Beyond Sensor Networks: Crystal-Free Wireless Audio Streaming

Paul de La Sayette
Inria
Paris, France
paul.de-la-sayette@inria.fr

Prabal Dutta
UC Berkeley
Berkeley, United States
prabal@berkeley.edu

Filip Maksimovic
Inria
Paris, France
filip.maksimovic@inria.fr

Abstract

Crystal-free radios integrate an entire wireless radio onto a single die, eliminating off-chip frequency references to reduce size, weight, power, and cost. Prior work has demonstrated crystal-free communication in low-rate wireless sensor networks, but higher-throughput consumer applications have received far less attention. To explore the viability of such applications, this paper designs, implements, and tests a crystal-free wireless audio system: a Single Chip Micro Mote (SCμM) transmitter broadcasts 8 kHz stereo audio over an IEEE 802.15.4 link to crystal-free SCμM receivers, which recover a stable audio signal and drive a DAC over bit-banged I^2S—with *no* crystal on either end of the link. A transmitter and two receivers discover each other via a frequency-sweep rendezvous, per-packet housekeeping keeps the link locked despite oscillator drift, and a packet-loss concealment scheme smooths over dropped packets. The broadcast link architecture supports two simultaneous receivers, as in a pair of earbuds. On ITU-T P.501 speech the system achieves a mean ViSQOL score of 3.75 at the analog output.

CCS Concepts

• **Hardware** → **Wireless devices**; **Sound-based input / output**; • **Networks** → **Wireless personal area networks**; • **Computer systems organization** → **Embedded systems**.

Keywords

Crystal-free radio, SCμM, 802.15.4, audio streaming, low-power wireless, IoT

1 Introduction

As demand for lower power and physically smaller wireless electronics grows, the prospect of monolithically integrated chip-scale radios becomes more appealing. The removal of off-chip passive components is an effective way of achieving size, weight, and power goals. This has the additional benefit of reducing the cost of these wireless components while maintaining functionality and adherence to wireless standards. However, removing off-chip components can have negative impacts on performance. On-chip passives for matching have high on-resistance and therefore a lower quality factor than their off-chip counterparts, resulting in inferior filtering, higher loss and noise. On-chip clocks and frequency synthesizers also have lower performance, resulting in higher jitter, phase noise, and higher susceptibility to environmental conditions including variation in supply voltage and temperature. However, significant reduction in power, especially from the removal of higher frequency off-chip crystal oscillators, can have a significant impact on device battery life in wireless sensor networks and personal area network applications.

Recent innovations in hardware have demonstrated that crystal-free radios are capable of communicating with off-the-shelf hardware using low-power wireless area network protocols such as Bluetooth Low Energy (BLE) and IEEE 802.15.4 [2, 6]. Physical layer modifications make the radio resilient to environmental variation by using both packet timing and carrier offset tracking [25]. The co-design of customized protocols, especially at the MAC layer [4] and minor modifications to existing time-synchronized channel-hopping networks [3], further allows robust networking without the use of external crystals.

The prior work in crystal-free communication focuses on wireless sensor networks. These applications favor low-power consumption and low data rates. Personal area networks are also well-suited to other higher-throughput applications in consumer electronics such as wireless computer peripherals or medical electronics such as instruments for diabetics [21]. This additional application motivates our investigation of crystal-free radio for audio streaming to wireless earbuds.

A distinct benefit of this application, as opposed to the wireless sensor network applications in previous literature, is that there is constant network traffic to stream audio with maximum fidelity. Previous work [18] has shown that as the packet rate in a network is reduced, the guard time must be proportionally increased to offset the effect of accumulating jitter due to random walk. Furthermore, as crystal-free radios are susceptible to environmental variation, the increased packet rate in this network also makes the system more resilient provided that connectivity is maintained [5].

In order to test the original crystal-free radio hypothesis, the Single Chip Micro Mote (SCμM) platform was developed [19]. This chip was the first known crystal-free platform that could transmit and receive standards-compliant, IEEE 802.15.4-formatted packets with no external components besides a power source and an antenna. It is also capable of transmitting Bluetooth Low Energy (BLE) packets. The chip also includes all necessary on-chip bias generation, clock and timer generation, and a low-power Cortex M0 processor. SCμM is used as the wireless networking platform both in this work and in all of the previously cited network literature.

This paper asks whether a crystal-free radio can sustain the continuous traffic that consumer audio demands. We build the complete chain on SCμM, describe the rendezvous, housekeeping, and packet-loss-concealment mechanisms that keep the link usable despite oscillator drift, and evaluate the resulting audio quality and receiver power draw.

