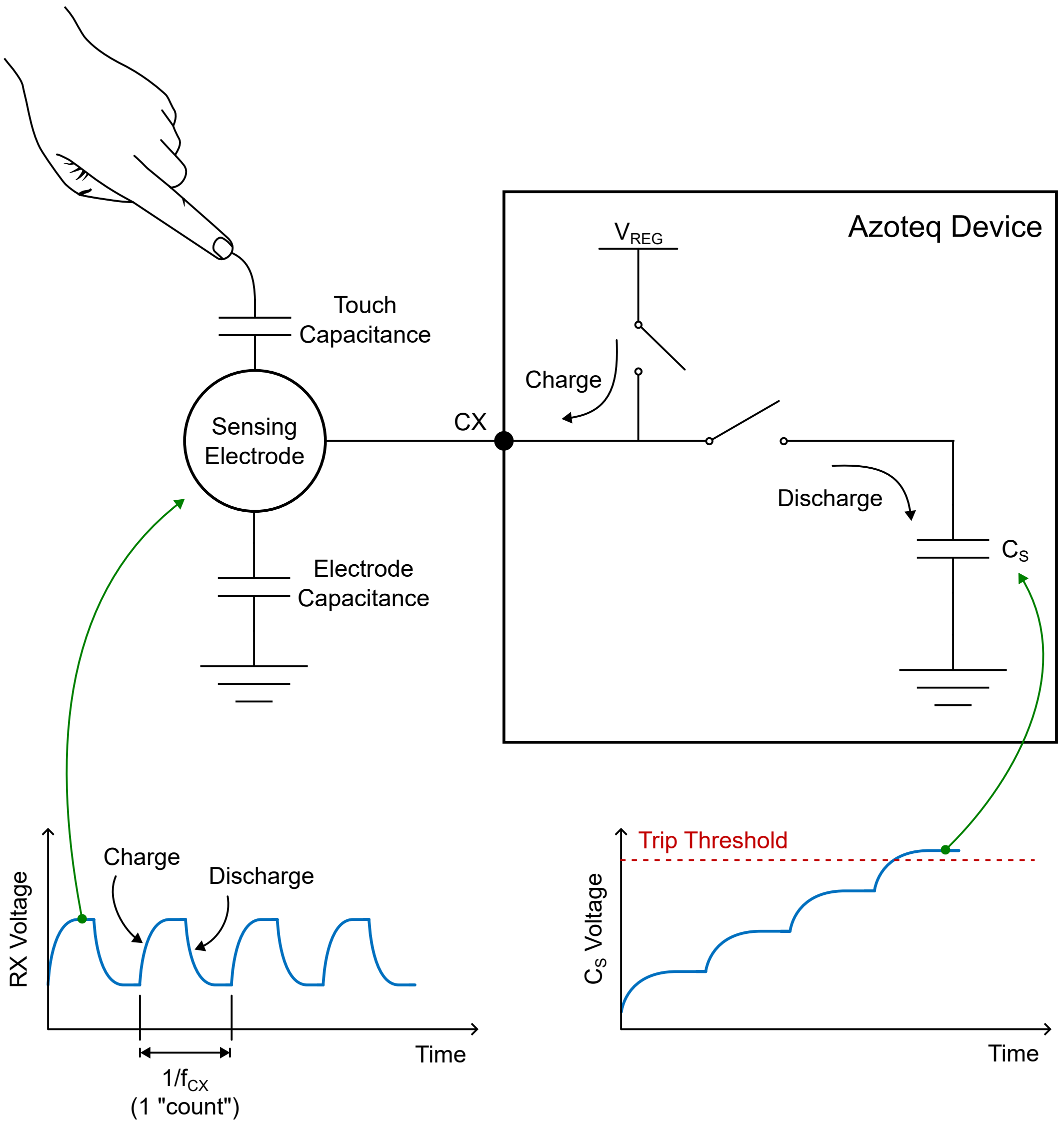

Figure 2.2: Charge Transfer Model

The capacitance of a sensing electrode is largely dominated by inherent “load” or “parasitic” system capacitance. Azoteq devices specialise in detecting minute changes in this capacitance.

Note that the rate at which the sensing electrode is charged and discharged is known as the *charge transfer frequency* or *conversion frequency*, f_{cx}. This has some implications on the accuracy of the capacitive measurement, as well as the duration of the sensing period and overall system current consumption.

More information on the charge transfer method and charge transfer cycles can be found in application note AZD102.

3 Azoteq Sensing Technologies

3.1 Self-Capacitive Sensing

Self-capacitance technology makes use of the parallel plate capacitor theory where

$$C = \frac{\epsilon_r \epsilon_o A}{d}.$$

The capacitance is measured between the electrode and the system's ground. It is dependent on the surface area A of the electrode and the approaching conductive object, as well as the distance d between them.

- As a finger (or other conductive object) approaches the electrode the distance (d) between the electrode and the earth decreases, effectively increasing the capacitance (C).
- From the formula $Q = CV$, as C increases, the charge (Q) per transfer increases too.
- This will decrease the number of transfers required to charge the reference capacitor C_s to a fixed voltage. Therefore, counts decrease when touching self-capacitance applications.

Typically, self-capacitive technology operates at a conversion frequency of around 250 kHz. As the capacitance of the electrode increases, it takes longer to be charged completely, and thus requires a lower conversion frequency to maintain accurate sensing.

Figure 3.1 shows an equivalent circuit model for a typical self-capacitive application.

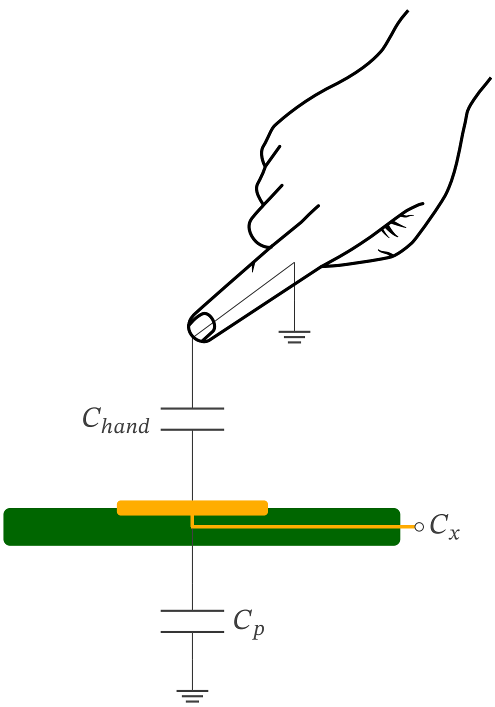

Figure 3.1: Self Capacitance Circuit

The term C_p represents the inherent capacitance of the electrode with respect to the ground potential. This is typically referred to as the parasitic capacitance. It is this capacitance, along with the capacitance coupling to the user $(C_p + C_{hand})$, that is measured at the C_x pin.

3.2 Mutual-Capacitive Sensing

Electrically charged conductive objects close to one another will form an electric field around them. *Mutual-capacitive* technology measures the change in capacitive coupling (mutual capacitance, C_m) between *two* electrodes. Mutual-capacitive sensing relies on two electrodes, the transmitter (CT_x) and receiver (CR_x). Figure 3.2 shows the capacitance C_m between the electrodes and the parasitic capacitances associated with each electrode.

- As a finger (conductive object) approaches the sensor, the electrodes couple more with the finger and effectively "steal" some of the charges. This causes the capacitance C_m between the electrodes to decrease.
- From the formula $Q = CV$, as C_m decreases, the charge (Q) per transfer will decrease too.
- This will increase the number of transfers required to charge the reference capacitor C_s to a fixed voltage. Therefore, counts go up when touching mutual-capacitive applications.

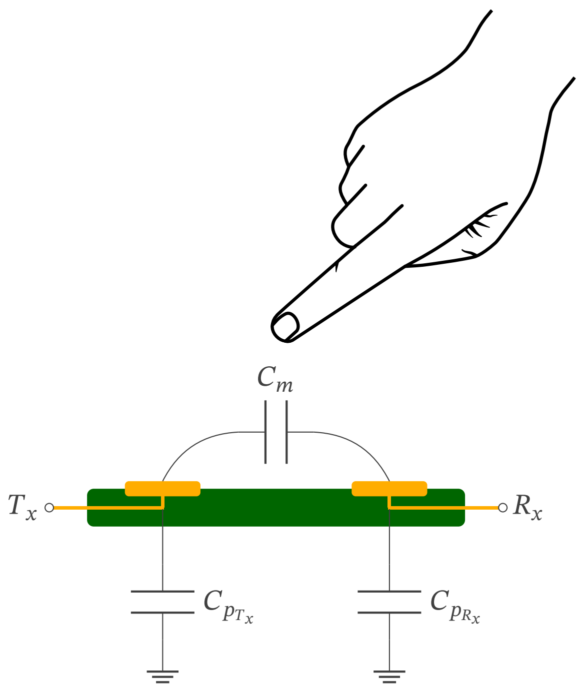

Figure 3.2: Mutual Capacitance Circuit

This explanation assumes that the finger is strongly coupled to the earth. In applications where the approaching conductive object is floating, C_m is typically found to increase. However, for a typical touch application, the user will be closely coupled to Earth.

A typical mutual-capacitive application operates at a conversion frequency of around 1 MHz.

Mutual-capacitive technology can be used in single-channel touch button applications as well as in keypads and trackpads. It is suited to these applications because the sensing electrodes are tightly coupled to one another, and therefore multiple electrodes can be placed close to one another with minimal inter-electrode interference.

3.3 Resonant Inductive Sensing

By placing a capacitor and inductor in parallel as shown in Figure 3.3, an *LC tank* is formed. This circuit has a resonant frequency f_{res}. The resonant frequency is dependent on the value of the inductor and capacitor. Thus, by keeping the capacitor C fixed, a change in the inductance L can be detected by measuring a shift in the resonant frequency. This is done by driving the T_x node close to the resonant frequency and measuring the amplitude of V_{tank}.

When a metal object approaches the inductor, eddy currents are formed in the object. This causes the frequency response of the *LC Tank* to shift and results in a decrease in the amplitude of V_{tank}. Azoteq's ProxFusion® and ProxSense® ICs drive the T_x node and measure the amplitude of V_{tank} at the R_x node to measure the change in the inductance L. In this way, the presence of a metal object near the inductor can be detected.

Typical applications for inductive sensors include waterproof snap-dome buttons and metal flex force sensors.

