LIÈGE université Library

https://lib.uliege.be https://matheo.uliege.be

Master thesis : Design Improvements and Optimisation of the OhmPi Watcher System for Large-Scale Autonomous Deep ERT Measurements

Auteur : Eyen, Guillaume
Promoteur(s) : Redouté, Jean-Michel; Nguyen, Frédéric
Faculté : Faculté des Sciences appliquées
Diplôme : Master : ingénieur civil électricien, à finalité spécialisée en "electronic systems and devices"
Année académique : 2025-2026
URI/URL : https://gitlab.com/ohmpi/ohmpi-watcher; http://hdl.handle.net/2268.2/27087

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.

UNIVERSITE DE LIEGE
FACULTE DES SCIENCES APPLIQUEES

Design Improvements and Optimisation of the OhmPi Watcher System for Large-Scale Autonomous Deep ERT Measurements

EYEN Guillaume

s182398

Thesis presented to obtain the degree of:

Master of Science in Electrical Engineering
Electronic Systems and Devices

Academic supervisors:

Prof. Jean-Michel Redouté
Prof. Frédéric Nguyen

Scientific advisor:
Dr Guillaume Blanchy

Academic year: **2025–2026**

Institutional and Legal Declarations

Intellectual property and attribution.

This dissertation combines original work carried out by the author with material from scientific publications, manufacturer documentation and open-source projects. External figures, technical information, software and hardware designs have been cited or acknowledged where applicable.

The developments presented here build upon the existing OhmPi and OhmPi Watcher projects and should therefore be understood as extensions of these platforms rather than claims of authorship over the original designs. Every effort has been made to acknowledge all relevant sources and any remaining omission would be unintentional.

Open-source material and licensing.

The Watcher is developed as an open-source project. The software source code associated with this work is distributed under the GNU Affero General Public License version 3 (AGPLv3), while the hardware design files are distributed under the CERN Open Hardware Licence Version 2 – Permissive (CERN-OHL-P-2.0). The complete source files and applicable licence texts are provided in the project GitLab repository [1].

Use of artificial intelligence tools.

The author acknowledges the use of generative artificial intelligence tools during the preparation of this work. These tools were used primarily to improve the clarity, fluency and structure of the written text, as well as to assist with the refinement of code and data-processing scripts.

Artificial intelligence tools were used in a supporting role and not as a substitute for experimental work or engineering judgement. The hardware and software design choices, experimental work, data selection, technical interpretation and conclusions presented in this dissertation were developed by the author, who remains fully responsible for the final content of the work.

Abstract

Deep electrical resistivity tomography (deep-ERT) extends conventional resistivity imaging to kilometre-scale profiles and greater investigation depths, but it requires large numbers of autonomous receiving stations that must record weak ground potentials, remain time-synchronised and stay remotely accessible over multi-day campaigns. Commercial distributed receivers meet these needs at a cost of several thousand euros per unit, which limits the achievable spatial coverage. This work addresses that gap by redesigning the OhmPi Watcher, an open-source autonomous receiver derived from the OhmPi platform, for a reference deployment of 100 stations spread over more than 10 km for at least five days.

The redesign targets four complementary aspects: hardware power architecture, analogue acquisition electronics, embedded software structure and wireless communication. The revised hardware integrates the power, acquisition, switching and protection circuitry on a single 100 × 160 mm board, extends the analogue front end from two to four independently conditioned electrode channels, and disconnects the bipolar analogue supply through a high-side switch when it is not required. The firmware was reorganised around an explicit finite-state machine managing recording, idle and deep-sleep operation, RTC-based wake-up and arbitration of the SPI bus shared by the SD card and the LoRa radio. Communication was improved through a revised 433-MHz antenna integration, multi-hop packet forwarding with duplicate suppression, and an autonomous STM32/SX1268 portable gateway able to bypass locally obstructed links.

Laboratory validation shows a two-channel recording current reduced from 174.6 mA to 77.6 mA (-56%), an estimated five-day consumption of about 4.16 Ah, an input-referred noise of 0.31 mV with channel gain errors below 0.07%, one hour of continuous logging at 99.964 Hz without record loss, and full recovery of an obstructed radio link from 0% to 100% packet reception via the gateway, for an estimated hardware cost of about 250 € per unit. Four-channel acquisition at 100 Hz and validation under real deep-ERT field conditions remain the main open points.

Acknowledgements

I would first like to express my sincere gratitude to my supervisors, Prof. Jean-Michel Redouté and Prof. Frédéric Nguyen, for giving me the opportunity to carry out this work at the intersection of electronics and geophysics. I am particularly grateful for the freedom they gave me to explore the different directions taken throughout this project.

I would also like to warmly thank Dr Guillaume Blanchy, who developed the original OhmPi Watcher prototype on which this work is based. His availability, responsiveness and technical insight were extremely valuable throughout the project. I am especially grateful for the time he devoted to discussing the system, reviewing this manuscript and providing detailed feedback, while at the same time giving me considerable freedom in the solutions I chose to investigate and implement.

I would like to thank my mother for the stability and encouragement she has given me throughout the years. Her support has helped me stay motivated and keep going during more demanding periods, and was also a key factor in successfully completing the air traffic controller selection process. I am equally grateful for the many messages of encouragement she sent before every university exam and all other important moments.

I would also like to thank my father for his support and for all the good moments we have shared over the past years, particularly through mountaineering, hiking and ski touring. These activities always provided a welcome break from my studies and helped me return to them with a clearer mind.

A very special thank you goes to my girlfriend Lucie for her constant support, patience, encouragement and for always believing in me throughout these years of study. As a geological and mining engineer, she also influenced my academic path by introducing me to geosciences before I eventually specialised in electronics, which ultimately led to a thesis combining both fields. I am also very grateful to her family for their warm welcome and the supportive environment they provided throughout my studies.

I would also like to thank my colleagues at the Liège Airport air traffic control tower. Over the past five years, they have accepted an impressive number of shift swaps and schedule changes so that I could attend compulsory university activities. Combining my studies with my professional activity would have been much more difficult without their support.

Finally, I would like to thank my friends for the many different ways they have supported me throughout these years, whether through academic help, sport, shared experiences or simply much-needed moments of fun and distraction.

Completing this degree has been a long journey, and I am grateful to everyone who contributed to it, whether academically, professionally or personally.

Contents

1 Introduction **1**

2 Literature review **3**

2.1 Electrical Resistivity Tomography and Induced Polarisation . . . 3
2.1.1 Physical principle . . . 3
2.1.2 Data acquisition . . . 5
2.1.3 Data inversion . . . 7
2.1.4 Deep ERT . . . 9
2.2 Autonomous geoelectrical acquisition systems . . . 11
2.2.1 From manual surveys to autonomous monitoring . . . 11
2.2.2 Distributed full-wave acquisition systems . . . 11
2.3 The OhmPi platform and Watcher prototype . . . 12
2.3.1 The OhmPi open-source platform . . . 12
2.3.2 The original OhmPi Watcher prototype . . . 13
2.4 Low-power design for autonomous sensor nodes . . . 16
2.4.1 Power budgeting and operating states . . . 16
2.4.2 Dynamic power management and duty cycling . . . 16
2.4.3 Power-supply efficiency and hardware integration . . . 17
2.4.4 Energy harvesting and application constraints . . . 18
2.5 Time synchronisation in distributed measurement systems . . . 18
2.5.1 Clock errors and synchronisation requirements . . . 18
2.5.2 Synchronisation strategies for distributed nodes . . . 19
2.5.3 GNSS timing, PPS and RTC-based local timekeeping . . . 19
2.6 Embedded data acquisition and local storage . . . 20
2.6.1 Analogue signal conditioning and digitisation . . . 21
2.6.2 Buffered storage and data integrity . . . 22
2.7 Low-power wireless communication and LoRa . . . 23
2.7.1 Wireless communication requirements for autonomous sensor networks . . . 23
2.7.2 LoRa modulation principle . . . 24
2.7.3 LoRa configuration and performance trade-offs . . . 27
2.8 Synthesis and design requirements . . . 27

3 Hardware development **29**

3.1 Microcontroller choice . . . 29
3.2 Power supply redesign . . . 32
3.2.1 Power budget and converter sizing . . . 32
3.2.2 Switchable power source . . . 35
3.3 Analogue acquisition chain redesign . . . 36
3.3.1 Operational-amplifier integration . . . 37

3.3.2 Analogue filtering and ADC input protection . . . 38
3.3.3 Electrode switching and fail-safe states . . . 40
3.3.4 ADC requirements and device comparison . . . 40
3.3.5 ADC architecture and channel expansion . . . 42
3.4 Electrode contact-resistance circuit redesign . . . 43
3.4.1 Relay driving circuit . . . 43
3.4.2 Current sensing and shunt protection . . . 44
3.5 PCB integration and component-level improvements . . . 46
3.6 Final hardware architecture . . . 47

4 Software development 48
4.1 Runtime architecture and resource management . . . 49
4.1.1 Original control flow and limitations . . . 49
4.1.2 Finite-state machine architecture and event handling . . . 49
4.1.3 SD-card and LoRa SPI arbitration . . . 51
4.2 Data acquisition and logging . . . 52
4.2.1 Acquisition scheduler and channel handling . . . 52
4.2.2 Diagnostic and self-check routines . . . 53
4.2.3 Binary record format and storage requirements . . . 53
4.2.4 SD-card buffering and file management . . . 55
4.3 Time reference and synchronisation . . . 57
4.3.1 GNSS, RTC and local time reference . . . 57
4.3.2 Pulse-per-second (PPS) based drift estimation . . . 58
4.4 Low-power scheduling and wake-up . . . 59
4.4.1 Daily operating schedule and peripheral shutdown . . . 59
4.4.2 RTC-alarm wake-up and deep sleep . . . 60

5 Wireless communication 62
5.1 Communication architecture and design objectives . . . 62
5.2 Operating frequency and regulatory constraints . . . 65
5.3 Hardware implementation . . . 65
5.3.1 Watcher radio and antenna integration . . . 65
5.3.2 Gateway radio platform . . . 66
5.3.3 Gateway power supply and physical integration . . . 68
5.4 Software implementation . . . 70
5.4.1 Watcher radio driver and interrupt-based reception . . . 70
5.4.2 STM32 HAL gateway firmware . . . 71
5.4.3 Gateway user interface . . . 74
5.5 Packet format and communication protocol . . . 75
5.5.1 Packet format, addressing and commands . . . 75
5.5.2 Store-and-forward packet propagation . . . 76
5.5.3 Duplicate suppression and repeated-packet handling . . . 77
5.5.4 Response handling and communication failure detection . . . 78
5.6 Communication cost and scalability . . . 79
5.7 Protocol limitations and future network developments . . . 80
5.7.1 Proposed routing extension for a mobile remote controller . . . 80
5.7.2 Topology-aware routing . . . 82
5.7.3 Network supervision GUI . . . 83

6 Experimental validation and results 84
6.1 Power consumption and autonomy . . . 85

6.1.1 Watcher operating-state consumption . . . 85
6.1.2 Subsystem power-consumption analysis . . . 85
6.1.3 Comparison with the original Watcher . . . 89
6.1.4 Low-power operation and estimated autonomy . . . 90
6.1.5 Gateway consumption and battery-voltage monitoring . . . 91
6.2 Acquisition and data-logging performance . . . 94
6.2.1 Acquisition timing . . . 94
6.2.2 SD-card write timing and buffering . . . 95
6.2.3 Power-loss integrity . . . 96
6.2.4 Continuous logging integrity . . . 97
6.3 Analogue front-end validation . . . 98
6.3.1 Electronic offset and noise . . . 98
6.3.2 DC gain and linearity . . . 99
6.3.3 Analogue front-end power-up settling . . . 101
6.3.4 Low-pass filter frequency response . . . 102
6.3.5 Contact-resistance measurement accuracy . . . 103
6.4 Wireless communication performance . . . 104
6.4.1 Antenna positioning . . . 104
6.4.2 Radio interoperability and direct-link performance . . . 105
6.4.3 Transmission and forwarding time . . . 106
6.4.4 Gateway-assisted link recovery . . . 108
6.4.5 Duplicate suppression . . . 108
6.5 Implementation cost and scalability . . . 110
6.6 Validation summary and remaining limitations . . . 111

7 Conclusion and future work **113**
7.1 Main outcomes . . . 113
7.2 Limitations and future improvements . . . 114

Bibliography **117**

A Complete electrical schematic **123**

B E22-400MBL-SC evaluation board schematic **124**

C Adafruit Feather M0 RFM9x pinout **125**

D Detailed firmware execution diagrams **126**

List of Figures

2.1 Theoretical representation of a four-electrode electrical resistivity measurement: a) current-flow lines and equipotential lines in the subsurface; b) theoretical current injection; c) ideal potential response including the chargeability effect; d) noisy potential measurement; and e) noisy measurement affected by spontaneous potential (SP). Reproduced from [6]. . . . 4
2.2 ABEM Terrameter LS 2 data acquisition system. Reproduced from [10]. 5
2.3 Generic electrode arrangement using the ABEM Terrameter LS 2 data acquisition system. Four 100 m cable spools are deployed here to obtain a profile length of approximately 400 m. Reproduced from [10]. 6
2.4 Example of a 3D survey composed of several parallel 2D profiles. Reproduced from [11]. . 6
2.5 Schematic diagrams of electrode arrays and their sensitivity patterns for 2D resistivity surveys. C1 and C2 are positive and negative current electrodes, respectively. P1 and P2 are two potential electrodes. The italic letters *a*, *n* and *s* are the array parameters that control the spacing and maximum separation of the arrays in a multi-electrode layout. Reproduced from [12] 7
2.6 Simplified workflow of ERT-IP data inversion. The response predicted from an initial subsurface model is compared with the measured data. The resulting residual is used to update the model until an acceptable fit is obtained. 8
2.7 Example of a resistivity model obtained from a deep-ERT survey, illustrating the kilometre-scale lateral extent and the large investigation depth achievable with this approach. Reproduced from Balasco et al. [2]. 9
2.8 IRIS Instruments V-FullWaver autonomous receiver. Reproduced from [5]. 12
2.9 Electronic signal conditioning circuit for Vmn from M/N front-end to the ADS1115 on the OhmPi. This measurement architecture is reused in the OhmPi Watcher. Reproduced from [26]. 13
2.10 Annotated view of the OhmPi Watcher v0.1.2, showing its main hardware components. . 14
2.11 Functional architecture of the original OhmPi Watcher prototype. The unit records the voltage across two receiving dipoles, synchronises the measurements using GNSS timing, stores the raw data locally and exchanges commands and status information through a LoRa radio link. 15
2.12 Comparison of voltage measurements obtained with the original OhmPi Watcher and the commercial V-FullWaver. Both instruments reproduce the same overall signal evolution, while the Watcher recording exhibits larger high-frequency fluctuations. Provided by Hadrien Michel, 2026. 15
2.13 Examples of NMEA sentences. Reproduced from [38]. 20
2.14 Delta-sigma ADC architecture. Reproduced from [44]. 21
2.15 Example of an upchirp and a downchirp in time domain (left) and frequency domain (right). Reproduced from [57]. 24

2.16 Illustration of eight symbols (SF=3, for didactic purpose only) obtained by cyclically shifting the initial frequency of the same up-chirp. All symbols have the same duration, bandwidth and sweep rate, but reach the upper frequency boundary at different instants before wrapping to the lower boundary. Reproduced from [60]. 25
2.17 Time–frequency representation of LoRa symbols from $SF = 7$ to $SF = 12$ at constant bandwidth. Each increase of one in the spreading factor doubles the symbol duration. Reproduced from [56]. 26

3.1 Physical comparison between the original OhmPi Watcher (v0.1.2) and the redesigned hardware developed in this work (v0.3.0). The new PCB integrates the revised power architecture, four electrode-conditioning channels and the associated control and protection circuitry into a more compact layout. 31
3.2 Schematic pin assignments of the Adafruit Feather M0 board in the redesigned Watcher. Modified from Appendix C 32
3.3 Power architecture of the redesigned Watcher. The 12 V battery supplies the continuously available low-voltage electronics through a dedicated 5 V converter, whereas the bipolar analogue supply is connected through a controllable high-side switch. The ±15 V domain can therefore be disconnected when the analogue acquisition circuitry is not required. . . . 33
3.4 High-side switching circuit used to disconnect the TEC 3-1223 from the 12 V battery. Q1 provides the high-side power switching, while Q2 allows the 3.3 V `REMOTE` signal from the Feather M0 to control its gate. 36
3.5 Analogue acquisition chain implemented for one receiving electrode in the redesigned Watcher. The input of the conditioning chain can be connected either to the electrode or to the analogue ground through the two DG413 switches. The resulting signal is then conditioned by the OPA4188, low-pass filtered, protected by the output clamping network and digitised by the ADS1115. 37
3.6 Close-up view of the analogue acquisition chain on the new PCB. DG413 switches not shown. 37
3.7 Electrode contact-resistance measurement circuit of the redesigned Watcher. The four G6L-1F relays select the electrode under test and are driven by the ULN2003A Darlington transistor array. The resulting test current flows through the series protection resistor and the current-sensing shunt, whose voltage is amplified by the INA282 before digitisation. . 43
3.8 Three-dimensional view of the empty PCB in KiCad 10. 46

4.1 Simplified comparison of the Watcher firmware execution structure before and after the software redesign. The original implementation coordinates acquisition, communication, timing, diagnostics and power management through multiple conditional paths, whereas the redesigned implementation separates these functions through explicit states and controlled transitions. Detailed execution diagrams of both implementations are provided in Appendix D. 50
4.2 Hardware modification connecting the `INT` output of the PCF8523 RTC on the Adalogger board to the interrupt-capable `A1` input of the Feather M0. This connection allows an RTC alarm to wake the microcontroller from deep sleep. 60

5.1 Controller-initiated multi-hop communication architecture of the Watcher network. Under normal conditions, a request issued by W_0 and the corresponding response are relayed between successive Watchers. A local propagation obstacle can interrupt this communication path and isolate the following receivers. A portable gateway can then be positioned to provide an alternative radio path across the problematic section. 64
5.2 Evolution of the Watcher antenna integration from the original internal wire antenna to two possible vertical integrations of the redesigned 433-MHz antenna. 66
5.3 E22-400MBL-SC evaluation board and main components. 67

5.4 Power-supply and battery-monitoring architecture of the portable gateway. The 18650 cell supplies the E22-400MBL-SC gateway board through the battery-protection and charging stage and the TPS63802 buck–boost converter. Battery voltage is monitored independently through a high-impedance voltage divider and low-pass filter connected to an STM32 ADC input. 68
5.5 Annotated views of the gateway power-supply perfboard. The board integrates the battery-protection and USB-C charging module, the TPS63802 buck–boost converter, and the battery-voltage monitoring circuit connected to the STM32 ADC input. 69
5.6 Complete autonomous gateway integrated into a weather-resistant enclosure. 70
5.7 STM32F103C8T6 pin and peripheral assignments used by the gateway firmware. 72
5.8 Example of the USB CDC debug output monitored through Tera Term during gateway development. The gateway parses the received Watcher packet, determines that it must be forwarded from left to right, and confirms the successful retransmission. 73
5.9 Gateway user interface. 75
5.10 Example of duplicate packet reception in the presence of a gateway. A packet transmitted by W_2 could occasionally reach W_3 directly while also being received and retransmitted by the gateway. Without duplicate suppression, W_3 could process and forward the same transaction twice. 77
5.11 Example of a mobile remote controller positioned next to an intermediate Watcher. Although the logical identifier of the remote remains W_0, its physical position along the ordered Watcher chain can change during field operation. 81

6.1 Supply-current evolution during the automatic subsystem characterisation. Each plateau corresponds to a predefined Watcher configuration maintained for 60 s. The initial switching sequence was used to align the firmware test with the multimeter recording. 86
6.2 Voltage measured across the 2.05 Ω supply-current shunt during a LoRa transmission at a configured output power of 20 dBm. The shunt-voltage increase relative to the continuous-reception baseline was approximately 59.2 mV. 87
6.3 Gateway-reported battery voltage as a function of the reference voltage measured with a digital multimeter. The dashed line represents the ideal relationship $V_{\text{gateway}} = V_{\text{DMM}}$, while the solid line represents the linear regression of the experimental data. 94
6.4 Measured DC transfer characteristics of the four analogue acquisition channels. The horizontal axis corresponds to the actual electrode voltage measured with a digital multimeter. The dashed line represents the ideal transfer function $V_{\text{ADC}} = 0.5V_{\text{in}} + 2.5$ V. 100
6.5 Measured settling of the four analogue acquisition channels after activation of the ±15 V supply. The inputs were connected to GNDA during the test. The dashed line indicates the nominal 2.5-V bias point. 101
6.6 Measured frequency response of the analogue low-pass filter compared with the theoretical first-order RC response and a first-order fit to the experimental data. 102
6.7 Measured contact resistance as a function of the reference resistance measured with a digital multimeter. Each point represents the mean of five measurements. The dashed line represents the ideal relationship $R_{\text{Watcher}} = R_{\text{DMM}}$. The four electrode channels produced very similar results, resulting in nearly overlapping curves. 103
6.8 Oscilloscope measurement of the Watcher LoRa transmission interval. The increased supply-current plateau associated with the transmission of a 17-byte packet lasted approximately 168.0 ms. 106
6.9 Serial output recorded during the duplicate-suppression validation. The first packet of each transaction is accepted, whereas subsequent identical copies are identified and suppressed. 109

List of Tables

2.1 Meaning of the main LoRa modulation parameters. 26
2.2 Main design requirements derived for the OhmPi Watcher. 28

3.1 Main hardware modifications introduced in the redesigned OhmPi Watcher. 30
3.2 Design power budget of the low-voltage domain. The 3.3 V peripherals are supplied through the on-board regulator of the Feather M0 and therefore also contribute to the current drawn from the 5 V rail. 34
3.3 Design power budget of the bipolar analogue domain. 34
3.4 Comparison of the ADS1115 and ADS1219 for four-channel acquisition at 100 Hz [43, 46]. 42

4.1 Main software modifications introduced in the redesigned Watcher firmware. 48
4.2 Main software libraries used in the final Watcher firmware. 48
4.3 Main runtime states of the redesigned Watcher firmware. 51
4.4 Main binary record types used by the Watcher firmware. Differences between record types 4, 6 and 7 are explained in Section 4.3.2. 54

5.1 Main wireless communication modifications introduced in this work. 62
5.2 LoRa parameters used during prototype development and characterisation. 71
5.3 Application-level LoRa packet format used by the Watcher network. 76
5.4 Main command codes currently used by the Watcher network. 76

6.1 Measured current consumption of the redesigned Watcher in its principal operating states. 85
6.2 Steady-state supply current measured during the automatic subsystem power-consumption characterisation. 86
6.3 Comparison of the measured current consumption of the original and redesigned Watchers at a 12.10 V supply voltage. 89
6.4 Measured gateway current consumption at a 3.7 V supply voltage, together with additional measurements of the EBYTE board alone. 91
6.5 Measured sequential ADC acquisition time. 94
6.6 Measured SD-card write and synchronisation timings at 100 Hz. 95
6.7 Summary of the 10 abrupt power-loss tests without periodic SD-card synchronisation. . . 96
6.8 Summary of the one-hour continuous logging test using the final fixed-grid sampling scheduler. 97
6.9 Measured offset and electronic noise of the four acquisition channels with their inputs connected to analogue ground. 98
6.10 Linear-regression results for the four analogue acquisition channels over the approximately -4 to $+4$ V input range. 100
6.11 Contact-resistance measurement accuracy for the four tested reference resistances. 104
6.12 Comparison of the three Watcher antenna configurations. 105
6.13 Direct-link interoperability between the Watcher RFM96 radio and the SX1268-based gateway radio. 105
6.14 Measured packet-transmission and gateway-forwarding times. 106

6.15 Gateway-assisted recovery of an obstructed radio link. 108
6.16 Packet identifiers used for the duplicate-suppression test. 109
6.17 Estimated hardware cost of one redesigned Watcher. The 100-unit estimate remains conservative because quantity discounts on electronic components are not included. . . . 110
6.18 Validation of the main Watcher design requirements defined in Table 2.2. 112

7.1 Main directions identified for future development of the OhmPi Watcher platform. 116

List of abbreviations

ADC	Analogue-to-Digital Converter
BOM	Bill of Materials
BW	Bandwidth
CDC	Communications Device Class
CMOS	Complementary Metal-Oxide-Semiconductor
CR	Coding Rate
CRC	Cyclic Redundancy Check
CSS	Chirp Spread Spectrum
DAC	Digital-to-Analogue Converter
DC	Direct Current
DMM	Digital Multimeter
ERP	Effective Radiated Power
ERT	Electrical Resistivity Tomography
ERT-IP	Electrical Resistivity Tomography – Induced Polarisation
FAT	File Allocation Table
GNDA	Analogue Ground
GND	Ground
GNSS	Global Navigation Satellite System
GPIO	General-Purpose Input/Output
GPS	Global Positioning System
GUI	Graphical User Interface
HAL	Hardware Abstraction Layer
I^2C	Inter-Integrated Circuit
IP	Induced Polarisation
IRQ	Interrupt Request
LSB	Least Significant Bit
LoRa	Long Range radio modulation
LoRaWAN	Long Range Wide Area Network

MAC	Medium Access Control
MCU	Microcontroller Unit
MOSFET	Metal-Oxide-Semiconductor Field-Effect Transistor
MQTT	Message Queuing Telemetry Transport
NC	Normally Closed
NMEA	National Marine Electronics Association
NO	Normally Open
NPN	Negative-Positive-Negative bipolar transistor
OSR	Oversampling Ratio
PCB	Printed Circuit Board
PGA	Programmable Gain Amplifier
PHY	Physical Layer
PPS	Pulse Per Second
PRR	Packet Reception Ratio
QSPI	Quad Serial Peripheral Interface
RAM	Random-Access Memory
RC	Resistor-Capacitor
RF	Radio Frequency
RMS	Root Mean Square
RSSI	Received Signal Strength Indicator
RTC	Real-Time Clock
RX	Receive
SAR	Successive-Approximation Register
SF	Spreading Factor
SNR	Signal-to-Noise Ratio
SP	Spontaneous Potential
SPI	Serial Peripheral Interface
SPS	Samples Per Second
SPST-NO	Single-Pole Single-Throw, Normally Open
SRD	Short-Range Device
TX	Transmit
UART	Universal Asynchronous Receiver-Transmitter
UTC	Coordinated Universal Time
VAT	Value-Added Tax

Chapter 1

Introduction

Electrical resistivity tomography (ERT) is a well-established geophysical imaging technique used to investigate groundwater systems, monitor environmental contamination, support mineral exploration, assess geotechnical stability, and characterise subsurface conditions for engineering projects. During an ERT survey, electrical current is injected into the ground through two electrodes, while the resulting potential differences are measured between pairs of receiving electrodes distributed along the surface. The measurements are sensitive to variations in lithology, water content, salinity and geological structure.

Deep electrical resistivity tomography (Deep-ERT) extends this principle to kilometre-scale profiles and greater depths, making it relevant to deep mineral and geothermal exploration, fault-zone mapping and the characterisation of sites for major underground infrastructure, including tunnels and projects such as the Einstein Telescope [2, 3, 4]. At these scales, the main challenges are largely operational: the current-injection system must be moved between widely spaced positions, while the receiving electrodes connected to numerous distributed stations must autonomously record weak voltage signals throughout multi-day campaigns. These measurements must be accurately time-synchronised and stored despite limited power availability and difficult field access.

Commercial distributed acquisition systems have already demonstrated the feasibility of deep-ERT surveys. The IRIS FullWaver system, for example, uses independent receivers that record complete voltage time series synchronised with the injected current [5]. However, the cost of several thousand euros per receiver limits the number of units that can realistically be deployed and therefore constrains the achievable spatial coverage and resolution. Its proprietary architecture also limits hardware modification and field repair by the user, while radio communication is provided as an optional rather than an inherent function of each receiver. These limitations motivate the development of a lower-cost, open-source and remotely controllable alternative.

In the reference application considered in this work, the deep-ERT survey extends over 10 km and involves 100 receiving stations deployed continuously for at least 5 days. Each station must record two receiving dipoles at a sampling interval of 10 ms, store the complete measurements locally and remain accessible for remote control and status monitoring. Electrode contact resistance must also be checked because cables or electrodes may be disturbed by animals, vegetation or human activity during unattended operation. Furthermore, the survey line may cross forested or topographically complex areas in which direct radio communication with every station cannot be guaranteed. Once installed, the stations are difficult and time-consuming to access, making autonomy, fault detection and reliable communication essential requirements. By comparison, conventional ERT profiles generally extend over only a few hundred metres and can often be acquired within a few hours with the equipment remaining directly accessible to the operators.

This work builds upon an existing measurement device, the OhmPi Watcher, developed to monitor ground potentials during deep-ERT acquisition. The initial OhmPi Watcher prototype is itself based on the

OhmPi open-source ERT measurement platform [6, 1], and has been modified to include GPS-based time synchronisation, local data storage, LoRa radio communication, electrode contact-resistance checking and a display for local status information. Compared to the original OhmPi platform, the OhmPi Watcher concept inherently introduced limitations related to size, weight, power consumption and communication reliability.

Power autonomy represents one of the main constraints of the system. Since the devices are deployed along the measurement profile without access to a fixed power supply, each unit must operate from an individual battery. The required battery capacity directly affects the cost, size and weight of the device. These effects become particularly significant when the system is replicated across 100 measurement units. Reducing the energy consumption of an individual Watcher can therefore decrease the total cost of a complete deployment, facilitate transportation and installation, and increase the possible duration of future measurement campaigns.

The primary objective of this work is thus to improve the energy efficiency and operational autonomy of the Watcher while preserving its measurement capabilities. This objective is pursued through four complementary approaches. First, the hardware architecture is revised to reduce unnecessary power losses, simplify the assembly and allow selected subsystems to be disconnected when they are not required. Second, the acquisition electronics are redesigned to improve their robustness, integration and scalability. In particular, the number of independently conditioned electrode channels is increased from two to four, while the analogue front end, switching, filtering and protection circuitry are integrated into a more compact and reproducible architecture. Third, the embedded software is refactored around clearly defined operating states so that measurement, communication and sleep periods can be controlled more efficiently. Finally, the wireless communication system is redesigned to provide reliable communication between spatially distributed stations and a remote controller, including the relaying of messages when direct radio links are obstructed by distance, vegetation or topography.

A further objective is to improve the maintainability and reproducibility of the platform. The hardware and software are therefore developed following an open-source approach, with an emphasis on clear documentation, modularity and ease of modification. The complete project files developed during this work are available in the project GitLab repository [1]. This is particularly important for an experimental platform that may be replicated, adapted to different field configurations or extended by future students or researchers.

The following chapters present the hardware development, software development, wireless communication improvements and experimental validation.

Chapter 2

Literature review

2.1 Electrical Resistivity Tomography and Induced Polarisation

Electrical Resistivity Tomography (ERT) and Induced Polarisation (IP) provide complementary information on the electrical properties of the subsurface by combining controlled current injection, voltage measurements, and numerical inversion. This section first introduces the physical principles governing resistivity and chargeability, then describes conventional multi-electrode acquisition systems and the main concepts used to reconstruct subsurface models. It finally presents deep ERT, in which distributed and synchronised receiving stations replace long multichannel cable spreads to enable investigations over kilometre-scale profiles and at greater depths, thereby establishing the context for the OhmPi Watcher developed in this work.

2.1.1 Physical principle

Electrical Resistivity Tomography (ERT) and Induced Polarisation (IP) are geophysical methods based on the electrical response of the subsurface to an externally applied current. In a conventional four-electrode configuration, a current I_{AB} is injected into the ground through two current electrodes, denoted A and B, while the resulting potential difference ΔV_{MN} is measured between two separate potential electrodes, denoted M and N (Fig. 2.1). The measured voltage depends not only on the electrical properties of the subsurface but also on the relative positions of the four electrodes [7, 8]. This geometrical influence is accounted for through a geometric factor K, allowing the apparent resistivity ρ_a to be calculated as

$$\rho_\mathrm{a} = KR = K\frac{\Delta V_{MN}}{I_{AB}}, \tag{2.1}$$

This value represents the equivalent resistivity of a homogeneous half-space medium that would produce the same electrical response for the considered electrode configuration [7, 9]. In heterogeneous ground, the apparent resistivity is therefore influenced by the spatial distribution of conductive and resistive materials within the investigated volume. By repeating the measurement for numerous electrode combinations and applying an inversion algorithm, a two- or three-dimensional model of the subsurface resistivity can be reconstructed (see Section 2.1.3).

In most soils and rocks, electrical conduction occurs predominantly through the interconnected pore fluid rather than through the solid silicate matrix. Bulk resistivity is therefore influenced by pore-fluid conductivity, water saturation, porosity, pore connectivity and temperature. Clay minerals can provide an additional surface-conduction mechanism, whereas electronically conductive minerals may strongly reduce the bulk resistivity [9, 7]. Consequently, resistivity contrasts may reflect changes in lithology, moisture content, salinity, clay content or fluid-filled fracture connectivity. However, the resistivity ranges associated

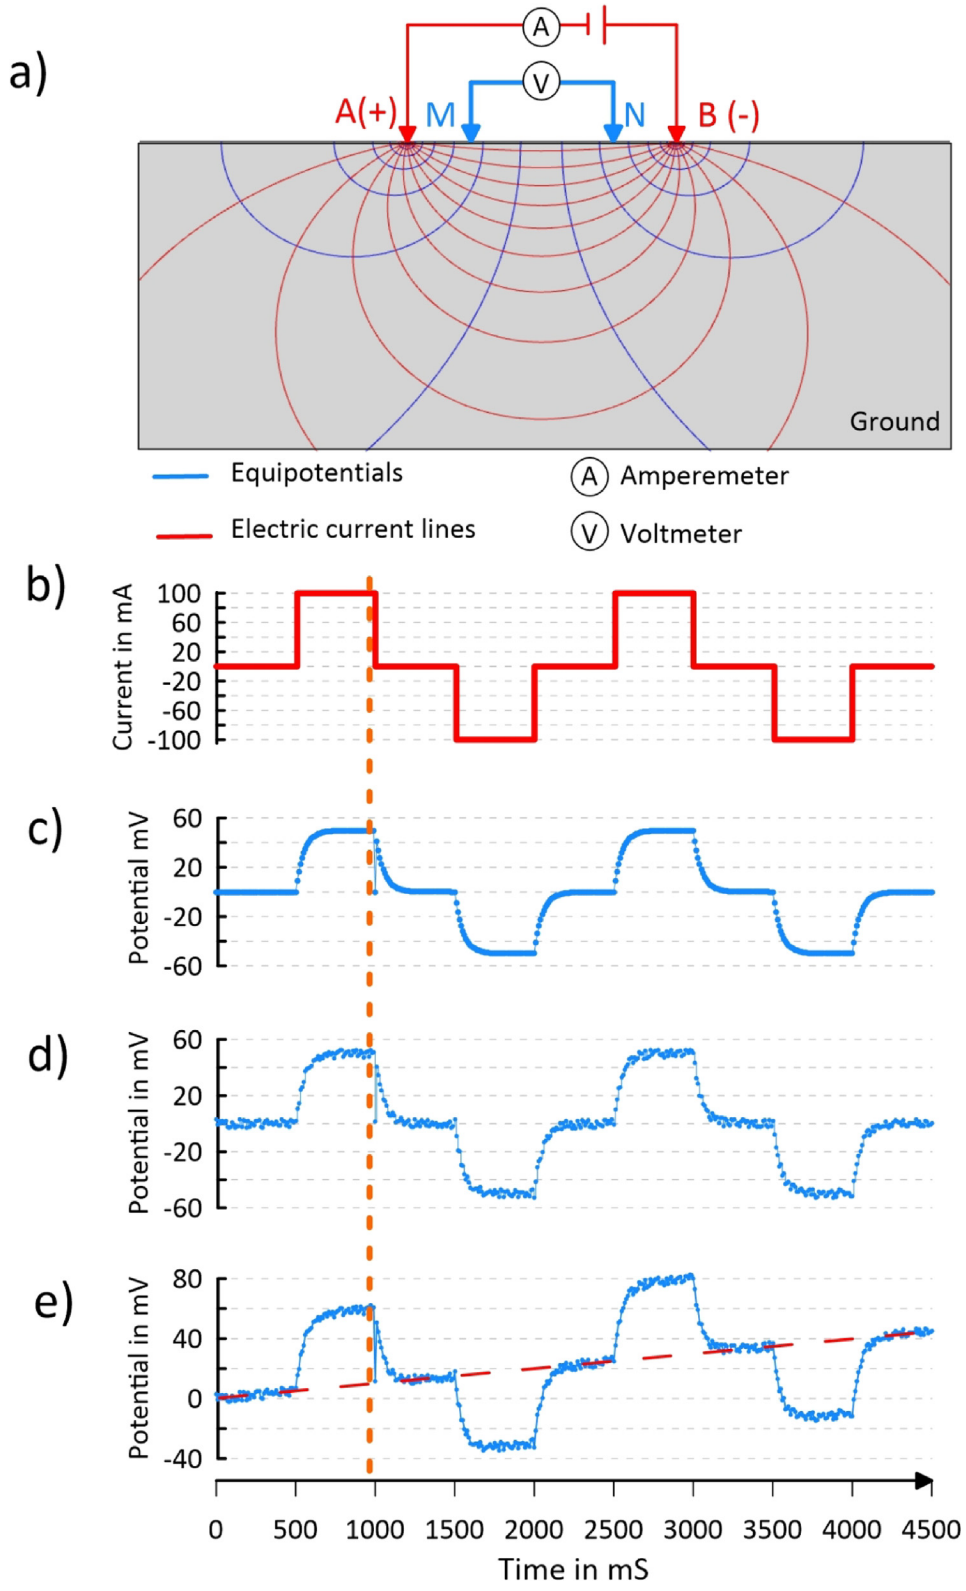

Figure 2.1: Theoretical representation of a four-electrode electrical resistivity measurement: a) current-flow lines and equipotential lines in the subsurface; b) theoretical current injection; c) ideal potential response including the chargeability effect; d) noisy potential measurement; and e) noisy measurement affected by spontaneous potential (SP). Reproduced from [6].

with different geological materials often overlap, and similar measured values may result from different subsurface conditions. The interpretation is therefore inherently non-unique, as discussed in Section 2.1.3.

Induced polarisation provides complementary information by characterising the delayed electrical response of the subsurface. Polarisation may result from electrochemical processes at interfaces between metallic mineral grains and the pore fluid, or from the redistribution of ions around clay particles and within partially restricted pore paths. When an electric field is applied, electrical charges progressively accumulate at these interfaces. When the applied current is interrupted, the charges return towards equilibrium and produce a secondary voltage that decays with time [9, 7].

In time-domain IP, this decay is measured during one or several time windows following current interruption. A commonly used quantity is the apparent chargeability m_a, defined as the average secondary voltage over a selected interval, normalised by the primary voltage:

$$m_a = \frac{1}{V_p(t_2 - t_1)} \int_{t_1}^{t_2} V_s(t)\,\mathrm{d}t, \tag{2.2}$$

where V_p is the primary voltage measured during current injection and $V_s(t)$ is the secondary voltage measured after the current is switched off. With this normalisation, chargeability is dimensionless and is commonly reported in mV/V. When the division by $(t_2 - t_1)$ is omitted, the quantity is referred to as integral chargeability and is commonly expressed in milliseconds [7].

ERT and IP are therefore complementary: resistivity primarily characterises the ability of the subsurface to conduct electrical current, whereas IP characterises its ability to temporarily store and release electrical charge. Their combined measurement may help distinguish materials that exhibit similar resistivities but different interfacial or electrochemical properties. The reliability of both measurements nevertheless depends on the injected-current amplitude, the input impedance and sensitivity of the receiver, electrode contact conditions, electromagnetic interference and acquisition timing [8, 7].

2.1.2 Data acquisition

Conventional ERT-IP acquisition systems are organised around a central resistivity meter connected to a series of electrodes through multiconductor cables (Fig. 2.3). Instruments such as the ABEM Terrameter (Fig. 2.2) combine a current transmitter, several voltage measurement channels, an automatic electrode switching system and a data acquisition unit [10]. For each measurement, two electrodes, denoted A and B, are selected for current injection, while two others, denoted M and N, measure the resulting potential difference. The injected current and measured voltage are recorded before the switching system automatically selects the next quadrupole combination defined by the acquisition protocol. Multichannel instruments can measure several potential dipoles during the same current injection, reducing the total acquisition time.



Figure 2.2: ABEM Terrameter LS 2 data acquisition system. Reproduced from [10].

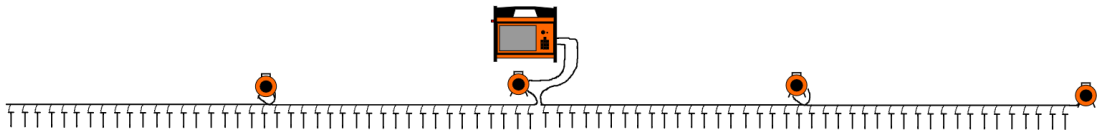

Figure 2.3: Generic electrode arrangement using the ABEM Terrameter LS 2 data acquisition system. Four 100 m cable spools are deployed here to obtain a profile length of approximately 400 m. Reproduced from [10].