2 System Design and Implementation

In this section, we discuss the design process, the resulting architecture, and its implementation on SCμM, showing that crystal-free radios are viable for the continuous audio-streaming use case.

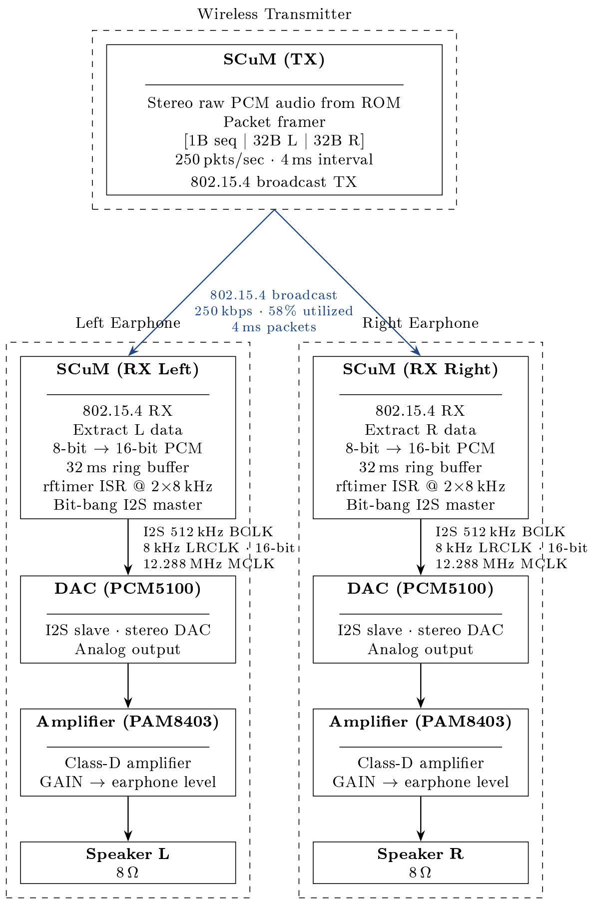

Figure 1: Complete system architecture block diagram.

2.1 Overview

We design a fully crystal-free audio chain: neither the transmitter nor the receivers use a crystal oscillator. Fig. 1 shows the hardware architecture of the system. One SCμM chip acts as a transmitter: it fetches audio data from memory and broadcasts audio packets. On each receiver, packets are received by SCμM's internal radio module. The data is extracted, checked, modified, and buffered before being converted to an analog signal at a precise sample rate and amplified to drive a speaker.

2.2 Radio Link

The wireless link leverages the 802.15.4 standard [12], adapted to the constraints of crystal-free operation and the limitations of SCμM. The link is connectionless: the transmitter broadcasts packets and receivers never transmit. While channel hopping is likely possible [4], we operate on a single frequency to minimize complexity.

2.2.1 Rendezvous Mechanism. When several SCμM chips power on, their RF oscillators can differ significantly in frequency, so a rendezvous is needed before they can communicate. The transmitter keeps its initial tuning settings fixed, while each receiver sweeps its own across the available range until it receives a packet with a correct cyclic redundancy check (CRC). The receiver then considers itself "locked": the two oscillators are close enough to sustain communication.

Figure 2: Packet structure: a 6 B MAC header, 1 B packet number, 32 B left-channel audio, 32 B right-channel audio, and a 2 B CRC.

Table 1: Candidate stereo PCM audio configurations versus the 802.15.4 link budget: The sound column is the raw audio bitrate (f_s× depth × 2 channels); the on-air column adds the 9 bytes of per-packet metadata and overhead, giving 73 bytes per packet; utilization is that rate over the 250 kbps budget. The 44.1 kHz row is the standard CD/studio rate, for reference.

Sample rate (kHz)	Depth (bit)	Sound (kbps)	On-air (kbps)	Link utilization (% of 250 kbps)
8.000	8	128.0	146.0	58.4 %
11.025	8	176.4	201.2	80.5 %
44.100	16	1411.2	1609.6	643.8 %

2.2.2 Housekeeping. Even once locked, the oscillators drift apart over time due to environmental variations. To compensate, we run a housekeeping routine after every received packet that fine-tunes the receiver's oscillator settings in real time. This routine uses intermediate frequency counting and the length of the packet to make a coarse estimate of carrier frequency offset. The oscillator is then tuned so that the offset remains fixed.

2.2.3 Payload. We settle on a packet with 64 B of audio data, since we observe that larger packets have a much higher failure rate. The packet structure is represented in Fig. 2: each packet carries both left and right audio samples together with a packet number, keeping the two channels inherently sample-aligned. The packet number lets each receiver detect lost or out-of-order packets; this is necessary for the packet-loss-concealment strategy described in Section 2.4.1.