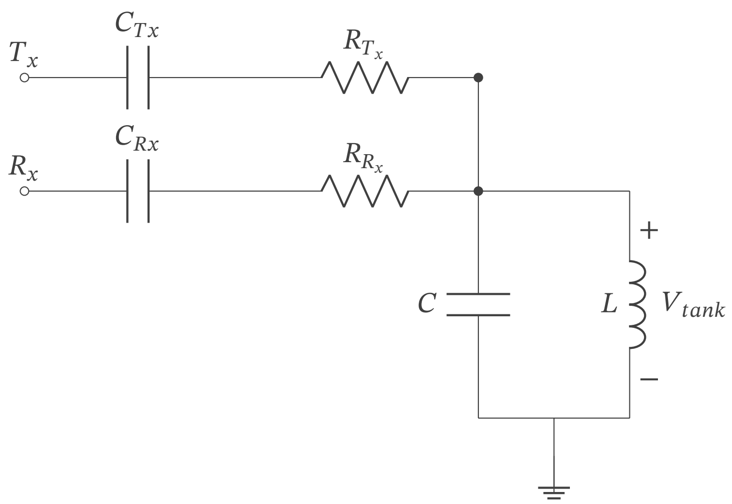

Figure 3.3: LC Tank Circuit for Inductive Sensing Mode 1

For more details on this sensing method please refer to AZD115.

3.4 Hall Effect Sensing

The Hall Effect is a phenomenon where moving charge carriers can be deflected in the presence of a magnetic field, producing a potential difference. Azoteq's hall sensing ICs use this phenomenon in combination with our ProxFusion® Technology to measure magnetic field strength using dedicated Hall plates within the IC.

Some Azoteq devices offer a single Hall-effect sensor that can be used as simple magnetic switches. This can be used in applications such as earbud docking detection.

Certain devices offer multiple Hall-effect sensors that can be used to calculate field differentials, which are used to determine the orientation of discrete magnets. These are used to create Hall-effect rotation encoders, typically used in applications such as mouse wheels and volume control dials.

For more details on Hall sensing and rotational encoders, please refer to AZD127.

4 Signal Conditioning

The basic charge transfer method faces several challenges that need to be addressed for effective sensing applications.

- **Limited Input Capacitance Range**: The range of input capacitance is constrained by the size of the internal reference capacitor. This limitation can restrict the device's ability to accurately detect larger capacitance changes or variations. The internal reference capacitor is typically on the order of 40 pF to 80 pF. However, some applications may require handling a load/parasitic of up to 200 pF on the sensing electrodes.
- **Low Sensitivity**: A typical capacitive sensing application requires the measurement of minute changes in capacitance. Compared to the typical load/parasitic capacitances on the sensing electrode, these small changes may be well below 1% of the total capacitance. This makes it challenging to detect touches reliably.
- **Sensor Drift**: Various factors may cause the measured capacitance to change over time. These factors include temperature and humidity changes, mechanical changes (physical changes such as components bending), and environmental changes (where the sensor is used in various positions, such as handheld vs tabletop).
- **Uncontrolled Measurements Resolution**: Without additional measures, the measurement resolution may not be controlled, leading to uncertainties in the sensed signals.

To overcome these challenges, Azoteq's devices have integrated signal conditioning circuitry that serve several purposes:

- **Wide range of input signals**: The sensing system can accommodate a wider range of input signals and external capacitances beyond what the basic charge transfer method can handle.
- **Improve sensitivity for small capacitive changes**: The sensing system compensates for parasitic capacitances in order to focus on detecting very small capacitive changes, which is often the target in sensing applications such as touch or proximity sensors.
- **Automatic tuning**: By monitoring slowly-changing environmental conditions, the sensing system can automatically compensate for changes in parasitic capacitance, and use the integrated signal conditioning circuits to maintain sensitivity.

These on-chip compensation circuits enhance the capabilities of Azoteq's ProxFusion® and ProxSense® ICs, enabling them to overcome the limitations of the basic charge transfer method and perform effectively in various sensing applications. This compensation eliminates the need for additional external components, complex PCB design guidelines, or limitations on PCB thicknesses, materials, and overlays.

The signal conditioning circuitry consists of two main structures: "multipliers" to control the range/resolution of the input signals and "offset subtraction", which reduces the effects of the parasitic capacitance and DC offsets. See Figure 4.1 for a visual representation of the internal analogue components.

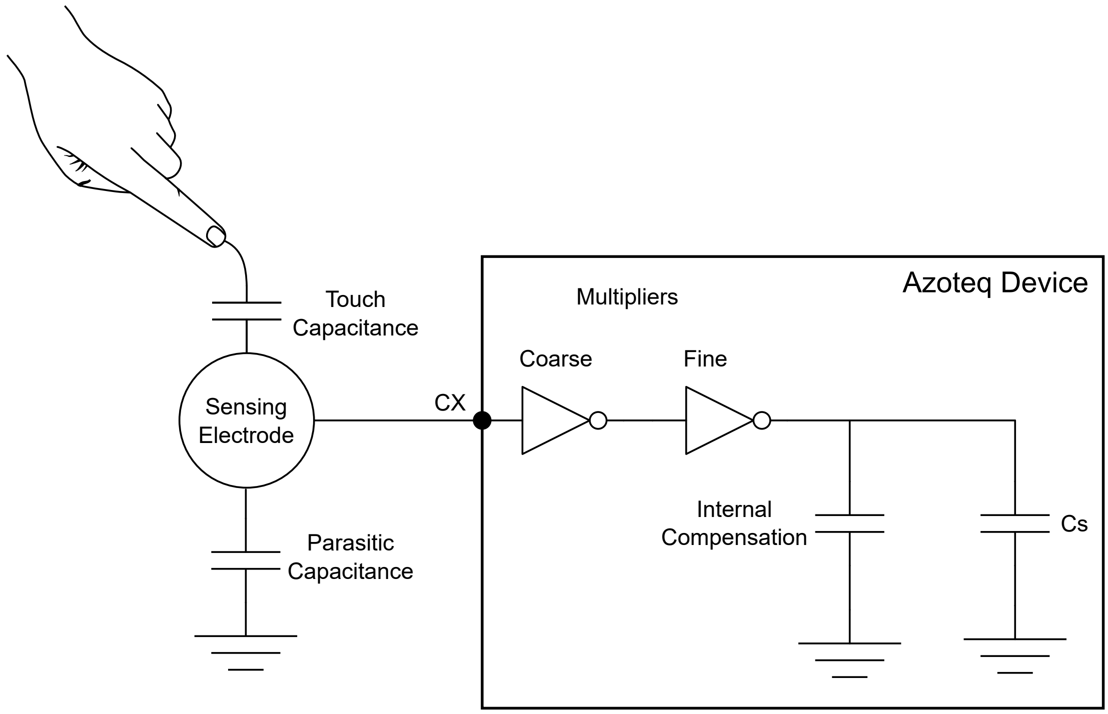

Figure 4.1: ProxSense® IC's Charge Transfer Overview

4.1 Resolution Control (Multipliers)

As the internal reference capacitor is relatively small, the charge entering the IC from the external electrode is divided to increase the number of charge transfer cycles needed to fully charge the reference capacitor. This division effectively "multiplies" the reported counts on the sensor, and is performed by a set of analogue circuits referred to as the "Coarse and Fine Multipliers".

Multiplying the counts increases the measurement resolution of the sensor. For example, consider a 10% change on 40 counts, which would provide a delta of 4 counts, compared to a 10% change on 1000 counts, which would provide a delta of 100 counts.

4.2 Offset Subtraction (Compensation)

Offset subtraction reduces the effect of static parasitic capacitance on a sensing electrode, since these parasitics can significantly reduce the sensitivity of the capacitive sensor. By subtracting the unwanted parasitics, the sensor becomes more sensitive to the variable capacitance caused by user interaction. In other words, it removes a portion of the DC component from the signal. Eliminating part of the DC signal further reduces the amount of charge transferred to the reference capacitor, increasing the number of counts.

Offset subtraction is achieved using an adjustable capacitor bank to subtract some fixed amount of charge during the charge transfer process. The amount of charge subtracted is represented by the "Compensation Value". Increasing the compensation value subtracts more parasitic capacitance, increasing the sensitivity of the sensor.

Compensation is performed after the two multiplier stages.

4.3 ATI Overview

ATI (Automatic Tuning Implementation) employs an advanced signal processing algorithm to optimise the hardware sensing circuits automatically. Automatic tuning allows designers to overcome design challenges such as parasitic capacitance, environmental effects, and manufacturing tolerances.

ATI uses a user-defined "Base" and "Target" value to calibrate the multiplier and compensation for the optimal operating range and sensitivity. ATI can be performed automatically during runtime to maintain optimal sensitivity across varying environmental conditions.

ATI offers several advantages:

- **Increased sensitivity**: ATI enhances sensitivity, allowing for improved detection of both touch and proximity events on the same sensor.
- **Automatic sensitivity adjustment**: ATI automatically adjusts sensitivity based on environmental conditions or changes in parasitic capacitance, ensuring optimal performance without manual intervention.
- **Design flexibility**: ATI reduces the conventional constraints on PCB layout, limits to overlay thicknesses and materials, and simplifies or negates the need for calibration during manufacturing. This makes integration into new designs simpler.
- **No external components or programming required**: ATI operates internally within the IC, eliminating the need for additional external components or programming to adjust sensitivity.

4.4 ATI Stages

The ATI routine in Azoteq's ICs is divided into two stages in order to calibrate the multipliers and offset compensation.

1. **Multiplier Calibration**:
   - The multipliers are calibrated to achieve some desired number of "Base" counts.
   - The ATI Base value represents the amount of charge to divide away by the coarse and fine multipliers.
   - The Base value denotes the desired number of charge transfer cycles necessary to fill the internal reference capacitor, after the multiplier stage but without any compensation enabled.
   - The Base value is specified by the designer, and the ATI algorithm modifies the multipliers until the resulting measured counts is as close to the Base value as possible.
   - Typical values for the Base value range from 100 to 200 counts.
2. **Compensation Calibration**:
   - The offset compensation is calibrated to reach a "Target" value.
   - The ATI target value corresponds to the number of charge cycles required to charge the internal reference capacitor with both multipliers and offset compensation enabled.
   - This step occurs after the multipliers have already been selected. The ATI routine adjusts the compensation value until the sensor reaches the Target counts.
   - Typical values for the Target value range from 500 to 1000 counts.

ATI is thus primarily controlled by these two parameters; the ATI Base and the ATI Target. These two parameters affect the sensitivity with the following relationship:

$$\text{Sensitivity} \propto \frac{\text{Target}}{\text{Base}}$$

Note that increasing the sensor's sensitivity by increasing compensation can result in more noise on the sensor counts.

