

Introduction

The Biosignals PRO architecture supports multiple SmartSensors connected to a single HUB while preserving a common temporal reference across all acquired channels. The SmartSensors are not synchronised by distributing one physical ADC sampling clock to every sensor. Instead, each SmartSensor records the timing of its own acquisition events and the HUB maps those events onto a single global time base with microsecond resolution.

This technical note describes the SmartSensor-to-HUB timestamp reconciliation mechanism, the measured timing uncertainty, the conservative theoretical relative timing bound between two independent SmartSensors, and the expected behaviour across continuous acquisitions and separate acquisition sessions.

Synchronisation Architecture

1. Acquisition timing inside each SmartSensor

Each SmartSensor digitises its signal using either a specialised analogue front-end (AFE) or an ADC combined with a discrete analogue front-end. When the AFE or ADC generates a new-sample interrupt, the SmartSensor records the timing of that acquisition event using its internal high-resolution timer.

The timing information is therefore associated with the actual data-acquisition event and not with the later transfer of the samples to the HUB. This separation is important because communication and buffering latency do not determine the timestamp assigned to the acquired signal.

2. Packet formation and HUB notification

Samples are buffered locally inside the SmartSensor. A transfer is initiated when either the configured packet length is reached or the configured timeout expires. The timeout mechanism limits latency for sensors operating at lower sampling frequencies. A packet can contain up to 250 samples.

When a packet is ready, the SmartSensor raises its dedicated interrupt line to notify the HUB. In the same local time domain, the SmartSensor knows both the interrupt time and the acquisition time of the first sample in the packet. It therefore includes in the data frame a signed time offset, in microseconds, between these two events.

3. HUB global timestamp and time reconciliation

The SmartSensor interrupt lines are monitored by the HUB through an IO expander. Interrupt processing has very high priority. When an interrupt is recognised, the HUB records its own timestamp from a global time base that is common to all ports and provides resolution down to the microsecond.

The HUB may retrieve the sensor data afterwards. The transfer delay does not change the acquisition timing because the packet carries the relative offset needed to place the first sample on the HUB timeline.

$$T_{\text{first_sample_global}} = T_{\text{interrupt_global}} + \Delta_t$$

Here, Δ_t is the signed offset provided by the SmartSensor between the interrupt event and the first sample in the packet. In normal operation, the first sample predates the packet-ready interrupt, so the offset represents that earlier position in time.

4. Timing of the remaining samples

Once the first sample has been placed on the HUB global time axis, the remaining samples in the packet are reconstructed from the configured sampling period. For a sampling period T_s , the timing of sample n is given by:

$$T_{\text{sample}(n)} = T_{\text{first_sample_global}} + n \times T_s$$

This process is applied independently to every SmartSensor and every HUB port. As a result, signals acquired at different sampling frequencies can still be directly correlated because their samples are all expressed relative to the same global time base. Packet length and communication latency affect data delivery, but they do not define the acquisition time.

Timing Accuracy and Tolerance

5. Measured single-SmartSensor timing uncertainty

Internal timing measurements were performed to characterise the delay between the SmartSensor interrupt reaching the HUB and the corresponding timestamp being stored. The measured HUB contribution, including the I2C transactions required to read and clear the IO-expander interrupt status, is approximately 120 us to 130 us.

The SmartSensor has an inherent maximum uncertainty of 5 us between the AFE/ADC producing a sample and the local sample timestamp being stored. Combining both contributions gives an expected total range of approximately 120 us to 135 us. The 135 us value is therefore used as the conservative upper bound for one SmartSensor timing chain.