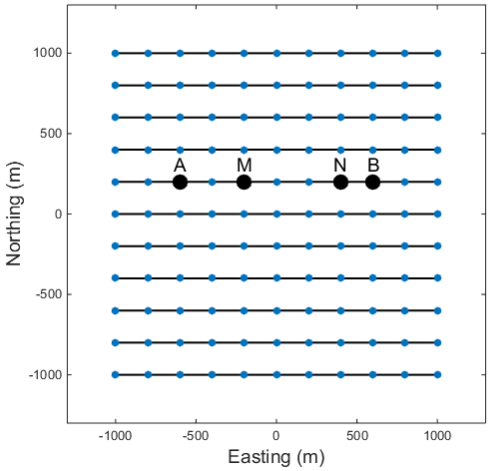

Figure 2.4: Example of a 3D survey composed of several parallel 2D profiles. Reproduced from [11].

The reference case considered in this work is a straight 2D surface profile, in which regularly spaced electrodes are installed along a single line. This geometry is commonly used to reconstruct a vertical resistivity section beneath the profile. When the profile length exceeds the available cable spread, a roll-along procedure can be used, in which part of the electrode layout is progressively moved forward while maintaining an overlap between successive acquisitions [10]. However, ERT acquisition is not restricted to straight surface lines. Three-dimensional surveys can be performed using several parallel profiles or a grid of electrodes (Fig. 2.4), while surface-to-borehole and borehole-to-borehole configurations can improve the sensitivity at depth [11]. Electrodes may also be installed in underground excavations such as tunnels, although this geometry is not considered in the present work [11].

The successive combinations of current and potential electrodes are defined by the selected electrode array. Common configurations include Wenner, Schlumberger, pole-dipole, pole-pole, dipole-dipole and gradient arrays (Fig. 2.5). Their geometry is controlled by the basic electrode spacing a and, for some arrays, by the separation factor n. Increasing these parameters generally increases the investigated depth but reduces the measured voltage and therefore the signal-to-noise ratio [12, 11].

The choice of array affects the spatial sensitivity, resolution, data density, acquisition time and sensitivity to noise. Wenner arrays generally provide strong and comparatively noise-resistant signals, whereas dipole-dipole and pole-dipole configurations can provide better resolution of complex lateral structures at the cost of greater noise sensitivity. Gradient arrays are particularly suitable for multichannel instruments because several potential measurements can be acquired during a single current injection [12].

Before acquisition, electrode contact resistances and cable continuity are generally checked. In the current-injection circuit, the dominant contribution is often the contact resistance at the metal–ground interfaces, where charge transport passes from electronic conduction in the electrodes to ionic conduction in the soil or rock. These contact resistances commonly reach several kilohms, whereas the resistance of the investigated ground volume may be only a few hundred ohms. For each selected quadrupole, the natural potential is

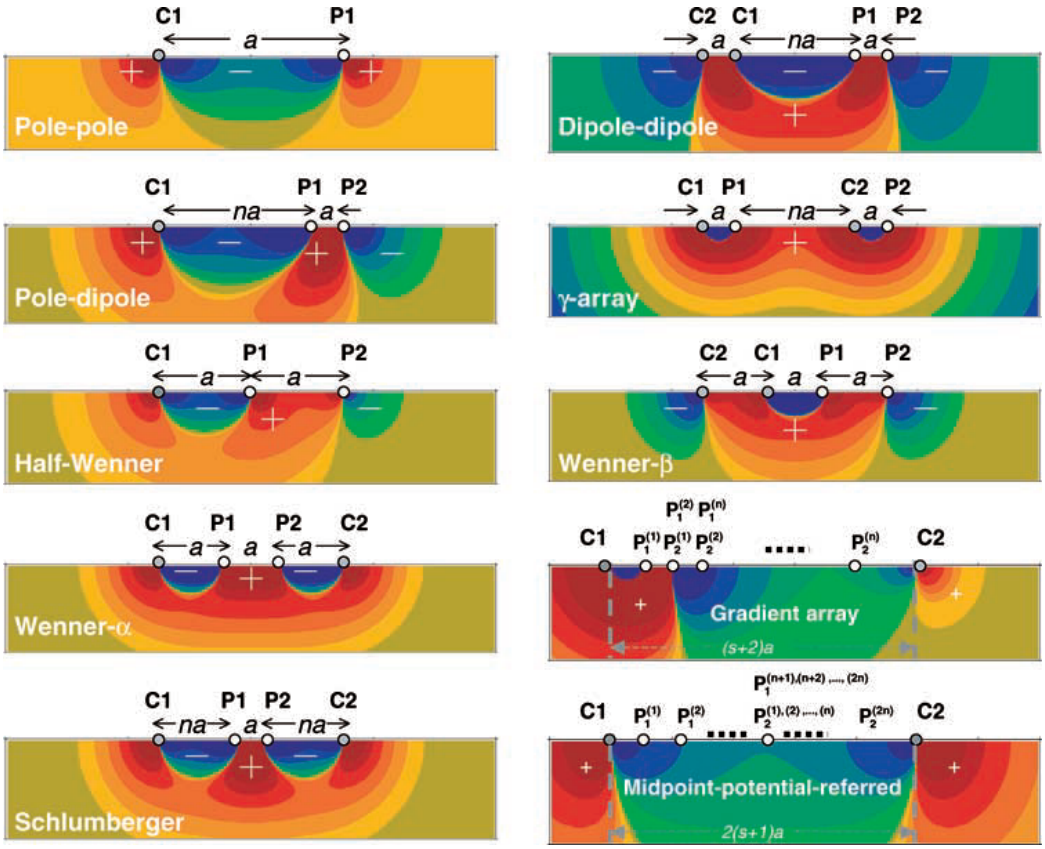

Figure 2.5: Schematic diagrams of electrode arrays and their sensitivity patterns for 2D resistivity surveys. C1 and C2 are positive and negative current electrodes, respectively. P1 and P2 are two potential electrodes. The italic letters *a*, *n* and *s* are the array parameters that control the spacing and maximum separation of the arrays in a multi-electrode layout. Reproduced from [12]

first measured, after which a known current is injected and the resulting voltage is recorded. As illustrated in Fig. 2.1, the current polarity is reversed over successive pulses to reduce offset and electrode-polarisation effects. In time-domain IP, the voltage decay following each current interruption is also recorded over predefined time windows.

Several measurement cycles may be repeated and averaged to improve data stability. The main acquisition parameters include the injected current, pulse duration, sampling rate and IP measurement windows. Data quality may be reduced by poor electrode contact, unstable electrodes, damaged cables, spontaneous potential (SP) variations and electrical noise from nearby infrastructure [13].

2.1.3 Data inversion

The apparent resistivity and IP data measured at the surface do not directly represent the true distribution of electrical properties in the subsurface but rather an averaged measurement over a given subsurface volume. Data inversion aims to estimate this distribution by dividing the subsurface into a numerical mesh composed of individual cells. The electrical properties assigned to these cells form a model vector $\mathbf{m}$. For a given model, the forward problem predicts the data that would be measured at the surface:

$$\mathbf{d}^* = f(\mathbf{m}), \tag{2.3}$$

where $\mathbf{d}^*$ represents the modelled data and f is the forward operator. The inverse problem consists of finding a model whose predicted response is sufficiently close to the field measurements $\mathbf{d}$.

The inversion is performed iteratively. It starts from an initial model, often a homogeneous subsurface, and calculates its predicted response. The difference between the measured and modelled data, called the residual, is then used to modify the model (Fig. 2.6). A sensitivity matrix describes how a change in the electrical property of each model cell would affect the predicted measurements. This process is repeated until the data misfit no longer decreases significantly or reaches a predefined acceptable level. Common ERT inversion algorithms use smoothness-constrained least-squares methods based on Gauss-Newton or quasi-Newton optimisation [14].

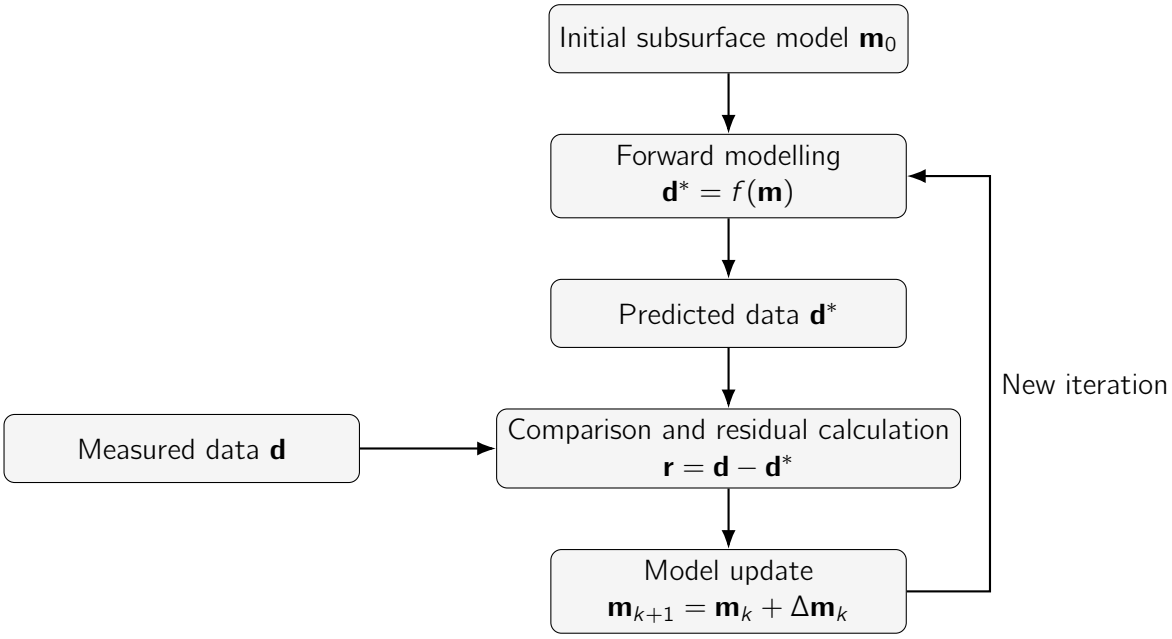

Figure 2.6: Simplified workflow of ERT-IP data inversion. The response predicted from an initial subsurface model is compared with the measured data. The resulting residual is used to update the model until an acceptable fit is obtained.

However, geophysical inversion is an ill-posed problem. The available measurements are finite, indirect and affected by noise, whereas the subsurface model may contain a large number of unknown cell values. Consequently, several different resistivity distributions can produce very similar surface measurements. Small variations or errors in the data may also cause large variations in the reconstructed model, particularly in regions with poor data coverage or low sensitivity. Nguyen et al. showed that decreasing sensitivity with depth can lead to significant discrepancies between the true subsurface properties and those recovered from ERT inversion [15]. A low data misfit therefore does not demonstrate that the obtained model is the unique or exact representation of the subsurface [16, 17].

Regularisation is introduced to stabilise the inversion and select one physically reasonable model among the many possible solutions. The objective function generally combines a data-misfit term Φ_d and a model-regularisation term Φ_m:

$$\Phi(\mathbf{m}) = \Phi_\mathrm{d}(\mathbf{m}) + \lambda\Phi_\mathrm{m}(\mathbf{m}), \tag{2.4}$$

where λ controls the balance between reproducing the measured data and maintaining a stable model. The regularisation term commonly penalises abrupt variations between neighbouring cells or deviations from a reference model. Excessive regularisation produces an overly smooth image, whereas insufficient regularisation may fit measurement noise and generate inversion artifacts. The final result must therefore be interpreted together with the residuals, data coverage, sensitivity and available geological information. In addition, a 2D inversion assumes that variations perpendicular to the profile are limited; significant

three-dimensional structures or topography can otherwise produce artifacts in the reconstructed section [18].

These principles are implemented in several geophysical inversion software packages. One example is pyGIMLi, an open-source Python library providing finite-element forward modelling and a generalised Gauss-Newton inversion framework with flexible regularisation strategies for ERT and other geophysical methods [19]. Nevertheless, inversion lies outside the scope of the present work, which focuses exclusively on the design and improvement of the ERT-IP data acquisition system.

2.1.4 Deep ERT

Most conventional ERT applications investigate the first 0-200 m of the subsurface using a central resistivity meter connected to a multi-electrode array through long multichannel cables and a switching system [2]. Surveys reaching depths between approximately 200 and 500 m can still be performed using an extension of this conventional architecture. However, the increasing cable length, electrode spacing and acquisition time make this approach difficult to apply at greater depths [2, 3].

The term deep electrical resistivity tomography has not always been used consistently in the literature. Earlier studies described investigations deeper than approximately 200 m as deep ERT [20]. More recently, Balasco et al. proposed a threshold investigation depth of 500 m to distinguish deep ERT from conventional and intermediate-depth surveys [2]. This definition is adopted throughout the present work.

Figure 2.7 illustrates the scale of a representative deep-ERT resistivity model. In contrast with conventional ERT profiles, which typically extend over hundreds of metres, deep-ERT surveys can cover several kilometres laterally while imaging structures at depths of several hundred metres to approximately one kilometre.

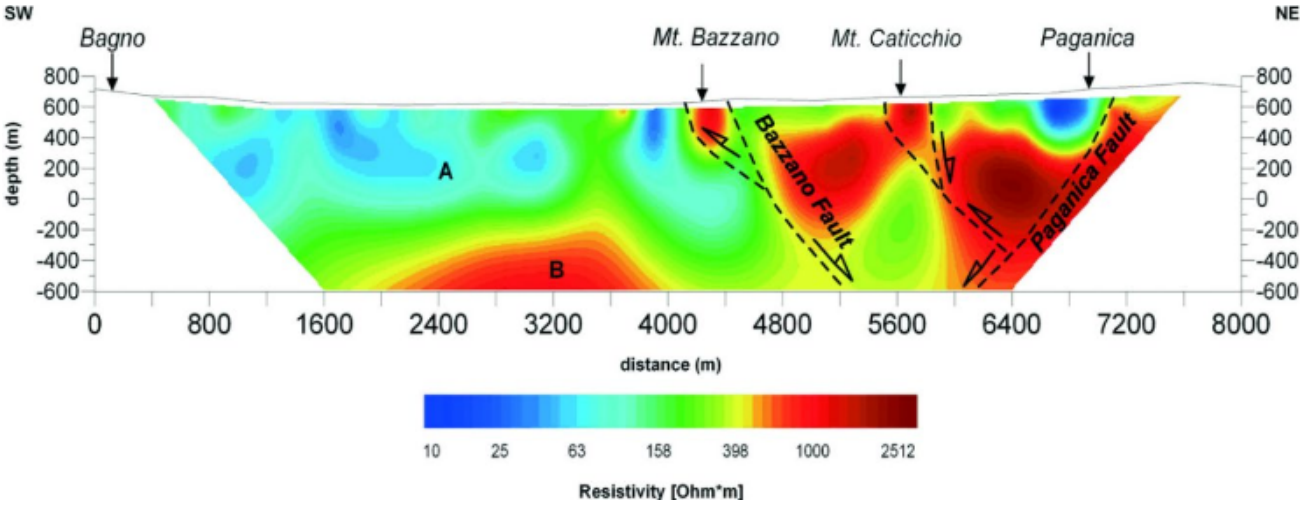

Figure 2.7: Example of a resistivity model obtained from a deep-ERT survey, illustrating the kilometre-scale lateral extent and the large investigation depth achievable with this approach. Reproduced from Balasco et al. [2].

Deep ERT relies on the physical separation of the current injection system from the voltage measurement system [3]. A dipole-dipole configuration is commonly used, with one current dipole, denoted AB, and several receiving dipoles, denoted MN (see Section 2.1.2). The receiving stations are autonomous data logging devices and are distributed along the survey profile. This distributed architecture avoids the need for prohibitively long multichannel cables connecting all electrodes and allows the distance between the transmitting and receiving dipoles to be progressively increased [2, 20, 3].

The transmitting system, composed of an electrical generator and a current measuring device, injects a low-frequency square-wave current into the ground. Its polarity is periodically reversed to limit electrode-polarisation effects. The receiving stations must be synchronised with the transmitter so that the measured

voltage cycles can be associated with the corresponding current polarity during the data processing. In modern systems, this synchronisation can be provided by GNSS timing [2, 3].

The main limitation of deep ERT is the rapid decrease in the measured voltage as the distance between the transmitting and receiving dipoles increases. For a dipole-dipole configuration, the useful voltage signal decreases approximately as $1/r^3$, where r is the distance between the two dipoles [2]. Deep ERT systems must therefore measure signals ranging approximately from 10^{-3} to 10^{-6} V in the presence of natural and anthropogenic electrical noise. Currents between 5 and 20 A are typically injected to improve the signal amplitude, although the achievable current is limited by the available power source and the electrode contact resistance [2].

The recorded voltage can be represented as

$$v(t) = s(t) + n(t), \tag{2.5}$$

where $s(t)$ is the useful periodic response to the injected current and $n(t)$ represents the electrical noise. This noise may originate from telluric currents, power lines, railways, pipelines, industrial equipment or poor electrode contacts. It is generally neither stationary nor Gaussian, making the extraction of the useful signal one of the main challenges of deep ERT [2].

Before the apparent resistivity can be calculated, the voltage records must therefore be processed to recover the useful signal. The usual processing sequence includes the removal of outliers, de-trending, filtering and the estimation of the signal amplitude using stacking, Fourier analysis or periodogram-based methods [2, 20, 3].

Stacking consists of synchronising several recorded cycles with the injected square wave, reversing the sign of the measurements acquired during negative-current half-cycles, and averaging the resulting cycles. The useful response is coherent with the current injection and is therefore reinforced by the averaging process, whereas uncorrelated noise tends to cancel out. Increasing the number of recorded cycles improves the estimation of the useful signal, but also increases the duration of each measurement and of the complete field campaign. Alternatively, Fourier or periodogram analysis can be used to identify the amplitude of the voltage component at the current-injection frequency, even when it is difficult to distinguish directly in the time-domain recording [2, 20, 3].

A representative application was carried out in the High Agri Valley in southern Italy. The survey consisted of a 20.8-km-long profile with an electrode spacing of 400 m and reached an investigation depth of approximately 1 km [20]. A maximum distance of 4 km was used between the transmitting and receiving stations. For each current position, voltage records lasting between 10 and 20 minutes were acquired to provide enough square-wave cycles for signal extraction. The complete campaign produced 351 voltage recordings and required approximately one month of field work with a team of five to six operators [20].

Deep ERT has also been applied to a wide range of large-scale geological problems, including the characterisation of fault-controlled sedimentary basins, geothermal and volcanic systems, groundwater circulation and large landslides [2]. Published examples include the imaging of fault and fluid-bearing structures near the German Continental Deep Drilling Project and the investigation of the Sèchilienne unstable slope [2]. Carrier et al. demonstrated the relevance of the method for geothermal exploration in the Greater Geneva Basin, where a 4.5-km-long profile composed of 25 independent receiving stations was used to investigate structures down to approximately 1 km in urbanised areas affected by strong electromagnetic noise [3]. The resulting resistivity model was compared with reflection seismic, gravity and borehole data, and identified a conductive structure potentially associated with deep groundwater circulation.

A more local application is currently being developed by the University of Liège within the framework of the Einstein Telescope project. In the Val-Dieu area of the Euregio Meuse-Rhine, deep ERT profiles extending over approximately 7 km have been acquired to investigate geological structures down to about 600 m.

These measurements are combined with microgravity data, field observations and existing geological models to improve the structural interpretation of the subsurface [21]. Such kilometre-scale surveys represent the target future application for the OhmPi Watcher, as autonomous receiving units could facilitate the deployment of denser, more flexible and more cost-effective measurement networks.

2.2 Autonomous geoelectrical acquisition systems

In autonomous geoelectrical monitoring, instruments must operate with limited human intervention and under constrained power and communication conditions. The following section reviews the evolution from centralised autonomous monitoring systems to distributed full-wave acquisition architectures.

2.2.1 From manual surveys to autonomous monitoring

Time-lapse ERT was initially performed through repeated manual surveys, in which an operator returned to the site to connect the instrument, configure the acquisition sequence and retrieve the measurements. Although this approach is suitable for occasional surveys, it becomes costly and logistically demanding when measurements must be repeated frequently or over long periods. LaBrecque et al. noted that, for long-term monitoring, field labour may eventually exceed the cost of both the acquisition system and the permanent electrode installation. Manual operation also introduces idle periods between acquisitions and limits the temporal resolution that can be achieved, particularly at remote or difficult-to-access sites [22]. Autonomous monitoring was therefore developed to reduce human intervention while enabling measurements to be collected regularly and under consistent acquisition conditions.

An autonomous geoelectrical monitoring system typically combines permanent electrodes, programmable acquisition, local storage and an independent power supply. Measurements are scheduled automatically, with the instrument returning to a low-power state between surveys. Telemetry may enable remote reconfiguration and data transfer. The ALERT system, for example, combines acquisition of data, communication, quality control, inversion and web-based delivery of results [23]. This improves temporal resolution, but the number of electrode configurations must remain compatible with the rate of the monitored process. If the subsurface changes significantly during the acquisition of a complete data set, measurements corresponding to different states are combined and the reconstructed image may be distorted [23, 24].

Reliable autonomous operation requires more than scheduled measurements. The system must store data safely, monitor measurement quality and power availability, and recover from non-critical faults without waiting for an operator. These requirements become especially important in isolated or harsh environments, where maintenance visits may be infrequent or impossible. Recent autonomous ERT systems developed for permafrost monitoring therefore combine rugged enclosures, low-temperature components, local storage, contact-resistance diagnostics and advanced power management. In one such system, a timer was added to disconnect the resistivity meter between acquisitions because its standby consumption, even if it was low, still represented a substantial fraction of the long-term energy demand [25]. Autonomous geophysical acquisition consequently involves a compromise between measurement frequency, data quality, system complexity, communication capability and available energy.

2.2.2 Distributed full-wave acquisition systems

While the systems described above automate repeated measurements using a central instrument connected to a permanent electrode network, distributed acquisition systems separate current injection from voltage recording by deploying independent receivers throughout the study area. A representative commercial system is the IRIS FullWaver shown on Fig.2.8, which combines an I-FullWaver unit recording the injected current with several autonomous V-FullWaver receivers [5]. Each receiver records the complete voltage

waveform on two dipoles at 100 Hz and synchronises its samples using an internal GNSS receiver with a PPS signal. The raw (full-wave) time series, covering the successive current pulses and current-off intervals, are stored locally and processed afterwards to calculate resistivity, chargeability and signal-quality parameters, while enabling filtering and stacking to improve the signal-to-noise ratio [5, 3]. Because the receivers are independent, their positions are not constrained by multichannel cables or a predefined electrode array. Carrier et al. demonstrated this architecture using 25 receiving stations along a 4.5-km profile, substantially reducing the amount and weight of cable required compared with a conventional system [3].



(a) V-FullWaver receiver.



(b) V-FullWaver during field deployment.

Figure 2.8: IRIS Instruments V-FullWaver autonomous receiver. Reproduced from [5].

The OhmPi Watcher applies a similar distributed full-wave acquisition principle through a lower-cost and open-source implementation [1]. The system therefore provides an alternative for deploying multiple independent receivers. Its origin, detailed architecture and operational limitations are presented in Section 2.3.

2.3 The OhmPi platform and Watcher prototype

The previous section positioned the OhmPi Watcher among existing distributed acquisition systems. This section describes the open-source platform from which it was derived and the prototype that formed the starting point of this work.

2.3.1 The OhmPi open-source platform

OhmPi was developed as a low-cost and open-source alternative to commercial resistivity meters, such as the ABEM Terrameter (Fig. 2.2), whose cost and limited adaptability can restrict their use in research, education and dedicated monitoring applications. The original system combines current injection, voltage and current measurement, automatic electrode switching, time stamping and data storage. Its hardware and software design files, construction instructions and bill of materials were published to facilitate replication and modification by other users [6].

The first OhmPi version published in 2020 uses a Raspberry Pi 3, which manages the current-injection sequence, data acquisition and recording. Current is supplied by a 12 V battery and its polarity is controlled using mechanical relays. The injected current and the potential difference between the receiving electrodes are measured using an ADS1115 16-bit analogue-to-digital converter associated with an analogue signal conditioning circuit shown in Fig. 2.9. Four relay-based multiplexers allow the current and potential terminals, denoted A, B, M and N, to be connected automatically to a network of up to 32 electrodes. This architecture therefore forms a complete small-scale resistivity meter capable of successively measuring several predefined sequences of quadrupoles [6].

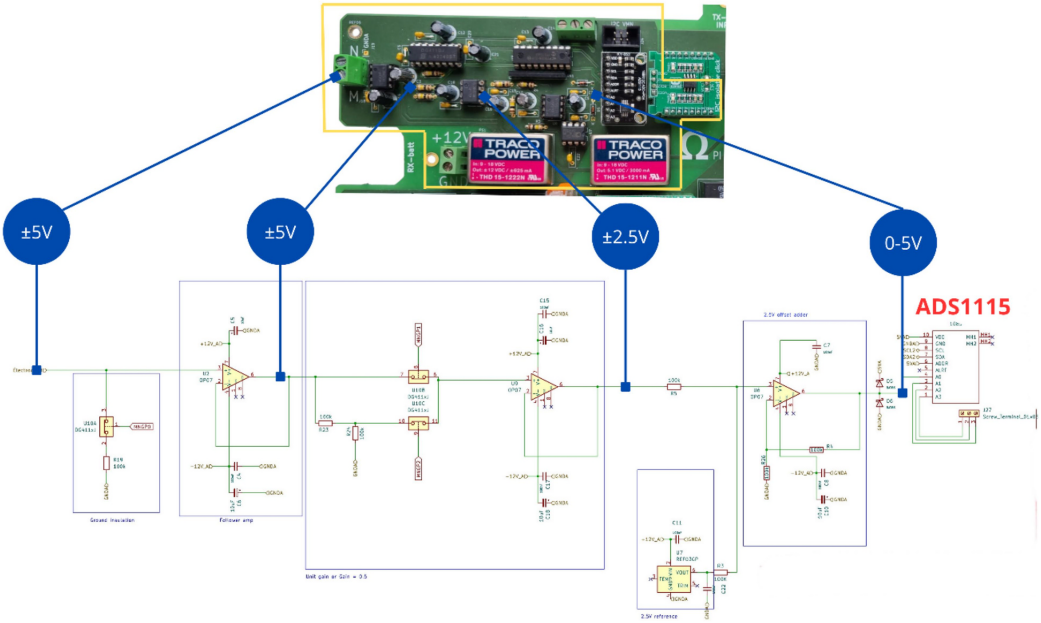

Figure 2.9: Electronic signal conditioning circuit for Vmn from M/N front-end to the ADS1115 on the OhmPi. This measurement architecture is reused in the OhmPi Watcher. Reproduced from [26].

Since its initial publication, the OhmPi has undergone a substantial hardware and software redesign. The Raspberry Pi 3 used in the original system was replaced by a Raspberry Pi 4, while the commercial relay modules and extensive point-to-point wiring were replaced by dedicated measurement and multiplexer boards organised within a configurable modular architecture. The new generation also supports higher and adjustable injection voltages, increases the input impedance of the potential-measurement channel, and provides more flexible electrode multiplexing. In parallel, the original Python acquisition script was replaced by an object-oriented software architecture supporting full-waveform acquisition, hardware configuration, system logging, MQTT communication and web-based remote operation [6, 26].

Laboratory validation of the original OhmPi v2020 system demonstrated that it could measure soil resistances between approximately 0.2 and 1000 Ω for electrode contact resistances ranging from 100 to 5000 Ω. Field measurements also showed good repeatability and close agreement with a commercial resistivity meter. Beyond its measurement performance, the main relevance of OhmPi to the present work lies in its open-source hardware architecture, which provided a basis for adapting the measurement electronics to a distributed acquisition concept [6].

2.3.2 The original OhmPi Watcher prototype

The OhmPi Watcher was developed by adapting the OhmPi measurement electronics to the specific requirements of deep-ERT surveys. Whereas the original OhmPi is a complete resistivity meter designed to inject current, switch between electrode combinations and acquire successive quadrupole measurements, the Watcher is intended to operate only as an autonomous distributed receiver. Each unit remains connected to two receiving dipoles and continuously records the corresponding ground-potential variations while current is injected by a physically separate transmitting system. By deploying many Watchers along a kilometre-scale profile, deep-ERT measurements can therefore be performed without the long multiconductor cables required to connect all receiving electrodes to a central resistivity meter [1].

The prototype (see Fig. 2.10) is built around an Adafruit Feather M0 microcontroller programmed in the Arduino environment and several Adafruit FeatherWing expansion boards. Its analogue measurement chain, adapted from the OhmPi platform, conditions the signals from the two receiving electrodes before digitisation. The resulting voltage time series are stored locally on an SD card for subsequent processing.

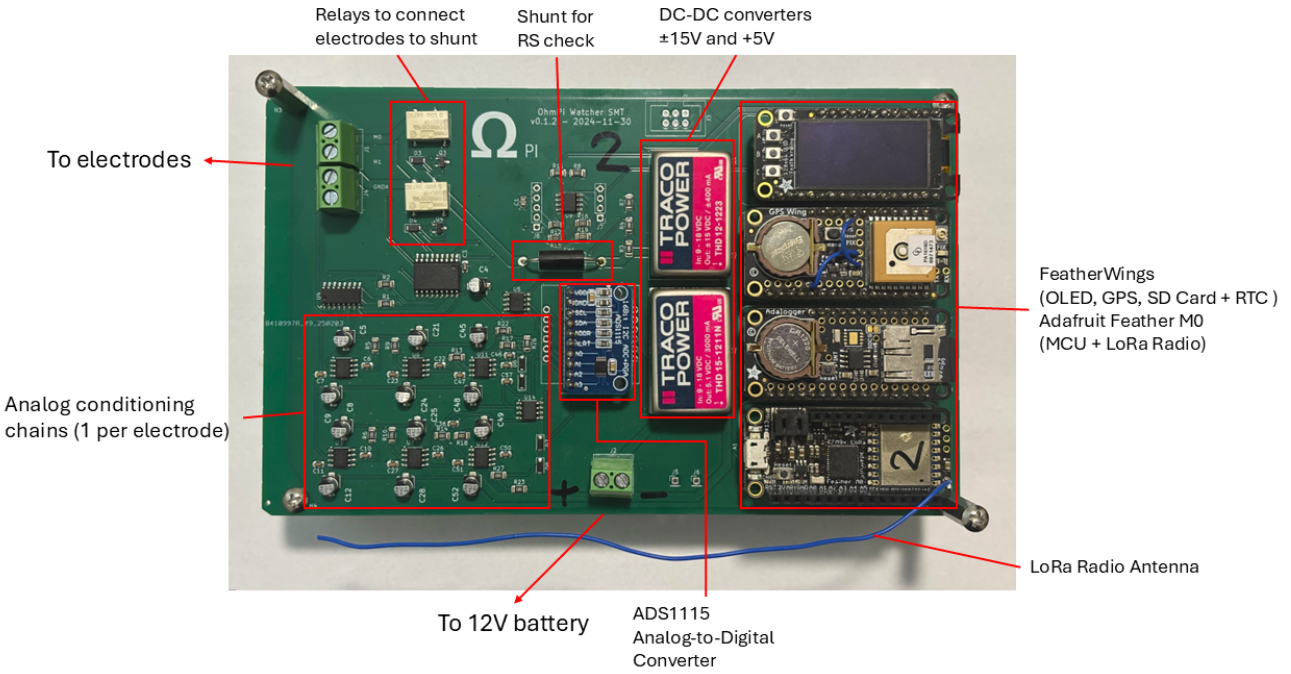

Figure 2.10: Annotated view of the OhmPi Watcher v0.1.2, showing its main hardware components.

Since the transmitter and receivers operate independently, a GNSS receiver provides the common time reference required to associate each recorded sample with the corresponding phase and polarity of the injected-current waveform (Section 2.5).

LoRa communication is used for low-data-rate exchanges such as remote control and status information, rather than for transferring the complete measurement time series, which would be incompatible with the limited throughput and energy budget of the radio link. The prototype also includes an OLED display for local status information and a contact-resistance test for checking the quality of the electrode-to-ground connection.

The different subsystems are powered from a 12 V, 12 Ah $LiFePO_4$ battery. A first DC–DC converter generates the ± 15 V rails required by the analogue signal-conditioning circuitry and by the analogue switches controlling whether the receiving electrodes are connected to ground or to the measurement chain. A second converter provides a regulated 5 V rail for the Feather-based control electronics and the low-voltage peripheral boards, including the ADC, GNSS receiver, data-logging and display subsystems. A built-in regulator on the Adafruit Feather M0 board generates 3.3 V where required.

The prototype was designed to combine continuous measurement, accurate timing, local data storage, remote interaction and basic field diagnostics over several days of autonomous operation. Figure 2.11 summarises these principal functional blocks and their interactions.

A comparison between measurements obtained with the original Watcher and the commercial V-FullWaver is shown in Fig. 2.12. Both systems reproduce the same overall voltage evolution, supporting the ability of the Watcher prototype to capture the useful geoelectrical signal. The Watcher recording nevertheless exhibits larger high-frequency fluctuations than the V-FullWaver measurement. One possible contributor to this difference is the absence of a dedicated analogue low-pass filter in the original Watcher acquisition chain, as discussed further in Section 3.3.2.

Although the prototype demonstrated the feasibility of an open-source and autonomous distributed receiver, its architecture was not yet optimised for the deployment of a large network of field units. In particular, its power consumption and wireless communication reliability required further investigation. These limitations form the starting point of the present work.

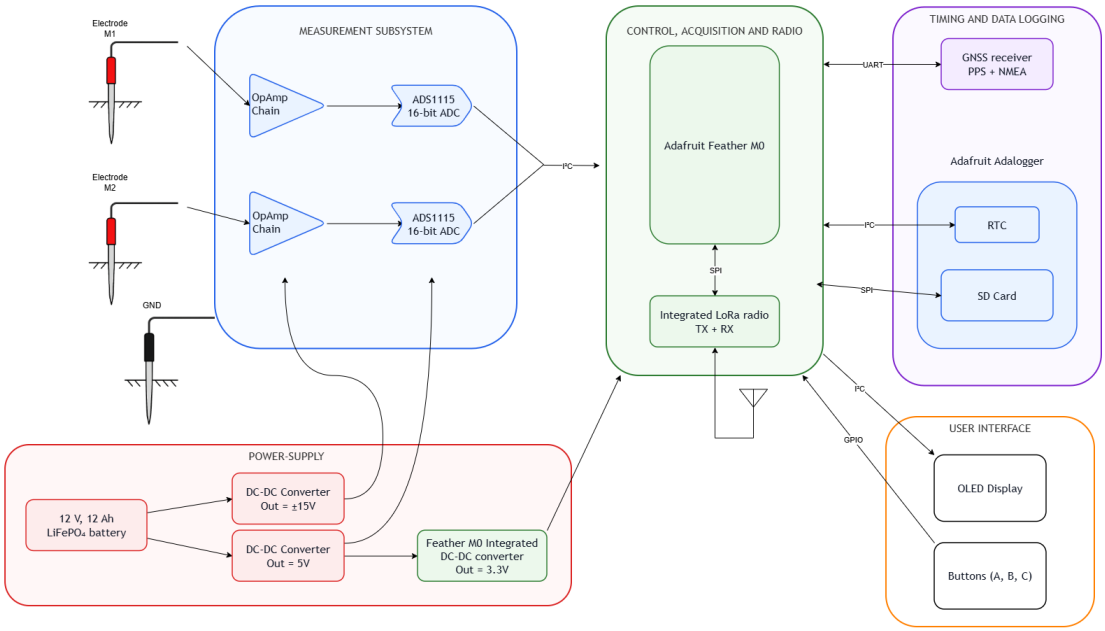

Figure 2.11: Functional architecture of the original OhmPi Watcher prototype. The unit records the voltage across two receiving dipoles, synchronises the measurements using GNSS timing, stores the raw data locally and exchanges commands and status information through a LoRa radio link.

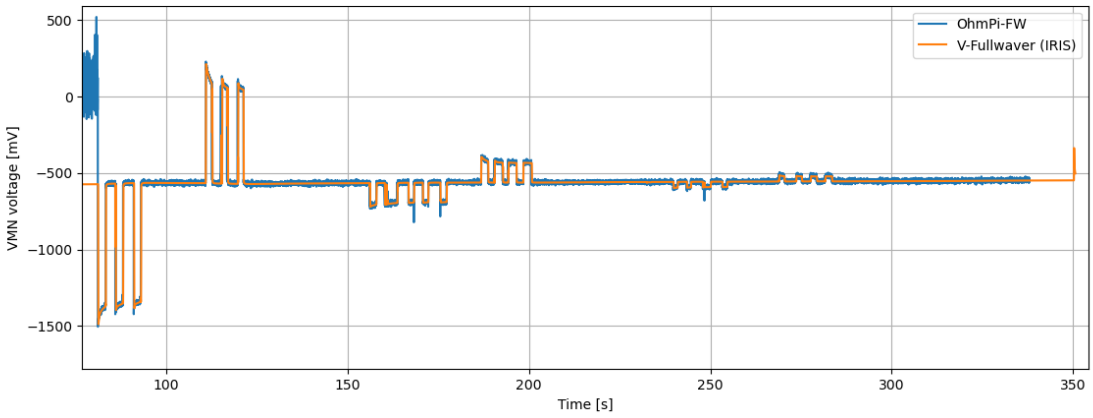

Figure 2.12: Comparison of voltage measurements obtained with the original OhmPi Watcher and the commercial V-FullWaver. Both instruments reproduce the same overall signal evolution, while the Watcher recording exhibits larger high-frequency fluctuations. Provided by Hadrien Michel, 2026.

2.4 Low-power design for autonomous sensor nodes

Autonomous wireless sensor nodes generally combine four main subsystems: sensing and analogue-to-digital conversion, data processing and storage, wireless communication, and power supply. Additional functions such as positioning, time synchronisation or local diagnostics may also be required depending on the application. Since such nodes are often deployed in large numbers and must operate unattended, their individual energy consumption directly affects the size, weight and cost of the power source, as well as the maintenance requirements of the complete network [27, 28]. These considerations are particularly important for distributed geophysical acquisition, where tens or hundreds of instruments may be installed over several kilometres in locations that are difficult to access.

No single low-power strategy is suitable for every sensor network. Energy optimisation must instead account for the requirements of the intended application, including sampling frequency, measurement accuracy, communication latency, coverage and robustness. Reducing the activity of a node may extend its lifetime, but it may also introduce missing measurements, delayed commands or temporary loss of network connectivity. Consequently, energy efficiency should be considered as a trade-off with the quality and availability of the service provided by the node, rather than as an independent design objective [29]. In a scientific instrument, this trade-off is further constrained by the need to preserve the fidelity, timing and completeness of the acquired data.

2.4.1 Power budgeting and operating states

A first step in low-power design is therefore to establish a power budget that distinguishes the different operating states of the system. If the node operates in n states, its average power consumption can be approximated as

$$P_{\text{avg}} = \sum_{k=1}^{n} d_k P_k, \tag{2.6}$$

where P_k is the power consumed in state k and d_k is the fraction of time spent in that state, with $\sum_{k=1}^{n} d_k = 1$. This representation shows that a component with a high active power may have a limited contribution on the overall energy budget if it is activated only briefly. Conversely, a small standby current can become significant when it is drawn continuously over a multi-day deployment. The energy and time associated with transitions between operating states must also be considered, since frequent transitions can offset part of the expected savings [28].

The contribution of each subsystem depends on the node architecture and its operating conditions. Wireless communication is often a major energy load, particularly when the receiver remains continuously active. However, in continuous or high-rate acquisition systems, the sensing chain, analogue front end, ADC and local storage may consume a comparable or even larger fraction of the available energy. A representative power budget must therefore include the complete acquisition chain together with the processor, radio and power-conversion stages [28, 30]. For multi-day autonomous operation, low-power design must consequently consider active and standby consumption, the duration of each operating state and the energy associated with transitions between states.

2.4.2 Dynamic power management and duty cycling

Dynamic power management reduces energy consumption by placing the processor or individual peripherals in low-power states whenever their functions are not required. For wireless communication, duty cycling reduces idle listening by alternating the radio between active and sleep periods. The achievable savings depend on the sleep current, wake-up time and transition energy of the selected hardware. They also

depend on the communication protocol: reducing the fraction of time during which the receiver is active decreases energy consumption, but increases the probability that an incoming message must wait until the next listening period. In a multi-hop network, this additional latency may accumulate over successive relays. Synchronised duty-cycling schemes can reduce idle listening and coordinate communication between neighbouring nodes, but require clock synchronisation, guard intervals and additional control exchanges [30, 31]. Duty cycling must therefore balance energy consumption against communication latency and network availability.

The Watcher imposes an additional constraint because the electrical potential must be sampled continuously at 100 Hz during scheduled deep-ERT acquisition periods. The complete waveform, including current reversals and current-off intervals, must be preserved for subsequent processing. Reducing the sampling rate or activating the measurement chain only in response to detected events could lower energy consumption, but could also remove information required for later analysis. The analogue front end, ADC, microcontroller and local storage must therefore remain operational throughout each acquisition period. Energy savings must consequently be sought primarily outside measurement periods and in auxiliary functions whose continuous operation is not required.