4.5 ATI in Capacitive Applications

In capacitive applications, ATI compensates for the following:

- The parasitic capacitance, C_p
- Changes in capacitance caused by the environment (temperature and humidity)
- Variations in the electrode layout and routing
- Changes in the overlay thickness or type of overlay used
- PCB substrate type (e.g. FR4, FPC, etc)

4.6 ATI in Inductive Applications

In inductive applications, ATI compensates for the following:

- Changes in capacitance caused by the environment (temperature and humidity)
- Variations in the coils during manufacturing
- Variations in the coil layout and routing
- PCB substrate type (e.g. FR4, FPC, etc)

4.7 ATI in Hall-Effect Applications

In Hall-effect applications, ATI compensates for the following:

- Magnetic field variation
- Changes in environmental temperature
- Variations in mechanical implementation

5 Filtering Overview

Touch and proximity events are detected by comparing the current counts to a reference value, known as the Long-Term Average (LTA). The difference between the counts and the LTA is known as the Delta, and this delta value is compared to a set threshold.

The LTA is slowly updated over time using a low-pass infinite impulse response (IIR) filter. This enables sensors to detect small relative changes in the measured signal while remaining optimally sensitive and adapted to minor changes in the system operating environment.

5.1 IIR (“Beta”) filter

The IIR filter employed by IQS devices can be described as a “low-pass single-pole IIR filter”. The filter operation is diagrammatically represented in Figure 5.1:

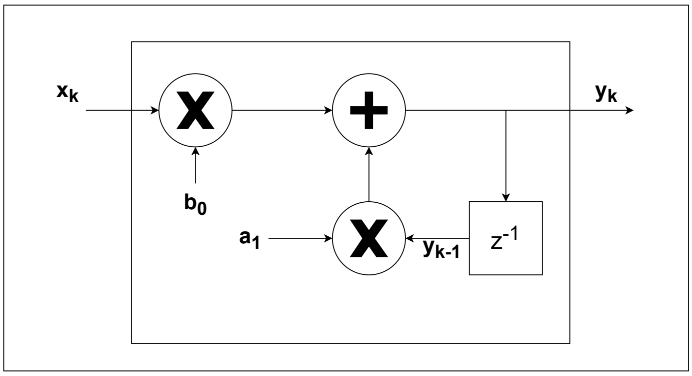

Figure 5.1: Beta Filter Diagram

In the above diagram,

- x_k is the new raw count sample,
- y_k is the filtered count output, and
- y_{k-1} is the previous filtered count value (first order time delayed feedback).

The filter coefficients are calculated with the same beta (β) parameter as:

$$a_1 = \frac{(2^\beta - 1)}{2^\beta} \qquad b_0 = \frac{1}{2^\beta}$$

The filter equation is thus given by:

$$y_k = a_1 y_{k-1} + b_0 x_k$$

And by substitution and simplification:

$$y_k = \left(1 - \frac{1}{2^\beta}\right) y_{k-1} + \frac{x_k}{2^\beta}$$

The “weight” of the new sample x_k is controlled by the beta value. Larger beta values place less weight on the new sample, causing the filter to respond more slowly to changes in the raw counts. This reduces the filter’s responsiveness but enhances its ability to reduce noise more effectively. Designers can configure the beta value to ensure that the noise reduction and responsiveness are optimal for the specific application.

5.2 Counts (“AC”) Filter

Raw measured samples are typically erratic and can exhibit some noise, which may cause false triggers or sporadic events. A beta filter is applied to the counts to reduce the noise and avoid false triggers. The result is known as the “Filtered Counts”.

Ideally, a fast-responding filter, with a low beta value, is used to acquire the filtered counts. Typical beta values used are between $\beta = 0$ (no filtering, rapid response) and $\beta = 4$ (slowly-following counts with some delayed response). This ensures that the filtered counts remain responsive to user interactions, preventing missed or delayed events. Figure 5.2 below shows the normalised step responses for these filter betas.

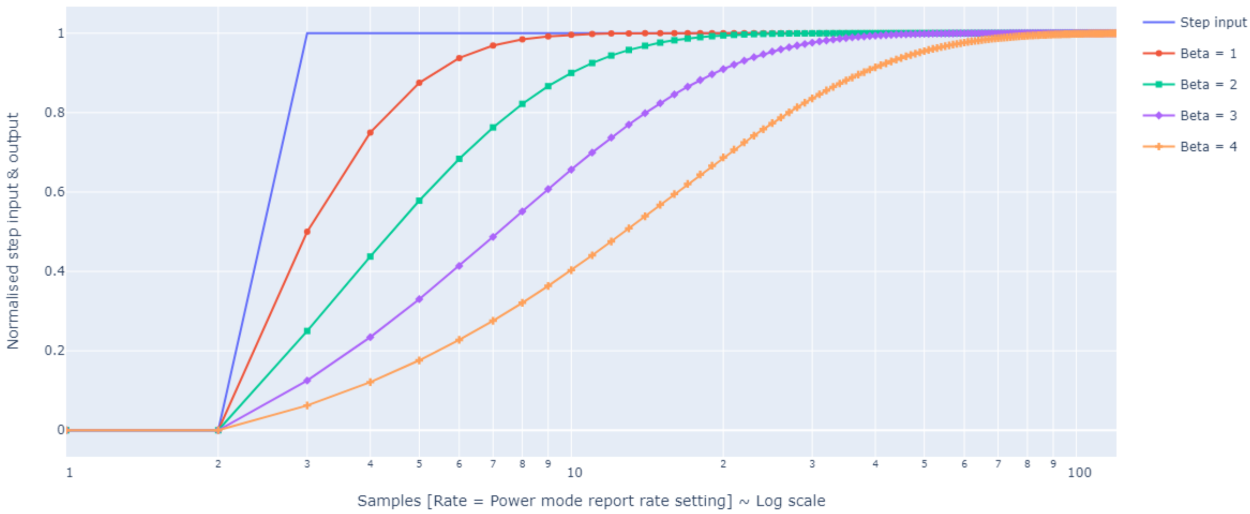

Figure 5.2: Counts Filter Beta Responses

5.3 Long Term Average (LTA)

The LTA is a filtered average of the measured counts. It allows ProxSense® and ProxFusion® ICs to intelligently track slow changes in the external environment.

Figure 5.3 below shows the normalised step responses for typical LTA beta values, which are between $\beta = 6$ (fast following) and $\beta = 10$ (slow following). These responses are notably slower than the counts filter. This is favourable in establishing a reference or LTA in a system with slowly varying environmental effects.

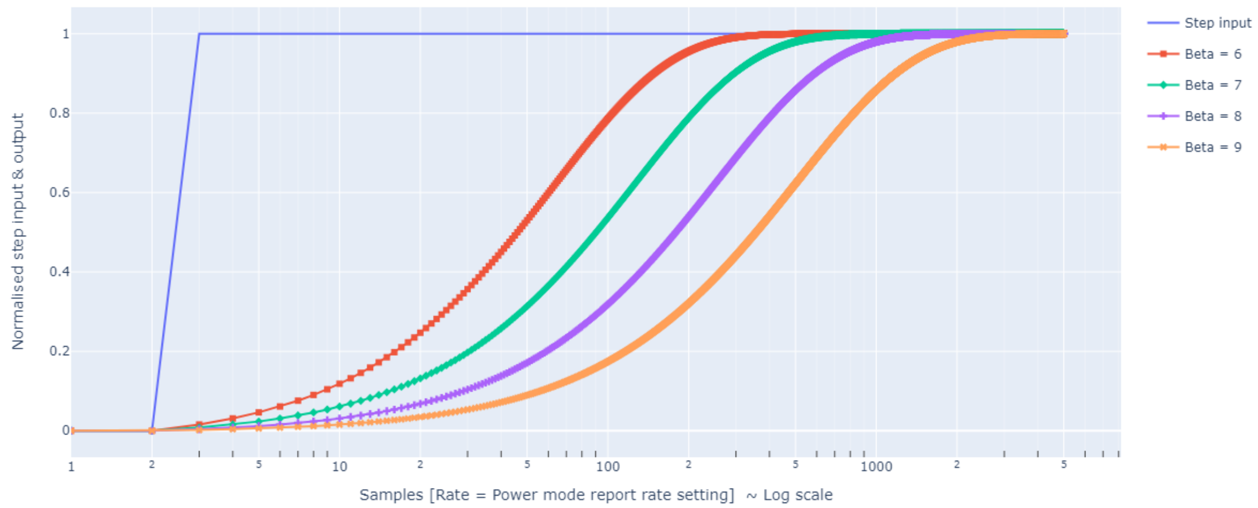

Figure 5.3: LTA Filter Beta Responses

A touch or proximity event is recorded if the current counts value and the LTA differ by more than a configurable threshold. During such an event, the LTA filter is halted. During this time the value of the LTA remains constant (i.e., not actively filtering). This prevents the user's interaction from affecting the LTA's environment tracking. Upon exiting a touch or proximity event, the LTA filter is resumed, and the LTA value is once again continuously updated.

As an example, see Figure 5.4. The counts value is represented by the blue line. The LTA value is represented by the red line. The Delta value is represented by the green line. The Delta is the difference between the counts and the LTA value. When the difference between the counts and LTA value is small, such as in areas A and C, the LTA follows the counts value. When the difference between the counts and LTA value is large, such as in area B, the LTA halts and does not follow the counts value.

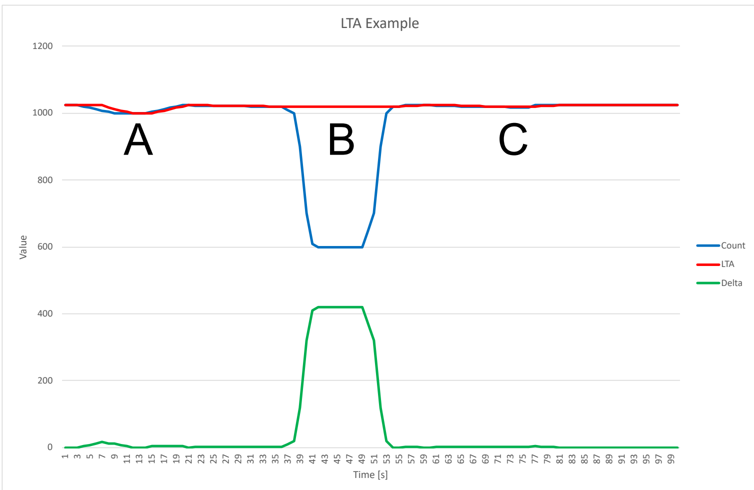

Figure 5.4: LTA follow Counts example for small variations

Halting the LTA may in some rare instances cause the sensor to get stuck in a proximity or touch event. This may happen if the environmental conditions change suddenly, or an additional load is added to a sensor. Therefore, the LTA filter is only halted for a maximum time, known as the filter halt period. If a touch or proximity event lasts longer than the halt time, the LTA will be adjusted (seeded) to be equal to the current counts. This clears the existing delta (zero difference between counts and LTA), thus clearing active states. This may also automatically trigger a re-ATI to recalibrate the sensor. The sensor LTA is now adjusted to the new environmental condition and can register capacitive changes of the same magnitude (as a relative deviation in counts) as before the object was introduced to the sensor's operating environment. The sensor is then calibrated to the new environmental condition and can register proximities to the same level of accuracy (sensitivity) as before.