2.2.4 Link Budget. The 2.4 GHz O-QPSK PHY offers a raw 250 kbps data rate budget. Table 1 investigates what that link budget can support for this application. We operate at an 8 kHz sampling frequency with a linear 8-bit depth per sample, at a comparable bitrate to the G.711 telephony standard [15]. We avoid saturating the channel to reduce collision risk with other devices in the same band.

2.3 Transmitter

The transmitter uses an interrupt at an 8 kHz sampling frequency to trigger sample generation. Depending on the mode, each sample is either fetched from internal memory, generated by a sine melody generator, or produced as a constant test pattern. Once enough samples are buffered, they are sent through the radio link.

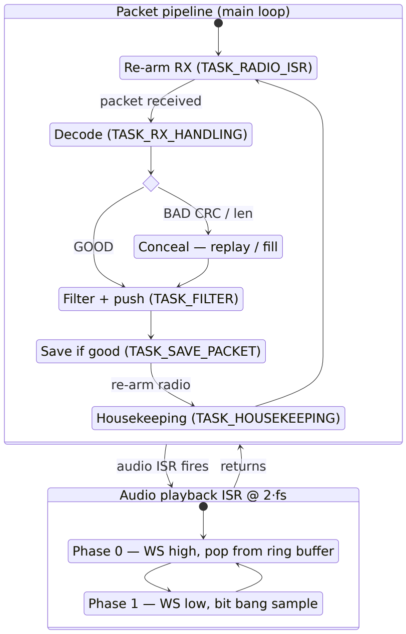

Figure 3: Finite state machine of the receiver's packet-processing loop and audio-playback interrupt.

2.4 Receiver

On the receiver, an incoming packet triggers a radio interrupt; the data is modified if necessary, and then stored in a ring buffer. Playback is handled by a separate interrupt, which sends each sample from the ring buffer to the DAC through I^2S, chosen because it is ubiquitous in simple audio systems. Fig. 3 gives an overview of the receiver's process states. We first describe how incoming packets are processed, then how samples are played back over I^2S, and finally how the two interrupts are scheduled to coexist.

2.4.1 Packet Loss Concealment (PLC). The radio link can suffer from a high packet loss rate. It is therefore paramount to implement an effective mitigation strategy for the different defects. To do so, we use a "seam" function that smooths the transition between discontinuous audio packets by applying slew-rate limiting. The packet-processing loop of Fig. 3 handles each case as follows:

- **Good packet:** send to the ring buffer with no modification and copy it to the last good packet buffer.
- **Bad CRC:** discard the data, send the last good packet to the ring buffer, and apply the seam function.
- **Good packet, but previous was incorrect or missed:** send to the ring buffer, but apply the seam function.

This strategy is loosely inspired by ITU-T G.711 Appendix 1 [16], but is significantly simplified because the G.711 algorithm would make it difficult to avoid race conditions on the buffer.

2.4.2 I^2S Overview. The Inter-IC Sound (I^2S) bus is a synchronous serial protocol for intra-system digital audio transfer. It uses a three-wire topology: a Serial Clock (SCK) for bit synchronization, a Word Select (WS) line for stereo channel multiplexing, and a Serial Data (SD) line transmitting two's complement, MSB-first audio samples. Implementations often add an optional high-frequency Master Clock (MCLK), a synchronized integer multiple of the sample rate, for oversampling on external DACs.

2.4.3 Clock Generation. To function properly with no internal clock reference, the DAC requires f_{MCLK} and f_{SCK} to be standard integer multiples of the sampling rate f_S. In our design, we use f_{MCLK}=12.288 MHz, f_{SCK}=512 kHz. To generate MCLK we use the SCμM HF_CLK oscillator divided by 2 through an internal clock divider; SCK is generated by further dividing by 24. WS is generated by a 16 kHz timer interrupt triggering at each edge. We use the same interrupt to start bit-banging the data sent to the DAC (Fig. 3). Because all these signals originate from the same HF_CLK oscillator, we know there will not be any drift between these references, which is necessary for proper DAC operation. Because of a SCμM hardware limitation, we have to use the same divider for the processor clock and MCLK. This constraint fixes the processor operating frequency, making it impossible to save energy by lowering the processor's clock frequency.

2.4.4 Digital Signal. SCμM does not include a hardware I^2S driver or DMA through its GPIO; the only way to send data on the bus is to use the processor via bit-banging.