2.4.3 Power-supply efficiency and hardware integration

The power-supply architecture also influences the achievable autonomy. DC–DC converters introduce conversion losses when supplying an active load and may continue to draw a quiescent current even when the downstream circuitry is inactive. Converter efficiency at the actual operating point is therefore more relevant than peak efficiency alone, particularly when the converter operates most of the time at light load [32, 33]. Shutdown current and restart time must also be considered when a converter is intended to be switched dynamically [28].

The original Watcher provides a practical example of these considerations. Its 12 V battery supplies separate DC–DC converters for the low-voltage electronics and the $\pm$15 V analogue domain (see Fig. 2.11). Although these converters provide sufficient power during active operation, their no-load and shutdown consumption remain relevant during the long inactive periods encountered in autonomous deployment. This illustrates why converter selection should consider not only nominal efficiency and output-power capability, but also quiescent consumption and the possibility of disconnecting inactive power domains [32, 33].

Independent power control of functional subsystems can provide greater savings than placing only the main processor in a low-power state. Peripheral circuits, sensors and power converters may otherwise continue to consume energy even when the processor is sleeping. Dividing the system into independently controllable power domains therefore allows functions that are not continuously required to be disconnected while maintaining the subsystems needed for sensing, timekeeping or communication. The achievable benefit depends on the standby consumption of each domain and on the frequency with which it must be reactivated [30, 25].

Hardware integration is another important design factor. Development boards facilitate prototyping but include regulators, interface circuits and peripheral components that were not necessarily selected for minimum standby consumption. When several such boards are combined, their individual quiescent currents and conversion losses accumulate. Replacing such development boards with a dedicated circuit board can reduce the number of continuously powered interfaces, eliminate unnecessary conversion stages and allow each functional block to be disconnected explicitly [6, 26]. However, this higher level of integration also requires substantially more design, validation and maintenance work and may make the system more difficult to understand, troubleshoot and modify. This can conflict with the open-source objectives of accessibility, repairability and ease of adaptation. The final architecture must therefore balance energy efficiency and compactness against modularity, transparency and future maintainability.

2.4.4 Energy harvesting and application constraints

Energy harvesting can extend the operating lifetime of autonomous sensor nodes. Solar energy is particularly attractive for outdoor deployments, but the available power varies with illumination, panel orientation and surrounding obstacles. Raghunathan et al. demonstrated with the Heliomote platform that photovoltaic harvesting can enable near-perpetual sensor operation, provided that energy collection, storage and node activity are managed according to the local solar conditions [34]. Other sources may also be considered depending on the deployment environment. Torah et al., for example, developed a vibration-powered sensor that adapted its measurement and transmission interval to the available harvested energy [35].

For the Watcher, however, solar harvesting cannot be applied uniformly along a deep-ERT profile. Receivers installed in open fields may benefit from sufficient sunlight, whereas those deployed in forests or shaded areas would receive substantially less energy. Moreover, the Watcher is inherently installed close to ground level, since it must remain near the receiving electrodes to avoid unnecessary cable lengths. This makes the equipment more exposed to theft, accidental damage and wildlife interference. Adding a sufficiently large solar panel would further increase the size and visibility of the installation, whereas the system should remain as compact and discreet as possible.

An additional limitation is that solar charging may interfere with the sensitive voltage measurements. A photovoltaic charging circuit generally relies on switching power conversion, which can introduce noise. A mitigation measure would include disabling the charger during acquisition periods. However, restricting solar charging to periods without measurements would reduce the available charging time and may substantially limit the benefit of solar harvesting for long-duration daytime acquisitions.

Low-power operation of the Watcher must therefore be achieved through a combination of efficient hardware, switchable power domains and state-based software control. The measurement chain and SD storage must remain active whenever resistivity data are being acquired, whereas the GNSS receiver, display, electrode-contact test and radio can be activated according to their respective operational requirements. Such an application-specific approach reduces battery size and deployment cost without compromising the continuity or quality of the voltage records. Among the functions that must nevertheless retain sufficient accuracy during intermittent operation is the timing subsystem, which is discussed in Section 2.5.

2.5 Time synchronisation in distributed measurement systems

Distributed deep-ERT acquisition requires measurements recorded by physically separated receivers and by the current-injection system to be placed on a common time axis. This section first introduces the main sources of clock error and possible synchronisation strategies, before describing the complementary roles of GNSS timing, the pulse-per-second (PPS) signal, the RTC and the microcontroller local clock in the Watcher.

2.5.1 Clock errors and synchronisation requirements

In a distributed measurement system, each acquisition unit operates with its own processor and local clock. Unlike a centralised instrument, the different nodes do not inherently share a common time base. Synchronisation is therefore required whenever measurements recorded at different locations must be compared or combined. This is particularly important for deep-ERT, where voltage time series recorded by independent receivers must be associated with the current waveform generated by a physically separate transmitter. A timing error may shift the recorded signal relative to the current reversals, leading to an incorrect identification of the pulse polarity or of the current-off intervals used for induced-polarisation analysis. More generally, time synchronisation is considered a fundamental service in wireless sensor networks because data fusion and event ordering require timestamps expressed on a common timescale [36, 37].

As a first approximation, the local time reported by node i can be represented as

$$C_i(t) = \alpha_i t + \beta_i, \tag{2.7}$$

where t is the reference time and $C_i(t)$ is the time indicated by the node. The parameter β_i represents the clock offset at the chosen time origin, while α_i describes the rate at which the local clock advances. For an ideal clock, $\alpha_i = 1$. Any deviation from this value corresponds to a relative frequency error, commonly referred to as clock skew, and causes the local time to progressively diverge from the reference time. Even if two clocks are initially aligned, small differences between their oscillator frequencies cause their offsets to increase progressively. Variations in oscillator frequency with temperature, supply conditions and ageing introduce an additional time-dependent drift. Consequently, synchronising the nodes only once at the beginning of a multi-day acquisition is generally insufficient. Sundararaman et al. note that a frequency error of only 0.001% already corresponds to approximately one second of timing error per day [37].

2.5.2 Synchronisation strategies for distributed nodes

Several approaches can be used to establish and maintain a common time reference between distributed acquisition nodes. One possible approach is to exchange timestamped messages through a wireless network. By comparing the transmission and reception times of these messages, each node can estimate the offset between its local clock and that of a reference node. However, the accuracy of this method is affected by variable delays introduced by radio transmission, packet processing and routing. These delays become even more difficult to estimate in a multi-hop network, where a timing message may pass through several intermediate nodes before reaching its destination. In addition, periodic synchronisation exchanges increase radio activity and energy consumption, and they require the communication network to remain operational at all times.

An alternative is to equip each acquisition node with its own external time reference. Rather than synchronising directly with one another, all nodes independently relate their local clocks to the same global timescale. The Watcher follows this principle by using a GNSS receiver at each station. This approach avoids timing uncertainties associated with the LoRa communication network and allows each receiver to establish its timing reference independently. However, it introduces additional power consumption and requires sufficient satellite visibility. The timing information delivered by the GNSS receiver and its combination with the local clocks are described below.

2.5.3 GNSS timing, PPS and RTC-based local timekeeping

In addition to determining position, a Global Navigation Satellite System (GNSS) receiver can provide access to a common and accurate timescale. During position determination, the receiver estimates the offset of its internal clock relative to the satellite time reference. Low-cost receivers generally expose this timing information through two complementary outputs. First, serial NMEA messages (Fig. 2.13) provide the absolute date and UTC time, together with position and fix-quality information. Their arrival time is not itself a precise timing marker because the messages must be generated, transmitted through a UART interface and parsed by the microcontroller. Second, a pulse-per-second (PPS) output produces an electrical edge at the boundary between two consecutive seconds. The PPS signal therefore provides a precise second marker, but does not by itself indicate which UTC second the pulse represents. Combining the absolute time contained in the NMEA messages with the precise boundary indicated by the PPS signal identifies both the instant of the second transition and the corresponding absolute UTC time [38].

The PPS signal marks only one instant per second, whereas the Watcher records voltage samples every 10 ms. The timing of the samples between two successive PPS edges must therefore be derived from the microcontroller's local oscillator and timer. Successive PPS observations can in principle be used to estimate both the phase offset of this local clock and its relative frequency error. Fuhrmann et al. showed

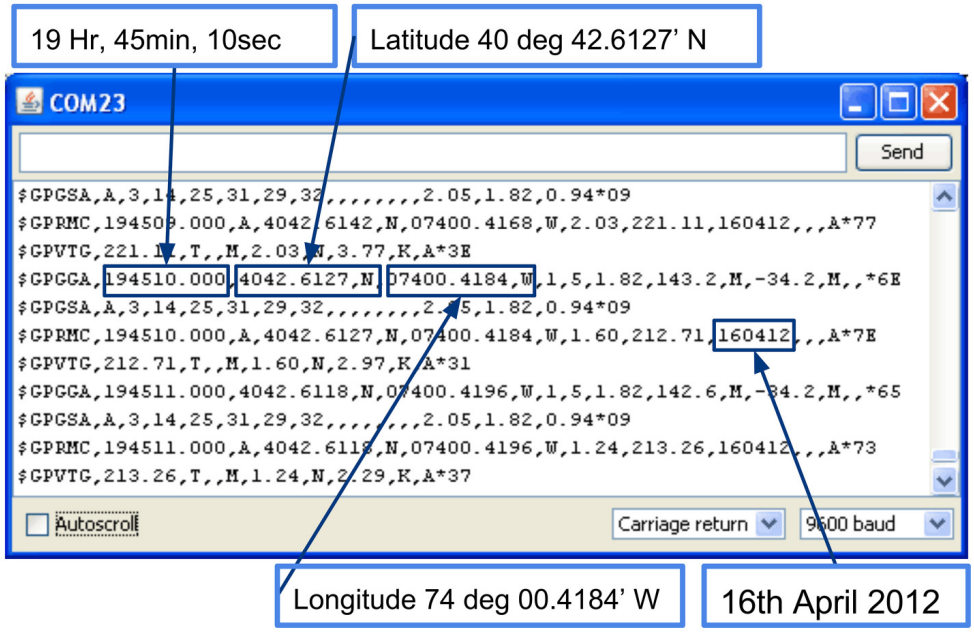

Figure 2.13: Examples of NMEA sentences. Reproduced from [38].

that maintaining synchronisation requires the correction of both quantities: correcting the offset aligns the clocks at a particular instant, while correcting the frequency error limits their subsequent divergence. Their results also illustrate that practical PPS signals may exhibit jitter or occasional outliers and that local oscillators remain affected by environmental variations and ageing [39].

Because the GNSS receiver may be temporarily unavailable or powered down to reduce energy consumption, a local time reference is also required between successive GNSS observations. A real-time clock (RTC) is a low-power integrated circuit that maintains calendar time independently of the main microcontroller. In the Watcher, this function is provided by a PCF8523 RTC integrated into the data-logging FeatherWing. The component uses a 32.768 kHz quartz oscillator and maintains calendar information from seconds to years. It communicates with the microcontroller through the I^2C bus and includes alarm, timer, oscillator-stop-detection and frequency-adjustment functions. Its backup supply allows timekeeping to continue when the main electronics are unpowered [40, 41]. Such an RTC is well suited to maintaining the date, controlling daily operating schedules and providing a fallback time reference. However, its one-second calendar resolution and the frequency variations of its quartz oscillator make it unsuitable as the sole reference for aligning samples acquired every 10 ms.

These complementary timing functions are combined in the Watcher using the PA1616D GNSS receiver, the PCF8523 RTC and the Feather M0 internal millisecond counter. Their practical integration and the modifications introduced in the redesigned firmware are presented in Section 4.3.

2.6 Embedded data acquisition and local storage

The common time reference established in the previous section is only useful if the corresponding voltage measurements are acquired and stored without interruption. Preserving these measurements requires a continuous data path from the receiving electrodes to the local storage medium. In the Watcher, the electrode voltages are first protected and adapted by an analogue front-end, converted into digital values by a precision ADC, timestamped and formatted by the microcontroller, and finally written to an SD card. The complete chain must sustain the required 100 Hz sampling rate while limiting noise, input loading, power consumption and the risk of data loss.

2.6.1 Analogue signal conditioning and digitisation

The Watcher is designed to measure receiving-electrode voltages within an expected range of ± 5 V. Since the ADS1115 analogue-to-digital converter (ADC) is powered from a 5 V supply and its analogue inputs must remain between ground and V_{DD}, the raw electrode voltage cannot be digitised directly. The analogue front end therefore uses three OP07 operational amplifiers to provide buffering, signal scaling and a 2.5 V level shift. Its ideal transfer function is

$$V_{ADC} = \frac{V_{electrode}}{2} + 2.5\ \text{V}, \qquad -5\ \text{V} \leq V_{electrode} \leq 5\ \text{V}. \tag{2.8}$$

The electrode voltage ranging from -5 to $+5$ V is consequently mapped onto the 0–5 V ADC input range. A diode-clamping network additionally protects the converter against unexpected voltage excursions. The OP07 provides low input offset voltage, low input bias current and high common-mode rejection, which are relevant for low-frequency precision measurements. However, the cumulative offset, drift, noise and power consumption of the three amplifier stages must also be considered [42, 43].

The original Watcher does not include a dedicated analogue low-pass filter before the ADC. Out-of-band interference can therefore reach the converter and, after sampling, contribute to components within the measurement band. An analogue low-pass filter can reduce this effect, provided that its cut-off frequency remains sufficiently high to preserve the voltage transitions and IP decay response. The design and experimental validation of such a filter for the redesigned Watcher are presented in Section 3.3.2.

Several ADC architectures can be considered for this application. Flash and pipeline ADCs are designed for much higher bandwidths and would provide little benefit at a 100 Hz acquisition rate. Integrating converters provide strong rejection of periodic interference but are generally too slow to preserve the complete waveform. Successive-approximation register (SAR) ADCs provide deterministic conversion timing and low latency, making them attractive for rapid multiplexing and simultaneous-channel systems. Their achievable precision nevertheless depends strongly on the external reference, input driver and analogue filtering. Delta-sigma ($\Delta\Sigma$) converters are particularly suited to the low-bandwidth, high-resolution measurements required by the Watcher because they trade bandwidth for improved noise performance and include digital filtering [44, 45].

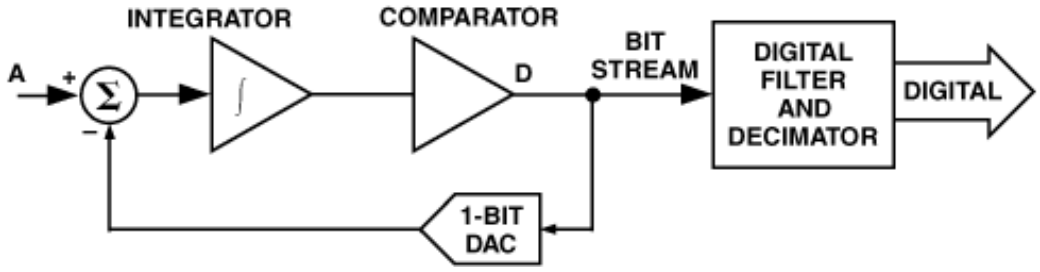

Figure 2.14: Delta-sigma ADC architecture. Reproduced from [44].

A delta-sigma ADC samples the input internally at a frequency much higher than its final output data rate. As illustrated in Figure 2.14, the analogue input is first compared with the feedback signal generated by a 1-bit DAC. The resulting error is accumulated by the integrator and quantised by the comparator, producing a high-frequency 1-bit stream. This bitstream is fed back through the DAC, forming a closed loop in which its time-averaged value tracks the input signal. At the same time, the modulator shapes a large part of the quantisation noise towards higher frequencies. A digital low-pass filter then suppresses this out-of-band noise, while the decimator reduces the sample rate and produces the final multi-bit conversion result [44].

The relationship between the internal modulator frequency and the output data rate is expressed by the oversampling ratio

$$\mathrm{OSR} = \frac{f_{\mathrm{mod}}}{f_{\mathrm{data}}}, \tag{2.9}$$

where f_{mod} is the modulator frequency and f_{data} is the output data rate. Increasing the OSR generally improves noise performance because more modulator samples contribute to each output result, but it also reduces the available signal bandwidth. The nominal bit depth must therefore not be interpreted as the effective measurement resolution, which also depends on the selected data rate, gain, voltage-reference stability and analogue-front-end noise [43, 46].

Digital filtering inside a delta-sigma converter does not eliminate the need for analogue input filtering. Although oversampling and the internal digital filter attenuate a large part of the out-of-band content, the filter response repeats around the modulator sampling frequency and its multiples. Frequency components located around these regions may therefore alias into the measurement band if they are not sufficiently attenuated before conversion. More generally, aliasing causes frequency components above the Nyquist frequency to be folded into the sampled frequency band, where they can no longer be distinguished from genuine in-band signals. An external analogue low-pass filter is therefore recommended when significant out-of-band interference is expected, with its bandwidth selected according to the useful signal spectrum [43].

The ADS1115 used in the original Watcher combines a delta-sigma converter, an input multiplexer, a programmable gain amplifier, an internal reference and an I^2C interface. When supplied at 5 V and configured with the ± 6.144 V full-scale range, its conversion range is compatible with the conditioned 0–5 V signal, although the absolute voltage applied to each input pin must still remain between ground and V_{DD} [43].

Its analogue inputs use a switched-capacitor sampling structure, which periodically draws small charging currents from the preceding stage. This dynamic loading can be represented by an effective input impedance whose value depends on the selected PGA range. For the ± 6.144 V range, the typical differential input impedance is approximately 22 MΩ [47, 43]. The final OP07 stage of the original Watcher therefore also acts as a low-impedance driver for the converter. Furthermore, although the ADS1115 can measure a negative differential voltage, neither individual input pin may be driven below ground. This constraint explains the level shift and external clamping network used in the original acquisition chain [43].

For a multiplexed acquisition system, ADC suitability cannot be assessed from nominal resolution alone. The relevant characteristics include the usable input range, effective resolution at the required data rate, conversion and multiplexing speed, source loading, power consumption and interaction with the analogue front end. These criteria are used in Section 3.3.4 to compare candidate converters and define the acquisition architecture of the redesigned Watcher.

2.6.2 Buffered storage and data integrity

Continuous high-rate measurements can produce more data than can be efficiently transferred through a low-power radio link. Werner-Allen et al. demonstrated this constraint during a 19-day deployment of 16 wireless nodes on an active volcano. Each node sampled seismic and acoustic signals at 100 Hz and retained the measurements in local flash memory because the available radio bandwidth was insufficient for continuous full-resolution transmission. The radio was instead used to report detected events and retrieve selected waveform intervals [48]. Although the Watcher records continuous electrical measurements rather than event-triggered seismic data, both systems require the scientific data to remain available independently of the wireless link.

Local non-volatile storage is generally more efficient when data are written sequentially in blocks rather than through numerous very small transactions. Flash memory is internally organised into pages and erase blocks, and repeated small accesses may therefore introduce additional execution and energy overhead [49].

Buffering several measurement records in volatile memory before committing them to non-volatile storage can reduce the frequency of physical storage operations. Mathur et al. showed that such buffering can reduce the energy cost of local flash accesses and that storing data locally may require substantially less energy than transmitting the same amount of information through a low-power radio [50].

Buffering nevertheless introduces a trade-off between storage efficiency and data integrity. Accumulating several records before committing them to non-volatile storage reduces the frequency of storage operations and their associated execution overhead, but increases the amount of recent data that may be lost following an unexpected reset or power interruption. More frequent commits reduce this vulnerable interval at the expense of additional storage activity. Power interruptions can also occur while the storage medium is internally programming or reorganising data. Experiments performed on raw flash memory have shown that interruptions during programme or erase operations may affect both the operation in progress and, under some conditions, previously written data [51]. Although these results cannot be transferred quantitatively to controller-managed SD cards, they illustrate the general need to consider power-loss behaviour when designing autonomous data-logging systems.

Local storage therefore decouples scientific acquisition from radio availability and allows the communication link to remain dedicated to low-data-rate operational exchanges. In the Watcher, LoRa is consequently used for commands, acknowledgements, diagnostics and status information rather than for continuous waveform transfer, as discussed in Section 2.7. The storage subsystem must, however, sustain continuous 100 Hz acquisition without blocking the measurement loop, while limiting both storage overhead and the amount of data vulnerable to an unexpected interruption. The record format, buffering policy and file-handling implementation adopted for the redesigned Watcher are presented in Section 4.2.

2.7 Low-power wireless communication and LoRa

This section first defines the communication requirements of autonomous sensor nodes, then introduces the LoRa physical layer and the main trade-offs between range, robustness, data rate and energy consumption.

2.7.1 Wireless communication requirements for autonomous sensor networks

Wireless communication enables autonomous sensor nodes to be monitored and controlled without requiring physical access to each device. However, the communication subsystem must operate within the limited energy, processing and memory resources available on battery-powered nodes. Radio operation can represent a significant fraction of the energy budget, particularly when the receiver remains active for long periods or when packets are repeatedly retransmitted. Communication protocols for wireless sensor networks must therefore balance network lifetime against competing requirements such as coverage, reliability, throughput and latency [27]. In the Watcher network, the radio is not intended to transfer the continuous high-resolution voltage series, which is stored locally, but rather to exchange commands, acknowledgements, diagnostic information and operational status. The required technology must consequently provide bidirectional communication over long distances with low average power consumption, while tolerating moderate latency and a comparatively low data rate.

These requirements correspond closely to the operating domain of low-power wide-area technologies. Short-range solutions such as Bluetooth or Zigbee generally provide insufficient coverage for extended outdoor deployments, whereas cellular technologies generally offer greater quality of service at the expense of higher power, infrastructure dependence and deployment costs, while also requiring adequate operator coverage in the monitored area [52]. LoRa provides an attractive intermediate solution by trading data rate for receiver sensitivity and communication range. It can be deployed as a private, self-contained radio network using only the sensor nodes and a locally installed gateway or control station, without relying on pre-existing cellular infrastructure or mobile-network coverage. This makes it particularly suitable for sparse off-grid deployments in which small amounts of information must be exchanged over large distances

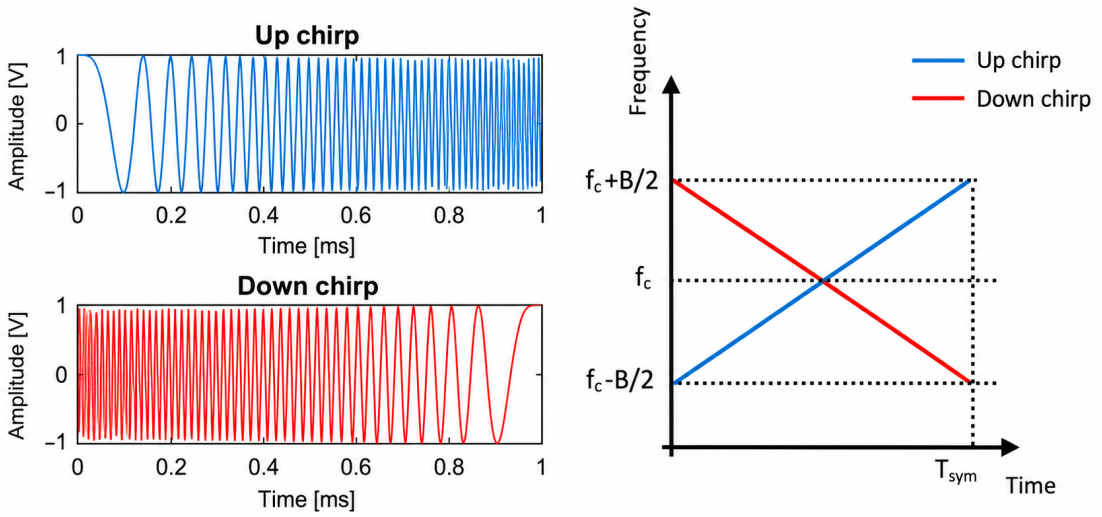

Figure 2.15: Example of an upchirp and a downchirp in time domain (left) and frequency domain (right). Reproduced from [57].

with a limited energy budget [53]. LoRa alone does not, however, guarantee a reliable link. Near-ground antennas, terrain, vegetation and other obstacles can introduce shadowing, multipath propagation and time-varying attenuation, while operation in an unlicensed band exposes the network to interference and regulatory duty-cycle constraints. Consequently, packet losses, collisions and restricted opportunities for acknowledgements or retransmissions must be considered when selecting the radio configuration and communication strategy [54]. The physical principles of LoRa modulation are introduced in the following section.

2.7.2 LoRa modulation principle

LoRa PHY and LoRaWAN refer to two distinct layers of the LoRa communication system. LoRa PHY defines the physical layer, including the radio waveform and the modulation and demodulation processes, whereas LoRaWAN is a higher-level MAC and networking protocol built on top of it. LoRaWAN specifies, among other elements, a star-of-stars architecture, device classes, channel-access procedures, addressing and security mechanisms. The LoRa physical layer can nevertheless be used independently with a custom point-to-point or multi-hop communication protocol. This distinction is important for the Watcher network, which uses the LoRa physical layer without implementing the complete LoRaWAN architecture. For simplicity, the term *LoRa* will therefore refer to the LoRa physical layer throughout the remainder of this document, unless LoRaWAN is explicitly mentioned [55, 56].

LoRa modulation is derived from Chirp Spread Spectrum (CSS). A *chirp* is a radio waveform whose instantaneous frequency sweeps progressively across a selected bandwidth BW. In LoRa, an *up-chirp* sweeps from the lower to the upper edge of this bandwidth, while a *down-chirp* does the opposite (Fig. 2.15).

The spreading factor SF determines the size of the LoRa modulation alphabet (Fig. 2.17). The number of possible symbol values is

$$M = 2^{SF}. \tag{2.10}$$

For example, $SF = 7$ gives $2^7 = 128$ possible symbols, each corresponding to one combination of seven bits. The symbol value $a \in \{0, \ldots, M-1\}$ is represented by a cyclic shift of a reference up-chirp. Using a frequency interval extending from 0 to BW, its instantaneous frequency can be written as

$$f(t;a) = \left[a\frac{BW}{M} + \frac{BW}{T_{\text{sym}}} t \right]_{\text{mod } BW}, \qquad 0 \leq t < T_{\text{sym}}. \tag{2.11}$$

The first term sets the starting frequency associated with the symbol, whereas the second term produces the linear frequency sweep. When the upper edge of the bandwidth is reached, the modulo operation returns the frequency to its lower edge. All LoRa symbols therefore have the same duration, bandwidth and sweep rate, but start at different frequencies and wrap around at different instants, as illustrated in Figure 2.16 [58, 59].

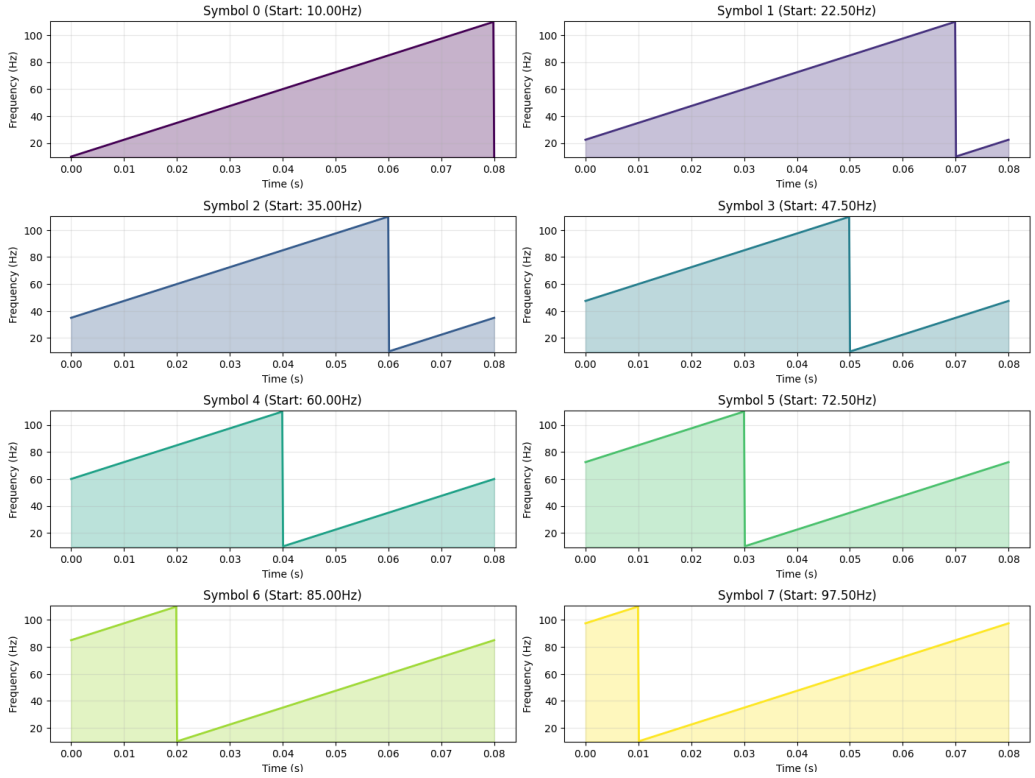

Figure 2.16: Illustration of eight symbols (SF=3, for didactic purpose only) obtained by cyclically shifting the initial frequency of the same up-chirp. All symbols have the same duration, bandwidth and sweep rate, but reach the upper frequency boundary at different instants before wrapping to the lower boundary. Reproduced from [60].

A LoRa symbol consists of one complete chirp divided into $M = 2^{SF}$ elementary time intervals called *chips*. A chip is not an information bit and is not a separate chirp; it is simply the smallest time interval used to describe the frequency sweep. The chip rate is equal to the selected bandwidth, so the duration of one chip is $T_{\text{chip}} = 1/BW$. For example, with $BW = 125$ kHz, one chip lasts 8 μs. Since one symbol contains 2^{SF} chips, its duration and symbol rate are

$$T_{\text{sym}} = \frac{2^{SF}}{BW}, \qquad R_{\text{sym}} = \frac{BW}{2^{SF}}. \tag{2.12}$$

At constant bandwidth, increasing SF by one doubles the number of chips and hence the symbol duration, as illustrated in Figure 2.17. A longer symbol gives the receiver more time to detect the chirp and generally allows weaker signals to be recovered. However, it also reduces the data rate and increases the time required to transmit a packet [55, 56].

LoRa also introduces forward-error-correction redundancy through the coding rate CR. The available code rates are 4/5, 4/6, 4/7 and 4/8. The numerator represents four information bits, while the denominator

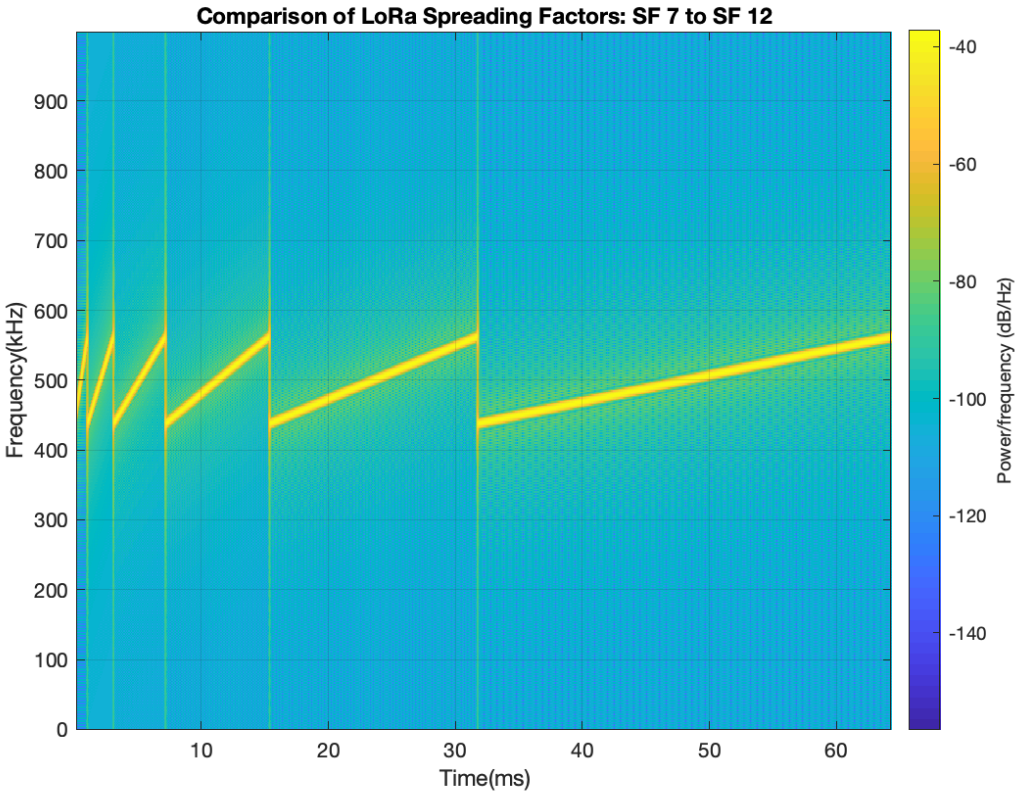

Figure 2.17: Time–frequency representation of LoRa symbols from $SF = 7$ to $SF = 12$ at constant bandwidth. Each increase of one in the spreading factor doubles the symbol duration. Reproduced from [56].

represents the total number of coded bits transmitted after adding redundancy. For example, a rate of $4/5$ adds one redundant bit for every four information bits, whereas $4/8$ adds four redundant bits. Additional redundancy increases the probability of recovering information affected by interference, but it also increases the number of transmitted bits and the packet duration. Using the coding-rate index $CR \in \{1, 2, 3, 4\}$, the corresponding nominal information rate can be approximated by

$$R_b = SF \frac{BW}{2^{SF}} \frac{4}{4 + CR}. \tag{2.13}$$

This expression accounts for the spreading factor, bandwidth and coding redundancy, but not for the preamble, header, CRC and other fields added to the radio packet. The effective application throughput is therefore lower than R_b [55, 53]. The respective roles of SF, BW and CR, together with their main effects on LoRa transmission, are summarised in Table 2.1.

Table 2.1: Meaning of the main LoRa modulation parameters.

Parameter	Physical meaning	Main effect when increased
Spreading factor SF	Determines the number 2^{SF} of chips and possible symbols	Longer symbols and lower data rate, but improved detection of weak signals
Bandwidth BW	Frequency range swept by each chirp and number of chips transmitted per second	Shorter symbols and higher data rate
Coding rate CR	Amount of redundant information added for error correction	Greater error tolerance, but longer transmissions

At the receiver, the symbol is recovered using a process known as *dechirping*. The received up-chirp is combined with a reference down-chirp having the opposite frequency sweep. The two sweeps cancel, leaving an approximately constant-frequency tone. The frequency of this tone depends on the original cyclic shift and therefore identifies the transmitted symbol. A discrete Fourier transform is commonly used to locate this tone and recover the corresponding symbol value. In this way, the energy initially spread over the complete bandwidth is concentrated during reception, allowing LoRa to detect signals that are partly buried in noise. Sweeping the complete bandwidth also provides useful resistance to narrowband interference and frequency-selective multipath propagation [55, 58, 59].

2.7.3 LoRa configuration and performance trade-offs

The LoRa parameters introduced above cannot be selected independently from the required communication range and energy consumption. Increasing the spreading factor improves the ability of the receiver to recover weak signals, while reducing the bandwidth lowers the amount of thermal noise integrated by the receiver. Both effects improve receiver sensitivity, but at the cost of a lower data rate and a longer packet time-on-air. Increasing the coding redundancy similarly improves robustness to corrupted symbols, while increasing the number of transmitted bits. Consequently, configurations providing the most robust radio link also generally require the longest transmission time [55, 56].

Receiver sensitivity can be related to the thermal noise floor by

$$P_{\text{sens}} = -174 + 10\log_{10}(BW) + NF + SNR_{\text{min}}, \tag{2.14}$$

where P_{sens} is expressed in dBm, BW is the receiver bandwidth in hertz, NF is the receiver noise figure in dB and SNR_{min} is the minimum signal-to-noise ratio required by the demodulator. The term -174 dBm/Hz corresponds to the thermal noise power density at room temperature. Reducing the bandwidth lowers the receiver noise floor, while increasing the spreading factor improves the receiver's ability to decode signals at lower SNR. For the SX1268, for example, the typical demodulator SNR decreases from approximately -2.5 dB at $SF = 5$ to -20 dB at $SF = 12$ [55].

The resulting sensitivity directly enters the radio link budget. Neglecting additional implementation losses, the received power can be approximated by

$$P_{\text{RX}} = P_{\text{TX}} + G_{\text{TX}} + G_{\text{RX}} - L_{\text{path}}, \tag{2.15}$$

where P_{TX} is the RF power delivered by the transmitter, G_{TX} and G_{RX} are the gains of the transmitting and receiving antennas, respectively, and L_{path} represents the propagation loss between both nodes. The powers are expressed in dBm, whereas antenna gains and propagation losses are expressed in dB. A reliable reception requires $P_{\text{RX}} \geq P_{\text{sens}}$. A configuration with high SF and low BW therefore increases the available link margin and potentially the communication range. However, the associated increase in time-on-air keeps the radio active for longer and increases the energy required per packet. For an autonomous system such as the Watcher, LoRa configuration must consequently balance link reliability against transmission delay and energy consumption rather than simply maximising communication range.

2.8 Synthesis and design requirements

The literature reviewed in this chapter shows that a distributed deep-ERT receiver must satisfy several requirements simultaneously. Unlike many low-power sensor nodes, the Watcher cannot reduce its energy consumption by decreasing the measurement rate during an acquisition period. The complete voltage waveform must remain available for subsequent resistivity and IP processing, requiring continuous 100 Hz acquisition, sufficiently low-noise analogue conditioning, accurate time stamping and reliable local storage.

Since the current transmitter and the receiving stations operate independently, all receivers must additionally remain referenced to a common time base, while retaining a local timing capability when GNSS reception is temporarily unavailable.

Energy optimisation must therefore focus primarily on functions that are not continuously required. The measurement chain, ADC, processing and storage must remain operational during acquisition, whereas the GNSS receiver, display, electrode-contact test and communication subsystem can be activated according to the operating state of the node. The power architecture should consequently provide low quiescent consumption and independent control of the principal functional domains. Local storage must remain the primary repository for the scientific data, while the wireless link is reserved for commands, acknowledgements, diagnostics and status information. This separation avoids imposing the throughput and energy requirements of continuous waveform transmission on the radio network.

Finally, deployment of a large number of receivers over kilometre-scale profiles requires each unit to remain compact, autonomous and sufficiently inexpensive to replicate. Communication cannot be assumed to remain reliable between every receiver and the control station because near-ground propagation, vegetation and terrain can strongly affect the radio link. The communication architecture must therefore provide a means of maintaining connectivity when a direct link is unavailable. Together, these considerations define the principal requirements summarised in Table 2.2 and provide the basis for the hardware and software developments presented in the following chapters.

Table 2.2: Main design requirements derived for the OhmPi Watcher.

Requirement	**Rationale**	**Design implication**
Continuous acquisition	Preserve the complete deep-ERT voltage waveform	Maintain the required electrode channels at 100 Hz during measurement periods
Acquisition flexibility	The original Watcher was limited to two electrode input channels, restricting the range of possible measurement configurations	Extend the acquisition electronics from two to four independently conditioned electrode channels while preserving the required two-channel acquisition performance
Measurement integrity	Deep-ERT signals may be weak and affected by electrical noise	Low-noise analogue conditioning, appropriate input protection and filtering
Time synchronisation	Independent receivers must be aligned with the injected-current waveform	GNSS-based reference combined with local timekeeping between synchronisation events
Reliable local storage	Raw measurements cannot depend on radio availability	Buffered SD-card storage with controlled data-loss behaviour
Low average consumption	Battery size and cost scale with the number of deployed units	Low-quiescent-current components, switchable power domains and state-based operation
Robust communication	Direct radio links may be degraded by terrain, vegetation and near-ground deployment	Low-data-rate command and status communication with tolerance to packet losses and indirect connectivity
Scalable architecture	Large deep-ERT surveys require many field units	Compact, reproducible and maintainable hardware with limited unit cost

Chapter 3

Hardware development

The hardware developed in this work is an evolution of the original OhmPi Watcher described in Section 2.3.2, rather than a complete redesign of its acquisition concept. The main functional architecture was retained: the system is supplied from a 12 V battery, an Adafruit Feather M0 provides embedded control and the LoRa interface, GNSS and RTC devices provide the timing references, and measurements are stored locally on an SD card. The analogue front end also retains the same fundamental transfer function, mapping the expected ±5 V electrode range onto the 0–5 V input range of the ADC.

The redesign instead focused on improving the suitability of the prototype for replication and long-duration autonomous deployment. The main hardware modifications are summarised in Table 3.1. They primarily concern the power architecture, the number and integration of analogue acquisition channels, the default electrical states of the electrode interfaces, and the protection and control circuitry.

These modifications were implemented while retaining the functions that had already proved suitable in the original prototype. Their integration into the redesigned hardware is illustrated in Fig. 3.1. The redesigned PCB consolidates the revised power, acquisition, switching and protection circuits into a more compact and structured layout. The following sections describe the modifications in greater detail and justify the main design choices associated with each of them.

3.1 Microcontroller choice

The Adafruit Feather M0 RFM96 LoRa development board used in the original Watcher was retained in the redesigned hardware. Replacing the microcontroller was considered, particularly with more recent low-power devices offering lower sleep consumption, additional GPIOs and greater peripheral flexibility. Such a change could potentially simplify some of the software constraints encountered later in the development, including the sharing of the SPI bus (see Section 4.1.3).