An example of such a persistent prox and touch state with a filter halt timeout is shown in Figure 5.5 below.

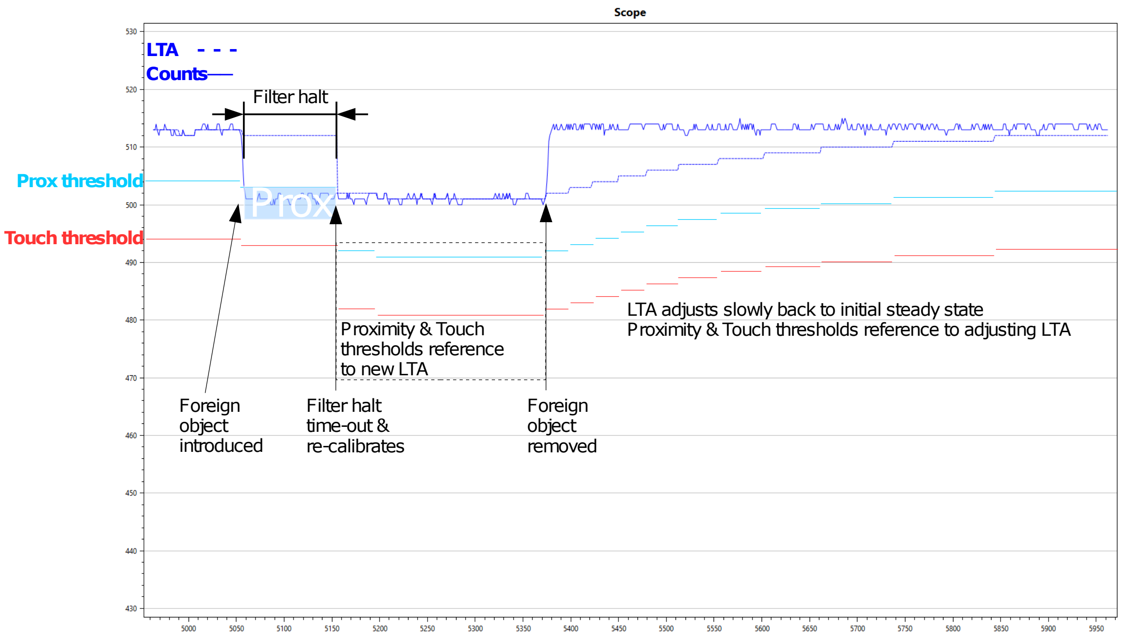

Figure 5.5: LTA Filter Halt for Proximity and Touch Detection

The LTA is also capable of readjusting upwards (away from the proximity threshold) as soon as the object or load causing the interference is removed. Removing the load causes a drop in capacitance, resulting in an increase in counts. Because user interaction causes a decrease in counts, an *increase* in counts above the LTA is unexpected behaviour. To account for this, the LTA filter is typically configured to follow the counts more quickly when the counts are above the LTA and slowly (normally) when the counts are below the LTA.

Under normal circumstances, the customer should not adjust LTA values. Normally, LTA registers are read-only and values are seeded after ATI completion or upon an explicit reseed command. Special cases can allow for writing custom LTA values.

6 Linearised Counts

"Counts" refers to the number of charge transfer cycles required to charge up the internal sampling capacitor C_S. As the capacitance of the capacitive-sensing electrode increases, fewer charge transfer cycles are required to charge the capacitor. Thus, as a user approaches a capacitive sensor, the measured capacitance increases, and the counts decrease. Counts are thus *inversely* proportional to the measured signal. This leads to non-linear relationships, which is not ideal for systems that require a linear output.

6.1 Counts Linearisation

Linearising the counts simply involves taking the inverse of the measured "raw counts".

$$\text{Signal} \propto \text{Linearised Counts} \propto \frac{1}{\text{Raw Counts}}$$

To avoid floating-point operations, Azoteq® devices use a large numerator in the division above. The exact numerator value used is dependent on the specific product. The relevant datasheet should be consulted.

Some devices offer counts linearisation as an optional feature that can be enabled or disabled as required. These devices typically use the following linearisation calculation:

$$\text{Linearised Counts} = \frac{\text{Target}^2}{\text{Raw Counts}} \tag{1}$$

This formula allows the counts to remain in the same range when enabling linearisation. With no activation, the raw counts and linearised counts will both be approximately equal to the ATI target.

Other Azoteq® devices perform linearisation as a fundamental step in the sensor processing. These devices make use of the following formula:

$$\text{Linearised Counts} = \frac{3276750}{\text{Raw Counts}} \tag{2}$$

This large numerator value maximises the resolution of the sensor. However, using linearised counts in this way tends to be less intuitive when it comes to ATI Base and Target selection.

6.2 Working With Base and Target

Devices that offer optional linearisation as shown in Equation (1) use Base and Target as normal, where the Base and Target are specified in terms of the "raw counts". For example, a typical capacitive sensor may use 100 counts for the Base and 500 counts for the Target.

Devices that linearise the counts as a core feature, as in Equation (2), specify the Base and Target in terms of linearised counts. For the previous example, the Base value would be chosen as $\frac{3276750}{100} = 32767$, and the Target would be chosen as $\frac{3276750}{500} = 6553$.

6.3 Evaluating Sensitivity With Linearised Counts

Sensitivity of a channel is typically performed by evaluating the size of the Delta during a touch or activation. With linearised counts, evaluating the size of the delta is less intuitive, as the delta is not representative of the change in raw counts. To calculate the delta in raw counts

$$\text{Delta Counts} = \frac{3276750}{\text{Linearised LTA}} - \frac{3276750}{\text{Linearised Counts}} \tag{3}$$

For example, consider a case with a delta of 500 linearised counts on an LTA of 10000. The delta in "raw counts" is

$$\begin{aligned}\text{Delta Counts} &= \frac{3276750}{10000} - \frac{3276750}{10500}\\ &= 327.67 - 312.07\\ &= 15.6\,\text{counts}\end{aligned}$$

7 Debounce and Hysteresis

7.1 Introduction

Debounce and Hysteresis are techniques used to limit false or overactive touch and release signals.

7.2 Debounce Operation

Debouncing is used when the counts initially crosses the proximity threshold. It forces the sensor to perform a number of measurements at a high rate, checking that the counts of each measurement exceeds the threshold. Once the device reads that a certain number of consecutive samples have all exceeded the threshold, the proximity event is set.

Debounce limits the information sent to a master device and prevents false triggers, as in Figure 7.1. It avoids unnecessary events that may occur due to noise or jitter on the sensor.

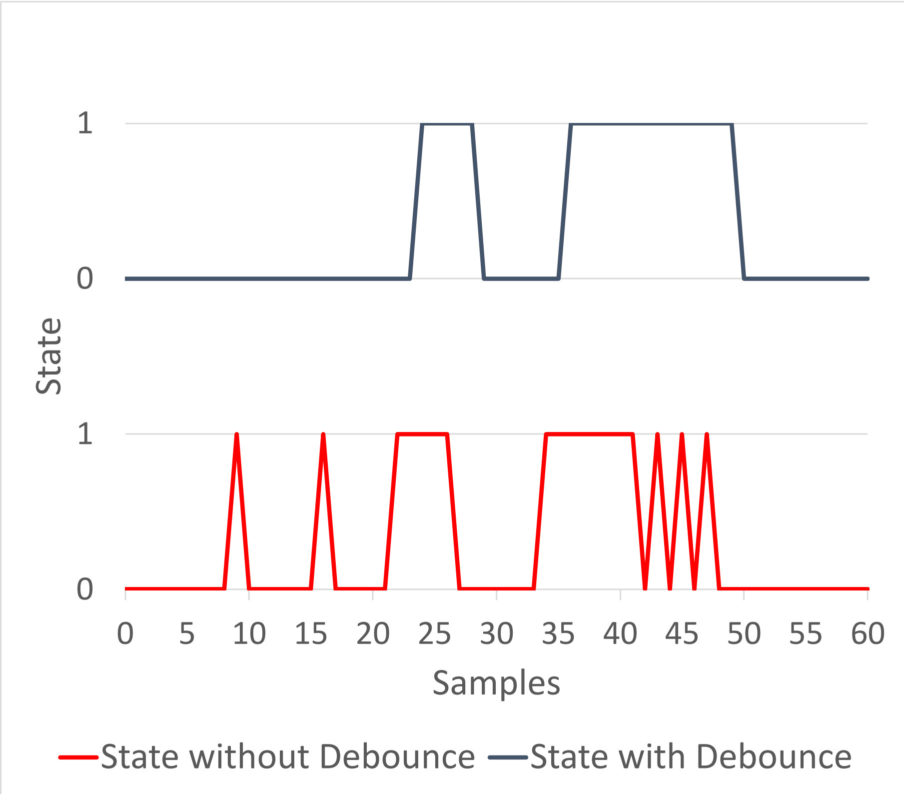

Figure 7.1: Debounce example

7.3 Hysteresis Operation

Hysteresis is applied to the touch threshold. Hysteresis allows a sensor to use different enter and exit thresholds for touch detection, with the exit threshold being lower than the enter threshold. This avoids jitter or bouncing of the touch event as the counts cross the touch threshold. The different enter and exit levels are illustrated in Figure 7.2.