Timing summary

Parameter	Timing value
HUB interrupt-to-port timestamp delay (measured)	120-130 us
SmartSensor internal sample timestamp uncertainty	up to 5 us
Single-SmartSensor conservative upper bound	135 us
Two-SmartSensor relative timing upper bound	approximately 270 us (theoretical, conservative)

6. Relative timing between two independent SmartSensors

For two independent SmartSensors connected to the same HUB, a conservative theoretical maximum relative timing uncertainty of approximately 270 us can be used. This value is obtained by applying the 135 us conservative bound to both independent timing chains.

$$\text{Maximum theoretical relative uncertainty} = 2 \times 135 \text{ us} = 270 \text{ us}$$

The 270 us value is intentionally conservative and should not be presented as a directly measured sensor-to-sensor skew specification. It is a theoretical upper bound derived from the measured single-sensor timing behaviour.

Adding further SmartSensors to the HUB does not introduce a proportional cumulative timing dilation. The relevant worst-case condition occurs when one SmartSensor raises its interrupt immediately after the HUB has already read the IO-expander interrupt status for another event. That interrupt can then be recognised in the following interrupt-processing cycle. In practice, the relative timing error is expected to be lower than the 270 us bound, but a dedicated common-input sensor-to-sensor validation would be required before publishing a tighter formal specification.

Acquisition Sessions and Repeatability

7. Timing consistency within one continuous acquisition

Within a single continuous acquisition, all SmartSensors remain associated with the same HUB global time base and the same timing context. The timestamp reconciliation mechanism therefore provides a common time axis throughout the acquisition, even when the connected sensors operate at different sampling frequencies.

For experiments in which relative timing must remain consistent between several experimental conditions, the current recommended workflow is to keep one continuous acquisition running and separate the different periods afterwards during data analysis.

8. Separate acquisition sessions and system restarts

The relative timing phase between independent SmartSensors is not guaranteed to be preserved across a complete stop/restart of the acquisition or a system restart. The SmartSensor timing state is reinitialised, so each new acquisition establishes a new synchronisation context.

This does not affect the ability to correlate sensors within each individual acquisition. It means only that a fixed phase relationship from a previous acquisition session should not be assumed to persist after a restart.

9. Pause/resume acquisition workflow

A pause/resume acquisition mechanism is planned for Biosignals Studio. The intended behaviour is to pause data recording while keeping the HUB and SmartSensors operational, so that the internal timing state and synchronisation context are maintained when recording resumes.

This functionality is not available in the current Biosignals Studio release. It is currently planned for availability by the end of Q1 2027. Until it is released, continuous acquisition followed by offline segmentation is the recommended approach when preservation of the same timing context is important.

Application Considerations

10. Timing-sensitive applications such as pulse transit time

For applications such as pulse transit time (PTT), where physiologically relevant changes can be only a few milliseconds, the conservative 270 μ s theoretical SmartSensor-to-SmartSensor timing bound is substantially smaller than a 5 ms to 10 ms effect.

However, timestamp synchronisation is only one contribution to the total end-to-end timing uncertainty. The sampling frequency, analogue and digital filter group delay, differences between sensor signal chains, and the algorithm used to identify the fiducial point of each waveform can also influence the measured delay.

Consequently, a 5 ms to 10 ms change should not be attributed exclusively to physiology solely from the synchronisation figure. For rigorous PTT validation, the complete measurement chain should be characterised using identical SmartSensors driven from a common precisely timed input and, where relevant, repeated across independent acquisition sessions.

Conclusions

Biosignals PRO synchronises SmartSensor data by mapping local acquisition timing to a common HUB global time base rather than by distributing one physical sampling clock to every sensor. Each packet carries the relative timing information required to reconstruct the first sample on the HUB timeline, while the remaining samples are positioned according to the configured sampling period.

The measured conservative upper bound for one SmartSensor timing chain is 135 us. For two independent SmartSensors, approximately 270 us can be used as a conservative theoretical maximum relative timing uncertainty, although this value has not yet been validated as a direct sensor-to-sensor skew specification.

Within one continuous acquisition, the timing relationship remains referenced to the same HUB time base. A fixed relative phase is not guaranteed across complete acquisition restarts. The current Biosignals Studio workflow for timing-critical experiments is therefore a continuous acquisition with offline segmentation; pause/resume operation that preserves the timing context is planned for the end of Q1 2027.

END OF DOCUMENT