However, changing the controller platform would also have required a substantial redesign of the existing embedded software and the revalidation of the complete acquisition and communication system. The Feather M0 was already directly compatible with the GPS, OLED, RTC and Adalogger FeatherWing modules used by the Watcher, while its integrated LoRa module allowed the existing radio hardware and software interface to be retained. Maintaining the same platform also preserved compatibility with the Arduino-based code developed for the original prototype.

The microcontroller itself was also not identified as the dominant contributor to the Watcher's energy consumption during normal operation. The redesign effort was therefore prioritised towards subsystems offering larger expected gains. Retaining the Feather M0 consequently reduced development risk and implementation time. A migration to a more highly integrated low-power microcontroller remains a possible future improvement.

Table 3.1: Main hardware modifications introduced in the redesigned OhmPi Watcher.

Function	**Original Watcher**	**Redesigned Watcher**
Power conversion	Oversized and expensive DC–DC converters	Converters selected from the actual rail power requirements
±15 V rail switching	Current-consuming remote shutdown	Low-loss high-side MOSFET switching
Electrode acquisition	2 electrode channels	Up to 4 electrode channels
Signal conditioning	Three discrete OP07 amplifiers per channel	One OPA4188 quad amplifier per channel
Analogue filtering	No dedicated input low-pass filter	RC low-pass filtering
ADC architecture	One ADS1115 shared by measurement and auxiliary functions	One ADS1115 dedicated to the electrode channels and one auxiliary ADS1115
GPIO expansion	MCP23008	MCP23017
Analogue switching	DG411 analogue switches	DG413 switches with more suitable default states
Relay driving	Discrete transistor and flyback-diode stages	Darlington transistor array
Contact-resistance circuit	2 Ω current-sensing shunt	2 Ω shunt with an additional 100 Ω series protection resistor
PCB design	Functional, but long traces impairing the continuity of the bottom GND plane and not annotated	More compact functional placement, improved ground-plane continuity, and additional silkscreen and test features

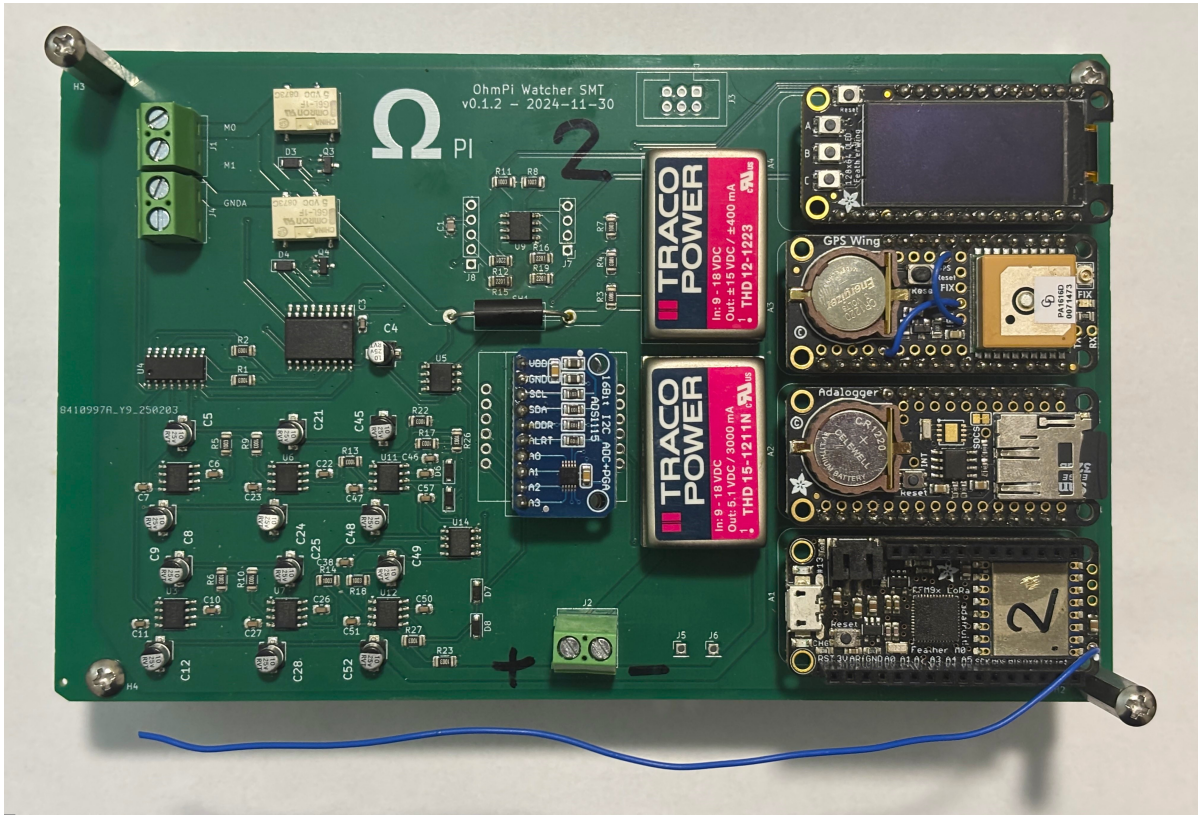

(a) Original Watcher prototype. 180×110 mm.

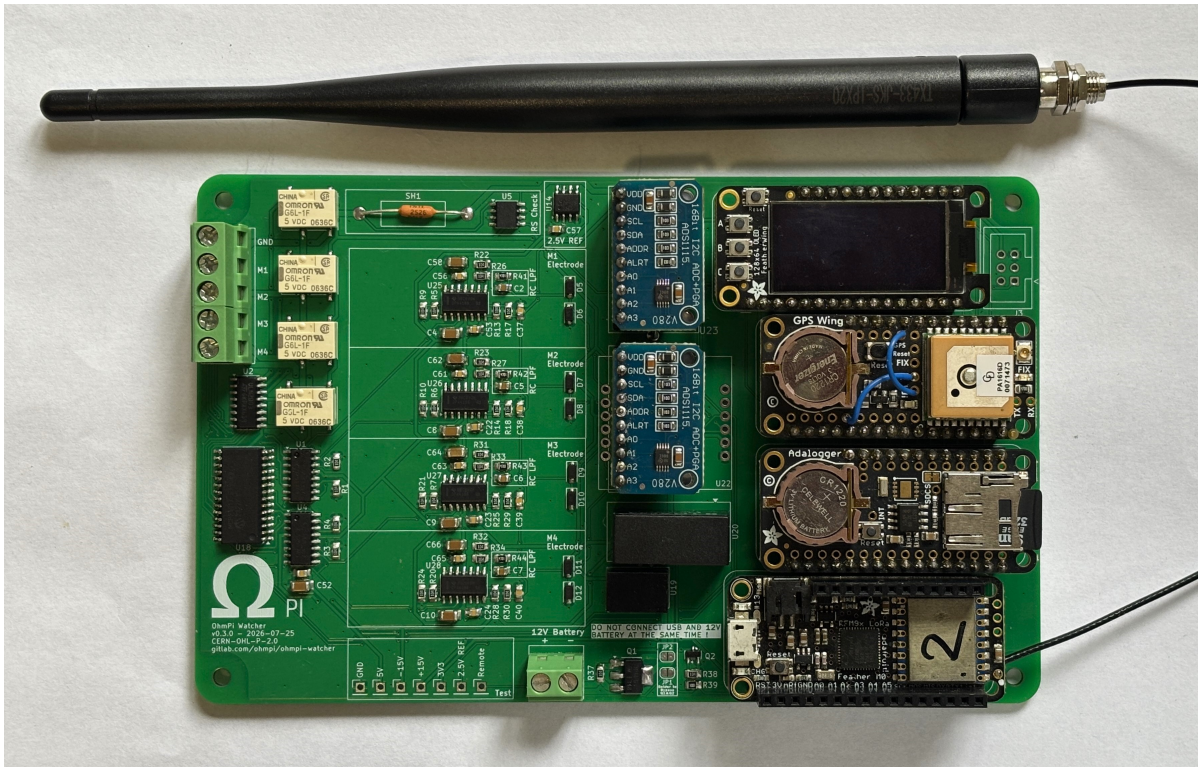

(b) Redesigned Watcher along with updated LoRa antenna. 160×100 mm.

Figure 3.1: Physical comparison between the original OhmPi Watcher (v0.1.2) and the redesigned hardware developed in this work (v0.3.0). The new PCB integrates the revised power architecture, four electrode-conditioning channels and the associated control and protection circuitry into a more compact layout.

The main pin assignments of the retained Feather M0 in the redesigned Watcher are summarised in Fig. 3.2. In addition to the pins internally assigned to the RFM96 LoRa module, the controller uses dedicated GPIOs for the GPS timing signal, OLED push-buttons, SD-card selection and analogue-power control, while most external peripherals share the I^2C bus.

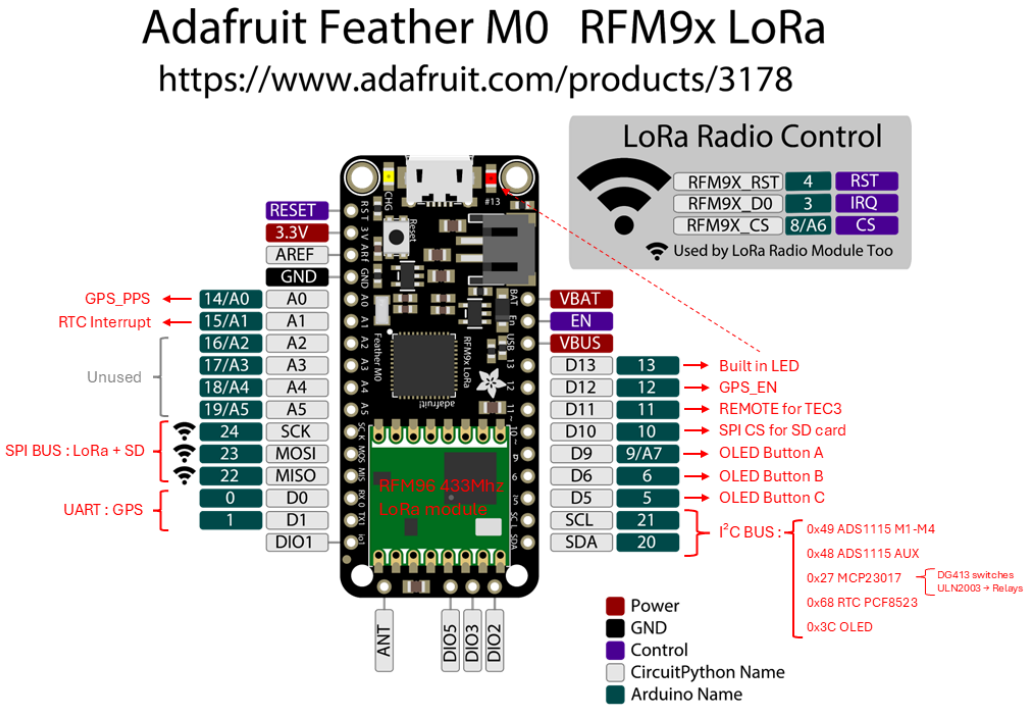

Figure 3.2: Schematic pin assignments of the Adafruit Feather M0 board in the redesigned Watcher. Modified from Appendix C

.

3.2 Power supply redesign

The power architecture of the original Watcher was retained at a functional level: a 12 V battery supplies a 5 V rail for the digital electronics and a ±15 V supply for the analogue front end. However, the converters used in the prototype provided substantially more output power than required by the actual loads. As discussed in Section 2.4, converter sizing is particularly important in an autonomous system because efficiency, no-load consumption and shutdown behaviour at the actual operating point may be more relevant than the nominal maximum output power.

The redesigned power architecture therefore follows two complementary principles. First, the converters were resized according to the electrical requirements of the loads connected to each rail. Second, the bipolar analogue supply was implemented as a separately switchable power domain so that the converter and the associated analogue electronics can be completely disconnected from the 12 V source when voltage acquisition is not required. The resulting architecture is summarised in Fig. 3.3.

3.2.1 Power budget and converter sizing

A rail-level power budget was established before selecting the replacement converters. The objective was to provide sufficient margin for the maximum simultaneous load while avoiding the large excess capacity of the original implementation. The loads were therefore grouped according to their supply rail and to the combinations of functions that can actually operate simultaneously.

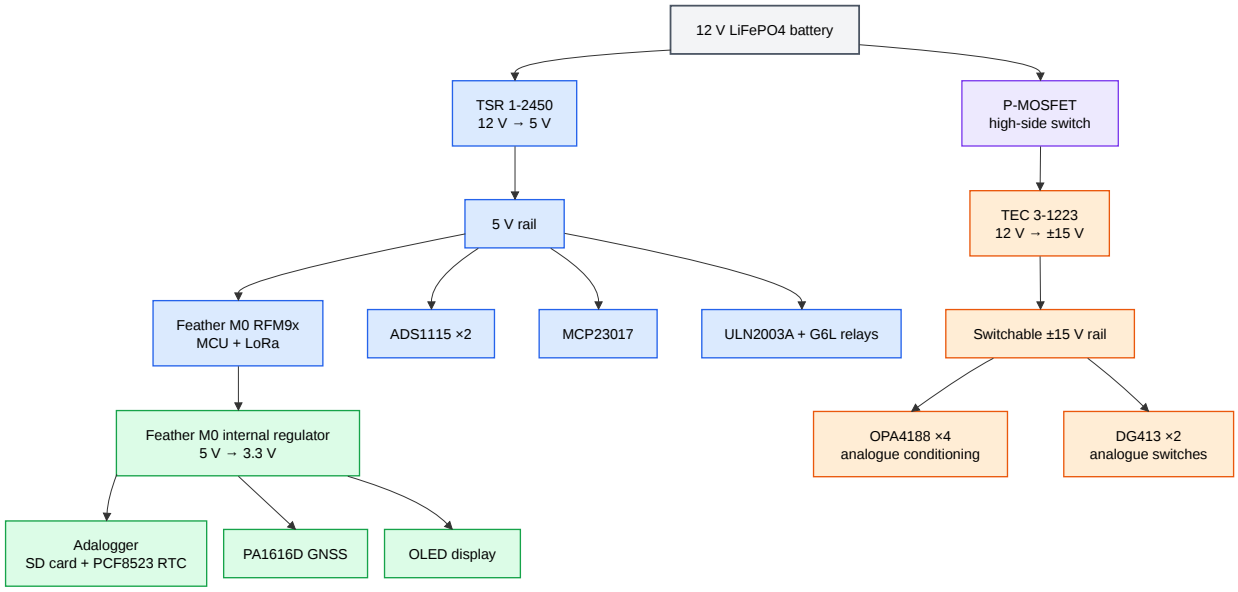

Figure 3.3: Power architecture of the redesigned Watcher. The 12 V battery supplies the continuously available low-voltage electronics through a dedicated 5 V converter, whereas the bipolar analogue supply is connected through a controllable high-side switch. The ±15 V domain can therefore be disconnected when the analogue acquisition circuitry is not required.

5 V domain

The low-voltage domain contains both 5 V and 3.3 V loads. Only the 5 V rail is generated directly from the 12 V battery, while the on-board regulator of the Feather M0 subsequently generates the 3.3 V supply used by the microcontroller, radio and FeatherWing peripherals. Consequently, the Adalogger, GPS FeatherWing and OLED display also contribute to the current drawn from the 5 V converter. The G6L relays were considered according to their actual operating sequence: only one is switched at a time during the successive electrode resistance checks. The resulting worst-case simultaneous load is summarised in Table 3.2.

The component current values were taken from the corresponding manufacturer datasheets [61, 41, 62, 43, 63, 64], except for the OLED peak current, which was measured experimentally with all pixels enabled.

The loads supplied through the Feather M0 3.3 V regulator represent a combined peak current of approximately 305 mA. This remains below the 500 mA peak output capability specified for the on-board regulator [61].

The estimated peak current of 342 mA is considerably below the capability of the converter used in the original Watcher. The THD 15-1211N provides a 5.1 V output rated at 3 A [32], almost nine times the estimated worst-case current requirement of the redesigned receiver. Moreover, the THD 15N series is no longer recommended for new designs by the manufacturer.

It was therefore replaced by a Traco TSR 1-2450 step-down switching regulator. The TSR 1-2450 provides a regulated 5 V output rated at 1 A, leaving sufficient margin for component tolerances and short current transients without retaining the excessive power capability of the original converter. It also has a typical no-load input current of approximately 1 mA, compared with 8 mA for the THD 15-1211N, and is significantly smaller and lighter.

Unlike the THD 15-1211N, the TSR 1-2450 is non-isolated. Although the original converter provided galvanic isolation, this function was not required by the Watcher architecture as the digital and acquisition circuits already shared a common electrical reference. No separation between independently grounded

Table 3.2: Design power budget of the low-voltage domain. The 3.3 V peripherals are supplied through the on-board regulator of the Feather M0 and therefore also contribute to the current drawn from the 5 V rail.

Load	Qty.	Supply	Peak current[1]
Feather M0 + LoRa radio (TX)	1	3.3 V	130 mA
Adalogger (RTC + SD card)	1	3.3 V	30 mA
GPS FeatherWing	1	3.3 V	35 mA
OLED display	1	3.3 V	110 mA
ADS1115	2	5 V	0.60 mA
MCP23017	1	5 V	1 mA
G6L relay	1 active at a time	5 V	36 mA
Estimated 5 V current			**342.6 mA**

[1] Approximate current contribution referred to the 5 V supply.

systems is therefore required. The isolated converter could consequently be replaced by the simpler non-isolated regulator without removing a required system function [65, 32].

±15 V domain

The bipolar domain was evaluated separately because it supplies only the analogue acquisition circuitry. It includes the four OPA4188 devices, the DG413 analogue switches and the 2.5 V reference circuit. Unlike most of the low-voltage electronics, these components are not required continuously and can therefore form a separately switchable power domain.

For the OPA4188 devices, the budget includes both the worst-case quiescent current of the complete quad amplifier and a conservative estimate of the current delivered by its active outputs. For one conditioning chain, this gives approximately 3.42 mA, of which 2.28 mA corresponds to the worst-case quiescent consumption and approximately 1.14 mA to the maximum estimated output load. The resulting bipolar-rail budget is given in Table 3.3.

The component current values were derived from the OPA4188, DG413 and REF03 datasheets together with the calculated worst-case output loading of the conditioning stages [66, 67, 68].

Table 3.3: Design power budget of the bipolar analogue domain.

Load	Qty.	Supply	Peak current (per device)
OPA4188	4	±15 V	2.28 mA quiescent + 1.14 mA output load
DG413	2	±15 V	5 μA
REF03	1	+15 V	1.4 mA
Estimated ±15 V current			**15.09 mA**[1]

[1] Per-rail current, computed here for +15 V as REF03 is only supplied by +15 V.

The original bipolar supply was provided by a THD 12-1223, a 12 W isolated converter capable of supplying up to 400 mA on each ±15 V output. The estimated requirement of approximately 15 mA per rail therefore used only a small fraction of its available current capability [33].

The THD 12-1223 was consequently replaced by a TEC 3-1223. The latter provides regulated ±15 V outputs rated at 100 mA per rail. This still provides substantial operating margin. The TEC 3-1223 also uses a smaller SIP-8 package and weighs 4.5 g, compared with 18 g for the THD 12 series. Its typical no-load input current is 30 mA, compared with 20 mA for the THD 12-1223 [69, 33].

The two replacement converters therefore remain comfortably above the estimated peak requirements while being more closely matched to their respective loads. In addition, the bipolar analogue domain can be completely disconnected from the battery when voltage acquisition is not required, as described in Section 3.2.2. Since the budgets presented above combine datasheet values, circuit-level estimates and conservative peak-current assumptions, the actual current consumption of the assembled Watcher was subsequently measured in its different operating states, as presented in Section 6.1.

The replacement also substantially reduces component cost. At the single-unit prices used during the design phase, the THD 15-1211N and THD 12-1223 cost approximately 41.78 € and 36.92 €, respectively, whereas the TSR 1-2450 and TEC 3-1223 cost approximately 3.10 € and 12.47 €. The combined cost is therefore reduced from approximately 78.70 € to 15.57 € (Mouser Belgium, excl. VAT, accessed August 2026). A complete cost comparison of the redesigned Watcher is presented separately in Section 6.5.

3.2.2 Switchable power source

In the original Watcher, the ±15 V DC–DC converter was disabled during inactive periods using the manufacturer's remote-control input. For the THD 12-1223, this input is voltage controlled and reduces the converter input current to approximately 2.5 mA when the device is disabled [33]. Although this already reduces the consumption of the analogue domain substantially compared with leaving the converter active, the converter remains permanently connected to the 12 V battery.

The TEC 3-1223 also provides a remote-shutdown function. However, the converter remains electrically connected to the battery in this state and specifies a typical off-state input current of approximately 2.5 mA. In addition, activating the remote-off function requires a continuous control current of approximately 2–4 mA through the REMOTE input [69]. The shutdown state therefore still involves a non-negligible continuous current consumption. Since the analogue domain may remain inactive for several hours, for example during the night, this residual current consumption was considered undesirable. The redesigned circuit therefore disconnects the TEC 3-1223 directly from the 12 V supply instead of relying on its internal remote-shutdown function.

A latching relay was first considered as an alternative. Such a device would provide physical galvanic disconnection and would require no holding current once switched. However, it would occupy more PCB area, introduce mechanically wearing contacts and, more importantly, could remain in an undesired state after a loss of power or controller reset. A solid-state high-side MOSFET switch was therefore preferred. It provides negligible steady-state control consumption, requires little PCB area, is less expensive than a relay and contains no mechanically wearing contacts, while ensuring that the analogue supply defaults to the OFF state when the controller is unpowered.

The resulting high-side switching circuit is shown in Fig. 3.4. A P-channel MOSFET (Q1) is placed between the 12 V battery rail and the positive input of the TEC 3-1223. Its gate is controlled through a 2N7002 N-channel MOSFET (Q2), allowing the 3.3 V logic output of the Feather M0 to control the 12 V high-side switch. When the `REMOTE` signal is low, Q2 is off and R37 pulls the gate of Q1 towards its source. The resulting gate-to-source voltage of Q1 is approximately zero, keeping the P-channel MOSFET off and disconnecting the TEC 3-1223 from the battery. In this state, no continuous bias current is required. The remaining battery current is limited to the leakage currents of the two MOSFETs. At 25 °C, both devices specify a maximum off-state drain leakage current of 1 μA under test voltages considerably higher than the 12 V used here, giving a conservative datasheet-based upper bound of approximately 2 μA for the complete switching circuit [70, 71].

When REMOTE is driven high, Q2 conducts and pulls the gate of Q1 towards ground. The resulting V_{GS} of approximately -12 V turns Q1 fully on and applies the battery voltage to the input of the TEC 3-1223. The 100 kΩ pull-up resistor R37 then draws only approximately 120 μA through Q2. The resulting gate voltage also remains within the ± 20 V absolute maximum rating of Q1 [70].

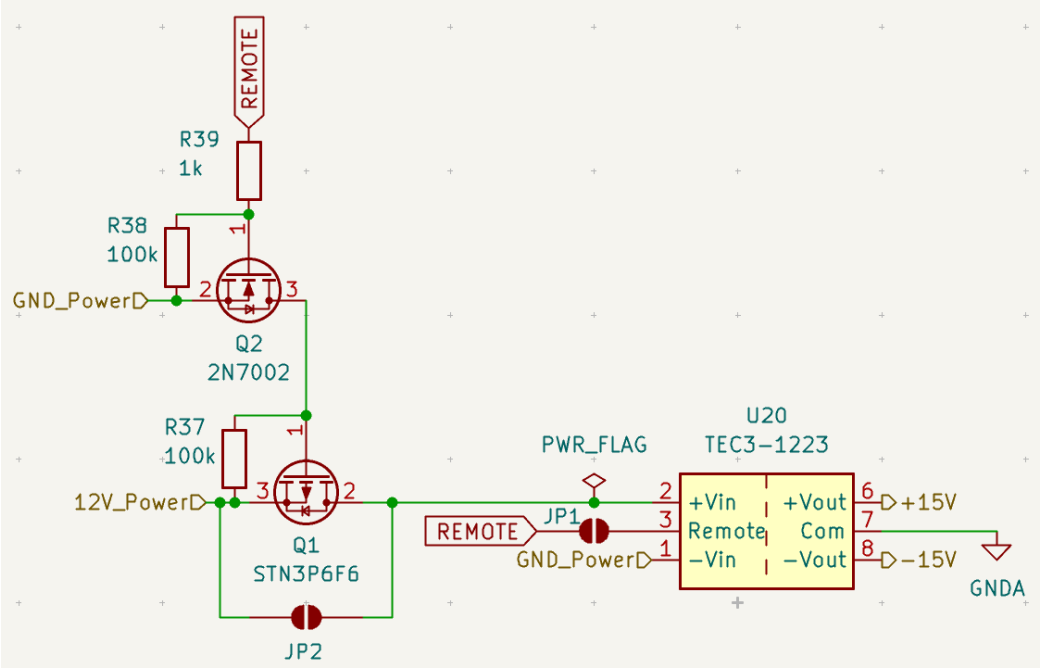

Figure 3.4: High-side switching circuit used to disconnect the TEC 3-1223 from the 12 V battery. Q1 provides the high-side power switching, while Q2 allows the 3.3 V REMOTE signal from the Feather M0 to control its gate.

As this switching architecture had not yet been experimentally validated when the PCB was designed, two configuration jumpers were added as a fallback solution. One jumper allows the high-side switch to be bypassed so that the TEC 3-1223 can be supplied directly from the 12 V rail, while the second allows the Feather M0 control signal to be connected to the converter's internal remote input. If the external switching circuit proved unsuitable, the board could therefore revert to the manufacturer-supported remote-control configuration without requiring a PCB revision. This alternative would retain the additional off-state consumption of the TEC 3-1223, but would ensure that the prototype remained operational during validation of the new power architecture.

The high-side switching circuit was first verified by LTspice simulation and subsequently validated experimentally on the assembled PCB. Both tests confirmed the expected behaviour, with the TEC 3-1223 supplied from the 12 V rail only when the control signal was enabled. The fallback jumpers were therefore not required during normal operation.

This result also shows that the remote-control function does not need to be integrated in the converter itself. Since the high-side MOSFET circuit can switch the converter supply externally, a simpler and less expensive non-isolated converter without a remote-control input could be used in a future revision, provided that it meets the required voltage, current and efficiency specifications.

3.3 Analogue acquisition chain redesign

The analogue acquisition chain was redesigned while preserving the measurement principle of the original Watcher. In particular, the same ideal transfer function was retained,

$$V_{\mathrm{ADC}} = \frac{V_{\mathrm{electrode}}}{2} + 2.5\ \mathrm{V}, \tag{3.1}$$

so that an electrode voltage ranging between −5 and +5 V is mapped onto the 0–5 V input range of the ADC. The 2.5 V offset required by this transfer function is derived from the +15 V analogue supply through a REF03 precision voltage reference, providing a stable bias for the level-shifting stage [68]. Retaining this transfer function preserved compatibility with the existing acquisition concept while allowing the implementation itself to be improved in terms of power consumption, integration, filtering, protection and number of measurement channels.

The conditioning circuit schematic for one electrode channel is shown in Fig. 3.5. The physical implementation of the circuit is shown in Fig. 3.6. Four identical conditioning chains are implemented on the redesigned PCB, one for each receiving electrode.

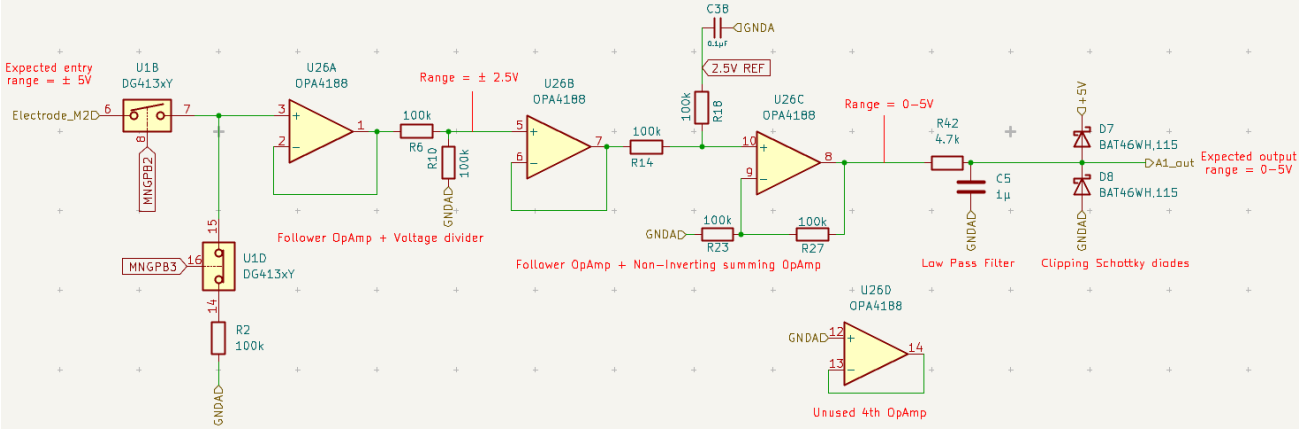

Figure 3.5: Analogue acquisition chain implemented for one receiving electrode in the redesigned Watcher. The input of the conditioning chain can be connected either to the electrode or to the analogue ground through the two DG413 switches. The resulting signal is then conditioned by the OPA4188, low-pass filtered, protected by the output clamping network and digitised by the ADS1115.



Figure 3.6: Close-up view of the analogue acquisition chain on the new PCB. DG413 switches not shown.

3.3.1 Operational-amplifier integration

The original Watcher implements each electrode-conditioning chain using three separate OP07C operational amplifiers. Extending this architecture directly to four electrodes would therefore require 12 individual OP07C packages. In the redesigned circuit, these three stages are integrated into a single OPA4188 quad operational amplifier for each electrode. Three of its four internal amplifiers reproduce the original

conditioning stages, while the fourth is unused. To avoid leaving its inputs floating, this unused amplifier is configured as a unity-gain buffer, with its non-inverting input connected to analogue ground and its inverting input connected to its output. This keeps the amplifier in a defined operating condition and avoids possible saturation or oscillation, although its intrinsic quiescent current remains.

The OPA4188 was selected not only for this higher level of integration, but also for its improved electrical performance. At 25 °C, the OP07C specifies an input offset voltage of 60 μV typical and 150 μV maximum, whereas the OPA4188 specifies 6 μV typical and 25 μV maximum. Its offset-voltage drift is also considerably lower, with a maximum of 0.085 μV/°C compared with 1.8 μV/°C for the OP07C. These characteristics are particularly relevant for the Watcher because the measured signals contain slowly varying and quasi-DC components for which offset and temperature drift directly affect the measurement [42, 66].

The OPA4188 also presents a very high input impedance, which is advantageous for voltage measurements from potentially high-impedance electrodes. The datasheet specifies a common-mode input impedance of $6 \times 10^{12}\ \Omega \parallel 9.5$ pF and a differential input impedance of 100 M$\Omega \parallel 6$ pF. The resistive common-mode input impedance is therefore approximately 6 TΩ at low frequency [66].

The replacement also substantially reduces the quiescent consumption of the conditioning circuitry. At $\pm$15 V, the OP07C specifies a no-load power consumption of 75 mW typical and 120 mW maximum, corresponding respectively to approximately 2.5 and 4 mA per device. A four-electrode implementation based on twelve OP07C devices would therefore require approximately

$$I_{\mathrm{OP07C,typ}} = 12 \times 2.5\ \mathrm{mA} = 30\ \mathrm{mA}, \tag{3.2}$$

with a datasheet maximum of approximately

$$I_{\mathrm{OP07C,max}} = 12 \times 4\ \mathrm{mA} = 48\ \mathrm{mA}. \tag{3.3}$$

For the OPA4188, the quiescent current at the same supply voltage is approximately 415 μA typical and 500 μA maximum per internal amplifier at 25 °C. Four complete quad devices therefore represent approximately 6.64 mA typical and 8 mA maximum. Even though one amplifier remains unused in each package, the expected quiescent consumption is consequently reduced by approximately 78% based on typical values, and by approximately 83% when comparing the 25 °C datasheet maxima [42, 66]. These values exclude the additional output current delivered to the surrounding passive components.

A three-package implementation was initially considered to further reduce component count and cost. In this configuration, each OPA4188 would implement one conditioning stage for all four electrodes, thereby using all twelve available amplifiers. However, the signals from each electrode would then have to travel between several amplifier packages distributed across the PCB. This would result in longer and more interleaved analogue traces, make the circuit less intuitive to inspect, and make it more difficult to keep the four measurement channels physically separated. The final layout therefore uses one OPA4188 per electrode. The small additional component cost was considered justified by the shorter routing, and the clearer correspondence between each physical channel and its amplifier package.

Integration also reduces the number of separate supply-decoupling locations. Instead of supplying and locally decoupling three individual OP07C packages for each electrode, only one quad package must be decoupled. The OPA4188 datasheet recommends local supply bypassing and indicates that 0.1 μF capacitors are adequate in most applications [66]. This further simplifies the PCB around each conditioning chain.

3.3.2 Analogue filtering and ADC input protection

A first-order passive low-pass filter was added between the final conditioning stage and the ADS1115 input. The original Watcher did not include a dedicated analogue filter at this point in the acquisition chain. In

the redesigned circuit, the filter uses $R = 4.7$ kΩ and $C = 1$ μF, giving a time constant of

$$\tau = RC = 4.7 \text{ ms}, \tag{3.4}$$

and a theoretical cut-off frequency of

$$f_c = \frac{1}{2\pi RC} = \frac{1}{2\pi(4.7 \times 10^3)(1 \times 10^{-6})} \approx 33.9 \text{ Hz}. \tag{3.5}$$

The purpose of this filter is to limit the analogue bandwidth before digitisation and to attenuate higher-frequency electrical interference. This is particularly relevant for the 100 Hz acquisition rate used by the Watcher, for which the Nyquist frequency is 50 Hz. The filter is intentionally simple and does not constitute a steep anti-aliasing filter. For a first-order low-pass response,

$$|H(f)| = \frac{1}{\sqrt{1 + (f/f_c)^2}}, \tag{3.6}$$

which gives approximately

$$|H(50 \text{ Hz})| \approx 0.56, \tag{3.7}$$

or about -5.0 dB at 50 Hz. The circuit therefore provides only moderate attenuation close to the Nyquist limit. This is particularly relevant to mains-related interference. With a sampling frequency of 100 Hz, a 50 Hz component lies exactly at the Nyquist frequency, while frequency components above 50 Hz are folded back into the sampled bandwidth. In particular, a 100 Hz component, such as the second harmonic of the mains frequency, aliases towards DC when sampled at exactly 100 Hz. Components slightly above 50 Hz similarly appear at lower frequencies after sampling. Once aliased, these components cannot be removed unambiguously by subsequent digital filtering. Sufficient analogue attenuation before the ADC is therefore desirable.

The selected RC values also represent a compromise between noise attenuation and transient response. With a time constant of 4.7 ms, the filter responds rapidly compared with the much longer current injection periods used during deep-ERT acquisition, while reducing high-frequency disturbances before conversion. The implemented frequency response is evaluated experimentally in Section 6.3.4.

The unused fourth amplifier in each OPA4188 package also provides a straightforward upgrade path for the analogue filter. In a future revision, it could be used to implement a second-order active low-pass filter, for example using a Sallen–Key architecture. Such a filter would provide a roll-off of 40 dB/decade, compared with 20 dB/decade for the present first-order network, and would therefore provide stronger attenuation of components around and above the Nyquist frequency.

The optimum cut-off frequency should nevertheless be determined from representative field measurements. Reducing f_c would improve rejection of 50 Hz, 100 Hz and other high-frequency interference, but could also distort useful voltage transients and the early part of the induced-polarisation response. Field recordings should therefore be analysed to determine the lowest cut-off frequency that provides effective anti-aliasing while preserving the geoelectrical information required for subsequent processing.

In addition to filtering, the ADC input is protected against voltage excursions outside its permissible range. Although the ADS1115 is configured with a nominal full-scale range of ±6.144 V, its analogue input pins must still remain within the supply rails during normal operation [43]. A pair of external Schottky diodes is therefore connected at each ADC input to clamp positive and negative excursions towards the 5 V and ground rails, while the series resistance already present in the low-pass network limits the

associated transient current. This arrangement follows the manufacturer's recommendation for protecting the ADS1115 inputs when the applied voltage may exceed the permitted range [43].

3.3.3 Electrode switching and fail-safe states

The analogue switching stage was also revised. In the original Watcher, DG411 analogue switches were used to connect or isolate the receiving electrodes from the conditioning circuitry. In the redesigned board, these devices were replaced by DG413 switches. Unlike the DG411, which contains four normally-closed (NC) switches, the DG413 combines two normally-open (NO) and two normally-closed (NC) channels [67]. As shown on the left-hand side of Fig. 3.5, one NC and one NO channel are used for each electrode-conditioning chain.

Three routing configurations are required:

- **Normal acquisition:** the electrode is connected to its conditioning chain, while the corresponding GNDA connection remains open.
- **Electronics noise and offset measurement:** the electrode is disconnected and the conditioning-chain input is connected to GNDA, providing a known zero-input condition.
- **Contact-resistance measurement:** both paths are opened, isolating the analogue conditioning chain while the external test voltage is applied to the selected electrode.

The mixed NO and NC configuration of the DG413 allows these three routing states to be implemented while providing a defined safe configuration if controller control is lost or during a reset.

The DG413 is a CMOS voltage-controlled analogue switch. Changing or maintaining its state therefore does not require a holding current such as the coil current required by an electromechanical relay. The device still draws its specified static supply and logic-input currents, but these currents remain small and are not associated with continuously energising a particular switch state. Consequently, changing from the DG411 to the DG413 does not introduce any state-dependent power penalty [67].

Removing the analogue switches entirely was also considered. In this alternative, the electrode conditioning circuitry would have been isolated during the contact-resistance test only by switching off the TEC 3-1223 and therefore removing the $\pm$15 V supply from the operational amplifiers. This would reduce component count and cost, but the approach was rejected for two reasons. First, an unpowered operational amplifier cannot be assumed to provide a defined high-voltage isolation barrier. The OPA4188 inputs contain protection structures referenced to its supply rails, and applying an external voltage while the device is unpowered could therefore create unintended current paths rather than reliably blocking the 12 V resistance-test excitation [66]. Second, removing the switches would eliminate the possibility of disconnecting the field electrodes while deliberately connecting the analogue-chain inputs to GNDA. This configuration provides a known zero-input condition in which the complete conditioning chain and ADC can be characterised without contributions from electrode or environmental noise. The mean measured ADC voltage can then be used to evaluate the actual 2.5 V offset of each channel, while its standard deviation provides an estimate of the electronic noise of the acquisition chain. Retaining the analogue switches therefore provides both protection during the contact-resistance test and a useful self-diagnostic and calibration configuration.

3.3.4 ADC requirements and device comparison

The ADC used for electrode acquisition must satisfy several requirements imposed by the redesigned Watcher. Because the redesigned Watcher is intended to support four electrode channels, the converter must provide at least four single-ended inputs, each with a target acquisition rate of 100 Hz, while remaining compatible with the 0–5 V output range of the existing analogue conditioning chain. Low power consumption is also desirable because the receiver is intended for long-duration autonomous operation.

Although the ADS1115 used in the original Watcher satisfies these basic requirements, a higher-resolution alternative was investigated. This was partly motivated by comparison with the commercial V-FullWaver, which records at the same 100 Hz sampling rate while specifying a voltage resolution of 1 μV [5]. By comparison, with the $\pm$6.144 V full-scale range used in the Watcher, the ADS1115 has a nominal LSB size of 187.5 μV, corresponding to approximately 375 μV when referred back to the electrode through the 0.5 analogue-chain gain [43]. This comparison highlighted the potential benefit of investigating a higher-resolution converter. Therefore, the ADS1219 was analysed as a possible higher-resolution alternative. It was attractive from a prototyping perspective, since inexpensive commercial breakout boards were readily available. This made the device easy to evaluate experimentally without requiring a dedicated PCB revision, and would also simplify replacement or comparison with other ADC options during development.

The ADS1219 is a 24-bit delta-sigma converter providing four single-ended inputs, rail-to-rail input buffers, selectable gains of 1 and 4 and data rates up to 1000 SPS. Its nominal resolution is therefore substantially higher than that of the 16-bit ADS1115 [46, 43]. However, nominal bit depth alone is not sufficient to compare the two devices, because the four electrode channels must each be acquired at 100 Hz. Since both converters multiplex their inputs sequentially, at least 400 conversions per second are required. For the ADS1115, the 475 SPS setting would nominally require

$$4 \times \frac{1}{475} \approx 8.42 \text{ ms}, \tag{3.8}$$

leaving only 1.58 ms of the 10 ms acquisition period for I^2C communication, channel reconfiguration, timestamping and storage. The 860 SPS setting was therefore selected, reducing the nominal conversion time for four channels to

$$4 \times \frac{1}{860} \approx 4.65 \text{ ms}. \tag{3.9}$$

For the ADS1219, the 330 SPS setting is insufficient because four conversions would require approximately 12.1 ms. Operation at 1000 SPS would instead require approximately 4 ms for four conversions. These values represent conversion times only and do not include communication or firmware overhead; the practical acquisition timing is evaluated later in Section 6.2.1.