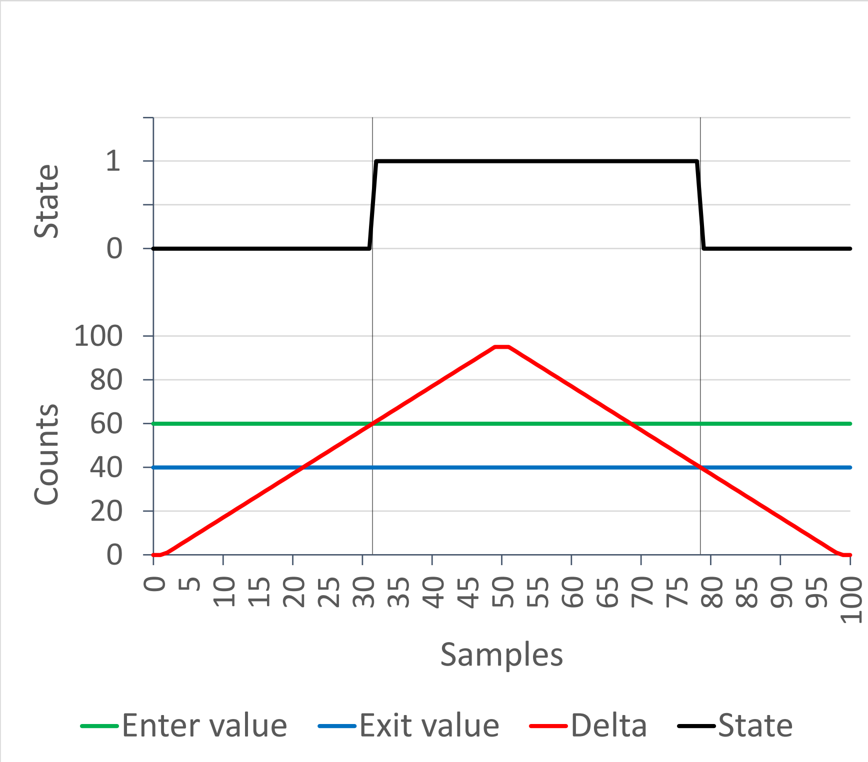

Figure 7.2: Hysteresis example

The red line above is a simple example of the delta measured on a sensor. The enter level (indicated by the green line) is the touch threshold of the sensor. As soon as the delta exceeds this entry threshold, the touch state is set. If the signal drops below the enter level, it does not yet set the state low. Only after the signal drops below the exit level (blue line) does the state go low.

It is recommended to add hysteresis to touch levels whenever possible to ensure a reliable and clean output.

8 Power Modes

Azoteq devices have a configurable report rate that determines how often the device should sample. For example, a device may be configured to sample every 10 milliseconds, or at 100 Hz. Therefore, once every 10 ms, the device will wake up and perform a measurement on each of its sensors. The device may then report the measurement results over I^2C. Once measurement and communication is complete, the device will go into a sleep mode for the remainder of the 10 ms, until the next measurement is scheduled. This process is shown in Figure 8.1.

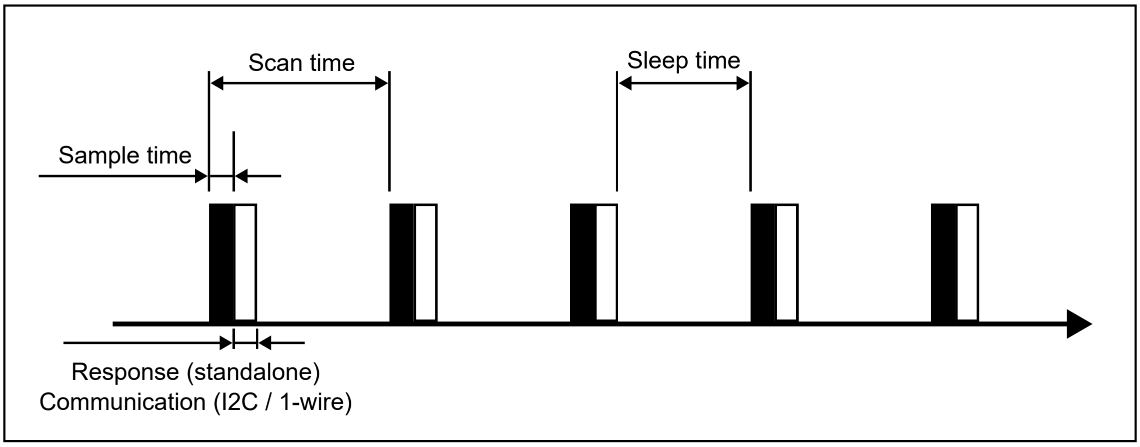

Figure 8.1: Sample, scan, sleep, and communication time diagram

Sampling and communication is often a power-intensive process. However, the average current consumption of the device can be lowered by increasing the sleep time between samples. This can be done by increasing the report period (to 50 ms for example), thereby reducing the overall sampling frequency of the device. This is especially useful in applications where low current consumption is required, such as in battery-powered devices.

However, increasing the report period may lead to sluggish behaviour, where the device takes longer to respond to user inputs. There is a trade-off between current consumption and system responsiveness.

Azoteq devices feature multiple power modes, where each power mode may be configured with a different report period.

- **Normal power** (NP) mode typically has a high report rate to perform samples in quick succession, ensuring the system is suitably responsive to user inputs. Typical report periods may be between 10 ms and 50 ms. This is often the power mode with the highest current consumption.
- **Low power** (LP) mode is configured with a slower report rate, with a typical period between 50 ms and 200 ms. This slower sampling rate can significantly reduce the current consumption.
- **Ultra-low power** (ULP) mode is configured with the slowest report rate, often between 100 ms and 500 ms. ULP can often also have additional power-saving features in order to achieve the lowest current consumption possible. For example, on multi-channel devices, ULP will only sample on a single channel, keeping all other channels dormant. This feature is sometimes referred to as “AutoProx”.

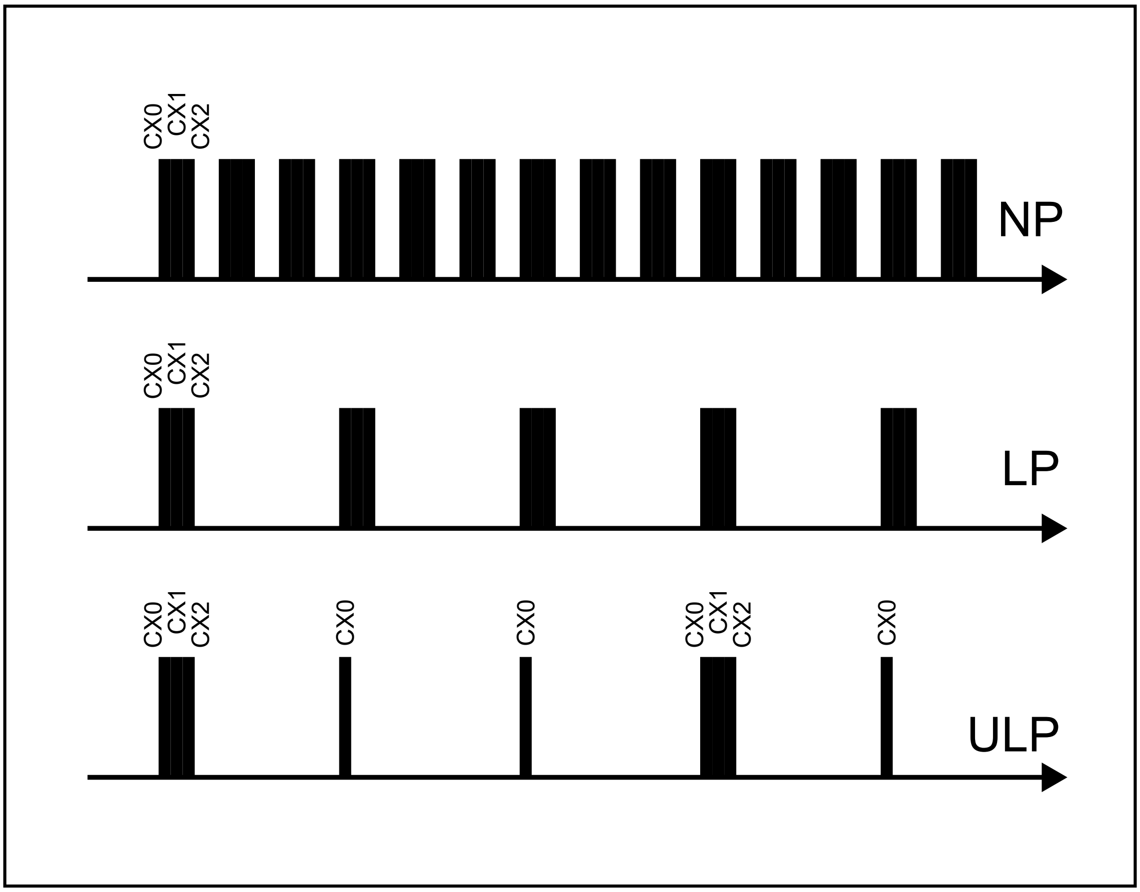

Figure 8.2: Power Modes Comparison

Some devices may include additional application-specific power modes. For example, Hall-rotation and encoder sensors include a “High-Accuracy” mode, which is often configured to sample as fast as possible to avoid aliasing.

8.1 Automatic Power Mode Switching

A compromise between low current consumption and fast response rate can be achieved by intelligently switching between different power modes. Current consumption is decreased during long periods of no user interaction by transitioning to a lower power mode. When a sensor detects some user input, the device can automatically transition to a higher power mode for better responsiveness. This functionality is present on most Azoteq devices.

9 Parasitic Capacitance

In any capacitive application, there exists some parasitic capacitance C_p to ground. This encompasses stray capacitance that exists between the electrode and PCB ground pours, PCB traces, other components, and other nearby conductive materials. These parasitics are illustrated in Figure 9.1.

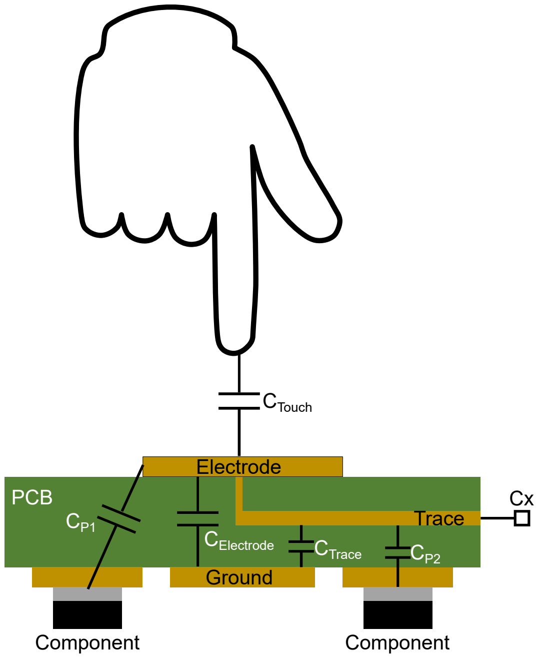

Figure 9.1: Parasitic Capacitance example

In the above image, $C_{Electrode}$, C_{Trace}, C_{P1}, and C_{P2} represent the various sources of parasitic capacitance and can be approximated as parallel plate capacitances. Parasitic capacitance tends to decrease the overall sensitivity of the capacitive sensor, as Azoteq devices try to detect minute changes in this capacitance caused by user interaction. Larger parasitic capacitances result in a smaller relative capacitance change when user interaction (C_{touch}) is introduced.