This 16 kHz Word Select interrupt accesses the ring buffer every other time and stores the retrieved sample at a static memory address. When the interrupt does not perform this access, we send the audio data by shifting the sample onto a GPIO using a hybrid C/assembly code to ensure timing accuracy: the time between bit shifts must be exactly the same as the time between two SCK positive edges. This is only possible because SCK and the processor clock originate from the same oscillator. An alternative is to bit-bang both SCK and the digital data simultaneously. However, tests of this approach showed that the DAC does not function without a continuous SCK signal, and generating that signal through bit-banging would consume all available processor bandwidth.

2.4.5 Scheduling and State Management. To ensure that the interrupt for the audio playback is not late, we set it to the highest interrupt priority and disable the radio interrupt when we know it will not have time to finish before it has to trigger the next audio interrupt. We also only start the save operations on the ring buffer if the task has time to finish before the next audio interrupt that pulls from the ring buffer, ensuring there is no race condition. We do not formally prove that no race condition can occur, but no buffer corruption is observed in practice. Moreover, interrupt collisions are not observed, outside of the serial communication every 4 s, which we deem an acceptable tradeoff to get serial data.

2.5 Digital to Analog Conversion

From this point on, we use a standard I^2S analog audio chain, with one exception: the DAC's clock reference (MCLK) comes from the SCμM receiver rather than being generated internally by the DAC. We choose the PCM510x DAC [24] because of its wide adoption in the prototyping space and its ability to be configured entirely in hardware, with no additional digital interface.

Table 2: List of components used in the experimental setup.

Role	Part
Power supply	Otii Ace Pro
SCμM programmer	nRF52840-DK
UART bridge	SH-U09C5
SCμM development board	Sulu v2.2 [22]
DAC	PCM5100 [24]
Amplifier	PAM8403 [8]
Logic analyzer	Saleae Logic 8

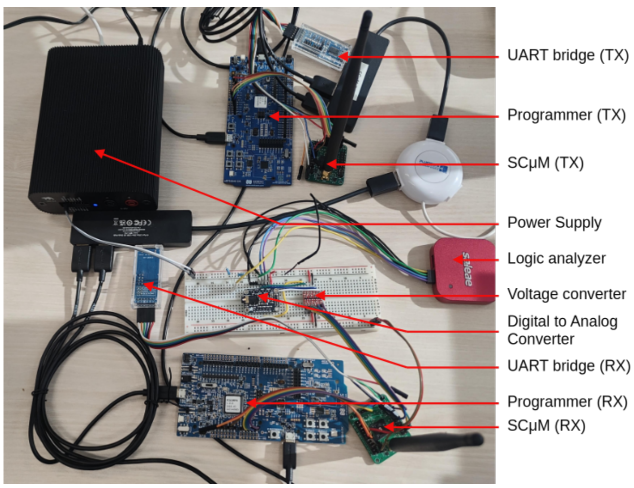

Figure 4: Experimental setup for a single-channel test: SCμM transmitter and receiver, DAC, and logic analyzer, along with nRF boards for programming and UART bridges for logging.

3 Evaluation

This section evaluates key performance criteria including clock drift, sound quality, and power draw.

3.1 Experimental Setup

Table 2 lists the hardware components used throughout the experimental setup. Fig. 4 shows the experimental setup for a single-channel experiment. The logic analyzer is running the I^{2}S/PCM analyzer extension [23]. It logs I^{2}S data, analog sound output, and GPIO flags.

The firmware is flashed onto SCμM with an nRF development kit and an open-source SDK [1]. We perform most debugging using GPIO flags measured through the logic analyzer.

3.2 Clock Drift

Fig. 5 shows the processor-clock frequency error over time for each of the SCμM chips relative to its average frequency. While the clocks stay fairly consistent over time, we measure significant differences in nominal frequency between the three chips, with deviations from −3277 ppm to +1111 ppm. This is due to a limitation of the calibration procedure of the SCμM HF_CLK: each fine calibration step is on the order of 5000 ppm. These chip-to-chip

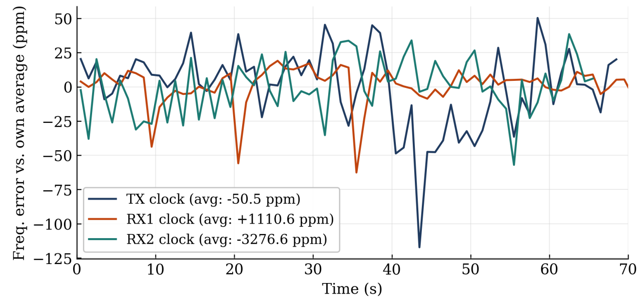

Figure 5: Processor (H_CLK) frequency versus average chip frequency, measured over time on the transmitter and both receivers. The average frequency deviation of each clock from the intended 12.288 *MHz* **is also shown. The observed drift over time is small enough to likely not affect audio quality.**

variations are significant enough to affect audio playback. The rate at which the transmitter sends data to the receiver and the rate at which the receiver sends data to the DAC differ; we therefore either drop samples or insert filler samples when the buffer is empty. Implementing a skew compensation algorithm [11] would mitigate this; we leave it as future work.