Table 3.4 summarises the characteristics relevant to the present application.

The comparison illustrates the difference between nominal and usable resolution. With its internal 2.048 V reference and gain 1, the ADS1219 has a nominal LSB size of approximately 0.244 μV. However, at 1000 SPS the datasheet specifies 36.98 μV RMS and 221.61 μV peak-to-peak input-referred noise, corresponding to 16.76 effective bits and only 14.17 noise-free bits. The nominal 24-bit output must therefore not be interpreted as 24 bits of usable measurement resolution at the data rate required here [46].

The ADS1219 input buffers would also provide limited benefit in the Watcher because the converter is already driven by the low output-impedance OPA4188 conditioning stage. More importantly, the existing conditioning circuit produces a 0–5 V signal. With its internal 2.048 V reference, the ADS1219 differential measurement range at gain 1 is limited to $\pm$2.048 V, and the analogue front end would therefore have to be redesigned to make efficient use of the converter range. The ADS1115, by contrast, directly accommodates the existing 0–5 V signal when configured for the $\pm$6.144 V range.

The ADS1219 also draws approximately 315 μA during conversion, compared with approximately 150 μA for the ADS1115. At the required data rates, the ADS1219 does not provide a higher noise-free resolution when expressed in bits, although its absolute input-referred peak-to-peak noise remains lower than that of the ADS1115 under the respective datasheet test conditions. The ADS1115 was therefore retained because it combines lower power consumption, direct compatibility with the existing analogue range and a

Table 3.4: Comparison of the ADS1115 and ADS1219 for four-channel acquisition at 100 Hz [43, 46].

Characteristic	**ADS1115**	**ADS1219**
Nominal resolution	16 bits	24 bits
Input configurations	Four single-ended or two differential	Four single-ended or two differential
Available data rates	8, 16, 32, 64, 128, 250, 475, 860 SPS	20, 90, 330, 1000 SPS
Selected data rate	860 SPS	1000 SPS
Nominal conversion time	1.16 ms	1.00 ms
Four-conversion time	4.65 ms	4.00 ms
Time remaining in a 10 ms period	5.35 ms	6.00 ms
Effective resolution at selected rate	16 bits	16.76 bits
Noise-free resolution at selected rate	14.8 bits	14.17 bits
Selected / relevant full-scale range	$\pm$6.144 V	$\pm$2.048 V, gain 1
Nominal LSB at this range	187.5 μV	0.244 μV
Input-referred peak-to-peak noise	430.06 μV	221.61 μV
Input stage	Switched-capacitor input	Rail-to-rail input buffers
Typical conversion current	150 μA	315 μA

nominal conversion rate compatible with the four-channel acquisition target at the hardware level, without requiring a redesign of the conditioning chain. The practical implementation and timing limitations are evaluated in Section 6.2.1.

A higher-resolution converter could nevertheless become advantageous in a future revision if combined with an acquisition architecture better suited to the required 100 Hz output rate. One possible candidate is the Analog Devices AD4111, a 24-bit sigma–delta ADC integrating a basic analogue front end for high-impedance bipolar voltage inputs up to $\pm$10 V. Such a device could potentially simplify the present external conditioning chain while providing a higher-resolution acquisition architecture, and should therefore be investigated in future work.

3.3.5 ADC architecture and channel expansion

Following this comparison, the ADS1115 was retained for the redesigned Watcher. The acquisition architecture was nevertheless modified to separate electrode-voltage acquisition from auxiliary measurements. Whereas the original Watcher used a single ADS1115 for both functions, the redesigned hardware uses two ADS1115 converters. One device is dedicated to the receiving electrodes, with its four single-ended inputs connected to M1–M4, while the second ADC is dedicated to auxiliary measurements, including the 12 V battery voltage and the contact-resistance current monitoring.

Expanding the receiver from two to four electrode channels provides greater flexibility in the geometrical configuration of the measurements. For example, four electrodes can be distributed along two perpendicular directions at two different distances, providing a short and a longer measurement baseline in each direction. However, using all four inputs is not required for operation. The redesigned Watcher can still be used with only two electrode channels, thereby preserving compatibility with the acquisition configuration of the original prototype.

The extension to four electrode channels also increased the number of digital control signals required by the analogue switching and contact-resistance circuits. Each electrode requires two control lines for the DG413 switches and one additional line for its resistance-check relay, resulting in twelve control outputs in total. The 8-bit MCP23008 GPIO expander used in the original Watcher was therefore replaced by an MCP23017, which provides 16 GPIOs through the same I^2C interface [63]. This provides sufficient outputs for the four-electrode architecture without consuming additional GPIOs of the Feather M0.

3.4 Electrode contact-resistance circuit redesign

The contact-resistance measurement principle of the original Watcher was retained, but the associated switching, driving and protection circuitry was revised. As shown in Fig. 3.7, four relays independently connect the receiving electrodes to the resistance-check circuit. Their coils are controlled through a ULN2003A Darlington transistor array, while the test current is measured across a shunt resistor and amplified by the INA282 before digitisation. An additional series resistor was introduced to limit the current under low-resistance or short-circuit conditions. The following sections describe these modifications in greater detail.

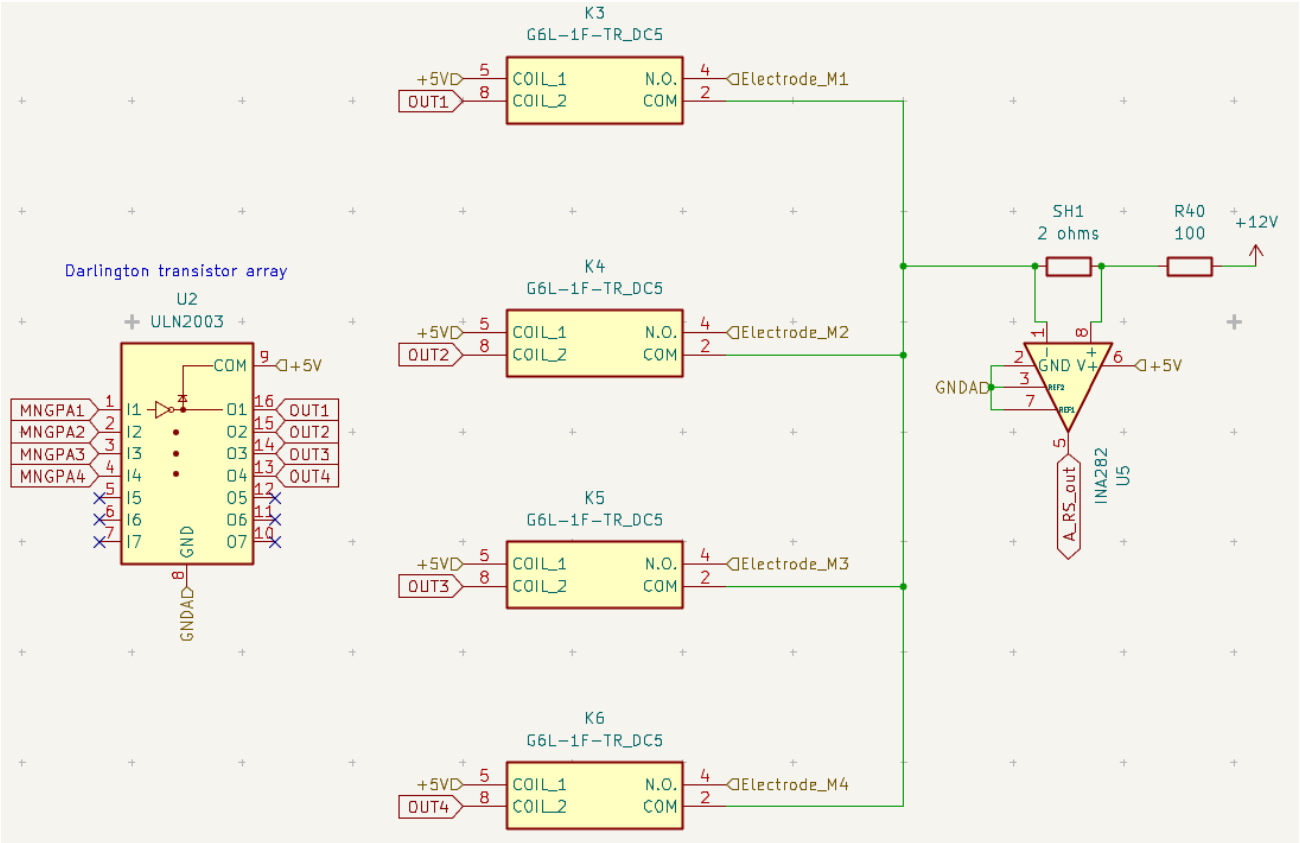

Figure 3.7: Electrode contact-resistance measurement circuit of the redesigned Watcher. The four G6L-1F relays select the electrode under test and are driven by the ULN2003A Darlington transistor array. The resulting test current flows through the series protection resistor and the current-sensing shunt, whose voltage is amplified by the INA282 before digitisation.

3.4.1 Relay driving circuit

The Omron G6L-1F relays used in the original Watcher were retained for the contact-resistance measurement. The selected 5 V version is a single-side-stable SPST-NO relay with a nominal coil current of 36 mA and a

coil power consumption of approximately 180 mW [64]. One relay is associated with each electrode and is energised only when that electrode is selected during the `rs_check()` procedure.

As shown in Fig. 3.7, the relay-driving circuit was however simplified. In the original implementation, each relay coil was controlled by a discrete transistor stage and protected by an external flyback diode. These components were replaced by a ULN2003A Darlington transistor array. The device integrates seven NPN Darlington pairs with open-collector outputs, each rated for collector currents up to 500 mA, largely exceeding the 36 mA required by one G6L-1F coil. Its high current gain allows the low-current logic signal from the MCP23017 to control the relay-coil current, while integrated clamp diodes suppress the inductive voltage generated when the coil is switched off [72]. Several discrete transistor and flyback-diode stages can therefore be replaced by a single package, reducing component count, PCB area and routing complexity.

Latching relays were not adopted despite their lower holding-power requirement. In the present application, only one relay is energised at a time and each remains active for approximately 300 ms, giving only about 1.2 s of cumulative coil activation for a complete four-electrode check. The resulting contribution to the overall energy consumption is therefore quite low. Moreover, the single-side-stable G6L-1F automatically returns to its normally open state when the coil is de-energised [64]. This provides a deterministic fail-safe configuration after a controller reset or loss of power, whereas a latching relay would retain its previous mechanical state. The G6L-1F was therefore retained while its driving circuitry was made more compact through the ULN2003A.

3.4.2 Current sensing and shunt protection

The contact-resistance measurement is based on applying the 12 V battery voltage to the selected electrode and estimating the resulting current from the voltage drop across a 2 Ω shunt resistor. This voltage is measured using an INA282 current-shunt monitor before being digitised by the auxiliary ADS1115, as shown in Fig. 3.7.

The INA282 provides a fixed gain of 50 V/V and is specifically designed to measure small differential voltages across current-sensing resistors at common-mode voltages much larger than its own supply voltage. Its specified common-mode input range extends from -14 to $+80$ V, making it suitable for the 12 V resistance-test circuit. The device also uses a zero-drift architecture and specifies an input offset voltage of 20 μV typical and 70 μV maximum [73]. With both reference inputs connected to ground, the output voltage is ideally

$$V_{\mathrm{INA}} = G_{\mathrm{INA}} V_{\mathrm{shunt}} = 50 I_{\mathrm{test}} R_{\mathrm{shunt}}, \tag{3.10}$$

With $R_{\mathrm{shunt}} = 2\ \Omega$, the test current can therefore be obtained from

$$I_{\mathrm{test}} = \frac{V_{\mathrm{INA}}}{50 R_{\mathrm{shunt}}} = \frac{V_{\mathrm{INA}}}{100\ \Omega}. \tag{3.11}$$

The total resistance of the test path is then obtained from the measured battery voltage and test current. Since the 100 Ω protection resistor and the 2 Ω shunt are also part of this path, their known contributions must be subtracted to estimate the electrode contact resistance:

$$R_{\mathrm{contact}} = \frac{V_{\mathrm{bat}}}{I_{\mathrm{test}}} - R_{\mathrm{protection}} - R_{\mathrm{shunt}}. \tag{3.12}$$

The measurable contact resistance range is determined by the usable output range of the INA282 together with the ADC input range. The ADS1115 is supplied from 5 V and its analogue input must remain within its

supply rails [43]. The INA282 is also supplied from 5 V and specifies an output swing up to approximately $V^+ - 0.4$ V and down to approximately $V_{GND} + 0.04$ V [73]. The INA282 therefore represents the limiting component, giving an approximately usable output range of 0.04–4.6 V.

The maximum current that can be measured without saturating the current monitor is consequently

$$I_{max} = \frac{4.6\ \text{V}}{50 \times 2\ \Omega} \approx 46\ \text{mA}. \tag{3.13}$$

For a nominal battery voltage of 12 V, the corresponding minimum measurable contact resistance is

$$R_{contact,min} = \frac{12\ \text{V}}{0.046\ \text{A}} - 100\ \Omega - 2\ \Omega \approx 159\ \Omega. \tag{3.14}$$

Conversely, using the conservative lower output-swing limit of approximately 40 mV as a bound for linear operation gives a corresponding current of

$$I_{min} = \frac{0.04\ \text{V}}{50 \times 2\ \Omega} = 0.40\ \text{mA}. \tag{3.15}$$

The corresponding maximum contact resistance is therefore

$$R_{contact,max} = \frac{12}{0.0004} - 100 - 2 \approx 29.9\ \text{k}\Omega. \tag{3.16}$$

The resulting datasheet-based conservative linear measurement range is therefore approximately 160 Ω to 30 kΩ. These limits remain theoretical and do not include the effects of component tolerances, relay and wiring resistances, battery-voltage variation or measurement noise.

The original Watcher used the 2 Ω shunt without an additional series current-limiting resistor. In the redesigned circuit, a 100 Ω resistor was added in series with the measurement path to limit the current in the event of a very low electrode resistance or an accidental short circuit. Neglecting the other small series resistances, the worst-case current with a 12 V supply becomes

$$I_{short} = \frac{12\ \text{V}}{100\ \Omega + 2\ \Omega} \approx 118\ \text{mA}. \tag{3.17}$$

Under this worst-case condition, the instantaneous power dissipated in the shunt is

$$P_{shunt} = I_{short}^2 R_{shunt} = (0.118)^2 \times 2 \approx 0.028\ \text{W}, \tag{3.18}$$

whereas the protection resistor dissipates approximately

$$P_{protection} = I_{short}^2 R_{protection} = (0.118)^2 \times 100 \approx 1.4\ \text{W}. \tag{3.19}$$

The selected 3 W protection resistor therefore remains within its nominal power rating even under this worst-case condition, which is applied for only approximately 300 ms. The added 100 Ω resistance consequently provides effective current limiting without altering the original 2 Ω shunt and INA282 current-sensing principle.

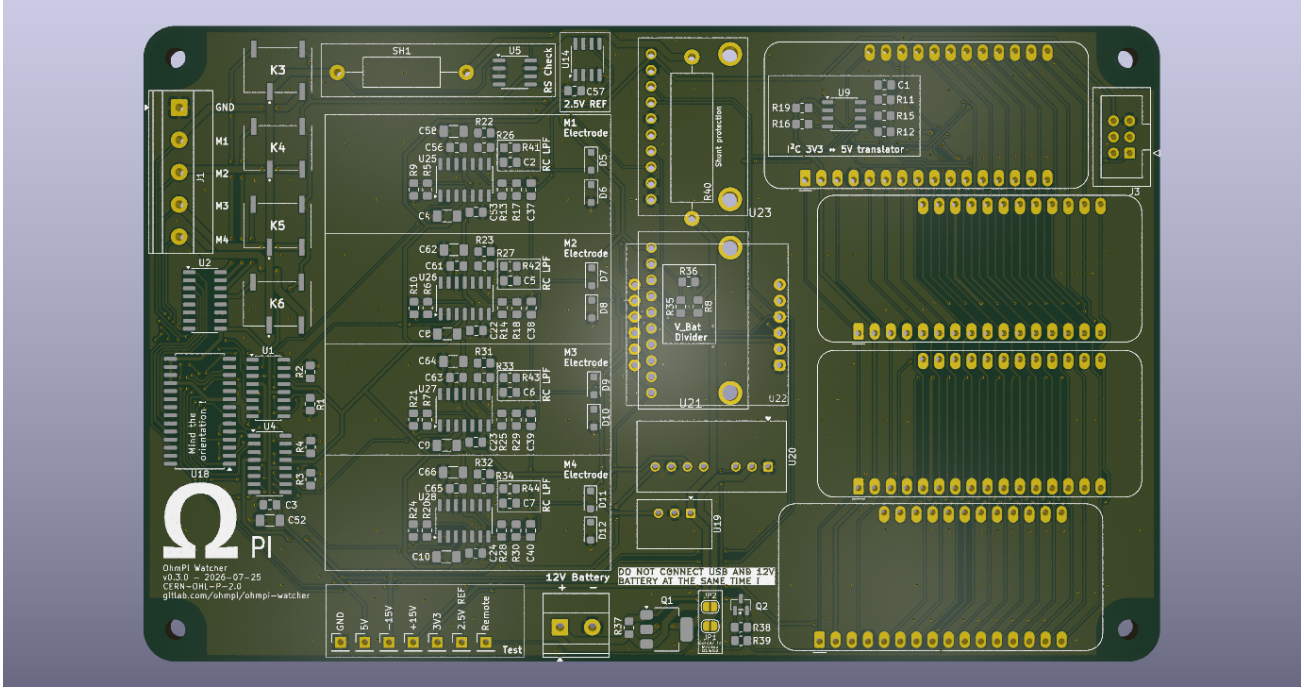

Figure 3.8: Three-dimensional view of the empty PCB in KiCad 10.

3.5 PCB integration and component-level improvements

The redesigned electronics were integrated on a two-layer PCB designed in KiCad 10 (Fig. 3.8 and Fig. 3.1b). Despite the increase from two to four electrode acquisition channels and the addition of the associated analogue and switching circuitry, the complete system could be fitted within a standard 100 mm × 160 mm Eurocard format. Component placement was organised by functional blocks, including the power-supply, analogue acquisition, contact-resistance measurement and digital control sections. Where mechanically possible, low-profile components were also placed underneath modules mounted on pin headers. This made use of otherwise unused PCB area and contributed to maintaining a compact overall footprint while still leaving some space available for future modifications.

Most of the redesigned circuitry was implemented using surface-mount components. Since the prototype was assembled manually, package selection also considered accessibility for hand soldering rather than minimising component size alone. Component package selection should account for the intended assembly method, as excessively small packages can unnecessarily increase assembly difficulty and development risk [74]. The use of integrated multi-channel devices and surface-mount passive components significantly reduced the space occupied by the circuitry compared with an equivalent implementation based on individual through-hole or discrete components.

Particular attention was also paid to grounding and current-return paths. Ground copper pours were implemented on both the front and back layers and connected through multiple stitching vias. Routing on the back layer was kept to a minimum where it would interrupt the ground plane, in order to maintain continuous and short return-current paths under the signal traces. A discontinuity in the reference plane forces return current to follow a longer path, increasing the loop inductance and potentially increasing switching noise and electromagnetic coupling [74].

For the same reason, signal paths were kept short where practical and local bypass capacitors were placed close to the supply pins of the corresponding integrated circuits. Short connections between an IC and its decoupling capacitor reduce the associated power-ground loop inductance and therefore the switching noise generated on the supply rails [74]. These considerations were particularly important around the analogue

acquisition circuitry, where coupling from the digital and power sections could directly affect the measured signals.

Several component-level changes were made to further simplify the PCB. In particular, the 10 μF electrolytic capacitors were replaced by X7R multilayer ceramic capacitors. Their smaller surface-mount packages facilitated the compact placement of the redesigned circuitry. MLCCs also provide low-inductance local decoupling when connected through short traces [74].

Some modularity was deliberately retained despite the increased level of PCB integration. In particular, the ADS1115 converters were kept on removable breakout boards connected through socket headers. This allows an ADC board to be replaced during development without desoldering the main PCB. During the design phase, an alternative footprint and the required connections for an ADS1219 were also included as a possible upgrade path. Although the ADS1219 was ultimately not selected, retaining this provision reduces the modifications required if a higher-resolution converter is investigated in a future revision.

The PCB was also designed to simplify assembly, bring-up and field diagnostics. As shown in Fig. 3.8, functional areas were identified directly on the silkscreen, and the four electrode inputs were labelled M1 to M4. Battery polarity was explicitly indicated at the 12 V input connector. A dedicated row of test points provides direct access to GND, 3.3 V, 5 V, +15 V, −15 V, the 2.5 V analogue reference and the analogue-supply REMOTE control signal. Providing clearly identified test points for important power rails and signals is recommended as a good practice because it facilitates measurements and reduces the risk of probing an incorrect node [74].

Finally, an operational warning was added directly to the silkscreen near the Feather M0 USB connector. Since the Watcher supplies the controller from its own external power architecture, simultaneous connection of the normal USB power line and the externally powered system is avoided during debugging. When USB serial communication is required while the Watcher remains battery-powered, a data-only USB connection can instead be used. This precaution prevents an unintended power path between the externally supplied board and the connected computer [61].

Overall, the PCB redesign did not simply accommodate the additional hardware functions, but also improved integration, accessibility and maintainability. Four electrode acquisition channels, the redesigned power supplies and the revised contact-resistance circuitry were integrated within the same Eurocard footprint while retaining modular elements and providing dedicated features for assembly, debugging and future hardware modifications.

3.6 Final hardware architecture

The complete electrical schematic of the redesigned Watcher is provided in Appendix A. It integrates the power architecture, the four analogue acquisition channels, the electrode switching and contact-resistance circuits, and the digital control interfaces described throughout this chapter.

The final hardware therefore retains the main acquisition, synchronisation, storage and communication functions of the original Watcher while integrating the modifications introduced in this work on a single 100 mm × 160 mm PCB.

Chapter 4

Software development

The Watcher firmware is implemented in C++ using the Arduino framework on the Feather M0. The software development focused on adapting the original firmware to the revised hardware while improving its structure, resource management and autonomous operation. The main acquisition, timing and diagnostic principles of the original implementation were retained, but the execution flow was reorganised to better separate acquisition, communication, peripheral reconfiguration and low-power operation. The firmware was also extended to support the additional analogue channels, revised diagnostic routines and RTC-based wake-up mechanism. Table 4.1 presents an overview of the main software modifications, while the following sections describe these modifications in more detail, as well as the main design choices introduced to improve the maintainability and reliability of the embedded software.

Table 4.1: Main software modifications introduced in the redesigned Watcher firmware.

Function	Original implementation	Redesigned implementation
Runtime control	Distributed conditional logic and multiple state flags	Explicit finite-state machine
Electrode handling	Two channel-specific routines	Indexed four-channel routines
SPI management	Shared peripheral access	State-dependent shared/dedicated SD access
Diagnostics	Temporary sample arrays	Online statistical processing
Night operation	Periodic watchdog wake-up	RTC-triggered alarm

For reproducibility, the main software libraries used in the final firmware are summarised in Table 4.2.

Table 4.2: Main software libraries used in the final Watcher firmware.

Library	Version
Adafruit SAMD Boards	1.7.17
SdFat	2.3.0
RadioLib	7.7.1
Adafruit ADS1X15	2.6.2
RTClib	2.1.4
Arduino Low Power	1.2.2
Adafruit MCP23017	2.3.2
Adafruit GPS Library	1.7.5

4.1 Runtime architecture and resource management

Beyond adaptation to the redesigned hardware, the runtime architecture was reorganised to explicitly separate the main operating modes and their transitions. The resulting implementation remains single-threaded: periodic acquisition is scheduled cooperatively from the main loop, while asynchronous external events are signalled through hardware interrupts. The current operating mode is represented by a finite-state machine.

4.1.1 Original control flow and limitations

The original firmware followed a conventional Arduino `setup()`–`loop()` structure in which the main loop successively evaluated numerous conditions related to acquisition, diagnostics, GNSS/PPS timing, user inputs and the daily operating schedule. The receiver state was therefore represented only implicitly through a combination of Boolean flags, time conditions and peripheral states rather than by a single explicit program state. A simplified representation of this execution structure is shown in Fig. 4.1a. A more detailed execution diagram of the original firmware is provided in Appendix D.

Several higher-level routines also grouped together functions with very different purposes. For example, `morningRoutine()` combined GNSS synchronisation, LoRa communication, SD-card configuration, file creation and diagnostic measurements, while `eveningRoutine()` handled communication, storage reconfiguration and the transition towards sleep. This structure was sufficient for the initial prototype, but became increasingly difficult to extend as SPI arbitration, additional hardware control and RTC-based deep sleep were introduced. The firmware was therefore reorganised around an explicit finite-state machine.

4.1.2 Finite-state machine architecture and event handling

In the redesigned firmware, the current operating mode of the Watcher is represented by the `WatcherState` string variable. The main `loop()` is organised around a `switch` statement, so that only the operations associated with the current state are executed. This provides a clear separation between continuous operating modes and the operations required to transition safely between them.

The main functional states are summarised in Table 4.3. Three states represent relatively stable operating modes: `IDLE_SHARED_SPI`, `RECORDING_DEDICATED_SPI` and `SLEEP`. The remaining states perform explicit transitions between these modes. This prevents resource reconfiguration, file handling and peripheral power switching from being distributed across unrelated sections of the main loop.

The corresponding call structure is shown in Fig. 4.1b. In contrast with the original implementation, peripheral-specific functions are now reached through state-dependent execution paths: voltage acquisition is confined to the recording state, while full radio-packet processing is performed only when the SPI bus is shared.

State transitions are triggered by events rather than by a fixed sequence of function calls. A remote measurement request can initiate the transition from idle to recording, a LoRa reception event detected during acquisition can request a return to a communication-capable state, and the daily operating schedule can initiate the transition towards sleep. Conversely, the RTC interrupt provides the event required to leave the deep-sleep state. Local push-buttons are handled in the same state-dependent manner and can, for example, request the start or stop of an acquisition.

Interrupts are deliberately limited to event signalling. A LoRa reception interrupt therefore only sets a volatile flag indicating that a packet is pending; the corresponding processing is deferred to the main execution flow. This keeps interrupt service routines short and avoids performing lengthy peripheral transactions asynchronously from the state-machine logic.

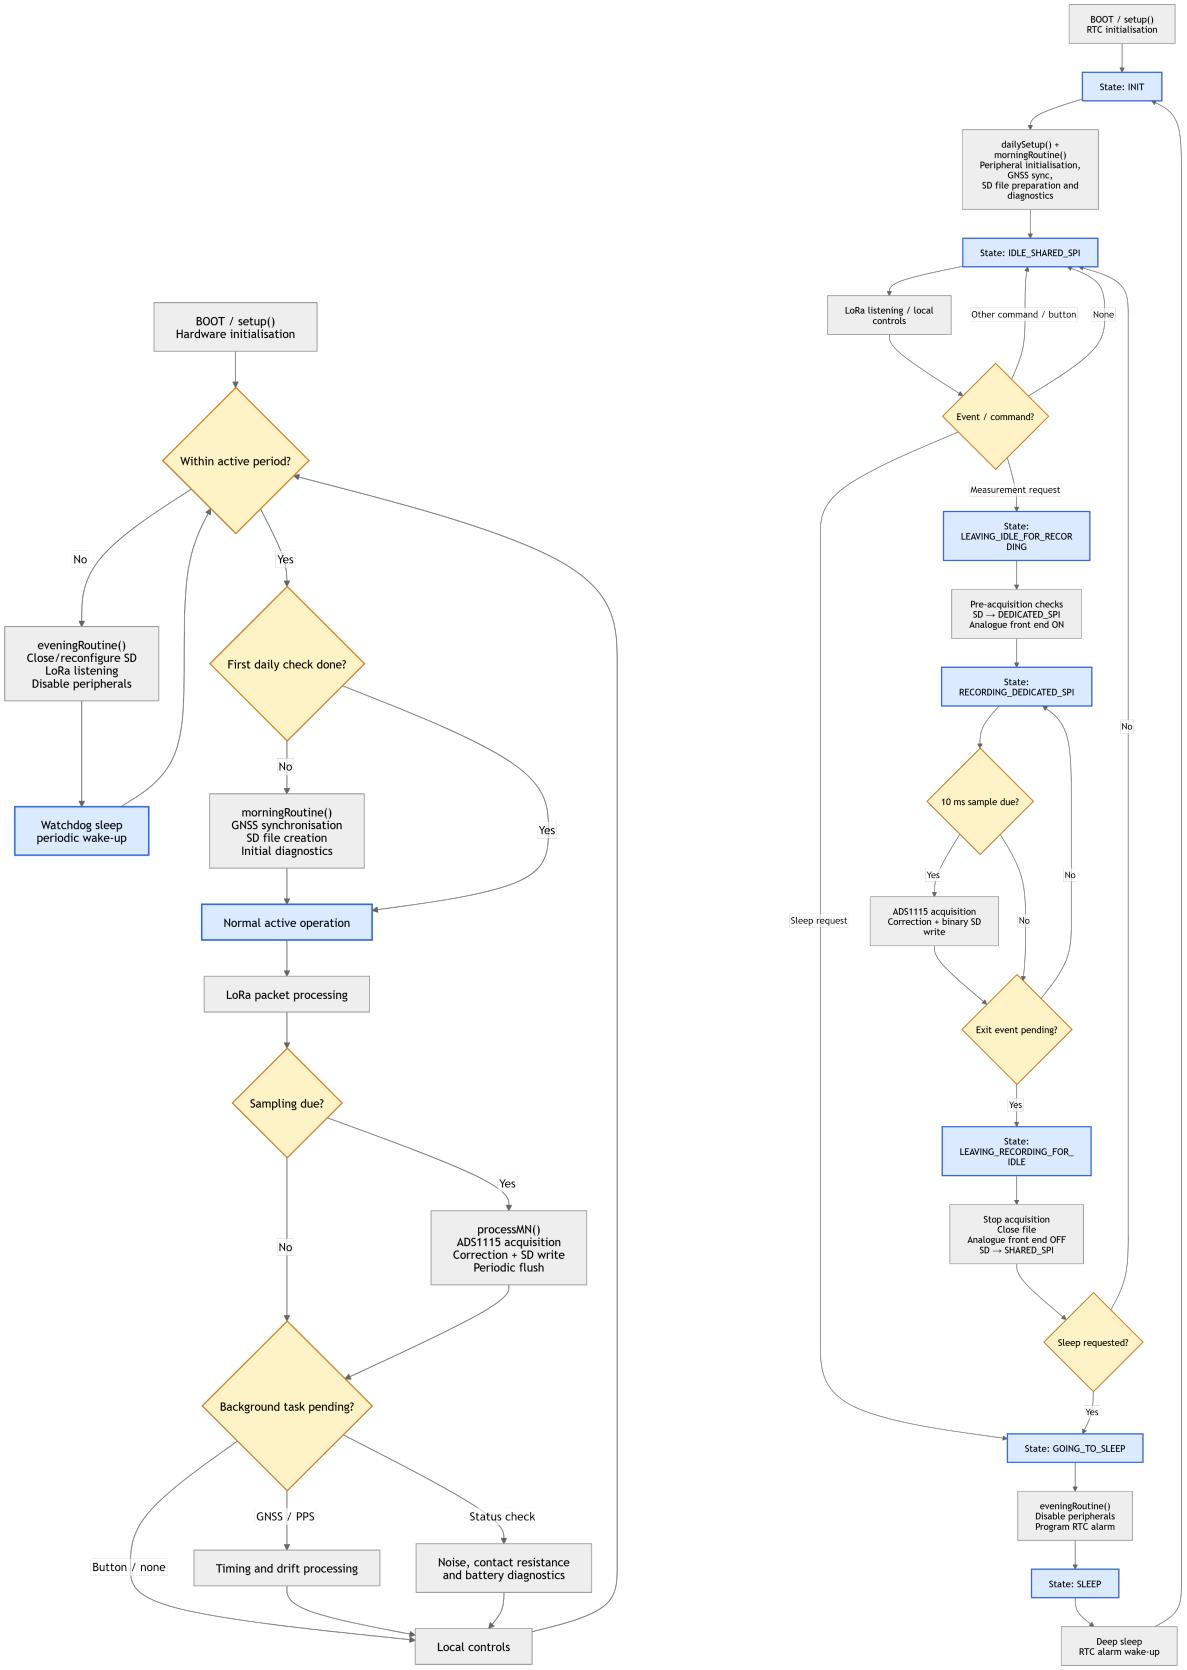

(a) Original firmware. **(b)** Redesigned firmware.

Figure 4.1: Simplified comparison of the Watcher firmware execution structure before and after the software redesign. The original implementation coordinates acquisition, communication, timing, diagnostics and power management through multiple conditional paths, whereas the redesigned implementation separates these functions through explicit states and controlled transitions. Detailed execution diagrams of both implementations are provided in Appendix D.

Table 4.3: Main runtime states of the redesigned Watcher firmware.

Watcher State	Main function
`INIT`	Initialises the peripherals required for daily operation and executes the start-up routine.
`IDLE_SHARED_SPI`	Keeps the receiver available for local and remote commands while allowing the SD card and LoRa radio to share the SPI bus.
`LEAVING_IDLE_FOR_RECORDING`	Prepares the analogue front end and reconfigures SD-card access before continuous acquisition.
`RECORDING_DEDICATED_SPI`	Prioritises periodic voltage acquisition and binary SD-card logging.
`LEAVING_RECORDING_FOR_IDLE`	Stops acquisition, returns the analogue circuitry to its inactive configuration and restores shared SPI access.
`GOING_TO_SLEEP`	Prepares the hardware for low-power operation and configures the RTC wake-up condition.
`SLEEP`	Places the microcontroller in deep sleep until an external RTC wake-up event occurs.

The refactoring was also extended to peripheral initialisation. The original `setup()` routine initialised most devices directly, whereas the redesigned firmware separates these operations into `initRTC()`, `initRadio()`, `initADC()`, `initMCP()` and `initGPS()`, coordinated by `dailySetup()`. This improves readability and allows individual peripherals to be reinitialised or modified with fewer dependencies on the remainder of the firmware.

4.1.3 SD-card and LoRa SPI arbitration

A major motivation for the explicit runtime states was the need to share the SPI bus between the SD card and the RFM96 LoRa transceiver. During normal idle operation, both devices must remain available: the SD file may be accessed for status or previously recorded information, while LoRa packets must be received and processed. In this mode, the SD card is therefore initialised using the shared-SPI configuration.

Continuous acquisition imposes a different requirement. Voltage records are generated at short and regular intervals, and the SD card becomes the dominant SPI peripheral during this period. Preliminary execution-time measurements showed that writing one measurement record required approximately 6000 μs with the SD card configured in shared SPI mode, compared with approximately 1000 μs in dedicated SPI mode. In addition, acquisition and processing of two electrode voltages required approximately 5400 μs. Shared-mode SD writing would therefore leave insufficient time to reliably complete the acquisition cycle within the target 10 ms sampling period.

Ideally, LoRa communication would remain fully available while voltage acquisition continues uninterrupted. This would allow incoming commands to be processed without introducing gaps in the recorded time series. However, with the existing Feather M0 peripheral allocation, the SD card and RFM96 transceiver share the same hardware SPI bus. Maintaining shared access to both peripherals was thus incompatible with the measured timing requirement. Under these hardware and timing constraints, prioritising the SD card during continuous acquisition was therefore considered the most practical compromise.

The firmware behaviour directly reflects this arbitration strategy. The system initialises the SD interface using the dedicated SPI configuration when entering `RECORDING_DEDICATED_SPI`. In this state, voltage acquisition and SD-card logging are given priority, while full LoRa packet processing is deferred. This does not prevent the receiver from detecting the arrival of a radio packet. The RFM96 `IRQ` line (see top right

of Fig. 3.2) remains active and is used to signal reception without immediately performing SPI transactions. The associated interrupt callback only sets the `loraPacketReceived` flag.

When this flag is detected, the state machine initiates a controlled transition out of the recording state. During `LEAVING_RECORDING_FOR_IDLE`, the measurement file is closed, the analogue acquisition chain is returned to its inactive configuration and the SD interface is restarted in shared mode. The pending LoRa packet can then be read and processed from `IDLE_SHARED_SPI`. The interrupt-based mechanism therefore does not provide truly simultaneous acquisition and communication, but it allows radio events to be detected during recording and serviced in a controlled manner without permitting asynchronous LoRa SPI accesses to disrupt the 10 ms acquisition cycle.

The resulting architecture deliberately favours deterministic acquisition over continuous command processing while measurements are being recorded. A controller providing independent SPI interfaces could in principle use separate buses for the SD card and LoRa transceiver and thereby allow radio transactions without reconfiguring the storage interface. However, as discussed in Section 3.1, migrating to a different controller platform would have required substantial firmware refactoring and revalidation of the acquisition and communication functions. The development effort was therefore prioritised towards other hardware and software improvements, and SPI arbitration was implemented within the constraints of the retained Feather M0 platform.

The SPI arbitration is treated primarily as a runtime-resource management problem rather than as part of the SD-card logging algorithm itself. The details of the binary record structure, buffering and SD-card write strategy are presented in Section 4.2, while the radio driver, packet structure and multi-hop communication behaviour are discussed separately in Chapter 5.

4.2 Data acquisition and logging

Once the runtime architecture has established which peripherals can be used simultaneously, the main acquisition constraint becomes the regular generation and storage of voltage samples within the 10 ms sampling period. This section describes the corresponding scheduling, diagnostic and storage mechanisms.

4.2.1 Acquisition scheduler and channel handling

Continuous voltage acquisition is organised around a sampling period of 10 ms, corresponding to 100 Hz. During `RECORDING_DEDICATED_SPI`, the firmware checks whether the next sampling instant has been reached and, when required, acquires the electrode voltages, applies the channel-specific offset and gain correction, constructs the corresponding binary record and writes it to the SD card. Using the dedicated-SPI configuration described in Section 4.1.3, this execution sequence is designed to fit within the 10 ms acquisition interval.

The acquisition routines were also generalised to match the four-electrode hardware architecture. Instead of handling individual channels through separate variables and repeated instructions, the redesigned firmware uses indexed arrays such as `mnChannel[4]`, `mgain[4]` and `moffset[4]`. Electrode voltages can therefore be acquired and converted using the same loop for all channels. The two ADS1115 devices are also separated according to their function: `adcMN` is dedicated to the electrode inputs M1–M4, while `adcAux` handles lower-rate auxiliary measurements such as battery voltage and contact-resistance current. This organisation reduces duplicated code and simplifies further changes to the number or configuration of measurement channels.

The sampling scheduler is designed so that the next acquisition instant is referenced to the intended sampling sequence rather than to the completion time of the previous measurement. Conceptually, $t_{k+1} = t_k + T_s$, with $T_s = 10$ ms, so that an occasional execution delay does not permanently shift all subsequent samples. Meeting this deadline depends on the combined ADC conversion, processing and SD-write times. The

practical execution-time margin, including the comparison between two- and four-channel acquisition, is evaluated experimentally in Section 6.2.1.

4.2.2 Diagnostic and self-check routines

In addition to continuous voltage acquisition, the firmware implements several self-check routines used to evaluate the condition of the acquisition electronics, electrode connections and power supply. These operations are executed outside the time-critical 10 ms sampling loop, for example during start-up or when a status check is requested. The higher-level `getStatus()` routine coordinates the analogue noise and offset measurement, the electrode contact-resistance check and the battery-voltage measurement, and can store the resulting diagnostic values on the SD card.

The analogue-chain diagnostic disconnects the field electrodes and connects the four conditioning-chain inputs to GNDA through the analogue switches. For each channel, `adc_noise()` acquires $N = 100$ samples and evaluates their mean and standard deviation. The mean estimates the output offset of the complete conditioning and conversion chain relative to its nominal 2.5 V operating point, while the standard deviation provides an indicator of the electronic noise. In the redesigned firmware, these statistics are computed incrementally using Welford's algorithm [75]. The mean and variance are updated for each new sample, avoiding the need to store the complete sample sequence in RAM. For 100 single-precision samples, this reduces the temporary storage requirement from 400 bytes to only a few running variables.

The contact-resistance routine uses a different analogue configuration. The measurement chains are first placed in a floating state so that the external resistance-test excitation is isolated from the voltage acquisition paths. The battery voltage is measured once, after which the four electrode relays are activated sequentially. For each electrode, the relay remains energised for the configured injection duration, the INA282 output is sampled through the auxiliary ADS1115, and the test current is calculated from the known shunt resistance and current-sense gain. The electrode contact resistance is then obtained after subtracting the known series protection and shunt resistances from the total measured loop resistance. Using a common indexed routine for the four relays and result values avoids the duplicated two-channel logic of the original firmware.

Together, the offset, noise, contact-resistance and battery measurements provide a compact assessment of the main system before or between acquisition periods. Their execution is kept separate from continuous voltage recording so that the comparatively long settling delays, multiple ADC conversions and relay activation times do not interfere with the regular sampling schedule.

4.2.3 Binary record format and storage requirements

Measurement and diagnostic data are stored in binary form rather than being converted to formatted text before each SD-card write. This reduces both the amount of data written and the processing required inside the acquisition loop, while providing a fixed record size that simplifies subsequent file parsing. The main data structure used by the firmware is defined as follows:

```
struct Record {
    uint32_t d;
    float v1;
    float v2;
    uint32_t t;
};
```

On the Feather M0, each of these fields occupies 4 bytes, giving a fixed record size of 16 bytes.