Consider two self-capacitive sensing setups, where one has a high parasitic capacitance of 100 pF, and the other has a lower parasitic capacitance of only 10 pF. A user interacts with the sensor, causing a change in the measured capacitance of C_{touch} = 1 pF.

- In the first design, the measured capacitance increases from 100 pF to 101 pF. There is a 1% increase in capacitance.
- In the second design, the measured capacitance increases from 10 pF to 11 pF. There is a 10% increase in capacitance.

The second design will be much more sensitive, as there is a much greater relative increase in the measured capacitance for the same user interaction. The on-chip compensation circuits can aid in reducing the effect of the parasitic capacitance and boosting the sensitivity. However, it is good practice to consider and reduce the parasitic capacitance of the electrode as much as possible when creating a new design. Guidance for PCB routing, overlay materials, and PCB substrate selection can be found in AZD125.

10 Capacitive Component Choices

It is important to carefully consider the minimum capacity constraints of µF-range reservoir capacitors on power supply inputs (V_{DD}) and regulator outputs (V_{REG}). These will be sensitive to derating and tolerance. In addition, inductive sensing circuits using LC resonators will be specifically sensitive to component tolerances.

10.1 Derating

The derating of critical power supply capacitors (in the µF range) should be taken into account, especially with ceramic surface-mount capacitors. A ceramic capacitor's capacitance value will decrease as the voltage across it approaches its rated voltage. This effect is also a function of the physical size of the capacitor, as smaller capacitors are more susceptible to this reduction in capacitance. The typical derating of various surface-mount capacitor sizes is shown in Figure 10.1.

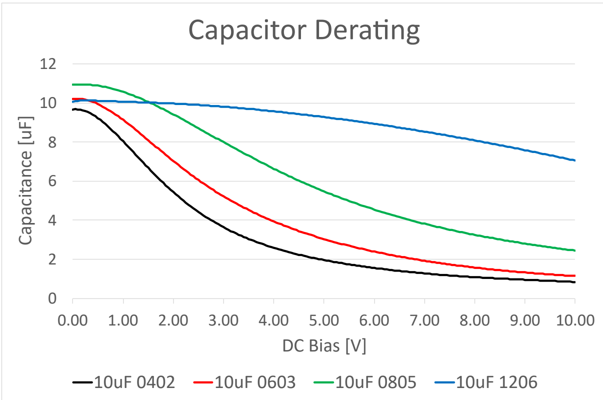

Figure 10.1: Capacitance Variation vs. DC Voltage for Selected 10 µF Capacitors

This figure shows that the 4 capacitors, which should be 10 µF each, behave differently when biased at a certain voltage. Azoteq ICs typically operate at 3.3 V, and the effects of derating should therefore be considered at 3.3 V in the graph above. For the 1206 component size, the 10 µF capacitor is still very close to the original value and should operate sufficiently. The 0805 capacitance changes to just below 8 µF, more than 20% lower than the desired value. The 0603 component drops to just below 6 µF, which is 40% lower than desired. Finally, the 0402 component size maintains a capacitance less than 4 µF, significantly less than the desired value.

This shows that a 10 µF 1206 capacitor operated at 3.3 V will still provide the required 10 µF capacitance. 0402 components may be considered to conserve board space; however, the 10 µF capacitor will only provide 4 µF capacitance, which may be insufficient.

Capacitors should be chosen with derating and datasheet limits in mind. For Azoteq capacitive sensing devices, this mainly applies to V_{REG} and V_{DD} requirements.

11 EMC

11.1 Introduction

In many sensing applications, the sensing electrode may be placed far from the IC, resulting in a long PCB trace or wire connecting the electrode to the IC. Noise may couple into the system via these long connections. Due to the extremely sensitive nature of Azoteq's sensing approach, this noise can influence the sensor and give false detections, resulting in faulty behaviour.

11.2 Relevant Standards

There are many EMC standards, and the most popular EMC standards will be introduced below.

EMC mainly has two concerns:

- Does the device cause other devices to malfunction? (Does it radiate too much EMC?)
- Does the device malfunction due to a lack of protection against received electromagnetic energy? (Is it immune to predefined levels of EMC?)

For the first question, emissions are measured to determine the energy radiated into the space around the device. Capacitive sensing emissions are typically in the nW to µW range, and thus compliance should not be a problem. For more details on this, please see AZD085 on the Azoteq website.

For the second question, much more care needs to be taken. Capacitive sensing circuits are very sensitive. Interfering sources can couple into these circuits, and thus thorough testing is advised. The main focus should be on noise sources that result in capacitive currents flowing to earth, thus focusing on common mode EMC tests.

Typical noise sources can be:

- lightning
- power supply voltage fluctuations
- 50 Hz or 60 Hz magnetic fields
- arcing due to breaks in inductive circuits
- radio transmitters
- electrostatic discharges
- switch mode power supplies

For more details on this, please see AZD085 on the Azoteq website.

12 I²C Communication

Azoteq® ICs support a standard two-wire I²C interface, complemented with a RDY (ready interrupt) line that indicates the availability of new data or important events. Azoteq® devices only function as a slave device on the I²C bus. The bus is controlled by a master device that generates the serial clock (SCL), controls bus access, and generates the START and STOP conditions. The serial clock and serial data (SDA) lines are open-drain and therefore must be pulled high to the operating voltage with a pull-up resistor (4.7 kΩ recommended).

12.1 I²C Address Range and Selection

The I²C addresses of Azoteq® devices fall within the range of 0x30 to 0x80. Certain Azoteq devices' I²C address can be changed at startup by pulling certain pins either high or low. Other devices may have unique order codes to provide alternative I²C addresses.

12.2 Clock Stretching

Clock stretching occurs when the slave device holds the SCL line low for some time between bytes, thereby pausing the I²C transaction. The transaction will only continue once the slave has released the SCL line, allowing the line to go high again.

When a slave device needs to receive or transmit bytes at a fast rate, it may need more time to store or prepare a byte of data to be transmitted back to the master. In these cases, the slave can hold the SCL line low after receiving and acknowledging a byte to force the master into a wait state until the slave is prepared to transfer the next byte. Figure 12.1 shows an example of a slave holding the SCL line low before the data is ready for transmission.

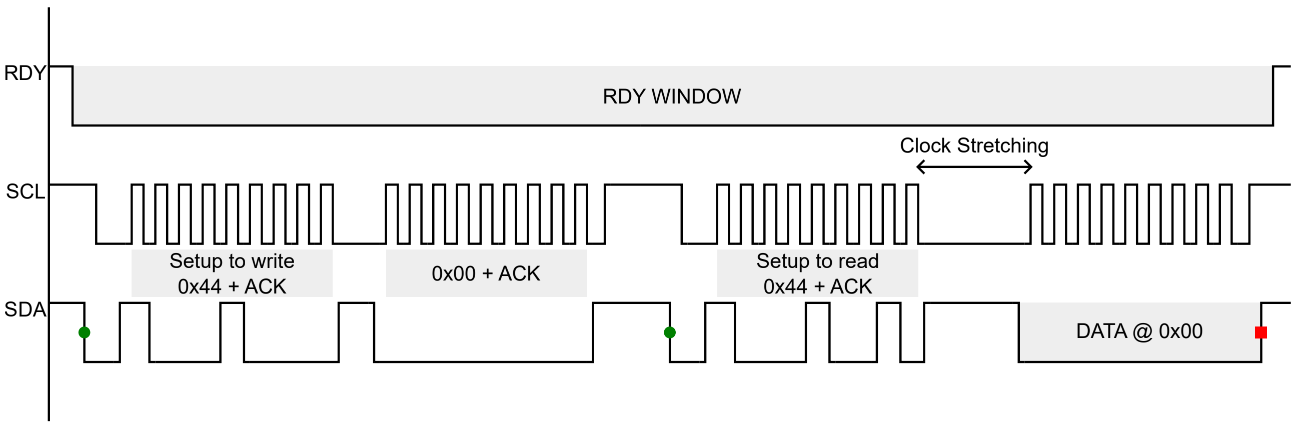

Figure 12.1: I²C Clock-Stretching

Due to this clock-stretching feature, Azoteq® devices are not compatible with an I³C bus.

12.3 RDY/IRQ

As mentioned previously, Azoteq® ICs support a standard two-wire I²C interface with the addition of a RDY (ready interrupt) line. The communication has an open-drain active low RDY signal to inform the master that new data is available and that it has opened a communication window. The RDY pin stays low for the duration of the I²C communication, until an I²C stop condition is detected. The RDY pin then goes high, and the master must wait for the next communication window in order to do comms again.

Some Azoteq devices allow communication at any time. However, it is recommended to perform I²C communication only while the ready pin is low, when the communication window is open. This reduces the total time spent performing I²C comms, reducing the current consumption of the system. It is optimal for the master to use this RDY pin as an interrupt input and obtain the data accordingly. It is also useful to allow the master MCU to enter low power or sleep mode, allowing wake-up from the touch device when user interaction is detected.

An example of an I²C read sequence can be seen in Figure 12.2. When the RDY signal goes low, a communication window is opened where data can be read. The communication window will remain open until an I²C stop is sent.

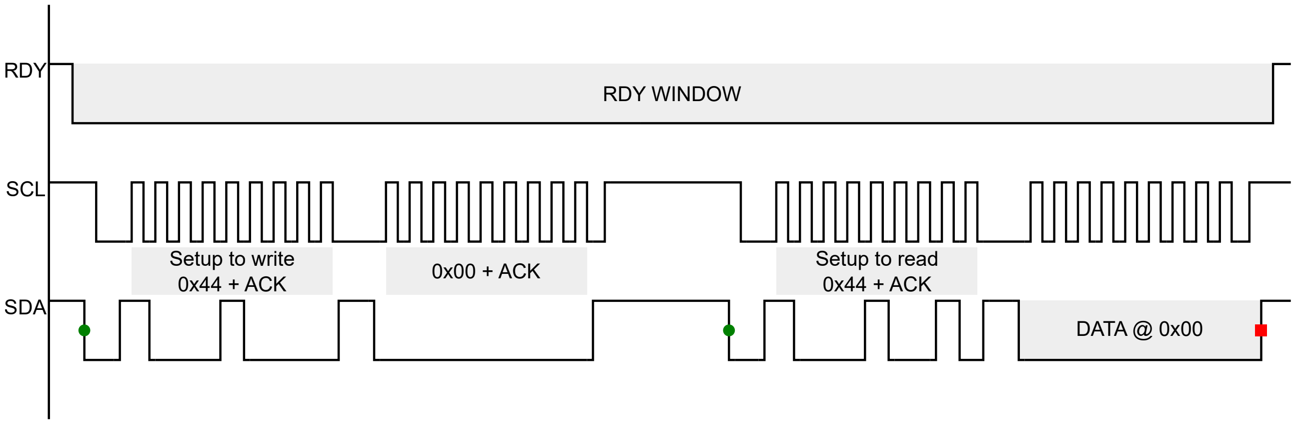

Figure 12.2: I²C Read Sequence

If the master attempts to read from the device outside of a communication window (i.e., while RDY is high), the device may respond with an invalid communication response (0xEE). The invalid communication response will also be sent by the device when attempting to read from a memory map register that does not exist.