3.3 Sound Quality

3.3.1 Measurement Method. Using the setup described in Section 3.1, we capture the analog voltage at the output of the audio DAC, and record the digital data sent to the DAC over I^{2}S. In addition, the SCμM receiver tracks CRC errors along with buffer-underflow and overflow events which are sent every 4 s to the computer using UART.

We load audio files from ITU-T Recommendation P.501 [14] into the transmitter memory. Audio from Annex E, meant for technical measurements, is used for Figs. 6 and 7. Samples from Annex D, meant for speech quality measurements, are used in Table 3.

3.3.2 Audio Output. Fig. 6 shows the digital and analog audio waveforms from the receiver compared with the reference sound. Because the analog output, digital data, and reference originate from independent clocks, each pair must be time-aligned before comparison. The digital stream is aligned with the analog signal via cross-correlation over the full displayed segment (approximately 2.7 s). The reference cannot be aligned in the same way because of clock drift. We therefore restrict the reference correlation search to a short window around the start, where the offset is locally constant. Since the three signals share no common physical unit (volts, raw PCM codes, normalized samples), each is independently z-scored.

3.3.3 Quality Scores. We evaluate listening-quality mean opinion scores (MOS) for the source audio (encoded to fit the link capacity), the digital trace, and the analog trace using two objective metrics, reported in Table 3. MOS ranges from 1 (bad) to 5 (excellent). We use test sentences in five languages drawn from ITU-T P.501 Annex D and evaluate them using PESQ (ITU-T P.862 [13]) and ViSQOL [10, 20]. ViSQOL has been shown to correlate well with subjective MOS

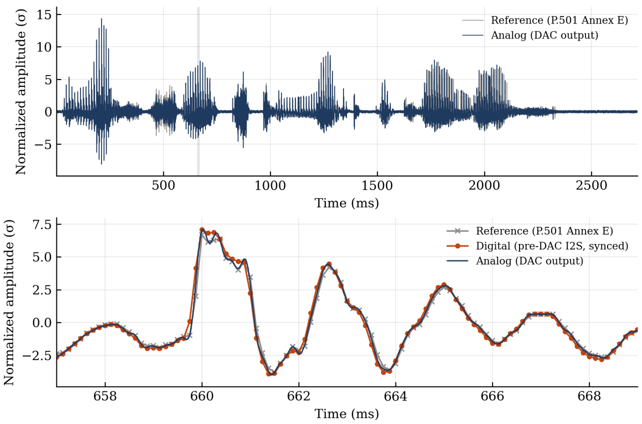

Figure 6: Audio output: waveform of the complete sentence "The last switch cannot be turned off" (top) and zoomed-in view (bottom).

Table 3: Objective estimates of mean opinion scores for 42 speech recordings from ITU-T P.501 Annex D, separated by language.

		PESQ (NB)	ViSQOL (speech)		
Language	n	Analog	Source	Digital	Analog
American English	7	2.13 ± 0.13	4.24	3.78 ± 0.08	3.94 ± 0.08
Mandarin Chinese	6	1.81 ± 0.16	3.75	3.34 ± 0.10	3.39 ± 0.11
French	8	2.13 ± 0.20	4.02	3.56 ± 0.12	3.70 ± 0.13
German	10	1.71 ± 0.15	4.27	3.61 ± 0.11	3.74 ± 0.10
Italian	11	1.72 ± 0.21	4.25	3.77 ± 0.19	3.89 ± 0.18
Average	42	1.88 ± 0.26	4.14	3.63 ± 0.20	3.75 ± 0.21

even in the presence of clock drift and jitter [9]; PESQ, in contrast, is known to underperform when clock drift is present between the reference and the captured signal, as is the case here. We therefore consider the ViSQOL score to be the more reliable estimate of perceived quality for our crystal-free link, and report PESQ mainly for reference.