The first field, `d`, identifies the type of information stored, while `v1` and `v2` contain the corresponding measurement or diagnostic values and `t` contains the local timestamp. A second structure, `TRecord`, uses

the same 16-byte layout but interprets the two central fields as 32-bit integer timing values rather than floating-point measurements. The main record categories generated by the firmware are summarised in Table 4.4.

Although the acquisition and diagnostic routines were generalised to four electrode channels, the binary storage format was initially kept compatible with the original two-channel implementation. The current `Record` structure therefore contains two measurement fields; four-channel storage would require an extension of this format and the corresponding post-processing routines.

Table 4.4: Main binary record types used by the Watcher firmware. Differences between record types 4, 6 and 7 are explained in Section 4.3.2.

`d`	**Record**	**First data field** `v1`	**Second data field** `v2`
0	Electrode voltage	M1 voltage	M2 voltage
1	Battery status	Battery voltage	Unused
2	Analogue noise	M1 standard deviation	M2 standard deviation
3	Contact resistance	M1 resistance	M2 resistance
4	GNSS timing	GNSS time	RTC time
5	Analogue offset	M1 offset	M2 offset
6	PPS timing	GNSS/PPS time	RTC time
7	Clock-drift reference	GNSS/PPS time	RTC time
10	Position	Latitude	Longitude

The storage requirement is dominated by type 0 voltage records because they are generated at the target rate of 100 Hz, whereas diagnostic and timing records are produced only occasionally. With a 16-byte record, the continuous voltage data rate is

$$\dot{D} = 16 \text{ bytes} \times 100 \text{ s}^{-1} = 1600 \text{ bytes}\,\text{s}^{-1}. \tag{4.1}$$

The amount of voltage data generated during an acquisition period of T hours is therefore

$$D(T) = 16 \times 100 \times 3600 \times T = 5.76\,T \text{ MB}, \tag{4.2}$$

This corresponds to approximately 5.76 MB for one hour, 34.56 MB for a 6-hour acquisition period and 69.12 MB for 12 hours. For a five-day field campaign containing 6 hours of recording per day, the voltage records alone would therefore occupy approximately

$$D_{\text{campaign}} = 5 \times 34.56 = 172.8 \text{ MB}. \tag{4.3}$$

The contribution of the lower-rate diagnostic, position and timing records is small relative to this continuous voltage data stream. These values also show that storage capacity is not a restrictive design parameter: even a modest-capacity SD card provides a large margin for a multi-day field campaign.

The current firmware preallocates approximately 65 MiB for each daily measurement file. This corresponds to

$$65 \times 2^{20} = 68\,157\,440 \text{ bytes}, \tag{4.4}$$

which, considering the continuous 1600 bytes s^{-1} voltage stream alone, represents approximately

$$T_{\text{prealloc}} = \frac{68\,157\,440}{1600 \times 3600} \approx 11.8 \text{ h} \tag{4.5}$$

of continuous two-channel acquisition. The selected preallocation is therefore sufficient for a typical daytime measurement period while avoiding file-system allocation operations during time-critical recording. The implementation and implications of file preallocation, sector buffering and file synchronisation are discussed in Section 4.2.4.

If all four available electrode voltages are eventually stored, a straightforward extension containing one 32-bit type field, four 32-bit floating-point voltage fields and one 32-bit timestamp would occupy

$$S_{4\text{ch}} = 4 + 4 \times 4 + 4 = 24 \text{ bytes.} \tag{4.6}$$

At 100 Hz, this would increase the continuous data rate to 2400 bytes s^{-1}, or approximately 8.64 MB/h. An 8-hour four-channel acquisition would therefore require approximately 69.12 MB. Storage capacity itself is consequently not a limiting factor for the four-channel extension; the more restrictive constraint remains the execution time required to perform four sequential ADC conversions within the 10 ms sampling interval.

4.2.4 SD-card buffering and file management

The SD-card interface is managed using version 2.3.0 of the SdFat library developed by Greiman [76].

Since each measurement record occupies only 16 bytes, writing every record as an independent 512-byte SD-card transaction would be inefficient. Partial-sector writes are therefore handled through the library's data cache. Once a complete sector has been assembled, SdFat forces the corresponding sector write to the SD card, while the following measurements begin filling the next sector [76].

Removable SD-card storage was retained for the redesigned Watcher because it combines sufficient write performance and capacity with particularly simple field data retrieval. SD cards are inexpensive, widely available and can be removed from the receiver and inserted directly into a computer for high-speed data transfer. This avoids transferring several hundred megabytes of campaign data through the Watcher's USB or serial interface and also allows cards to be exchanged easily between deployments.

Alternative embedded storage technologies could nevertheless be considered in a future revision. In particular, QSPI NOR flash could provide compact on-board storage with direct high-speed access from the microcontroller and would eliminate the removable-card connector. However, the recorded data would then need to be retrieved through the Watcher itself, requiring an additional connector or a short-range wireless data-transfer procedure. This could be facilitated by selecting a microcontroller platform with integrated Bluetooth or Wi-Fi connectivity in addition to the LoRa radio, such as an ESP32, allowing the recorded files to be transferred locally without removing the device from its enclosure. In the present implementation, however, removable SD-card storage provides a simpler solution, since the recorded data can be retrieved directly without requiring an additional communication interface.

A standard SD-card sector contains 512 bytes. With the 16-byte binary record format described in Section 4.2.3, one sector can contain

$$N_{\text{record/sector}} = \frac{512}{16} = 32 \tag{4.7}$$

measurement records. At the target acquisition rate of 100 Hz, one sector therefore represents

$$T_{\text{sector}} = 32 \times 10 \text{ ms} = 320 \text{ ms} \tag{4.8}$$

of voltage measurements. The buffering mechanism therefore groups the small 16-byte records into storage units that match the block-oriented organisation of the SD card, while limiting the amount of data held in the current incomplete sector.

File synchronisation must, however, be distinguished from normal record writing. In SdFat, `write()` transfers measurement data through the sector cache, while `flush()` synchronises pending cached information and file-system metadata with the storage device. The `close()` operation also performs this synchronisation before closing the file [76]. Preliminary measurements on the Watcher showed that an explicit `flush()` required approximately 20 ms under the tested configuration. This is twice the target 10 ms sampling period and would therefore introduce an unacceptable blocking operation into the continuous acquisition loop.

The redesigned firmware consequently performs no periodic `flush()` operation during continuous voltage recording. Measurements are successively passed to `write()`, and each completed 512-byte sector is forced out of the SdFat data cache as it becomes full. At most one incomplete 512-byte sector of measurement data remains only in microcontroller memory at a given time. This corresponds to up to 32 records, or approximately 0.32 s of acquisition. The possibility of losing this final incomplete group of measurements in the event of an unexpected power interruption was considered acceptable compared with the deterministic timing penalty that would be introduced by repeated file synchronisation. This behaviour was evaluated experimentally through controlled power-loss tests, as presented in Section 6.2.3.

When acquisition is stopped normally, sampling is first terminated and the file is then closed. Since `close()` performs the required synchronisation, an additional `flush()` immediately beforehand is unnecessary [76]. The potentially long synchronisation operation is therefore moved to a state transition, where it no longer competes with the 10 ms sampling deadline.

To further reduce file-system activity during acquisition, the measurement file is preallocated when it is created. SdFat provides the `preAllocate()` function to reserve contiguous storage for an initially empty file [76]. By allocating the required clusters before continuous recording begins, additional cluster allocation and FAT updates are avoided as the file progressively grows. The preallocation size is selected from the storage requirement derived in Section 4.2.3.

The use of preallocation also provides an advantage in the event of an uncontrolled interruption. Since the storage region has already been allocated before acquisition, complete measurement sectors written before the interruption remain located within the preallocated file region even if the normal finalisation procedure could not be executed. The valid portion of the binary data could therefore be identified during post-processing by scanning the fixed-size records up to the last consistent measurement. During normal operation, the firmware preserves the position of the last valid record when the file is temporarily closed and restores this position when acquisition resumes.

The resulting file-management strategy combines fixed-size binary records, sector buffering, contiguous file preallocation and deferred synchronisation. Together with the dedicated-SPI configuration described in Section 4.1.3, these choices minimise storage-management operations within the 10 ms acquisition window while retaining clean file finalisation during controlled state transitions.

Regular sampling and reliable storage are only useful if measurements recorded by independent receivers can subsequently be placed on a common time axis. The following section therefore describes how the Watcher establishes and maintains its temporal reference.

4.3 Time reference and synchronisation

Accurate timing is required to relate the voltage time series recorded by independent Watchers to the current-injection sequence. Since the receivers operate autonomously and do not share a physical clock, each unit must establish a common absolute time reference while maintaining a sufficiently regular local time base during acquisition. The firmware therefore combines three complementary timing sources: the GNSS receiver provides an absolute UTC reference, the RTC maintains calendar time when the GNSS module is inactive, and the microcontroller's `millis()` counter provides the high-resolution local time base used to timestamp individual measurements. The GNSS pulse-per-second (PPS) signal is additionally used to characterise the drift of this local oscillator.

The timing and synchronisation mechanism described in this section was inherited from the original Watcher firmware and reused without modification in the present work. It is nevertheless described here to clarify how measurement timestamps are generated and how timing drift is estimated, since these functions remain part of the final firmware architecture. The contribution of the present work related to the RTC concerns its additional use for low-power scheduling and wake-up from deep sleep, as described in Section 4.4.2.

4.3.1 GNSS, RTC and local time reference

At start-up, the GNSS receiver is used whenever possible to establish an absolute time reference based on NMEA strings, as explained in Section 2.5.3. Once a valid GNSS time has been obtained, the RTC can be updated accordingly and subsequently maintains the date and time while the GNSS receiver is switched off. This separation is important for power management: continuous GNSS operation is not required simply to retain calendar time. If a GNSS fix is temporarily unavailable, the RTC provides the reference required to initialise the acquisition until a new GNSS synchronisation can be performed. The RTC is therefore used as a persistent low-power time reference, whereas GNSS provides the external reference from which its accuracy can be checked and corrected.

In the intended field-operation strategy, GNSS resynchronisation would not be performed periodically except during startup and before going to sleep. Instead, it could be triggered on demand through a LoRa command sent by the injection team. This would allow a new absolute time reference to be acquired at convenient moments, for example during pauses in the injection sequence, while avoiding unnecessary GNSS activation during normal acquisition periods. The injection team could therefore decide when a timing update is required and initiate the synchronisation remotely.

Neither GNSS messages nor the RTC are used directly to timestamp every 100-Hz voltage sample. Instead, the firmware establishes `tstart`, which relates the microcontroller's `millis()` counter to the time of day expressed in milliseconds. When a valid GNSS time is obtained, `gpsData.gpsTime` contains the number of milliseconds elapsed since midnight. The corresponding relations are

$$\texttt{tstart} = \texttt{millis()} - \texttt{gpsData.gpsTime}, \tag{4.9}$$

$$\texttt{tstart} = \texttt{millis()} - \texttt{getTimeRTC()}, \tag{4.10}$$

$$\texttt{record.t} = \texttt{millis()} - \texttt{tstart}, \tag{4.11}$$

where Eq. 4.9 is used when a valid GNSS reference is available, Eq. 4.10 provides the RTC-based fallback when no GNSS fix is obtained before the configured timeout, and Eq. 4.11 is then applied to each recorded voltage sample.

The resulting `record.t` directly represents the number of milliseconds elapsed since midnight at the sampling instant. Its calculation requires only a subtraction inside the time-critical acquisition loop and therefore avoids repeated accesses to the GNSS receiver or RTC. The RTC-based initialisation has a lower

temporal resolution because the RTC provides time only to the nearest second, but it allows acquisition to proceed when a GNSS fix is unavailable.

4.3.2 Pulse-per-second (PPS) based drift estimation

Although `millis()` provides a convenient local time base for timestamping the voltage measurements, it is derived from the microcontroller clock and can progressively deviate from the GNSS time reference. The GNSS receiver therefore provides an additional pulse-per-second (PPS) signal on `GPS_PPS`. As discussed in Section 2.5.3, the PPS signal provides a more precise timing marker than the arrival of an NMEA message. The firmware uses this reference to evaluate the error accumulated by the local time reconstruction.

After a valid GNSS message has been processed, the firmware stores the three available time references as

$$\texttt{gpsTime} = \texttt{gpsData.gpsTime}, \tag{4.12}$$

$$\texttt{millisTime} = \texttt{millis()} - \texttt{tstart}, \tag{4.13}$$

and

$$\texttt{rtcTime} = \texttt{getTimeRTC()}. \tag{4.14}$$

These values are saved in a type 4 timing record. The firmware then sets the `checkPPS` flag and waits for the next rising edge of `GPS_PPS`. If the PPS signal is already high when the GNSS message is processed, `waitLowPPS` first forces the firmware to wait for the signal to return low, preventing the same pulse from being interpreted as the following second.

When the next PPS transition is detected, the GNSS time contained in the previously received message is rounded to the corresponding integer second. In the firmware, this reference is calculated as

$$\texttt{gpsTime} = \left\lceil \frac{\texttt{gpsData.gpsTime}}{1000} \right\rceil \times 1000, \tag{4.15}$$

with an additional 1000 ms correction if the first expected PPS pulse was missed. At the same instant, the local reconstructed time is obtained again as

$$\texttt{millisTime} = \texttt{millis()} - \texttt{tstart}. \tag{4.16}$$

The local timing error estimated from the PPS reference is then

$$\texttt{adrift} = \texttt{gpsTime} - \texttt{millisTime}. \tag{4.17}$$

The variable `drift` stores the latest accepted value of this difference. A consistency check is applied before updating it:

$$|\texttt{adrift} - \texttt{drift}| < 800 \text{ ms}. \tag{4.18}$$

This relatively large threshold is used to reject an erroneous association with the neighbouring PPS second, which would introduce an error close to 1 s. Accepted drift estimates are stored as type 7 records, while the PPS time references themselves are stored as type 6 records. The RTC value is recorded simultaneously,

allowing the behaviour of both the microcontroller time base and the RTC to be compared with the GNSS reference during post-processing.

The resulting `drift` value can subsequently be applied when a measurement is retrieved for a requested absolute time. In particular, the firmware corrects the requested timestamp before searching the binary file according to

$$t_{\text{search}} = t_{\text{requested}} + \texttt{drift}. \tag{4.19}$$

The PPS mechanism therefore does not modify the timestamp of every record during acquisition. Instead, it provides discrete reference points characterising the difference between the GNSS time and the local `millis()`-based time scale. This preserves the lightweight timestamping operation required by the 10 ms acquisition loop while retaining the information required to assess and correct timing offsets afterwards.

The RTC is not only used to preserve the time reference while the GNSS receiver is inactive. In the redesigned firmware, it also becomes the key component enabling the Watcher to remain in deep sleep between daily operating periods.

4.4 Low-power scheduling and wake-up

Multiple-day field deployment requires the Watcher to remain active only during periods in which measurements or remote interactions are expected. The redesigned firmware therefore integrates low-power operation directly into the runtime state machine rather than treating sleep as a delay inside the main loop. Two configuration parameters, `hour_start` and `hour_stop`, define the daily active period in UTC. Outside this interval, the receiver enters a dedicated shutdown sequence and places the microcontroller in deep sleep until the next programmed wake-up time.

4.4.1 Daily operating schedule and peripheral shutdown

The active period is evaluated using the RTC through the `isActiveHour()` function. For the usual case in which `hour_start < hour_stop`, the Watcher is considered active when

$$\texttt{hour_start} \leq h_{\text{RTC}} < \texttt{hour_stop}. \tag{4.20}$$

The implementation also supports operating periods spanning midnight by reversing the logical condition when `hour_start > hour_stop`. Since these parameters are read from the configuration file during `dailySetup()`, the operating schedule can be modified between deployments without recompiling the firmware. When the end of the active period is detected, the state machine transitions to `GOING_TO_SLEEP`. The same state can also be requested remotely through the LoRa command interface, so scheduled and remotely requested shutdowns share the same controlled sleep sequence.

Before entering deep sleep, `eveningRoutine()` places the peripherals in known low-power states. The analogue inputs are disconnected from the electrodes and connected to GNDA, the ± 15 V analogue supply is disabled, pending GNSS and PPS operations are cancelled and the GNSS receiver is switched off. The measurement file is closed and the SD interface is stopped, while the RFM96 transceiver is placed in sleep mode. Finally, the OLED display is switched off. Moving these operations to the explicit `GOING_TO_SLEEP` state ensures that storage and communication resources are finalised before the processor itself is suspended, rather than relying on a collection of independent power-management flags throughout the main loop.

4.4.2 RTC-alarm wake-up and deep sleep

The original firmware periodically woke the processor using the watchdog and checked whether the beginning of the next operating period had been reached. Although simple, this approach requires repeated processor wake-ups throughout the inactive period. The redesigned implementation instead uses the alarm function of the PCF8523 RTC to generate a single hardware wake-up event at the next occurrence of `hour_start`. When entering `GOING_TO_SLEEP`, the firmware first clears the `rtcWakeRequested` flag and calls

`setRTCWakeAlarm(hour_start, 0),`

which configures the RTC alarm for `hour_start:00` UTC. The minute and hour comparisons are enabled, whereas the day and weekday fields are ignored, allowing the same alarm configuration to be used from one day to the next.

The alarm output of the RTC must be connected to an interrupt-capable microcontroller input in order to wake the SAMD21 from deep sleep. This connection was not available through the original Watcher wiring and was therefore added manually by connecting the RTC `INT` output to pin `A1` of the Feather M0 using a jumper wire, as shown in Figure 4.2. In the firmware, this input is defined as `RTC_INT` and configured as an interrupt wake-up source:

```
LowPower.attachInterruptWakeup(
    RTC_INT,
    rtcWakeISR,
    FALLING
);
```

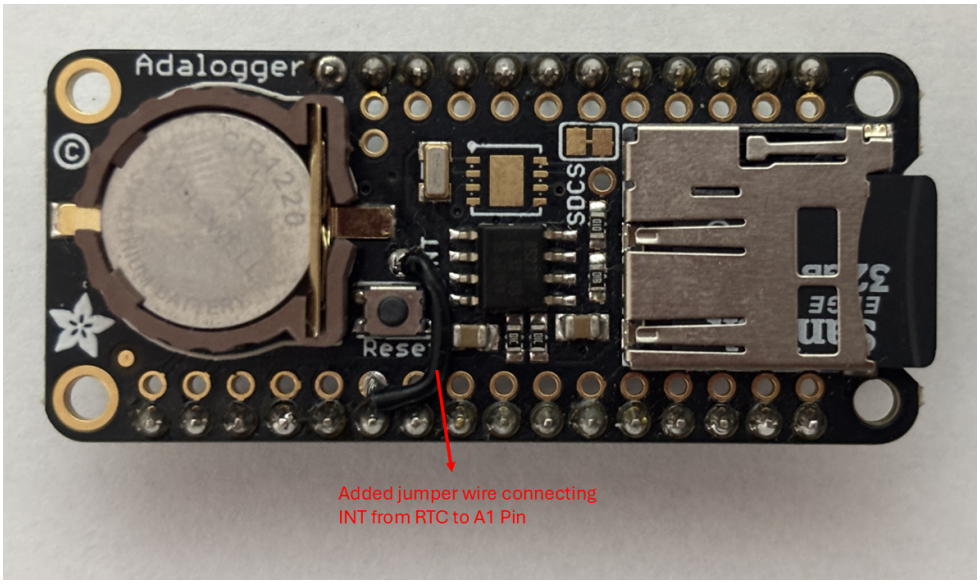

Figure 4.2: Hardware modification connecting the `INT` output of the PCF8523 RTC on the Adalogger board to the interrupt-capable `A1` input of the Feather M0. This connection allows an RTC alarm to wake the microcontroller from deep sleep.

Once all peripherals have been shut down, the state machine enters `SLEEP` and calls `LowPower.deepSleep()` from the `ArduinoLowPower` library. When the programmed RTC alarm occurs, the falling edge on `RTC_INT` executes the lightweight `rtcWakeISR()` callback, whose only operation is

`rtcWakeRequested = true.`

No peripheral communication is performed inside the interrupt service routine. Execution resumes after `deepSleep()`, the alarm is disabled and its pending flag is cleared using `disableRTCAlarm()`, and the runtime state is returned to `INIT`. The normal start-up sequence can then run again through `dailySetup()` and `morningRoutine()`, reinitialising the required peripherals, reloading the configuration, establishing the daily time reference and preparing a new measurement file.

This event-driven wake-up mechanism avoids the repeated watchdog wake-ups required by the previous implementation. During the inactive period, the microcontroller can therefore remain continuously in deep sleep while the low-power RTC independently maintains time and generates the next wake-up event. The measured reduction in Watcher current consumption obtained with this operating mode is evaluated in Section 6.1.1.

Chapter 5

Wireless communication

This chapter presents the improvements introduced to the Watcher's wireless communication system. The developments concern the radio link itself, the design of a portable relay gateway, the associated radio software and the packet-forwarding mechanisms used for multi-hop communication. The main modifications are summarised in Table 5.1, while the following sections describe their implementation and the associated design choices.

Table 5.1: Main wireless communication modifications introduced in this work.

Function	Original implementation	Redesigned implementation
Watcher antenna	Internal horizontal wire antenna	Dedicated 433-MHz antenna with vertical external placement
Link recovery	Communication dependent on successive Watcher-to-Watcher links	Portable gateway able to bridge a problematic radio link
Gateway platform	Not applicable	Autonomous STM32/SX1268-based relay gateway
Packet forwarding robustness	Original store-and-forward routing	Gateway-assisted forwarding with duplicate suppression

5.1 Communication architecture and design objectives

LoRa communication provides the control and coordination link between the remote controller W_0 operated by the injection team and the distributed Watchers. Communication follows a controller-initiated request–response model: W_0 sends a command or status request to a selected Watcher, the packet is relayed along the measurement profile until it reaches the addressed node, and the corresponding response is then forwarded back towards W_0. The Watchers do not autonomously query one another or initiate peer-to-peer exchanges as part of normal network operation.

The high-rate voltage measurements are not transmitted over LoRa but are stored locally on the SD card; the radio is therefore used primarily as a low-data-rate control and diagnostic channel. Reliable communication remains important because loss of this request–response path can prevent distant Watchers from receiving commands or returning their status during a field deployment.

The original Watcher network already relied on that linear multi-hop store-and-forward mechanism. Under normal conditions (Fig. 5.1a), a request issued by W_0 for a distant Watcher is retransmitted by successive intermediate Watchers until the final destination is reached. A response follows the same principle in the

reverse direction. Intermediate Watchers therefore act only as relays for an ongoing controller-initiated transaction, except for commands explicitly intended to be applied along the traversed section of the chain.

This architecture should not be confused with a mesh network. The current protocol does not perform neighbour discovery, dynamic route selection or arbitrary peer-to-peer communication between Watchers. Instead, packet propagation follows the known physical ordering of the receivers along the measurement line.

In practice, however, the continuity of this communication chain depends on the reliability of each successive radio link. The original prototype used a relatively unfavourable antenna arrangement, and the Watchers are typically installed close to the ground, where terrain, vegetation and local obstacles can strongly affect radio propagation. As illustrated in Fig. 5.1b, a single unreliable link between two neighbouring Watchers can therefore interrupt the chain and isolate all units located farther along the profile, even though the forwarding protocol itself is functioning as intended.

The objective of the wireless redesign was therefore primarily to improve the robustness of this existing architecture under realistic deployment conditions. Two complementary approaches were followed. First, the Watcher radio link was improved through a more suitable 433-MHz antenna and improved antenna placement. Second, a portable gateway was developed to provide an additional relay point where the normal Watcher-to-Watcher chain was interrupted, for example by topography or excessive distance. As shown in Fig. 5.1c, the gateway provides an alternative radio path across the problematic section and restores continuity of the communication chain.

In this work, the term *gateway* refers to this custom LoRa relay unit rather than to a LoRaWAN gateway.

The resulting system retains the original linear store-and-forward principle while providing additional means to maintain communication along difficult sections of the measurement profile. The following sections describe the corresponding hardware, radio software and multi-hop packet-handling mechanisms.

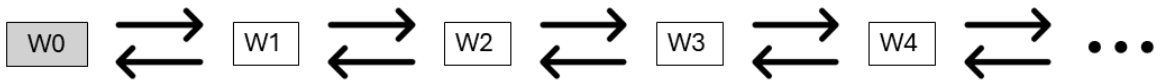

(a) Normal multi-hop communication along the Watcher chain.

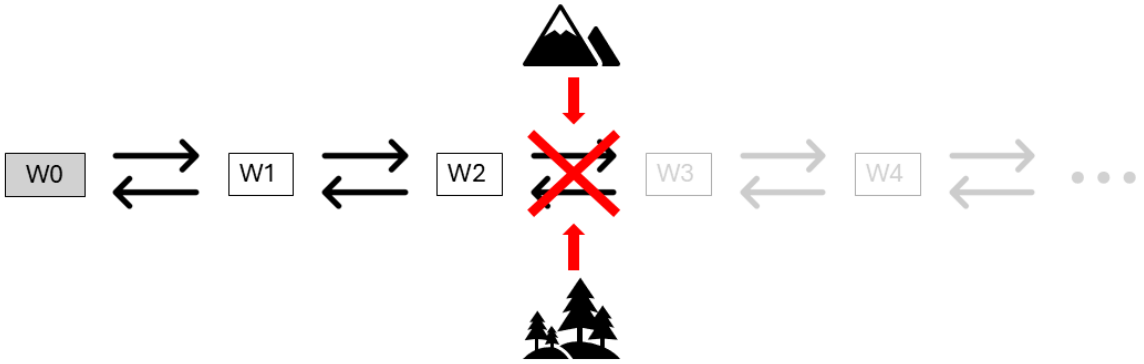

(b) Communication interruption caused by an obstacle.

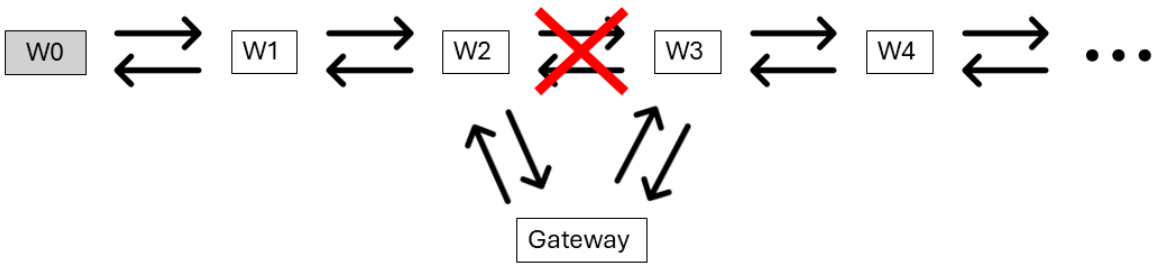

(c) Restoration of the communication path using a relaying gateway.

Figure 5.1: Controller-initiated multi-hop communication architecture of the Watcher network. Under normal conditions, a request issued by W_0 and the corresponding response are relayed between successive Watchers. A local propagation obstacle can interrupt this communication path and isolate the following receivers. A portable gateway can then be positioned to provide an alternative radio path across the problematic section.

5.2 Operating frequency and regulatory constraints

The original Watcher platform operates in the European 433-MHz short-range-device (SRD) frequency band, which was retained to preserve compatibility with the existing radio hardware and antenna system. Although LoRa can also operate around 868 MHz in Europe, changing band would have required a broader redesign of the existing wireless platform without directly addressing the main objectives of this work.

The lower carrier frequency is advantageous for the Watcher deployment environment, where antennas are installed close to the ground and radio links may be affected by terrain and vegetation. Compared with 868 MHz, operation around 433 MHz generally benefits from lower path loss, stronger diffraction around obstacles and reduced attenuation through vegetation. In free space, the difference in path loss is

$$20 \log_{10} \left(\frac{868}{433} \right) \approx 6.0 \text{ dB}. \tag{5.1}$$

The 868-MHz region nevertheless offers different regulatory options, including higher permitted transmit powers in some sub-bands, and could therefore remain relevant for a future radio redesign.

Operation in the 433-MHz band is subject to the regulatory conditions applicable to short-range devices. In Belgium, non-specific SRDs may operate in the 433.05–434.79 MHz band under general authorisation, with a maximum effective radiated power of 10 mW e.r.p. and a duty cycle below 10% [77]. ETSI EN 300 220 specifies the corresponding technical requirements for this class of equipment [78].

The radio settings used during development and characterisation should therefore not be interpreted as a field-deployment configuration. In particular, the 20 dBm Watcher and 22 dBm gateway transmit-power settings exceed the permitted 10 mW (10 dBm) e.r.p. limit. For field deployment, the configured transmit power must be reduced while also accounting for antenna gain and cable losses.

The 500 kHz LoRa bandwidth itself is compatible with the available 1.74 MHz band, but the complete transmitted spectrum must remain within its boundaries. A centre frequency of exactly 433.000 MHz, as used during prototype testing, would nominally extend approximately from 432.75 to 433.25 MHz. A field configuration should therefore move the carrier further inside the authorised band, for example to 433.92 MHz, while retaining the same bandwidth.

5.3 Hardware implementation

The wireless hardware modifications concern the replacement of the existing Watcher antenna and, most importantly, the development of the dedicated gateway introduced in Section 5.1.

5.3.1 Watcher radio and antenna integration

The Watcher retains the Adafruit Feather M0 LoRa platform and its integrated RFM96 transceiver. The modification therefore focused on the antenna rather than on the radio electronics themselves.

In the original prototype, a simple 16.5 cm-long wire antenna was used. This length is close to a quarter wavelength at 433 MHz and was therefore not intrinsically unsuitable for operation at the selected frequency. However, the wire was placed horizontally inside the enclosure, close to the electronics and to ground level, so that its effective RF behaviour depended strongly on the installation geometry. In the redesigned implementation, this was replaced by a dedicated 433-MHz antenna designed to be mounted vertically and positioned as high as practical on the enclosure. This provides a more reproducible antenna orientation and reduces the influence of the surrounding electronics and ground proximity. Vertical and exposed antenna placement is also consistent with the integration recommendations given for comparable sub-GHz LoRa modules [79].

The selected antenna is an Ebyte TX433-JKS-IPX20, specified for 433 MHz operation and connected to the RFM96 through an IPEX-1 connector. According to the manufacturer, the antenna has a nominal gain of 3 dBi, a characteristic impedance of 50 Ω, and a voltage standing-wave ratio not exceeding 1.5 at its intended operating frequency [80]. These specifications provide a more defined RF interface than the original wire antenna. Since antenna performance can also be influenced by its installation and surrounding environment, the influence of antenna placement was subsequently evaluated experimentally through packet-reception, RSSI and SNR measurements, as discussed in Section 6.4.1.



(a) Original Watcher antenna integration. A wire antenna (long blue wire, right side) is placed inside the waterproof enclosure and remains approximately horizontal when the enclosure is closed.

(b) Example of the redesigned antenna mounted vertically at the top of the enclosure, providing the greatest elevation above the box.

(c) Alternative vertical side-mounted integration. This arrangement is more compact and less exposed, but places the antenna slightly lower than the top-mounted configuration.

Figure 5.2: Evolution of the Watcher antenna integration from the original internal wire antenna to two possible vertical integrations of the redesigned 433-MHz antenna.

5.3.2 Gateway radio platform

The gateway is based on the Ebyte E22-400M22S radio module, which integrates a Semtech SX1268 transceiver and communicates with its host microcontroller through SPI. The module provides programmable transmit power up to 22 dBm and can transmit over 410–493 MHz, making it suitable for operation in the same 433-MHz band as the Watcher network [79]. In contrast with the SX1276-based RFM96 used on the Watcher, the SX1268 belongs to a newer generation of LoRa transceivers and uses additional control signals such as `BUSY` and `DIO1` alongside the SPI interface [81]. Both radio platforms can nevertheless exchange LoRa packets when configured with compatible physical-layer parameters, as discussed in Section 2.7.2 and verified in Section 6.4.2.

The 22-dBm hardware variant was deliberately selected instead of the 30- and 33-dBm versions available in the same E22 family. Under the reference test conditions specified by Ebyte—an open environment, a data rate of 2.4 kbps and antennas positioned 2 m above ground—the corresponding communication distances are approximately 6, 12 and 16 km, respectively [79]. These values represent manufacturer reference conditions and should not be interpreted as guaranteed field ranges.

For the intended licence-exempt field operation, the actual transmit power must be reduced according to Section 5.2. The higher-power variants therefore provide no regulatory advantage in this configuration. Even when considering their maximum hardware capabilities, the 22-dBm variant already offers a reference

range well beyond the expected distance between neighbouring Watchers.

Ebyte specifies a transmit current of approximately 100–140 mA for the 22-dBm module, compared with 500–620 mA at 30 dBm and 1.1–1.2 A at 33 dBm [79]. Semtech similarly specifies a typical transmit current of 107 mA for the SX1268 at 22 dBm and 490 MHz [81]. The 22-dBm version therefore provided the most appropriate compromise for prototype development.

Rather than relying on higher transmit power, communication reliability was improved through better antenna integration and the use of a gateway. Unlike the Watchers, which remain close to the ground at their measurement locations, the gateway can be positioned higher on a tree, pole or other support, reducing local obstruction and providing a clearer propagation path towards both sides of a problematic section.

For the prototype gateway, an E22-400MBL-SC evaluation board containing the E22-400M22S was initially selected only to evaluate the radio module and establish communication with the existing Watcher radios (see Fig. 5.3). During these tests, however, its integrated peripherals proved sufficient for the requirements of the relay unit itself. The evaluation platform includes an STM32F103C8T6 microcontroller, a 0.96-inch OLED display, three push-buttons, TX/RX indicators, USB connectivity and a SMA antenna connector [82]. It was therefore retained and repurposed as the core of the gateway rather than developing a separate controller board.

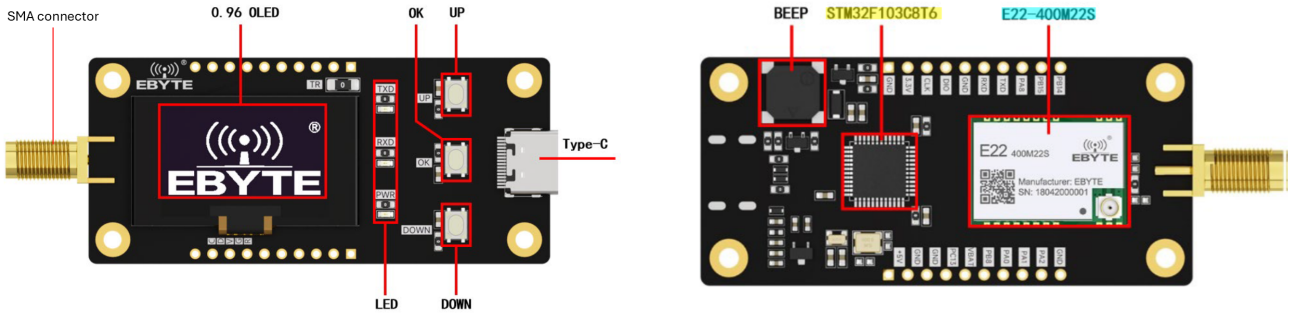

(a) Top view. **(b)** Bottom view.

Figure 5.3: Annotated views of the E22-400MBL-SC evaluation board, highlighting the main components and interfaces used in the gateway implementation.

Although well suited to the gateway, the evaluation board could not directly replace the Feather M0 in the Watcher. The STM32F103C8T6 itself provides sufficient processing and peripheral capabilities, but only a limited subset of its GPIOs is exposed by the E22 evaluation platform [83, 82]. These available connections are insufficient for the complete Watcher hardware, which must interface simultaneously with the acquisition electronics, relays, GNSS receiver, RTC, SD card, display and radio.

A future Watcher revision could instead integrate a microcontroller directly on the main PCB together with an E22-400M22S radio module. This would provide greater freedom in peripheral allocation and radio integration, but would reduce the modularity of the present Feather-based architecture and require a substantially larger hardware and firmware redesign.

Although the limited GPIO accessibility prevents the evaluation board from directly replacing the Feather M0 inside a Watcher, this limitation is much less restrictive for the remote controller, which requires considerably fewer external peripherals. The E22-400MBL-SC could therefore provide a basis for a future replacement of the current Feather-based remote. In the configuration used during this work, including the TX433-JKD-20P antenna, the complete platform was more than 50% less expensive than an Adafruit Feather M0 LoRa while already integrating the radio, display, push-buttons and USB interface required for remote operation. The main drawback would be the change in software environment, since the remote

firmware would have to be maintained using the STM32 HAL rather than the Arduino framework.

5.3.3 Gateway power supply and physical integration

To make this elevated and autonomous deployment practical, the E22 evaluation board was complemented by a battery supply assembled on a small perfboard and integrated into a compact weather-resistant enclosure.

The gateway is powered by a rechargeable 18650 lithium-ion cell. A combined battery-protection and USB-C charging module allows the cell to remain permanently integrated into the enclosure while providing a convenient charging interface. Since the battery voltage varies over its discharge cycle whereas the gateway electronics require a regulated 3.3-V supply, a TPS63802 buck–boost converter was added between the battery and the E22 evaluation board. The TPS63802 can operate from an input voltage between 1.3 and 5.5 V and automatically transitions between buck, buck–boost and boost operation depending on the relation between its input and output voltages. Its typical quiescent current is 11 μA, making it suitable for a battery-powered device [84].

For a 3.3-V output, the converter can provide up to 2 A when the input voltage is at least 2.3 V [84]. This provides a large current margin compared with the approximately 100–140 mA transmit current specified for the 22-dBm E22-400M22S module, while also supplying the microcontroller, display and remaining peripherals.

The complete power-supply and battery-monitoring architecture is summarised in Fig. 5.4.

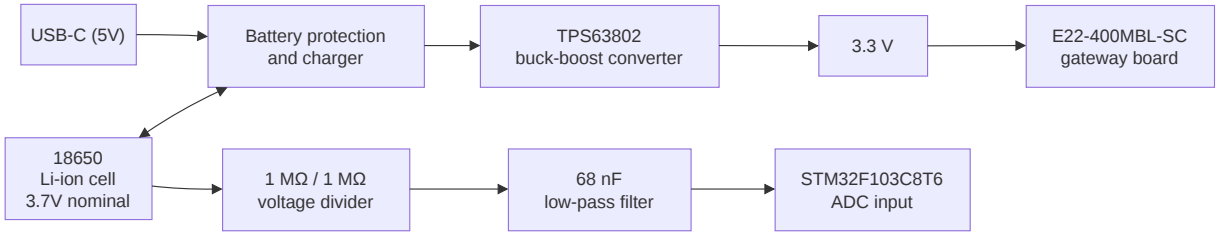

Figure 5.4: Power-supply and battery-monitoring architecture of the portable gateway. The 18650 cell supplies the E22-400MBL-SC gateway board through the battery-protection and charging stage and the TPS63802 buck–boost converter. Battery voltage is monitored independently through a high-impedance voltage divider and low-pass filter connected to an STM32 ADC input.

Battery-voltage monitoring was also integrated into the additional circuit using one of the 12-bit ADC inputs available on the STM32F103C8T6. Two 1 MΩ resistors form a 1:2 voltage divider between the battery and the ADC input. The relatively high resistance values were selected to minimise the permanent current drawn by the monitoring circuit. At the maximum 4.2-V voltage of a fully charged 18650 cell, the divider continuously draws only

$$I = \frac{4.2}{2 \times 10^6} \approx 2.1\ \mu\text{A}. \tag{5.2}$$

The battery voltage is reconstructed from the measured ADC voltage according to

$$V_{\text{bat}} = V_{\text{ADC}} \frac{R_1 + R_2}{R_2}. \tag{5.3}$$

Since $R_1 = R_2 = 1$ MΩ, this reduces to $V_{\text{bat}} = 2V_{\text{ADC}}$, while ensuring that the maximum voltage applied to the ADC remains well below the 3.3-V supply voltage.

The use of such high-value resistors also results in a relatively high source impedance at the ADC input. A 68 nF capacitor connected between the ADC input and ground therefore provides a local charge storage for the ADC sampling input while simultaneously low-pass filtering the battery measurement. Together with the 500 kΩ Thévenin resistance of the divider, this gives a time constant of approximately

$$\tau = R_{\text{th}}C = 500\ \text{k}\Omega \times 68\ \text{nF} \approx 34\ \text{ms}. \tag{5.4}$$

The power-supply circuit was implemented manually on a perfboard as a proof of concept rather than as a final hardware implementation. The objective at this stage was to integrate only the functions required for autonomous gateway operation and to validate the complete system with limited additional hardware. A more mature implementation could combine the charging and protection circuitry, voltage regulation and battery monitoring on a dedicated PCB, resulting in a more compact, reproducible and mechanically robust assembly. Such an optimisation was not prioritised in the present work, since the main development effort focused on the radio communication architecture and gateway firmware.

The resulting prototype perfboard therefore provides the three functions required for autonomous operation: battery protection and charging, generation of the regulated 3.3-V supply, and battery-voltage monitoring. Its physical implementation is shown in Fig. 5.5. Together with the E22-400MBL-SC evaluation board and the 18650 cell, it forms the self-contained gateway shown in Fig. 5.6, which was integrated into a compact weather-resistant enclosure. The complete manufacturer schematic of the E22-400MBL-SC evaluation board is provided in Appendix B.