If the communication window is not serviced within the I²C timeout period (in milliseconds), the session is ended (RDY goes HIGH), and processing continues as normal. This allows the system to continue and keep reference values up to date even if the master is not responsive. However, the corresponding event data and immediate values were missed, and this should therefore be avoided.

12.4 Streaming and Event Mode

I²C streaming mode refers to the continuous reporting of data at the relevant power mode report rate, as seen in Figure 12.3. The device will open a communication window after every measurement cycle, allowing the master to read the newest measurement data. This mode is useful for debugging or logging purposes, such as when configuring the device through Azoteq's graphical user interface (GUI). This gives real-time feedback to the user while evaluating the performance of a particular device, design, or configuration. However, the high amount of time spent doing communication has a severe impact on current consumption.

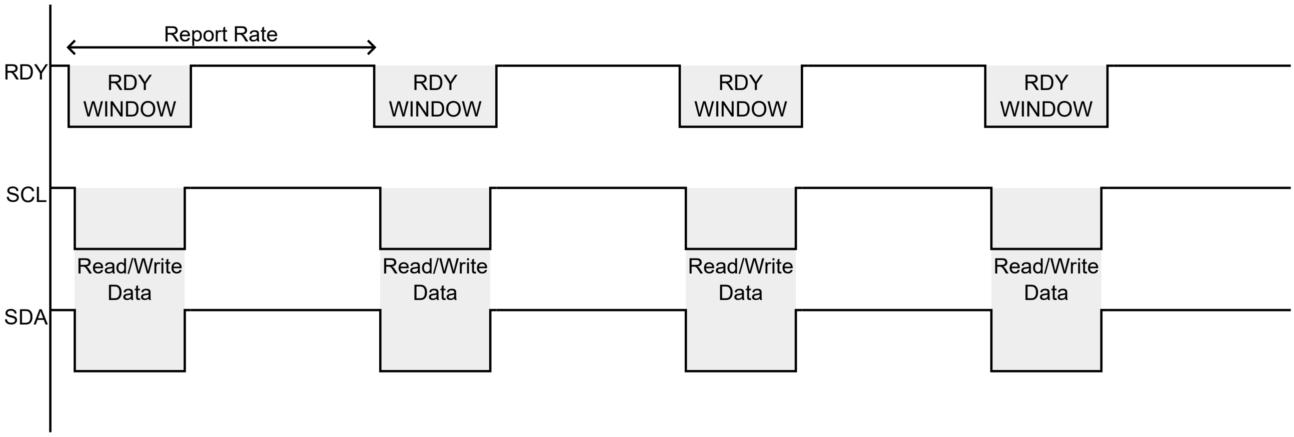

Figure 12.3: I²C Streaming Sequence

Event mode allows the device to bypass the communication window when no activity is sensed. Event mode is usually enabled since the master does not want to be interrupted unnecessarily during every cycle if no activity occurrs. The communication will resume (RDY will indicate available data) if an enabled event occurs, as seen in Figure 12.4.

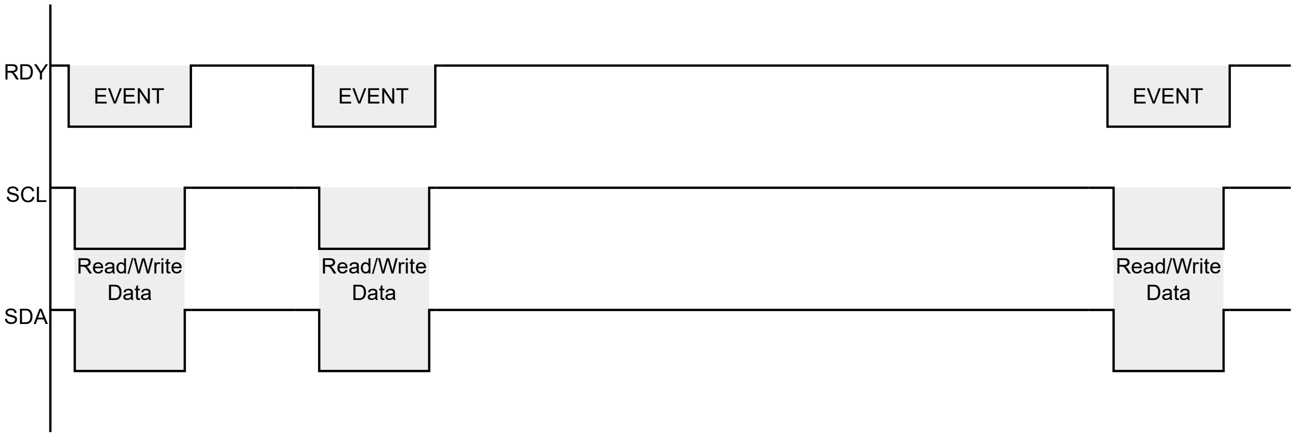

Figure 12.4: I²C Event Sequence

12.5 Terminate Communication

Once the RDY pin has indicated that a communication window is available, the window remains open until an I^{2}C stop condition is detected. The RDY pin then goes high, and the master must wait for the next communication window in order to do comms again. In order to perform multiple I^{2}C transactions, the master should use a repeated start condition between transactions to keep the window open. The final transaction should be terminated with an I^{2}C stop condition. This is shown in Figure 12.5.

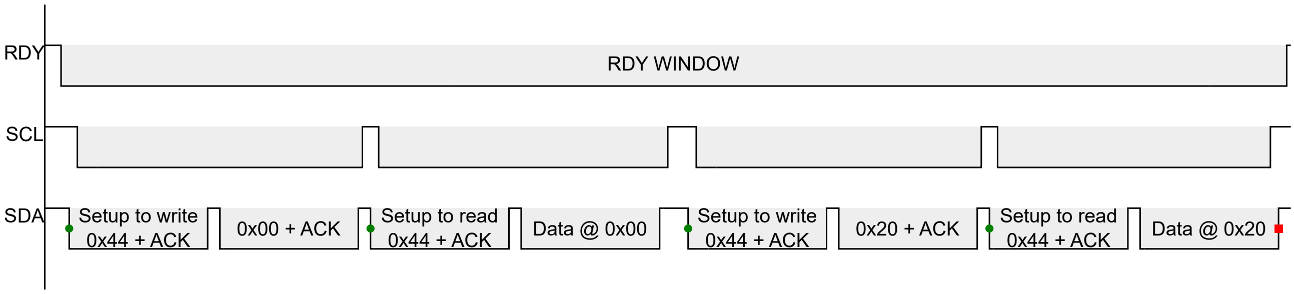

Figure 12.5: I²C Repeated Start Read Sequence

In some cases, the master controller cannot produce a repeated start condition. Therefore, communication will have to occur in multiple communication windows. This can become an issue when changing many settings or when attempting to read a large amount of data from the device.

The I²C stop bit check can be disabled on certain Azoteq® devices, which will prevent an I²C stop condition from closing the communication window. If multi-register transmission is required in only one communication window, the stop bit check can be disabled at the start of a communication window and then re-enabled at the end of the window. This is shown in Figure 12.6.

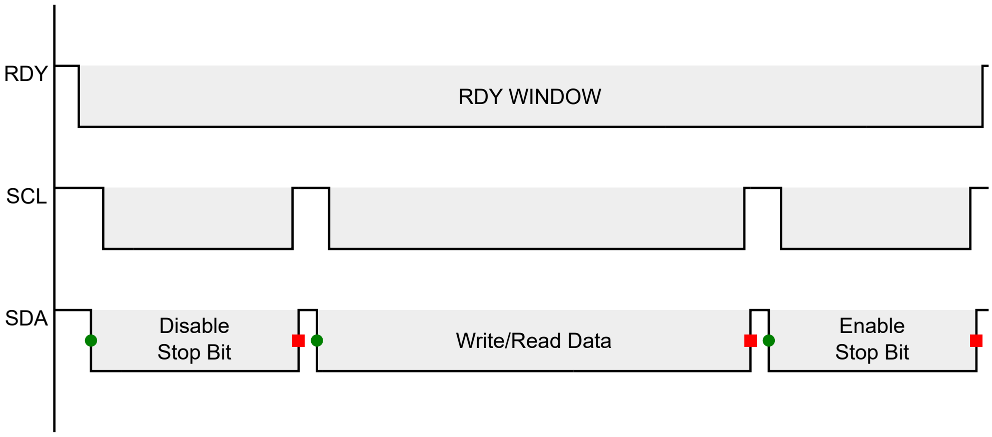

Figure 12.6: I²C Stop Bit Disable/Enable Sequence

Alternatively, with the stop bit check disabled, many Azoteq devices allow an end communication command (usually 0xFF) to be sent to close the window. An example of this is seen in Figure 12.7.

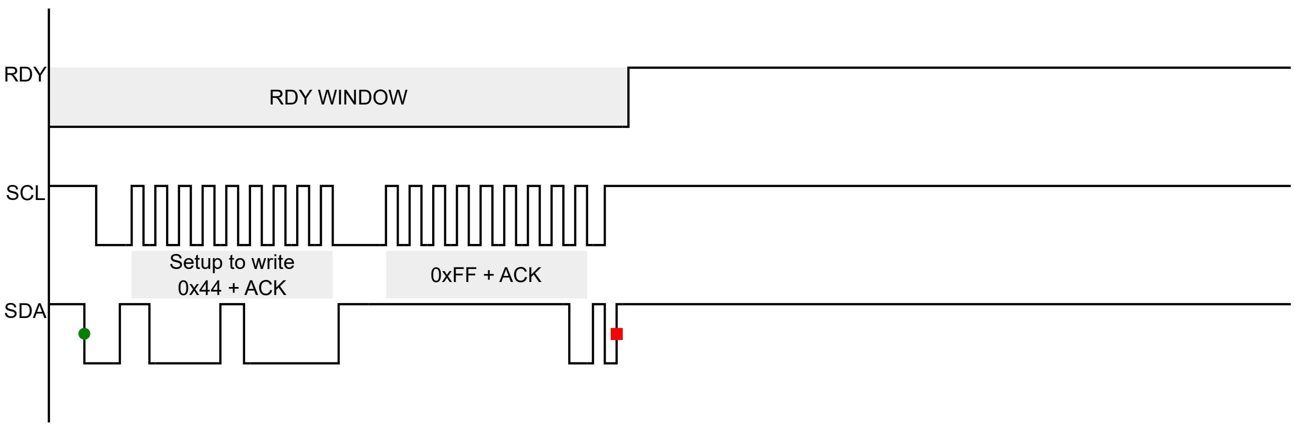

Figure 12.7: I²C Force Stop Communication Sequence

It should be noted that an end communication command will have to be sent in the same communication window that the stop bit is disabled.