Our mean analog-output ViSQOL score (3.75) is in line with values reported for other real-time voice systems under imperfect conditions: a large-scale study of ViSQOL scores collected from Google Hangouts Meet calls found scores of 4.21–4.28 under good network conditions, 4.04–4.16 under moderately impaired conditions, and 3.72–3.94 under severely challenging network conditions [7]. This places our crystal-free transmission on par with a VoIP call experiencing significant network impairment. For context, a drive-test study of a commercial mobile network reported a mean POLQA-based MOS of 3.57 for legacy GSM calls [17], though its different metric [9] and corpus make it only a broad reference point.

3.3.4 Quality Loss Breakdown. Table 3 lets us attribute the quality loss to each stage of the chain. The source audio, once encoded to fit the link capacity, already scores 4.14, so the encoding alone accounts for a loss of 0.86 points with respect to the 5.0 ceiling. The digital data recovered by the receiver drops the score by a further 0.51 points, while the filtering and oversampling performed by the DAC recover 0.12 points.

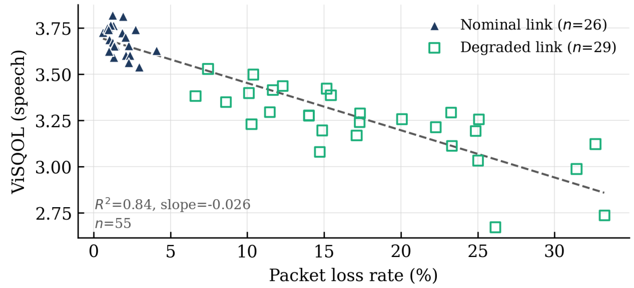

Figure 7: Audio quality versus packet loss, where a packet counts as lost when it is corrupted or never received.

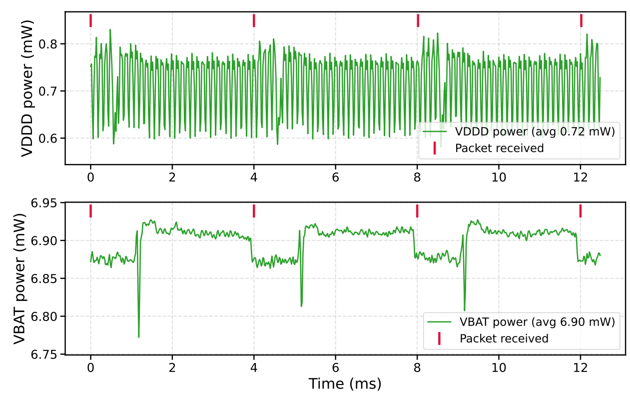

Figure 8: Measured power consumption versus time for the receiver on two domains: VDDD (Cortex M0 and RAM) and VBAT (global chip supply). Red lines indicate the start of a new packet. The measurements are taken in two different sessions and aligned on a packet arrival.

Fig. 7 shows the effect of packet loss on quality: the ViSQOL score deteriorates markedly as the radio link degrades and more packets are lost or received corrupted. Under a nominal radio link, however, packet losses remain small and their effect on audio quality is limited. We therefore attribute the 0.51 point drop from source to digital mostly to clock synchronization rather than to packet loss.

3.4 Power Consumption

This section focuses on the radio receiver power consumption. We assume that the transmitter is less power constrained than the receiver. We also do not measure the power usage of the DAC and amplifier as these components exist in any audio chain and are not specific to this configuration.

Measurement setup: SCμM is supplied with three different voltages: VBAT=1.8 V (global chip supply), VDDD=1.1 V (Cortex and RAM), and VDDIO=3.3 V (IO). The Otii Ace power supply supplies the three voltage rails and measures the currents drawn.

Results: Fig. 8 shows that the power used by VDDD (Cortex M0 core and RAM) has a strong periodicity at 8 kHz, which corresponds to the audio sampling rate. The higher consumption fits with the timing of accesses to the ring buffer. VBAT (global chip supply) notably powers the radio receiver. It shows a higher power consumption while listening for packets.The average total power excluding the IO domain is therefore 7.6 mW.

4 Conclusion

This work demonstrates continuous wireless audio streaming between fully crystal-free devices. A SCμM transmitter broadcasts 8 kHz stereo audio over IEEE 802.15.4 to SCμM receivers. Automatic frequency rendezvous allows the devices to communicate without an external frequency reference. The resulting analog audio is intelligible and achieves a mean objective ViSQOL estimate of 3.75. During streaming, the receiver consumes an average of 7.6 mW, excluding the I/O power domain.

The prototype nevertheless exhibits substantial audio-clock mismatch and packet corruption, requiring buffer management and packet-loss concealment. We attribute the audio-clock mismatch primarily to SCμM's limited clock-calibration resolution rather than to crystal-free operation itself. Moreover, the current implementation cannot enter a sleep state: the radio listens continuously, while the processor is responsible for generating I^2S in software and servicing the audio stream at 8 kHz. Therefore, the measured power draw reflects an always-active implementation rather than the minimum power required for crystal-free communications.