As discussed in Section 6.1.5, however, the energy available from a single typical 18650 cell is insufficient for five days of continuous operation. A future redesign would therefore require increased battery capacity.

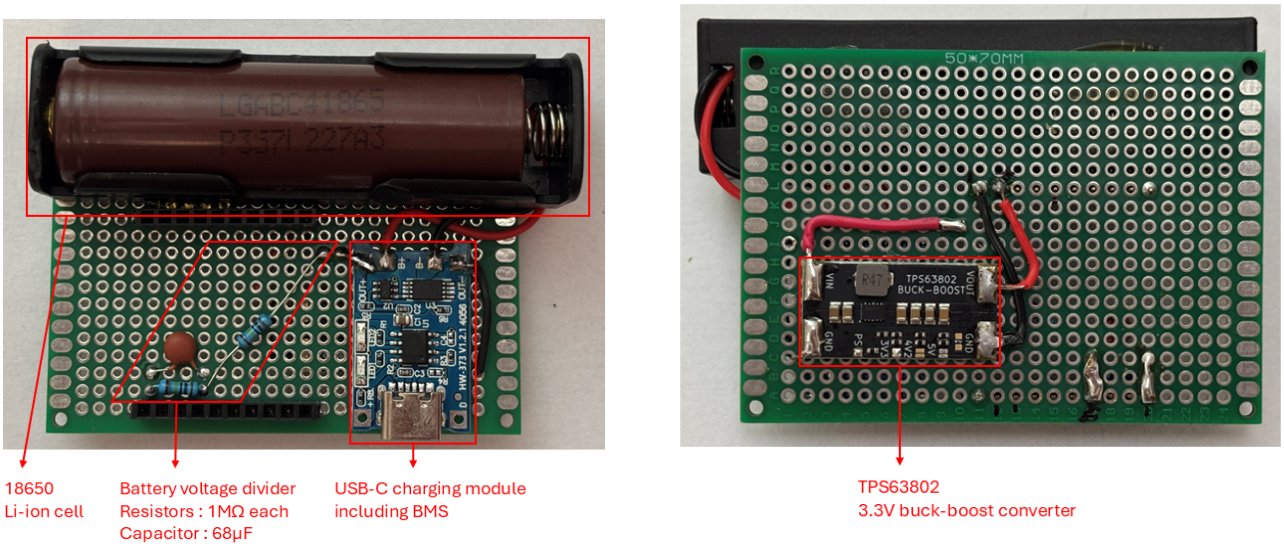

(a) Top view.

(b) Bottom view.

Figure 5.5: Annotated views of the gateway power-supply perfboard. The board integrates the battery-protection and USB-C charging module, the TPS63802 buck–boost converter, and the battery-voltage monitoring circuit connected to the STM32 ADC input.

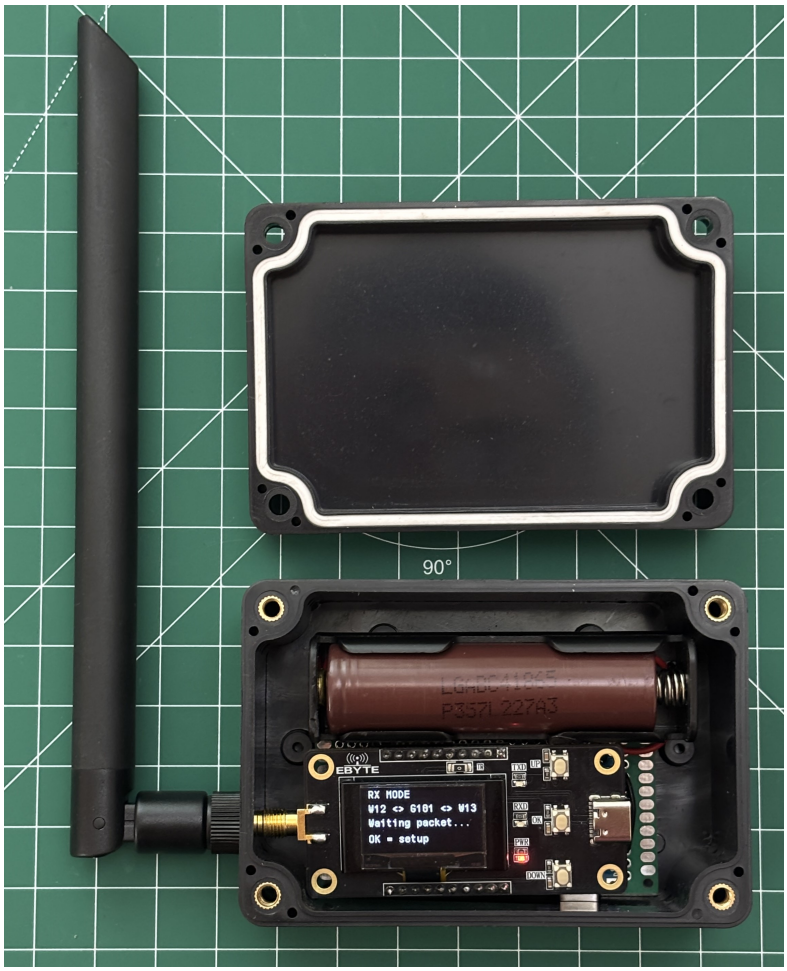

Figure 5.6: Complete autonomous gateway integrated into a weather-resistant enclosure.

5.4 Software implementation

The Watchers and the gateway use different microcontroller platforms and therefore require separate software implementations. The Watcher firmware remains based on the Arduino framework running on the Feather M0, whereas the gateway firmware was developed specifically for the STM32F103C8T6 of the E22-400MBL-SC evaluation board using the STM32 Hardware Abstraction Layer (HAL). Despite this difference, both implementations use compatible LoRa physical-layer parameters and the same application-level packet structure, allowing the gateway to behave as an additional node within the existing Watcher communication architecture.

5.4.1 Watcher radio driver and interrupt-based reception

On the Watcher, the RFM96 radio is controlled using the RadioLib library. The LoRa physical-layer parameters used by the Watcher network are summarised in Table 5.2. The gateway uses the same configuration to ensure interoperability between the SX1276-based RFM96 and the SX1268-based E22 module.

These values correspond to the prototype configuration used during development and experimental characterisation. The carrier frequency and transmit power must be adjusted for regulatory-compliant field operation as discussed in Section 5.2.

Packet reception is interrupt-driven, while the corresponding SPI access and packet processing are deferred

Table 5.2: LoRa parameters used during prototype development and characterisation.

Parameter	Value
Carrier frequency	433 MHz
Bandwidth	500 kHz
Spreading factor	11
Coding rate	4/5
Preamble length	8 symbols
Header mode	Explicit
Payload CRC	Enabled
Watcher transmit power	20 dBm
Gateway transmit power	22 dBm

to the main firmware execution. The implementation of this mechanism and its interaction with SD-card SPI arbitration are described in Section 4.1.3.

The radio driver is kept separate from the higher-level communication protocol. It handles transceiver configuration and packet transmission and reception, whereas addressing, forwarding, duplicate suppression and command interpretation are implemented at the application level and described in Section 5.5.

5.4.2 STM32 HAL gateway firmware

A separate firmware had to be developed for the gateway because the E22-400MBL-SC is based on an STM32F103C8T6 rather than on the Arduino-compatible SAMD21 used by the Watcher. The demonstration software originally programmed on the Ebyte evaluation board was useful for verifying the hardware, but its source firmware could not be recovered in a form that could be adapted to the present application. A new STM32 project was therefore created and the gateway application was progressively rebuilt around the peripherals available on the evaluation board.

The development environment consequently differs from that of the Watcher. The Arduino framework provides a relatively high-level abstraction of board initialisation and peripherals through libraries and the conventional `setup()` and `loop()` structure. The STM32 HAL remains a hardware abstraction layer, but requires the microcontroller peripherals to be configured more explicitly. In the gateway firmware, the system clock, GPIOs, SPI, I^2C, ADC and USB interface are initialised before the program enters its main `while(1)` execution loop. Peripheral access is then performed through HAL functions and peripheral handles generated and configured in STM32CubeIDE. The main microcontroller pin assignments are shown in Fig. 5.7.

The SX1268 software is divided into a generic radio driver and a platform-specific interface layer. The implementation uses the open-source LibDriver SX1268 driver, while a dedicated port was written to connect its generic functions to the STM32 HAL. This port implements the SPI transactions, radio reset and `BUSY` monitoring, and controls the `TXEN` and `RXEN` signals required by the RF front end of the E22 module. This separation limits the hardware-dependent code to a small interface layer while the higher-level SX1268 commands remain independent of the STM32 implementation.

The gateway LoRa configuration is summarised together with that of the Watcher in Table 5.2.

Radio events are signalled by the `DIO1` output of the SX1268, which is connected to an external-interrupt input of the STM32. The driver interprets the corresponding SX1268 interrupt status and distinguishes reception completion, transmission completion, timeout and CRC errors. Higher-level operations are then handled from the main application rather than being incorporated into the packet-forwarding protocol itself. After a packet has been processed or retransmitted, the radio is returned to receive mode.

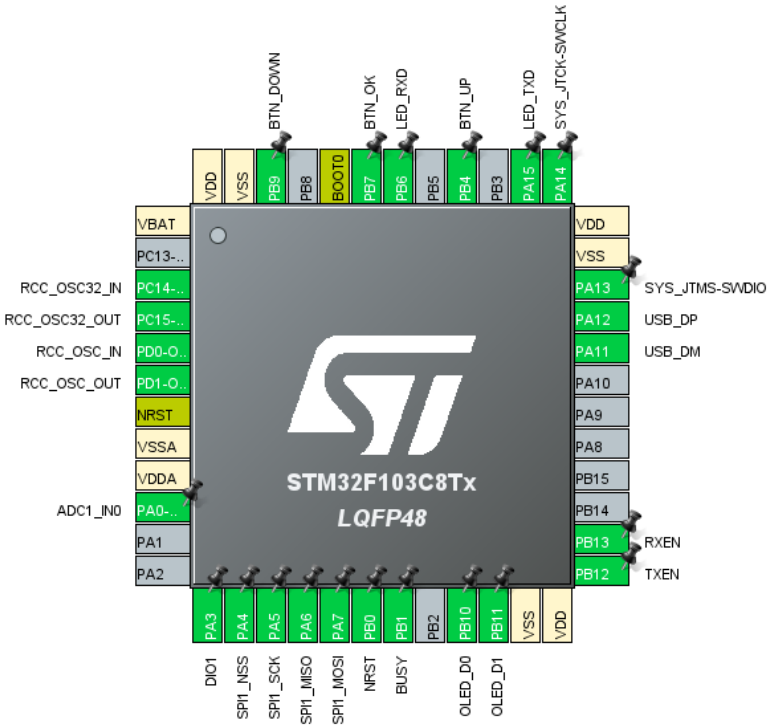

Figure 5.7: STM32F103C8T6 pin and peripheral assignments used by the gateway firmware.

For development and diagnostics, a USB CDC interface was also included. Debug messages are transmitted only when a USB serial connection is detected, avoiding unnecessary blocking operations when the gateway is used autonomously. The interface was extensively used during radio bring-up to inspect packet contents, forwarding decisions, radio errors and measured link parameters. During development, these messages were monitored on a PC using Tera Term 5 as the serial terminal. Figure 5.8 shows an example in which the gateway receives a Watcher-compatible packet, identifies that it must cross the configured gateway boundary from left to right, and retransmits it successfully.

Firmware programming is currently performed through the STM32 SWD interface using an ST-Link programmer. Although USB CDC communication is implemented, no USB bootloader was developed within the scope of this work. Such a bootloader would require additional Flash memory and start-up logic and could be considered in a future revision, subject to the remaining program-memory capacity.

Incremental radio validation

Gateway development followed a progressive hardware bring-up procedure. The board peripherals, such as LEDs, buttons and OLED display were first validated independently, after which SPI communication with the E22-400M22S was established. LoRa transmission and reception were then introduced and validated before the higher-level packet-handling functions were added.

Two E22-400MBL-SC boards were used during development, which allowed the transmission and reception functions to be tested independently of the Watcher firmware. A second STM32 project was developed specifically as a radio test transmitter. It generates packets using the same 17-byte format as the Watcher network and allows controlled packets to be sent between the two boards.

The transmission test firmware was designed to progressively validate the radio implementation. It allowed different commands to be generated and the packet fields to be modified deliberately, including the `seq` field used in the prototype packet format. This field was used during communication and duplicate-suppression

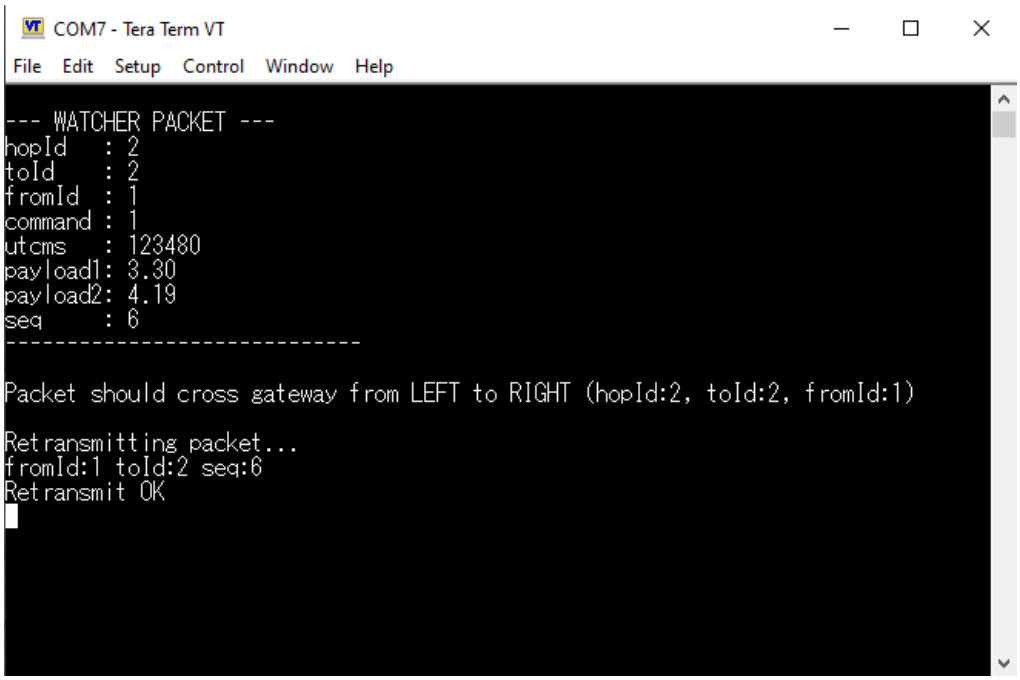

Figure 5.8: Example of the USB CDC debug output monitored through Tera Term during gateway development. The gateway parses the received Watcher packet, determines that it must be forwarded from left to right, and confirms the successful retransmission.

testing, although it was subsequently found to be unnecessary for normal protocol operation, as discussed in Section 5.5.1. The second E22-400MBL-SC board consequently also acted as a temporary remote controller during development, allowing Watcher-compatible commands to be generated without relying on the original Feather-based remote.

After basic indoor tests, preliminary outdoor tests were performed with the two boards separated over increasing distances. Packet Received-Signal Strength Indicator (RSSI) and Signal-to-Noise Ratio (SNR) were retrieved directly from the SX1268 after reception and used to assess the radio link. These preliminary tests confirmed the operation of the radio implementation and motivated the more systematic communication tests reported in Section 6.4.

The second E22-400MBL-SC used as the test transmitter is not dedicated permanently to this development role. Once the communication tests are completed, it can be reprogrammed with the same gateway firmware as the first board and assigned a different node identifier, thereby becoming a second fully functional gateway without any hardware modification. This reuse was one of the reasons for purchasing two identical evaluation boards: the pair could first provide an independent transmitter–receiver test platform and subsequently be deployed as two gateway units.

Gateway application behaviour

At the application level, the gateway operates primarily as a relay node. After receiving a valid packet, the firmware determines whether the packet is addressed to the gateway itself or must continue through the Watcher network. Packets addressed to the gateway are passed to a local command handler, whereas packets requiring retransmission are processed according to the forwarding rules described in Section 5.5.

Because the gateway does not contain any geophysical acquisition circuitry, most Watcher-specific commands are not applicable to it. The only local status function currently implemented is the battery-voltage request. When such a request is received, the firmware measures the 18650-cell voltage through the divider connected to the STM32 ADC and returns the result using the same response convention as

the Watchers. This allows the gateway battery state to be queried remotely without introducing a separate packet format or command interface.

Gateway power-management strategy

The power-management requirements of the gateway differ from those of the Watchers. A Watcher can enter a scheduled deep-sleep state when no measurement or remote interaction is expected, whereas the gateway may need to relay a packet at an arbitrary time. The radio must therefore remain available continuously during normal gateway operation.

For this reason, the present firmware keeps the SX1268 in continuous receive mode whenever the gateway is active. After a packet has been processed or retransmitted, the radio is immediately returned to receive mode. This strategy prioritises communication availability and low forwarding latency over receiver duty cycling. Introducing periodic receive windows would reduce radio consumption, but would also require additional synchronisation with the asynchronous Watcher traffic and could delay packet forwarding.

The OLED display is not required for packet forwarding once the gateway has been configured and can therefore be switched off during unattended operation to reduce its contribution to the power consumption. The measured current consumption of the gateway and the influence of the display and radio operating modes are presented in Section 6.1.5. These measurements also show that the energy available from a single typical 18650 cell is insufficient for five days of continuous gateway operation. Increasing the installed battery capacity is therefore required for a future field-ready implementation.

Further reductions could be investigated without duty cycling the radio. In particular, the STM32 could potentially enter a low-power state while the SX1268 remains in receive mode, with the radio `DIO1` signal used to request processor wake-up when a packet is received. This approach was not implemented in the present prototype and would require additional validation.

5.4.3 Gateway user interface

The OLED display and three push-buttons already present on the E22-400MBL-SC evaluation board were reused to provide a simple local interface shown in Fig. 5.9. Their main purpose is to make the gateway deployable without requiring a computer or modification of the firmware for each field position.

At start-up, the operator selects the identifiers of the Watchers located on the two sides of the gateway. The right-hand identifier is initially set to the next Watcher after the selected left-hand identifier, while both values can be changed using the buttons. Once confirmed, the gateway enters receive mode and the display indicates its position in the chain, for example

$$\mathrm{W12} \leftrightarrow \mathrm{G101} \leftrightarrow \mathrm{W13}.$$

This provides an immediate visual check that the gateway has been configured for the intended interrupted section of the profile. The configuration menu can also be reopened locally without reprogramming the STM32.

During firmware development, additional diagnostic screens were implemented to display received-packet information, including packet length, interrupt status, RSSI and SNR. These diagnostics were particularly useful during the paired E22 tests and preliminary range measurements. They are not essential to normal gateway operation and may therefore be simplified or disabled in the final field firmware.

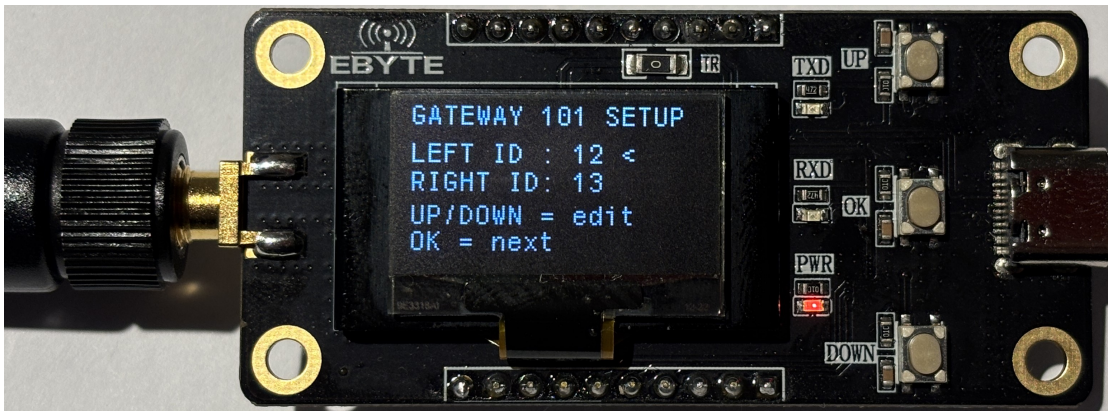

(a) Setup screen allowing the selection of the left and right neighbouring Watcher IDs.

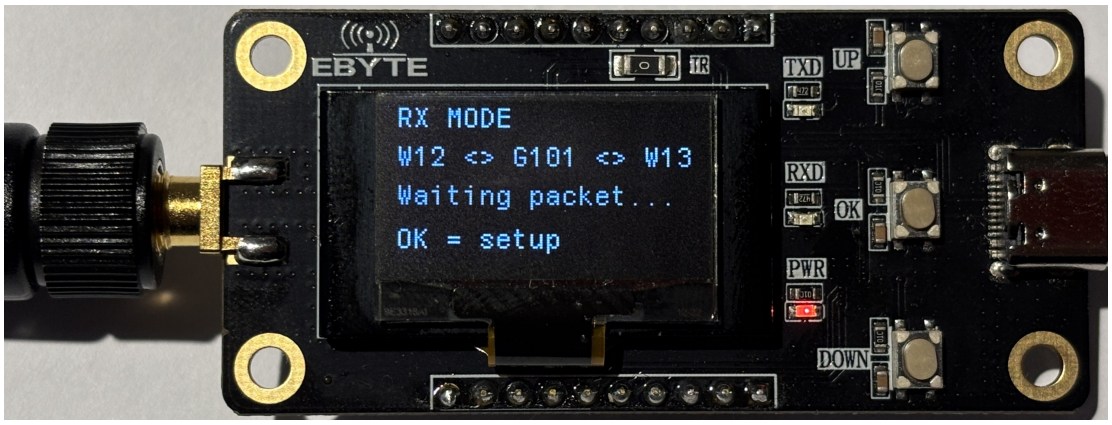

(b) Normal operating screen showing the configured position in the chain.

Figure 5.9: Gateway user interface.

5.5 Packet format and communication protocol

The radio interfaces described in the previous section provide the physical transfer of LoRa packets, while the application protocol determines how these packets are propagated through the distributed network. The protocol retains the store-and-forward principle of the original Watcher implementation, while this work introduces gateway-assisted forwarding and a duplicate-rejection mechanism.

5.5.1 Packet format, addressing and commands

Communication between the remote controller, Watchers and gateways uses a fixed 17-byte application packet. Its structure is summarised in Table 5.3.

The distinction between `hopId` and `toId` is central to the multi-hop mechanism. The former identifies only the next node that should act on the transmitted frame, whereas the latter remains equal to the final destination throughout the packet propagation. Similarly, `fromId` identifies the original source and is not replaced by the identifiers of intermediate relays.

The prototype packet also contains the 8-bit `seq` field. This field was used throughout the experimental validation and was therefore present in all measured 17-byte packets reported in this work. Duplicate detection initially relied entirely on it. However, the final duplicate-suppression logic can identify recent transactions using the combination of `fromId` and `utcms`, together with the 15-s recent-history lifetime described in Section 5.5.3. The `seq` field is therefore no longer required for normal protocol operation.

The two floating-point payload fields are intentionally generic so that the same packet structure can

Table 5.3: Application-level LoRa packet format used by the Watcher network.

Byte(s)	Field	Type	Description
0	`hopId`	`uint8_t`	Identifier of the next node expected to process the packet.
1	`toId`	`uint8_t`	Final destination of the packet.
2	`fromId`	`uint8_t`	Identifier of the node that originally generated the packet.
3	`command`	`uint8_t`	Command or response code.
4–7	`utcms`	`uint32_t`	Time or transaction value associated with the request.
8–11	`payload1`	`float`	First command-dependent payload value.
12–15	`payload2`	`float`	Second command-dependent payload value.
16	`seq`	`uint8_t`	Sequence number generated by the source node; no longer required in the proposed final protocol.

support several commands without defining a different frame format for each operation. Depending on the command, they can contain quantities such as battery voltage, electrode contact resistance, ADC statistics, measured voltage or geographical coordinates. The command set currently used by the Watcher firmware is summarised in Table 5.4.

Table 5.4: Main command codes currently used by the Watcher network.

Command	Function
1	Start voltage monitoring along the addressed section of the chain.
2	Stop voltage monitoring along the addressed section of the chain.
3	Start voltage monitoring on the addressed Watcher.
4	Stop voltage monitoring on the addressed Watcher.
5	Request electrode contact-resistance information.
6	Request ADC noise information.
7	Request ADC offset information.
8	Request battery voltage.
9	Request a voltage measurement corresponding to a given timestamp.
10	Request the Watcher to enter its sleep state.
11	Request the Watcher geographical position.

For commands requiring a response, the Watcher returns a packet using the same structure and the corresponding response code obtained by adding 100 to the request command. For example, command 8 produces response 108. The gateway follows the same convention for its battery-voltage response, allowing it to be queried through the same application protocol as the Watchers.

The 32-bit and floating-point fields are currently serialised by copying their four constituent bytes directly into the packet. This is compatible with the SAMD21 and STM32F103 platforms used in the present system. A future implementation intended to support more heterogeneous microcontroller architectures could instead explicitly define byte order and numerical representation to make the protocol independent of the host architecture.

5.5.2 Store-and-forward packet propagation

In the target application, Watchers are physically arranged along the measurement profile in increasing identifier order, $W_1, W_2, \ldots, W_N$. With a one-node forwarding step, a packet travelling towards a higher-numbered Watcher is therefore successively addressed to increasing `hopId` values, while a packet travelling in the opposite direction follows decreasing values.

A Watcher first verifies whether the received `hopId` corresponds to its own identifier. Packets intended for another immediate hop are ignored even if they are physically received by the radio. If the Watcher is also the final destination, i.e. if `toId` equals its identifier, the associated command is processed locally. Otherwise, the Watcher replaces only `hopId` by the identifier of the next Watcher and retransmits the packet. The final destination, original source, command, transaction-time value and payload are preserved. In the 17-byte prototype format used during testing, the `seq` field was also forwarded unchanged.

For example, a request generated by W_0 for W_5 normally propagates as

$$W_0 \rightarrow W_1 \rightarrow W_2 \rightarrow W_3 \rightarrow W_4 \rightarrow W_5,$$

with each intermediate Watcher changing only the immediate `hopId`. Commands 1 and 2 constitute particular cases because they apply not only to the final destination, but also to every Watcher encountered along the forwarding path. Command 1 requests the start of voltage monitoring on each traversed Watcher, whereas command 2 requests the corresponding stop operation. Each intermediate Watcher therefore updates its local monitoring state before forwarding the packet towards the final `toId`. These two commands allow a complete section of the measurement line to be started or stopped using a single request.

Gateway-assisted forwarding preserves the same store-and-forward mechanism. Consider a gateway positioned between W_4 and W_5 across an unavailable direct radio link. When W_4 transmits a packet with `hopId`=5, the gateway receives the frame and, if the packet must cross its configured boundary, retransmits it without modifying its contents. The `hopId` therefore remains equal to 5, allowing W_5 to process the relayed packet exactly as it would a packet received directly from W_4. The same mechanism operates in the reverse direction.

5.5.3 Duplicate suppression and repeated-packet handling

The introduction of the gateway creates a situation in which the same packet may reach a Watcher through two different radio paths: as illustrated in Fig. 5.10, if the direct Watcher-to-Watcher link remains marginally functional while the gateway also retransmits the same frame, the next Watcher may receive both the original packet and the relayed copy.

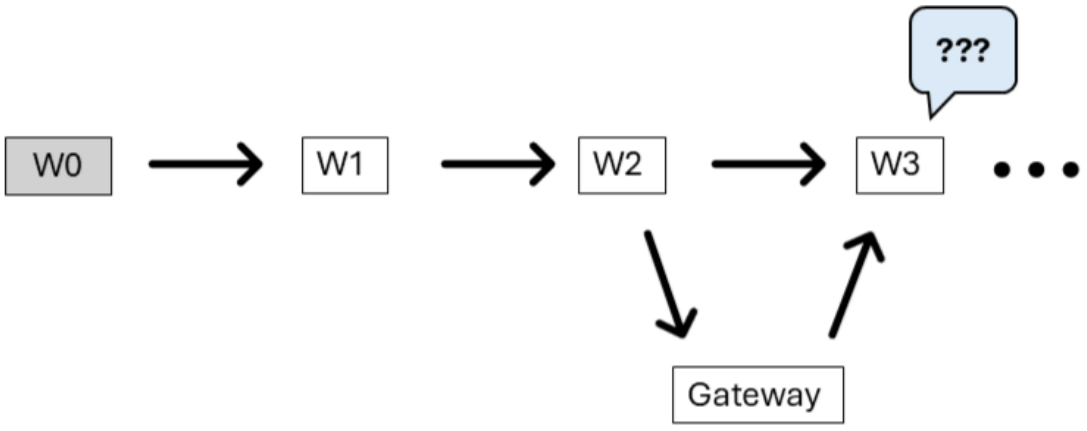

Figure 5.10: Example of duplicate packet reception in the presence of a gateway. A packet transmitted by W_2 could occasionally reach W_3 directly while also being received and retransmitted by the gateway. Without duplicate suppression, W_3 could process and forward the same transaction twice.

This behaviour is problematic as, without protection, both copies could be processed and retransmitted, creating duplicate commands and unnecessary traffic.

To prevent this, each Watcher keeps a short history of the transactions it has recently accepted. A transaction is identified by the combination

$$(\texttt{fromId}, \texttt{utcms}),$$

where `fromId` identifies the node that generated the packet and `utcms` acts as the transaction identifier. These values are preserved when the packet is forwarded by intermediate Watchers or retransmitted by the gateway. Consequently, a packet received directly and its copy received through the gateway carry the same identifier and can be recognised as duplicates.

When a packet is accepted for the first time, its `fromId`–`utcms` pair is stored in a recent-history table and the packet is processed normally. If another packet carrying the same pair is received while this entry is still present, it is considered to be another copy of the same transaction and is discarded before being processed or forwarded. In the example of Fig. 5.10, W_3 therefore accepts whichever copy arrives first and ignores the second one.

The recent-history table retains up to eight transactions for each `fromId`. Each entry has a time-to-live of 15 s, after which it is removed and the same identifier may be accepted again. This finite retention time is sufficient to cover the short delay between the direct and gateway-relayed copies without maintaining an indefinitely growing history of previous traffic.

Duplicate suppression consequently allows the gateway to remain useful even when the direct Watcher-to-Watcher link is intermittent rather than completely unavailable: the first copy of a transaction is accepted, whereas any additional copy received through the alternative path is ignored.

5.5.4 Response handling and communication failure detection

The current protocol uses end-to-end responses rather than a hop-by-hop acknowledgement after every retransmission. When a Watcher processes a request, it generates a response packet whose `fromId` corresponds to the responding node and whose `toId` corresponds to the remote controller. The request `utcms` value is preserved in the response, allowing the remote controller to associate an incoming packet with the transaction that generated it.

During development and communication testing, command 8 was used extensively as a simple end-to-end connectivity test. This command requests the battery voltage of the addressed node and therefore already provides a convenient request–response transaction without initiating a geophysical acquisition. Successful reception of the corresponding response confirmed that the complete communication path between the remote controller and the addressed node was operational. The same principle was also used when validating gateway-assisted forwarding.

This approach keeps the number of control frames low, but a timeout at the remote controller only indicates that the complete request–response transaction failed; it does not directly identify the radio link responsible for the failure. A lightweight link-diagnostic procedure could therefore be added for deployment and troubleshooting.

A possible future extension would therefore be an automated `PING`/`PONG` mechanism. Starting from the remote controller, Watchers could be queried successively along the measurement line. Each addressed Watcher would return a short response containing the transaction identifier together with the RSSI and SNR measured for the received `PING` packet, for example using the two generic payload fields. This would provide not only a binary indication of connectivity, but also a simple estimate of the quality of the corresponding radio link, without initiating any geophysical measurement or status acquisition.

If $W_1, \ldots, W_k$ repeatedly responded whereas W_{k+1} and all subsequent Watchers did not, the most likely communication discontinuity would be located between W_k and W_{k+1}. The reported RSSI and SNR values

could additionally help identify links that remain operational but are approaching unreliable conditions. A gateway could then be installed across the problematic section and the diagnostic sequence repeated to verify both restored connectivity and improved link quality.

Several retries should be performed before declaring a link unavailable, since the absence of a single response may result from transient interference rather than from a persistent propagation problem. In addition, this end-to-end procedure cannot by itself distinguish a failed radio link from a completely unpowered or defective W_{k+1}. Confirmation from the opposite side of the suspected break, or a temporary gateway placed near the boundary, would remove this ambiguity.

5.6 Communication cost and scalability

The store-and-forward architecture improves communication range without requiring every node to maintain a direct radio link with the remote controller, but this robustness comes at the cost of an increasing number of radio transmissions. If a source and destination are separated by H successive radio hops, a simple request requires H transmissions. In the absence of packet loss or additional retransmissions, a transaction consisting of one request and one response returning through the same path therefore requires

$$N_{\mathrm{TX}} = 2H. \tag{5.5}$$

For fixed LoRa parameters, packet length and transmit power, the corresponding transmission energy can be approximated as

$$E_{\mathrm{transaction}} \simeq 2H\, P_{\mathrm{TX}}\, t_{\mathrm{TX}}. \tag{5.6}$$

where P_{TX} is the electrical power consumed during transmission and t_{TX} is the duration of one packet transmission. This expression accounts only for transmission energy; receiver operation and microcontroller activity introduce additional consumption. It nevertheless illustrates why limiting the number of retransmissions becomes increasingly important as the number of hops in the communication path increases.

A gateway bridging an unavailable Watcher-to-Watcher link also increases the number of transmissions. A direct logical hop that would normally require one transmission becomes two physical transmissions, Watcher–gateway and gateway–Watcher. If the same gateway is crossed by both the request and its response, two additional transmissions are therefore introduced for the complete transaction. This additional cost is acceptable when the alternative is a complete loss of connectivity, but unnecessary gateway forwarding and duplicate retransmissions should be avoided whenever possible.

For the same reason, the present routing strategy favours directional forwarding rather than systematic flooding. Retransmitting every packet in both directions could increase robustness when the network topology is uncertain, but would also increase channel occupancy, packet collisions and energy consumption throughout the network. A more efficient compromise would be to use the expected direction under normal conditions and trigger additional or alternative forwarding only after a communication failure or timeout has been detected.

Duty cycle and channel occupancy at network scale

The regulatory 10% duty-cycle constraint introduced in Section 5.2 applies to each individual transmitter rather than to the cumulative transmissions of the complete Watcher network. In the present linear forwarding architecture, a command propagated to the end of the chain causes each intermediate Watcher to transmit twice: once when forwarding the request and once when forwarding the corresponding response towards W_0.

If the state of the entire network is changed every 5 min, each Watcher performs 24 transmissions per hour. For the 12-h active period considered in the reference operating scenario, this corresponds to 288 transmissions per day per Watcher. Using the measured packet transmission duration of approximately 167.6 ms, the cumulative transmission time of one Watcher is

$$T_{\mathrm{TX,h}} = 24 \times 0.1676 \approx 4.02 \text{ s/h}, \tag{5.7}$$

corresponding to a duty cycle of

$$D = \frac{4.02}{3600} \times 100 \approx 0.11\%. \tag{5.8}$$

The number of Watchers does not increase this per-node duty cycle for network-wide state-change commands, since each intermediate Watcher performs only one request forwarding and one response forwarding per transaction. The resulting duty cycle therefore remains well below the 10% regulatory limit applicable to the considered 433-MHz SRD band.

Sequential polling produces a less uniform forwarding load. During a complete status poll of 100 Watchers, the node closest to the remote controller, W_1, performs 199 transmissions: one for its own response and two for each of the 99 more distant Watchers. Using the measured packet duration of 167.579 ms, this corresponds to

$$T_{\mathrm{TX},W_1} = 199 \times 0.167579 \approx 33.35 \text{ s}, \tag{5.9}$$

or approximately 0.93% of one hour. The individual regulatory duty-cycle limit would therefore still not be approached by a single complete polling cycle. However, as shown experimentally in Section 6.4.3, the cumulative transmissions required to poll all 100 Watchers occupy approximately 28.2 min of radio time. Sequential polling is therefore primarily limited by network latency and channel occupancy rather than by the per-node duty-cycle constraint.

5.7 Protocol limitations and future network developments

Although the implemented store-and-forward protocol provides a simple and predictable communication path along the ordered Watcher chain, several limitations remain for deployment at the intended network scale. In particular, the current routing logic still assumes a fixed topological position of the remote controller, while the systematic use of adjacent Watchers as relays does not necessarily provide the fastest or most energy-efficient path. In addition, operating a large network would benefit from a higher-level supervision interface. These aspects are discussed below as possible extensions of the present system.

5.7.1 Proposed routing extension for a mobile remote controller

The Watchers $W_1, W_2, \ldots, W_N$ are deliberately assigned increasing identifiers according to their physical order along the measurement profile. This ordering provides a simple routing rule for communication between Watchers: packets travelling towards a higher-numbered node are forwarded towards increasing identifiers, while packets travelling in the opposite direction follow decreasing identifiers.

The current implementation assumes that the remote controller W_0 is located at the low-ID end of the chain, before W_1. Under this assumption, a response addressed to `toId` $= 0$ can always be returned by successively decreasing the Watcher identifier. This assumption is not well suited to the intended field operation, because the remote controller may need to be repositioned along the measurement profile.

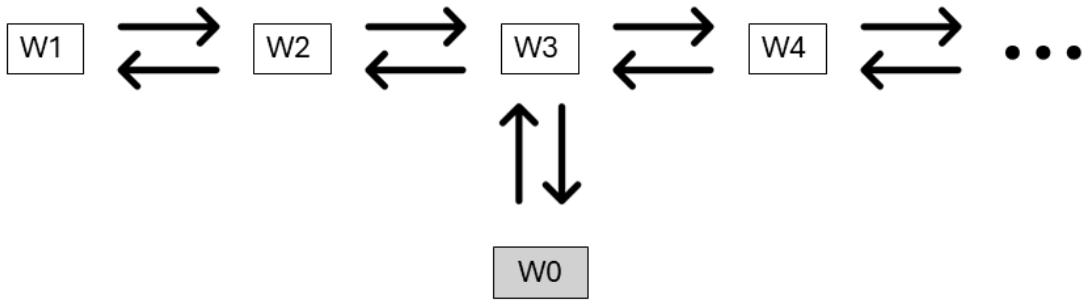

Figure 5.11: Example of a mobile remote controller positioned next to an intermediate Watcher. Although the logical identifier of the remote remains W_0, its physical position along the ordered Watcher chain can change during field operation.

In the configuration shown in Fig. 5.11, responses originating from Watchers located to the left of W_3 must propagate towards increasing identifiers, whereas responses originating from Watchers located to its right must propagate towards decreasing identifiers. The numerical destination $\texttt{toId} = 0$ therefore no longer contains sufficient information to determine the forwarding direction.

Since the remote controller is intended to remain physically close to one Watcher when operated, its topological position can be represented by a single *anchor Watcher*. In the example of Fig. 5.11, the remote controller is associated with W_3, corresponding to

$$\texttt{remoteAnchorId} = 3.$$

A proposed extension of the application packet is therefore to include this anchor identifier in every transaction initiated by W_0. The value would be preserved while the request propagates through the network and copied into the response generated by the destination. Each Watcher could then determine the direction of a response addressed to W_0 by comparing its own identifier with `remoteAnchorId`, rather than with $\texttt{toId} = 0$.

A straightforward implementation would reuse the byte occupied by the `seq` field during the experimental work. The sequence number is no longer required as duplicate suppression relies on `fromId` and `utcms`. The proposed routing extension could therefore replace `seq` by a `uint8_t remoteAnchorId` without increasing the application packet size. All radio timing, interoperability, forwarding and energy measurements reported in this work would thus stay the same.

The remote-controller interface would also need to allow the operator to select the Watcher next to which the controller is currently located. This could be implemented using a simple identifier-selection interface Watchers of the gateway in Section 5.4.3. The selected value would then be inserted automatically as `remoteAnchorId` in subsequently generated transactions.

For a response addressed to the remote controller, the forwarding rule would become

$$\begin{cases} \text{forward towards } W_{i+1}, & \text{if } i < \texttt{remoteAnchorId}, \\ \text{forward towards } W_{i-1}, & \text{if } i > \texttt{remoteAnchorId}, \\ \text{forward directly to } W_0, & \text{if } i = \texttt{remoteAnchorId}. \end{cases}$$

For example, with $\texttt{remoteAnchorId} = 3$, a response generated by W_1 would follow

$$W_1 \rightarrow W_2 \rightarrow W_3 \rightarrow W_0,$$

whereas a response generated by W_6 would follow

$$W_6 \rightarrow W_5 \rightarrow W_4 \rightarrow W_3 \rightarrow W_0.$$

The same principle would remain compatible with gateway-assisted forwarding. For packets addressed to W_0, a gateway could use `remoteAnchorId` as the effective topological destination when determining on which side of its configured `leftId`/`rightId` boundary the packet must propagate. For packets addressed to a normal Watcher, the final `toId` would continue to define the topological destination.