12.6 Force Communication

It is best to initiate communication with the device only during a communication window. However, a communication request will force a communication window to open. In event mode, communication windows are only provided when an event is reported, and a communication window must be requested to write or read settings outside of this window. A force communication command is initiated by writing 0xFF to the device. An example of this is seen in Figure 12.8 below.

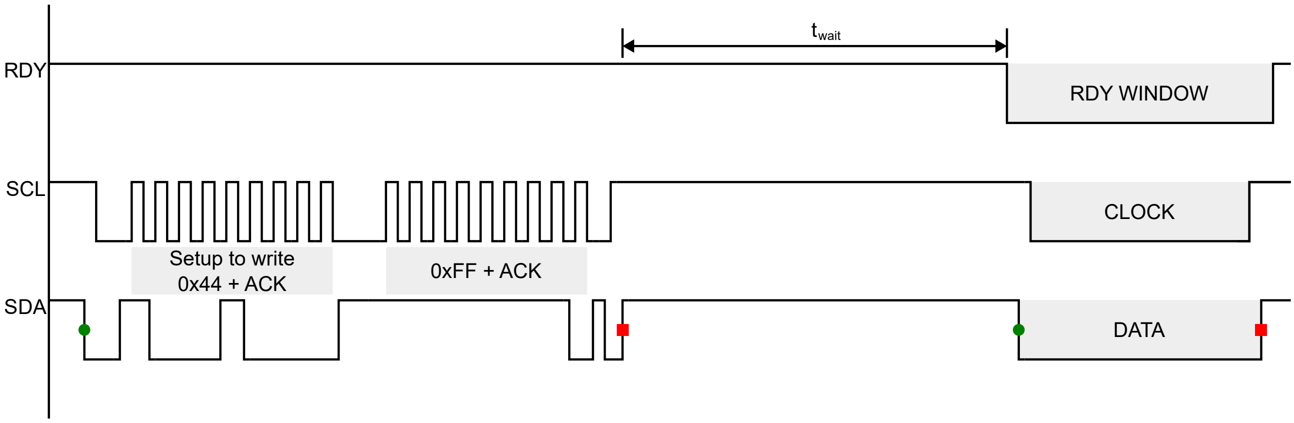

Figure 12.8: Force Communication Sequence

In the above example, 0xFF is written to device address 0x44 to force a communication window. After a certain amount of time (t_{wait}), RDY will go low to indicate that a communication window has opened. The t_{wait} is product-specific and is usually dependent on the sampling and processing duration of the device.

A force communication request should be avoided while RDY is in the active state. If a communication request is sent at the exact moment when an event causes RDY to go low, the window will close again after sending the I^{2}C STOP signal. In such a scenario, the device will provide an invalid communication response (0xEE) because the master is attempting to read from the device outside of a communication window (i.e., while RDY is high). To prevent this issue, it is recommended to read the product number during each ready window to ensure that the response received is valid.

12.7 Program Flow

Although it is recommended to run the device in event mode to avoid interrupting the master MCU, the basic program flow for device communication, as shown in Figure 12.9, still applies for I^{2}C communication.

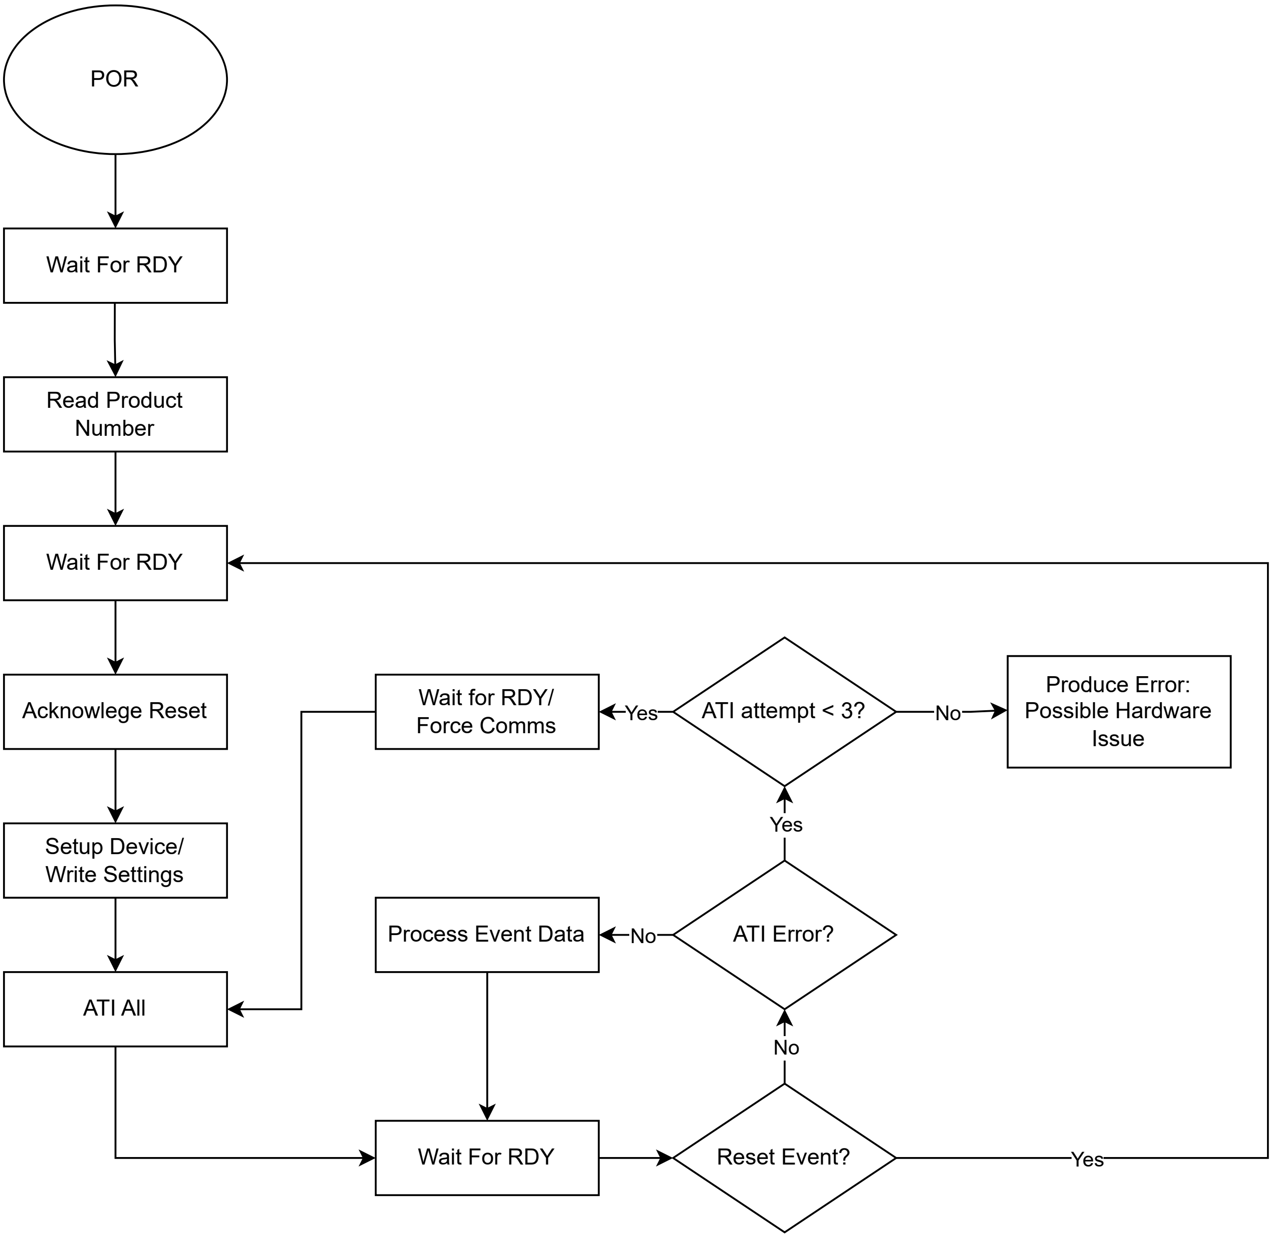

Figure 12.9: Program Flow Diagram

Once the device has been powered:

1. Wait for the first ready window.
2. Read the product number to ensure that the correct device is being used.
3. Wait for the next available ready window.
4. Acknowledge the device reset.
5. Write all the necessary settings to the device.
6. Initiate ATI All command.
7. Wait for a ready window caused by an event and read the status flags.
8. If a reset event occurred, then repeat step 3 to 8.
9. If an ATI error occurred, then check if number of ReATI attempts are below 3.
10. If the number of ReATI attempts are below 3, then force a communication window and ATI all.
11. If the number of ReATI attempts are 3 or more, then produce a hardware error.
12. If no ATI error occurred, then process the event data. Repeat from step 7.

The acknowledge reset and ATI all commands can be applied when writing the device settings. An ATI will be initiated once the window has been closed.

To avoid constant errors due to interaction with the sensor during an ATI, it is recommended to add a delay between ReATI attempts.

Contact Information

	South Africa (Headquarters)	China
Physical Address	1 Bergsig Avenue Paarl 7646 South Africa	Room 501A, Block A T-Share International Centre Taoyuan Road, Nanshan District Shenzhen, Guangdong, PRC
Tel	+27 21 863 0033	+86 755 8303 5294 ext 808
Email	info@azoteq.com	info@azoteq.com

	USA	Taiwan
Physical Address	7000 North Mopac Expressway Suite 200 Austin TX 78731 USA	Xintai 5th Road, Sec. 1 No. 99, 9F-12C Xizhi District 221001 New Taipei City Taiwan
Tel	+1 512 538 1995	+886 932 219 444
Email	info@azoteq.com	info@azoteq.com

Visit www.azoteq.com
for a list of distributors and worldwide representation.