A future iteration of SCμM with finer clock-frequency control could reduce the audio-clock mismatch. Hardware I^2S and DMA support would also allow greater margin for executing more sophisticated compression and packet-loss-concealment algorithms without compromising audio timing, and would allow the system to sleep between radio and audio events, further reducing the receiver's power draw.

Acknowledgments

This work was funded by the French National Research Agency (ANR-22-PEFT-0007) as part of France 2030 and the NF-FITNESS project. We thank Alfonso Cortes and Thomas Watteyne for their guidance and editorial work. Large language model (LLM) tools were used to accelerate code development and assist with data visualization and text formatting.

References

[1] Alexandre Abadie, Titan Yuan, Thomas Watteyne, and Filip Maksimovic. 2025. An Open-Source Multi-Platform SDK for Programming SCuM. In *The 2025 Workshop on Crystal-Free/-Less Radio and System-based Research for IoT (CrystalFreeIoT 2025) - EWSN 2025*. Leuven, Belgium. https://inria.hal.science/hal-05281393

[2] Abdullah Alghaihab, Xing Chen, Yao Shi, Daniel S. Truesdell, Benton H. Calhoun, and David D. Wentzloff. 2020. 30.7 A Crystal-Less BLE Transmitter with -86dBm Frequency-Hopping Back-Channel WRX and Over-the-Air Clock Recovery from a GFSK-Modulated BLE Packet. In *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*. IEEE, San Francisco, CA, USA, 472–474. doi:10.1109/ISSCC19947.2020.9062935

[3] Tengfei Chang, Timothy Claeys, Mališa Vučinić, Xavi Vilajosana, Titan Yuan, Brad Wheeler, Filip Maksimovic, David Burnett, Brian Kilberg, Kris Pister, and Thomas Watteyne. 2020. Industrial IoT with Crystal-Free Mote-on-Chip. In *2020 IEEE Symposium on VLSI Circuits*. IEEE, Honolulu, HI, USA, 1–2. ISSN: 2158-5636. doi:10.1109/VLSICircuits18222.2020.9162981

[4] Tengfei Chang, Thomas Watteyne, Filip Maksimovic, Brad Wheeler, David C. Burnett, Titan Yuan, Xavier Vilajosana, and Kristofer S. J. Pister. 2021. QuickCal: Assisted Calibration for Crystal-Free Micromotes. *IEEE Internet of Things Journal* 8, 3 (2021), 1846–1858. doi:10.1109/JIOT.2020.3015725

[5] Tengfei Chang, Thomas Watteyne, Brad Wheeler, Filip Maksimovic, David C. Burnett, and Kris Pister. 2022. Surviving the Hair Dryer: Continuous Calibration of a Crystal-Free Mote-on-Chip. *IEEE Internet of Things Journal* 9, 6 (March 2022), 4737–4747. doi:10.1109/JIOT.2021.3108077

[6] Xing Chen, Abdullah Alghaihab, Yao Shi, Daniel S. Truesdell, Benton H. Calhoun, and David D. Wentzloff. 2021. A Crystal-Less BLE Transmitter With Clock Recovery From GFSK-Modulated BLE Packets. *IEEE Journal of Solid-State Circuits* 56, 7 (2021), 1963–1974. doi:10.1109/JSSC.2020.3046610

[7] Michael Chinen, Felicia S. C. Lim, Jan Skoglund, Nikita Gureev, Feargus O'Gorman, and Andrew Hines. 2020. ViSQOL v3: An Open Source Production Ready Objective Speech and Audio Metric. arXiv:2004.09584 [eess.AS]. doi:10.48550/arXiv.2004.09584

[8] Diodes Incorporated. 2013. *PAM8403: 3W Filterless Class-D Stereo Audio Amplifier*. Diodes Incorporated. https://www.diodes.com/assets/Datasheets/PAM8403.pdf

[9] Andrew Hines, Jan Skoglund, Anil Kokaram, and Naomi Harte. 2013. Robustness of speech quality metrics to background noise and network degradations: Comparing ViSQOL, PESQ and POLQA. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, Vancouver, BC, Canada, 3697–3701. doi:10.1109/ICASSP.2013.6638348

[10] Andrew Hines, Jan Skoglund, Anil C. Kokaram, and Naomi Harte. 2015. ViSQOL: an objective speech quality model. *EURASIP Journal on Audio, Speech, and Music Processing* 2015, 1 (May 2015), 13. doi:10.1186/s13636-015-0054-9

[11] O. Hodson, C. Perkins, and V. Hardman. 2000. Skew detection and compensation for Internet audio applications. In *2000 IEEE International Conference on Multimedia and Expo. ICME2000. Proceedings. Latest Advances in the Fast Changing World of Multimedia (Cat. No.00TH8532)*, Vol. 3. IEEE, New York, NY, USA, 1687–1690 vol.3. doi:10.1109/ICME.2000.871096

[12] IEEE 2024. *IEEE Standard for Low-Rate Wireless Networks*. IEEE. https://standards.ieee.org/standard/802_15_4-2024.html

[13] International Telecommunication Union. 2001. *Perceptual evaluation of speech quality (PESQ): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs*. ITU-T Recommendation P.862. International Telecommunication Union. https://www.itu.int/rec/T-REC-P.862

[14] International Telecommunication Union. 2025. *Test signals for use in telephonometry*. ITU-T Recommendation P.501. International Telecommunication Union. https://www.itu.int/rec/T-REC-P.501-202504-I/en

[15] ITU-T. 1988. *Pulse Code Modulation (PCM) of Voice Frequencies*. Recommendation G.711. International Telecommunication Union. https://www.itu.int/rec/T-REC-G.711

[16] ITU-T. 1999. *A High Quality Low-Complexity Algorithm for Packet Loss Concealment with G.711*. Recommendation G.711, Appendix I. International Telecommunication Union. https://www.itu.int/rec/T-REC-G.711

[17] Sorin Leonte, Andra Pastrav, Ciprian Zamfirescu, and Emanuel Puschita. 2024. Voice Quality Evaluation in a Mobile Cellular Network: In Situ Mean Opinion Score Measurements. *Sensors* 24, 20 (Jan. 2024), 6630. doi:10.3390/s24206630

[18] Filip Maksimovic, Austin Patel, David Burnett, Thomas Watteyne, and Kristofer S.J. Pister. 2023. A Time Synchronized Multi-Hop Mesh Network with Crystal-Free Nodes. In *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*. IEEE, Kuala Lumpur, Malaysia, 4158–4163. ISSN: 2576-6813. doi:10.1109/GLOBECOM54140.2023.10437167

[19] Filip Maksimovic, Brad Wheeler, David C. Burnett, Osama Khan, Lydia Lee, Ioana Suciu, Lydia Lee, Alex Moreno, Arvind Sundararajan, Bob Zhou, Rachel Zoll, Andrew H. Ng, Tengfei Chang, Xavier Villajosana, Thomas Watteyne, Ali M. Niknejad, and Kristofer S. J. Pister. 2019. A Crystal-Free Single-Chip Micro Mote with Integrated 802.15.4 Compatible Transceiver, sub-mW BLE Compatible Beacon Transmitter, and Cortex M0. *Symposium on VLSI Circuits* 2019 (2019), C88–C89. doi:10.23919/VLSIC.2019.8777971

[20] Pranay Manocha, Adam Finkelstein, Richard Zhang, Nicholas J. Bryan, Gautham J. Mysore, and Zeyu Jin. 2020. A Differentiable Perceptual Audio Metric Learned from Just Noticeable Differences. In *Interspeech 2020*. ISCA, Shanghai, China, 2852–2856. doi:10.21437/Interspeech.2020-1191

[21] Novo Nordisk. 2026. *NovoPen 6 & NovoPen Echo Plus*. https://www.novonordisk.com/our-products/smart-pens/novopen-6.html

[22] PisterLab. 2026. Sulu Development Board for SCμM. https://github.com/PisterLab/scum-dev-board. Accessed: 2026-07-19.

[23] Saleae, Inc. 2026. *I2S / PCM Protocol Analyzer*. Saleae, Inc. Logic 2 built-in analyzer extension. https://support.saleae.com/protocol-analyzers/analyzer-user-guides/using-i2s-pcm

[24] Texas Instruments. 2012. *PCM510x 2VRMS DirectPath, 112dB Audio Stereo DAC with 32-bit, 384kHz PCM Interface*. Texas Instruments. https://www.ti.com/lit/gpn/PCM5102

[25] Brad Wheeler, Filip Maksimovic, Nima Baniasadi, Sahar Mesri, Osama Khan, David Burnett, Ali Niknejad, and Kris Pister. 2017. Crystal-free narrow-band radios for low-cost IoT. In *2017 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. IEEE, Honolulu, HI, USA, 228–231. https://ieeexplore.ieee.org/abstract/document/7969059/