Directly addressing a gateway from a mobile remote controller introduces an additional limitation. Gateway identifiers are independent of their physical position in the ordered Watcher chain, so a destination such as G_{101} does not indicate whether the gateway is located to the left or right of the current remote position. A more complete implementation would therefore require the remote controller to maintain a mapping between each gateway identifier and its position in the Watcher chain, for example through the corresponding `leftId`/`rightId` boundary. Such information could later be maintained by the supervision interface discussed in Section 5.7.3. A simpler but less efficient fallback would be to propagate a gateway-discovery request in both directions until the addressed gateway is encountered.

The `remoteAnchorId` mechanism was not implemented in the firmware submitted with this work and should therefore be considered a proposed extension rather than part of the validated communication protocol.

5.7.2 Topology-aware routing

The present store-and-forward protocol deliberately uses neighbouring Watchers as successive relays. This provides predictable routing with little network knowledge, but it is not necessarily the most energy-efficient or the fastest path. In favourable radio conditions, a Watcher may be able to communicate directly with another node located several positions farther along the profile. Relaying every packet through all intermediate Watchers would then introduce transmissions that are not physically required.

A more advanced architecture could represent the communication network as a graph in which each node maintains information about the other nodes that it can reach reliably. Routing could then favour paths containing fewer transmissions rather than systematically following consecutive Watcher identifiers. For example, if W_1 can communicate reliably with W_4, a packet could bypass W_2 and W_3, reducing both latency and the number of radio transmissions.

Gateways could play an even more important role in such an architecture. Because they can be positioned significantly higher than the Watchers, a gateway may have line of sight to a large portion of the measurement profile rather than only to the two nodes surrounding a local communication break. It could therefore act as a higher-level relay, receiving packets from several sections of the line and forwarding them towards the appropriate destination through a shorter communication path.

Implementing such topology-aware or mesh-like routing would, however, require a substantial redesign of the present protocol. Nodes would need to discover or maintain a list of reachable neighbours, link quality would have to be characterised using quantities such as packet reception rate, RSSI or SNR, and a routing metric would be required to select an appropriate path. The lowest-hop route would not necessarily always be the lowest-energy route, since a longer radio link may be less reliable and therefore require additional retransmissions. A future adaptive implementation could also select a different transmit power or a more robust LoRa configuration for difficult links.

Such an architecture is beyond the scope of the present work but represents a particularly promising continuation of the Watcher network. The gateway developed here already introduces several elements required for this evolution, including independent node addressing, elevated relay placement and the ability to inspect radio-link quality.

A comparable principle is used in Meshtastic, an open-source LoRa mesh network in which packet forwarding is not tied to a fixed ordering of node identifiers [85]. Nodes form a decentralised mesh, with packets carrying destination and hop-limit information; forwarding can rely on controlled rebroadcasting, while directed traffic can also benefit from learned next-hop information. Such an approach illustrates how a future Watcher network could progressively replace the present ID-based linear routing with topology-aware forwarding, allowing the remote controller and intermediate gateways to be positioned more freely along the measurement profile.

5.7.3 Network supervision GUI

A further useful development would be a graphical supervision interface providing an overview of the complete Watcher network. With up to approximately 100 receivers distributed over several kilometres, checking each unit individually from the remote controller would become increasingly impractical.

The interface could periodically retrieve the status information already available from each Watcher, including battery voltage and electrode contact-resistance checks, and associate these values with the corresponding position along the profile. The time of the most recent successful communication could additionally provide a simple indication of node connectivity.

Nodes requiring operator intervention could then be highlighted immediately. A Watcher reporting a low battery voltage, abnormal electrode contact resistance or repeated communication failures could, for example, be displayed in a warning state, allowing the operator to identify both the nature and approximate location of the problem before travelling to the corresponding position along the profile.

Such a supervision interface was not developed within the scope of the present work. It would nevertheless become particularly valuable for operation of the network at its intended scale and could also provide the topology information required by future routing extensions, such as the physical location of mobile remote controllers and gateways.

The preceding sections describe the implemented communication architecture, its expected communication cost and its main remaining limitations. Its practical performance also depends on packet duration, forwarding delay, receiver consumption and radio-link quality. These quantities are therefore characterised experimentally in Section 6.4, together with antenna performance, RSSI, SNR and packet-reception measurements.

Chapter 6

Experimental validation and results

The experimental validation presented in this chapter evaluates the performance of the redesigned OhmPi Watcher and of the LoRa gateway developed in this work. Most experiments were performed on the Watcher v0.3.0, while the original v0.1.2 prototype was retained where a direct comparison was required. The gateway was evaluated both as an independent low-power radio node and as an intermediate relay between Watchers.

The main measurement equipment used throughout the experimental validation was:

- a UNI-T UT61E+ digital multimeter [86], used for voltage, resistance and supply-current measurements. Its PC logging software was used to record current continuously during automated power-consumption sequences;
- a laboratory DC power supply, used to supply the Watcher and gateway under controlled conditions and to generate reference DC voltages for analogue and battery-monitoring measurements;
- a Tektronix TDS1001C-EDU digital oscilloscope, used for time-domain measurements of the LoRa transmission current profile and for frequency-response measurements of the analogue low-pass filter;
- a Thurlby Thandar Instruments TG315 function generator, used to generate sinusoidal test signals for the low-pass-filter frequency-response measurements;

During Watcher current measurements, serial communication was maintained through a USB power-blocking adapter so that the USB connection did not provide an additional supply path. Firmware-level execution times were measured using the Feather M0 `micros()` timer. SD-card experiments were evaluated from the recovered binary files using dedicated Python post-processing scripts.

Radio experiments used numbered packets so that losses, duplicates and forwarding behaviour could be identified unambiguously. Wireless-link quality was evaluated using the packet reception ratio (PRR), received signal strength indicator (RSSI) and signal-to-noise ratio (SNR) reported by the radio transceivers.

The validation is organised around the main design objectives of the work. The chapter first evaluates power consumption and the resulting autonomy of both the Watcher and the gateway. It then examines acquisition timing and SD-card logging performance, followed by the electrical behaviour of the redesigned analogue front end. Finally, the wireless subsystem is characterised in terms of antenna behaviour, radio interoperability, transmission and forwarding delay, gateway-assisted link recovery and duplicate suppression. The chapter concludes by summarising the achieved performance, the remaining limitations for large-scale field deployment, and the estimated implementation cost per unit for the redesigned system.

6.1 Power consumption and autonomy

6.1.1 Watcher operating-state consumption

The supply-current consumption of the redesigned Watcher was first characterised in the principal operating states defined by the firmware state machine. The objective of this first measurement was to quantify the consumption experienced during normal field operation, before investigating the contribution of individual hardware subsystems in Section 6.1.2.

Measurements were performed at 12.10 V. Five readings were recorded after the current had stabilised in each configuration. The three long-duration firmware states considered were `IDLE_SHARED_SPI`, `RECORDING_DEDICATED_SPI`, and `SLEEP`.

Short transition states such as `LEAVING_IDLE_FOR_RECORDING`, `LEAVING_RECORDING_FOR_IDLE`, and `GOING_TO_SLEEP` were not characterised separately because their duration is short compared with the main operating periods.

Table 6.1: Measured current consumption of the redesigned Watcher in its principal operating states.

Firmware state	Current (mA)	Std. dev. (mA)	Power (W)
`IDLE_SHARED_SPI`	25.742	0.008	0.311
`RECORDING_DEDICATED_SPI`	77.578	0.164	0.939
`SLEEP`	17.680	0.072	0.214

The `IDLE_SHARED_SPI` state, in which the Watcher remains available for LoRa communication while the analogue front end is disabled, required approximately 25.7 mA.

During `RECORDING_DEDICATED_SPI`, the analogue front end is powered, two electrode channels are acquired at 100 Hz, and the binary measurement records are written to the SD card. This state required approximately 77.6 mA, corresponding to an increase of about 51.8 mA relative to the idle communication state.

Finally, the `SLEEP` state required 17.68 mA, approximately 31% less than `IDLE_SHARED_SPI`. The remaining consumption in this state shows that a substantial baseline load remains supplied even when the microcontroller and radio are placed in their low-power modes.

6.1.2 Subsystem power-consumption analysis

The operating-state measurements characterise the complete Watcher, but do not directly identify the contribution of its individual subsystems. An additional automated sequence was therefore performed in which the firmware successively enabled or disabled selected hardware functions while the supply current was recorded continuously using the UNI-T UT61E+ logging software.

Each configuration was maintained for 60 s. The first 5 s following each transition were excluded from the steady-state analysis. At the beginning of the sequence, the analogue supply was switched on and off three times for 3 s to generate a distinctive current pattern that could be used to align the firmware sequence with the multimeter recording. The resulting current profile as measured in the logging software is shown in Fig. 6.1.

The mean current measured for each configuration is summarised in Table 6.2. The values in Tables 6.2 and 6.1 come from separate measurement campaigns, which explains the small differences between comparable states.

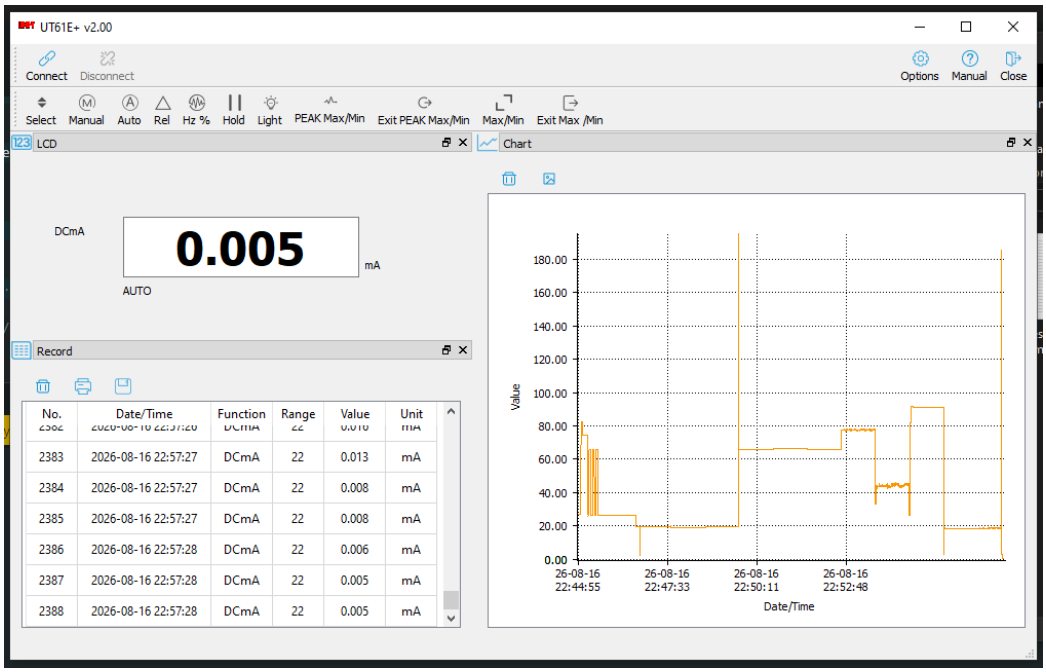

Figure 6.1: Supply-current evolution during the automatic subsystem characterisation. Each plateau corresponds to a predefined Watcher configuration maintained for 60 s. The initial switching sequence was used to align the firmware test with the multimeter recording.

Table 6.2: Steady-state supply current measured during the automatic subsystem power-consumption characterisation.

Configuration	Mean current (mA)
Idle / LoRa RX	25.973
Idle / radio sleep	19.60
Idle / radio sleep / OLED off	18.871
Analogue ON / inputs connected to GNDA	65.805
Analogue ON / electrodes connected	66.144
Analogue ON / inputs floating	65.841
Recording, 2 channels at 100 Hz	77.551
GPS active	44.122
RS check active, M1 with 988 Ω	91.109

LoRa radio.

Placing the LoRa transceiver in sleep mode reduced the idle current from 25.97 to 19.60 mA. Continuous LoRa reception therefore accounts for approximately

$$\Delta I_{\mathrm{LoRa,RX}} \approx 6.4 \text{ mA}, \tag{6.1}$$

or about one quarter of the normal idle-state current. This contribution is significant because the radio remains continuously in receive mode during idle communication periods.

The transient consumption associated with LoRa transmission was characterised separately using a 2.05 Ω series shunt resistor, whose resistance was measured with the UT61E+ multimeter. The voltage across the shunt was recorded with the oscilloscope while the Watcher repeatedly transmitted a 17-byte packet at a configured output power of 20 dBm.

As shown in Fig. 6.2, the shunt-voltage increase measured between the continuous-reception baseline and the transmission plateau was 59.2 mV, which corresponds to an additional transmit current of

$$\Delta I_{\mathrm{LoRa,TX}} = \frac{59.2 \text{ mV}}{2.05\ \Omega} \approx 28.9 \text{ mA}. \tag{6.2}$$

At the 12.10 V supply voltage, this represents an additional input power of approximately 349 mW. The transmission duration measured independently in Section 6.4.3 was approximately 168 ms, resulting in an incremental electrical energy consumption of

$$E_{\mathrm{LoRa,TX}} \approx 58.7 \text{ mJ} \tag{6.3}$$

per transmitted packet relative to the continuous-reception baseline.

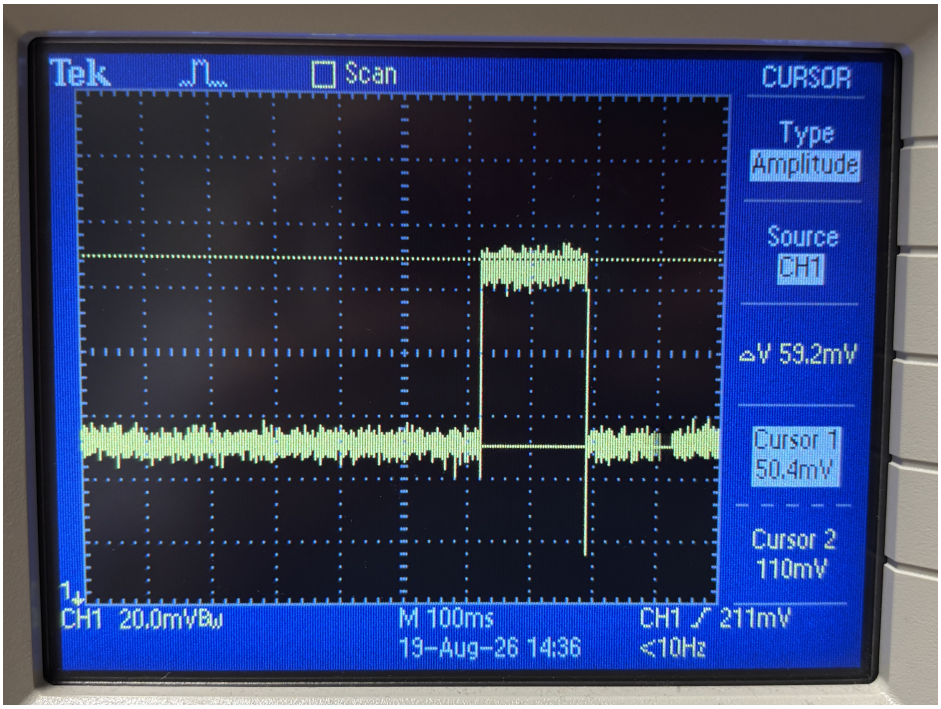

Figure 6.2: Voltage measured across the 2.05 Ω supply-current shunt during a LoRa transmission at a configured output power of 20 dBm. The shunt-voltage increase relative to the continuous-reception baseline was approximately 59.2 mV.

OLED display.

Disabling the OLED controller reduced the current by approximately 0.7 mA, indicating that its contribution to the overall Watcher consumption is comparatively small.

Analogue front end.

The largest additional load was associated with activation of the analogue domain. With the analogue inputs connected to GNDA, the supply current increased from 26.0 to 65.8 mA, corresponding to an increase of approximately 40 mA.

The configuration of the DG413 analogue switches had little influence on this consumption. Currents of 65.805, 66.144, and 65.841 mA were measured with the analogue inputs respectively connected to GNDA, connected to the electrode terminals, and left floating. The analogue-domain consumption is therefore primarily associated with the powered signal-conditioning and power-supply circuitry rather than with the particular analogue-switch state.

These measurements correspond to the fully populated four-channel version of the redesigned Watcher. In applications requiring only two receiving electrodes, the components dedicated to channels M3 and M4, including their OPA4188 devices and the second DG413, could be omitted. Although unused channels do not process a measurement signal, their powered components still contribute quiescent current to the $\pm$15-V analogue domain.

Such a partially populated implementation would reduce the op-amp contribution to the analogue-rail current, although the total analogue consumption would not decrease proportionally because the DC–DC converter, remaining conditioning stages, and switching circuitry represent fixed contributions. This two-channel hardware configuration was not characterised experimentally.

Acquisition and SD logging.

Starting continuous two-channel acquisition at 100 Hz together with binary SD-card logging increased the current from 66.14 to 77.55 mA. The combined acquisition and storage activity therefore contributes approximately

$$\Delta I_{\mathrm{acq+SD}} \approx 11.4 \text{ mA}. \tag{6.4}$$

This increase is substantially smaller than that caused by enabling the analogue domain itself. The powered analogue front end consequently remains the dominant additional load during the recording state.

GPS receiver.

Activation of the GPS receiver increased the current from approximately 26.0 to 44.1 mA, corresponding to an additional 18.1 mA. The GPS-active configuration also exhibited greater temporal variation than the other steady-state configurations.

Contact-resistance measurement

The contact-resistance circuit was characterised by maintaining the M1 injection relay active with a 988 Ω reference resistance connected between M1 and ground. The resulting supply current was 91.1 mA, approximately 25 mA above the corresponding powered analogue configuration.

This relatively high current is only drawn briefly during normal operation. With a 300 ms injection duration, four successive electrode checks represent approximately 1.2 s of relay activation per complete contact-resistance measurement and therefore contribute little to the daily energy budget.

Residual sleep consumption.

Finally, disabling the software-controllable peripherals while leaving the microcontroller awake resulted in a current of 18.87 mA. The deep-sleep current reported in Table 6.1 was 17.68 mA.

The relatively small difference indicates that the residual sleep consumption is not dominated by the SAMD21 itself. The exact contribution of the remaining continuously powered subsystems was not isolated in the present work. A more detailed hardware-level characterisation, in which individual peripherals and power domains are disconnected successively, would be required to identify the dominant contributors. Further reductions could then be achieved through selective power gating of the corresponding circuitry.

6.1.3 Comparison with the original Watcher

To quantify the improvement introduced by the redesigned hardware and power-management strategy, equivalent steady-state operating configurations were reproduced on the original Watcher v0.1.2 prototype. Both systems were supplied from the same 12.10 V laboratory source, and the current was measured at the 12-V input using the same digital multimeter and measurement procedure.

Unlike the preceding subsection, this comparison is not restricted to the formal firmware-state names of the redesigned Watcher, since the two implementations use different control architectures. Instead, the test conditions were selected to reproduce the same hardware functions as closely as possible on both systems. The comparison therefore includes the idle LoRa-receive condition, analogue-front-end activation, continuous two-channel recording, GPS operation, and sleep.

The original firmware behaviour was retained during the comparison rather than being replaced by the redesigned state machine, so the idle, recording, and sleep conditions correspond respectively to `IDLE_SHARED_SPI`, `RECORDING_DEDICATED_SPI`, and `SLEEP` of the redesigned Watcher.

Table 6.3 compares the mean current consumption measured for the corresponding operating configurations of both implementations.

Table 6.3: Comparison of the measured current consumption of the original and redesigned Watchers at a 12.10 V supply voltage.

Operating state	Original (mA)	Redesigned (mA)	Reduction (%)
Idle / LoRa RX	26.690	25.742	3.6
Analogue front-end ON	161.598	65.924	59.2
Recording, 2 channels	174.604	77.578	55.6
GPS active	48.696	44.012	9.6
Sleep	20.905	17.680	15.4

The idle consumption remained similar for the two systems, decreasing only from 26.69 mA to 25.74 mA. This indicates that the continuously powered digital subsystem represents a comparable baseline load in both implementations.

The largest improvement was observed when the analogue front end was enabled. The original Watcher consumed approximately 161.6 mA in this state, compared with only 65.9 mA for the redesigned system, corresponding to a reduction of approximately 59%. Since the acquisition and SD-logging activity added only approximately 13.0 mA and 11.7 mA above the analogue-on state for the original and redesigned systems, respectively, most of the reduction in recording consumption originates from the lower consumption of the redesigned analogue and power-supply architecture rather than from the data-logging process itself.

During continuous two-channel recording, the current decreased from approximately 174.6 mA to 77.6 mA,

equivalent to a reduction of approximately 56%. At the measured supply voltage, this corresponds to an input-power reduction from approximately 2.11 W to 0.94 W.

A smaller improvement was measured during GPS operation, for which the current decreased from approximately 48.7 mA to 44.0 mA. The GPS state therefore remained dominated by the receiver itself rather than by the modifications to the analogue subsystem.

The steady-state sleep current of the original Watcher was approximately 20.91 mA, compared with 17.68 mA for the redesigned implementation, corresponding to a reduction of about 15%. During the original watchdog-based sleep test, some multimeter readings temporarily increased to 25.01 mA while the status LED flashed, consistent with the periodic wake-up behaviour of the original firmware, which occurred every 16 s. This transient value was therefore excluded from the steady-state sleep average. In contrast, the redesigned Watcher remained in a stable deep-sleep state until an external wake-up event occurred.

To complement the state-by-state comparison, the two implementations were also evaluated under the autonomous fallback schedule used when the LoRa communication link cannot be restored. In this situation, the Watcher must remain capable of acquiring useful data without remote commands. The receiver therefore spends approximately 30 min in an idle communication state at the beginning and end of the day, records continuously for 12 h, and remains asleep for the remaining 11 h.

Applying the measured steady-state currents to this same 24-h duty cycle gives

$$Q_{\mathrm{day,original}} = 12I_{\mathrm{rec,original}} + 1I_{\mathrm{idle,original}} + 11I_{\mathrm{sleep,original}} \approx 2.35\ \mathrm{Ah/day}, \tag{6.5}$$

whereas the redesigned Watcher operated according to the same fallback schedule would require

$$Q_{\mathrm{day,new,same}} = 12I_{\mathrm{rec,new}} + 1I_{\mathrm{idle,new}} + 11I_{\mathrm{sleep,new}} \approx 1.15\ \mathrm{Ah/day}. \tag{6.6}$$

At 12.10 V, these values correspond to approximately 28.5 Wh/day and 13.9 Wh/day, respectively. The redesigned hardware therefore reduces the estimated daily energy consumption by approximately 51% even when both systems operate under this conservative fallback schedule.

This comparison isolates the benefit of the redesigned hardware under a communication-loss scenario. During normal field operation, however, the redesigned firmware can exploit LoRa commands to alternate between recording and lower-consumption communication states, allowing the analogue domain to be disabled whenever voltage acquisition is not required. The resulting autonomy under this intended operating strategy is evaluated in the following section.

6.1.4 Low-power operation and estimated autonomy

The impact of the measured operating-state consumption was evaluated for the reference deep-ERT deployment considered in this work. A typical measurement day was assumed to consist of 12 h of active operation, during which the Watcher alternates between voltage-recording periods and idle communication periods, each lasting on the order of a few minutes. For the energy-budget estimate, these alternating periods were approximated as 6 h of active voltage recording and 6 h in the idle communication state, followed by 12 h in deep sleep. The corresponding daily charge consumption is therefore

$$Q_{\mathrm{day}} = I_{\mathrm{rec}}t_{\mathrm{rec}} + I_{\mathrm{idle}}t_{\mathrm{idle}} + I_{\mathrm{sleep}}t_{\mathrm{sleep}}. \tag{6.7}$$

Using the experimentally measured currents of 77.58 mA, 25.74 mA, and 17.68 mA, respectively, gives

$$Q_{\mathrm{day}} \approx 832\ \mathrm{mAh}. \tag{6.8}$$

At the 12.10 V test voltage, this corresponds to an average daily energy consumption of approximately

$$E_{\text{day}} \approx 10.1 \text{ Wh}. \tag{6.9}$$

The equivalent average current over a complete 24-h period is therefore approximately 34.7 mA. Over the five-day acquisition scenario retained as the reference application, the corresponding consumption is approximately 4.16 Ah, or 50.3 Wh, before accounting for short-duration radio transmissions and state-transition events.

Continuous recording remains the dominant contributor to this energy budget, accounting for approximately 56% of the daily charge consumption. Despite occupying half of the day, the sleep state contributes approximately 25.5%, while the 6-h idle period accounts for approximately 18.6%.

Using the lower-bound communication scenario established in Section 5.6, an intermediate Watcher traversed by all state-change transactions performs approximately 288 LoRa transmissions during the 12-h active period.

The oscilloscope measurements presented in Section 6.1.2 showed that a 17-byte LoRa transmission at a configured output power of 20 dBm requires approximately 58.7 mJ of additional electrical energy relative to the continuous-reception baseline. For 288 transmissions per day, this corresponds to approximately 16.9 J, or 4.7 mWh per day. This represents less than 0.05% of the estimated 10.1 Wh daily Watcher energy consumption. The LoRa transmission energy is therefore negligible compared with the multi-hour recording, idle and sleep periods considered in the autonomy estimate.

6.1.5 Gateway consumption and battery-voltage monitoring

Gateway consumption

The gateway supply current was characterised at a 3.7 V input voltage using the same automated multimeter logging approach as for the Watcher. Since the gateway is intended to remain continuously available for packet reception, the LoRa RX state represents its principal operating condition. The influence of the OLED and of packet retransmission was therefore characterised separately.

Table 6.4: Measured gateway current consumption at a 3.7 V supply voltage, together with additional measurements of the EBYTE board alone.

Operating condition	Mean current (mA)
LoRa RX, OLED ON	52.77
LoRa RX, OLED OFF	50.12
Periodic LoRa transmission test	84.60
EBYTE board only, 5 V supply	23.78
EBYTE board only, direct 3.3 V supply	23.27

Disabling the OLED reduced the continuous RX current by approximately 2.65 mA, corresponding to about 5% of the total gateway consumption. The OLED can therefore be disabled after gateway configuration without affecting radio operation.

The EBYTE board was measured separately with the MCU awake and the radio in sleep mode. Its current was 23.78 mA from the 5-V input and 23.27 mA when supplied directly at 3.3 V, corresponding to approximately 118.9 mW and 76.8 mW, respectively. This supports the direct 3.3 V supply used in the gateway design.

The comparatively high 50.12 mA continuous-RX consumption should therefore not be attributed to the radio alone. The SX1268 itself is specified at only 4.2 mA in active receive mode [81], whereas the EBYTE board still consumed 23.27 mA at 3.3 V with the radio asleep and the STM32 awake. A substantial part of the gateway baseline consumption therefore originates from the evaluation-board implementation rather than from LoRa reception itself. A future dedicated gateway PCB could reduce this overhead, in particular by allowing the STM32 to enter a low-power mode while leaving the SX1268 in receive mode and using `DIO1` as a wake-up source.

The transmission cost was evaluated separately by transmitting one 17-byte LoRa packet every 500 ms for 60 s at the configured 22 dBm output power, while returning to receive mode between successive transmissions. The mean current measured during this interval was 84.60 mA, compared with 51.90 mA during the adjacent continuous-RX periods. Approximately 120 packets were transmitted during the test.

The additional charge associated with one RX–TX–RX cycle can therefore be estimated from the difference between the transmission-test and continuous-RX currents:

$$Q_{\mathrm{TX}} = \frac{(I_{\mathrm{TX,test}} - I_{\mathrm{RX}})T}{3600 N_{\mathrm{TX}}} = \frac{(84.60 - 51.90) \times 60}{3600 \times 120} \approx 4.54 \times 10^{-3}\ \mathrm{mAh}. \tag{6.10}$$

At a gateway supply voltage of 3.7 V, this corresponds to an additional energy of approximately

$$E_{\mathrm{TX}} = Q_{\mathrm{TX}} V \approx 0.0168\ \mathrm{mWh} \approx 60.5\ \mathrm{mJ} \tag{6.11}$$

per retransmission.

Using the same lower-bound traffic scenario established in Section 5.6, a gateway located on the traversed communication path would also retransmit twice per request–response transaction: once in each direction. If all 144 daily transactions cross the gateway, this therefore corresponds to 288 gateway retransmissions during the 12-h active period. The associated daily charge is

$$Q_{\mathrm{TX,day}} = 288 Q_{\mathrm{TX}} \approx 1.31\ \mathrm{mAh/day}. \tag{6.12}$$

This contribution remains very small compared with the continuous receive consumption of the gateway. Using the OLED-off RX current of 50.12 mA as the reference operating condition gives

$$Q_{\mathrm{RX,day}} = 50.12\ \mathrm{mA} \times 24\ \mathrm{h} \approx 1.203\ \mathrm{Ah/day}. \tag{6.13}$$

Including these estimated 288 daily retransmissions gives a total daily charge consumption of approximately

$$Q_{\mathrm{day}} \approx 1.203\ \mathrm{Ah} + 0.00131\ \mathrm{Ah} \approx 1.204\ \mathrm{Ah/day}. \tag{6.14}$$

The retransmission contribution therefore represents only about 0.1% of the total daily gateway charge consumption under this traffic assumption. Gateway autonomy is consequently dominated by continuous LoRa reception rather than by the short transmission events.

Over a five-day deployment, the corresponding charge requirement is

$$Q_{\mathrm{5d}} \approx 5 \times 1.204 \approx 6.02\ \mathrm{Ah}. \tag{6.15}$$

At the nominal 3.7 V voltage of a standard 18650 Li-ion cell, this requirement can also be expressed directly in terms of battery capacity. Assuming a nominal cell capacity of 2.5 Ah, the minimum number of cells required is

$$N_{\text{cell}} = \frac{6.02\ \text{Ah}}{2.5\ \text{Ah}} \approx 2.4. \tag{6.16}$$

A five-day deployment would therefore require at least three 18650 cells connected in parallel, corresponding to a nominal 7.5 Ah capacity at 3.7 V. This provides a nominal capacity margin of approximately 25% relative to the estimated 6.02 Ah requirement.

These values do not include battery-capacity derating, temperature effects, ageing or the design margin required for practical field deployment. They nevertheless show that the present gateway autonomy is dominated by the continuous baseline consumption of the evaluation-board implementation rather than by packet retransmission itself. Reducing this baseline through a dedicated low-power design could therefore substantially reduce the battery capacity required for future deployments.

Battery-voltage monitoring

The gateway battery voltage is monitored using an STM32 ADC connected through a 1 MΩ/1 MΩ resistive divider. The corresponding ideal relationships are

$$V_{\text{ADC}} = V_{\text{bat}} \frac{R_2}{R_1 + R_2} = \frac{V_{\text{bat}}}{2}, \qquad V_{\text{bat}} = 2V_{\text{ADC}}. \tag{6.17}$$

A 68 nF capacitor connected to the divider midpoint forms a low-pass filter. Since

$$R_{\text{th}} = R_1 \parallel R_2 = 500\ \text{k}\Omega, \tag{6.18}$$

its time constant is

$$\tau = R_{\text{th}} C = 500\ \text{k}\Omega \times 68\ \text{nF} \approx 34\ \text{ms}. \tag{6.19}$$

The measurement circuit was validated by replacing the battery with a laboratory DC supply and varying its voltage between approximately 2.5 and 4.2 V. The actual supply voltage was measured independently with a digital multimeter and compared with the voltage reported by the gateway, as shown in Fig. 6.3.

The measured relationship was

$$V_{\text{gateway}} = 0.9968\, V_{\text{DMM}} - 0.0394\ \text{V}, \tag{6.20}$$

with

$$R^2 = 0.99999. \tag{6.21}$$

The slope differs from unity by less than 0.4%, confirming a good linearity of the monitoring circuit over the tested range. The mean absolute voltage error was approximately 50 mV, with a maximum observed error of approximately 56 mV. The lowest successfully measured supply voltage was 2.496 V; at approximately 2.40 V, the gateway shut down.

Finally, the high divider resistance limits its continuous current consumption to

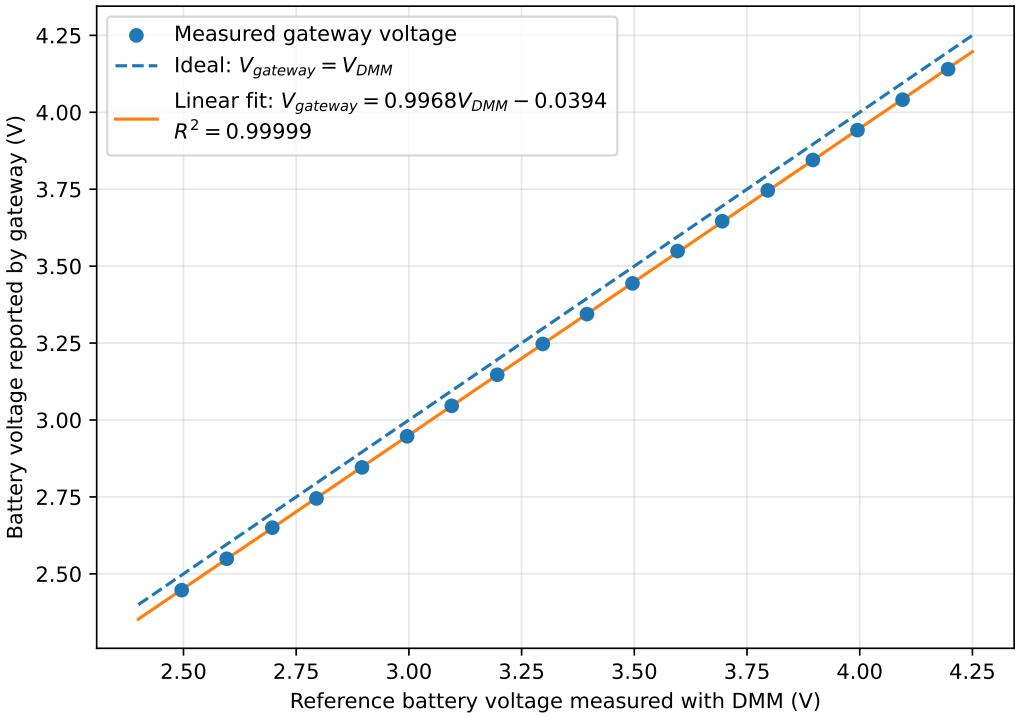

Figure 6.3: Gateway-reported battery voltage as a function of the reference voltage measured with a digital multimeter. The dashed line represents the ideal relationship $V_{\text{gateway}} = V_{\text{DMM}}$, while the solid line represents the linear regression of the experimental data.

$$I_{\text{divider}} = \frac{V_{\text{bat}}}{R_1 + R_2} \approx 2.1\ \mu\text{A} \tag{6.22}$$

at 4.2 V, which is negligible compared with the approximately 50 mA gateway operating current.

6.2 Acquisition and data-logging performance

6.2.1 Acquisition timing

The time required to acquire the electrode voltages was measured using the microcontroller `micros()` timer. The test was performed with the ADS1115 configured at its maximum data rate of 860 SPS. Two configurations were evaluated: sequential acquisition of two channels, corresponding to the original Watcher configuration, and sequential acquisition of all four channels available on the redesigned hardware. Each measurement was repeated 1000 times. The results of this test are shown in Table 6.5.

Table 6.5: Measured sequential ADC acquisition time.

Configuration	Mean	Maximum
2 channels	5.355 ms	5.362 ms
4 channels	10.709 ms	10.720 ms

Consequently, the current implementation already exceeds the 10-ms interval required for a 100-Hz four-channel acquisition rate before record construction and SD-card writing are considered.

This limitation does not originate directly from the nominal conversion rate of the ADS1115. At the selected 860-SPS setting, four conversions represent only

$$4 \times \frac{1}{860} \approx 4.65 \text{ ms}, \tag{6.23}$$

as discussed in Section 3.3.4. The difference between this theoretical conversion time and the measured 10.7 ms therefore results from the complete sequential acquisition procedure, which also includes input-multiplexer reconfiguration, I^2C transactions and software overhead associated with the present ADC driver.

However, the nominal ADS1115 data rate indicates that this is not necessarily a fundamental limitation of the converter itself. A more detailed investigation of the acquisition sequence would be required to determine whether lower-level ADC control, reduced communication overhead or a different conversion scheduling strategy could bring the complete four-channel acquisition within the available timing budget. For example, the electrode channels could be distributed between multiple converters to allow part of the acquisition to be processed in parallel.

6.2.2 SD-card write timing and buffering

The SD-card storage path was characterised to determine whether binary measurement records could be written continuously without disturbing the 10-ms acquisition period. Each measurement record occupies 16 bytes, while the SD-card sector size is 512 bytes. Consequently, 32 consecutive records correspond exactly to one sector.

A first controlled test was performed using the same SdFat configuration as the Watcher firmware, with the SD card operated in dedicated SPI mode at 12 MHz. A total of 3200 records were written at 100 Hz to a preallocated file. Most calls to `write()` required only about 10 μs. However, every 32nd call was systematically longer, with durations close to 1 ms. The periodicity of these longer operations coincides exactly with the completion of a 512-byte sector, which is consistent with the 16-byte records being first accumulated in a sector-sized buffer before being transferred to the SD card.

The influence of file preallocation was then evaluated directly in the complete Watcher firmware. Two file-handling configurations were compared: one in which the file was allowed to grow dynamically during acquisition, and one in which a 65-MiB region was reserved beforehand using `preAllocate()`. In the latter case, the current write position was preserved when changing SPI mode so that subsequent measurement records remained within the preallocated region.

The measured timings obtained with these different configurations are summarised in Table 6.6.

Table 6.6: Measured SD-card write and synchronisation timings at 100 Hz.

Operation and configuration	Minimum	Mean	Maximum
`write()`, controlled preallocated test	10 μs	41.1 μs	1.02 ms
`write()`, file growing during acquisition	11 μs	49.3 μs	24.0 ms
`write()`, preallocated Watcher file	10 μs	138 μs	4.30 ms
`flush()`, file growing during acquisition	15.6 ms	18.7 ms	21.7 ms
`flush()`, preallocated Watcher file	3.06 ms	3.58 ms	3.82 ms

The comparison shows that preallocation mainly improves the worst-case behaviour of the storage path. When the file was allowed to grow during acquisition, individual `write()` operations reached 24.0 ms, already exceeding the complete sampling period. When the measurement stream was instead kept inside the preallocated region, the maximum observed write duration decreased to 4.30 ms. The corrected file position was also verified explicitly: after 6000 records, corresponding to 96 000 bytes, the measured final position matched the expected position exactly.

Explicit file synchronisation was characterised by calling `flush()` every 1000 records, corresponding to one operation every 10 s at 100 Hz. Preallocation again reduced the observed latency: the mean flush duration decreased from 18.7 ms when the file was growing during acquisition to 3.58 ms when the reserved region was used correctly. Nevertheless, even this reduced latency remains significant relative to the 10-ms sampling interval.

In the two-electrode acquisition configuration, the two ADC conversions and associated processing require approximately 5.4 ms. Adding a typical flush duration of 3.6 ms therefore leaves only about 1 ms for the binary record write and the remaining acquisition tasks. Since individual `write()` calls were observed to require up to 4.30 ms, periodically forcing an SD-card synchronisation cannot be guaranteed to remain within the available 10-ms timing budget. Consistent with this observation, the acquisition iterations containing a `flush()` exceeded the nominal sampling interval during the Watcher-level timing test.

Periodic `flush()` operations were therefore not retained in the time-critical recording loop of the present implementation. Complete 512-byte sectors are transferred progressively as the SdFat buffer is filled, while explicit synchronisation is still performed when leaving the recording state under controlled conditions. The remaining design question is the amount of recently written data that may remain unrecoverable if power is removed before such a controlled synchronisation occurs. This aspect is evaluated separately through abrupt power-loss tests.

6.2.3 Power-loss integrity

To evaluate whether periodic synchronisation was necessary despite its timing cost, the consequence of an abrupt power interruption was characterised experimentally. Ten acquisitions were interrupted by removing the Watcher supply without calling either `flush()` or `close()`. A sequence number and run identifier were embedded in each 16-byte test record, allowing the last contiguous valid record to be identified after recovery of the SD card.

Across the ten tests, the confirmed unrecoverable interval ranged from 0 to 26 records, corresponding to a maximum observed loss of 0.26 s at 100 Hz. Since the diagnostic sequence number was reported through the serial interface only every ten records, the exact instant of the power interruption was known with an uncertainty of up to nine additional writes. The mean and median confirmed losses were both 0.09 s. The main results are summarised in Table 6.7.

Table 6.7: Summary of the 10 abrupt power-loss tests without periodic SD-card synchronisation.

Metric	Result
Mean confirmed data loss	0.09 s
Median confirmed data loss	0.09 s
Maximum observed confirmed loss	26 records (0.26 s)
Maximum expected loss	32 records (0.32 s)
Files ending on a 512-byte sector boundary	10/10

A consistent pattern was observed in all ten recovered files: the last valid record always coincided with a 512-byte sector boundary. Since each measurement record occupies 16 bytes, one sector contains 32 records. This behaviour is consistent with complete sectors being progressively transferred to the SD card while the final partially filled sector remains buffered. No loss of multiple complete sectors was observed during these experiments. This result should nevertheless be interpreted as an experimental observation rather than as a strict upper bound, since additional buffering may occur internally within the SD-card controller.

Considering this result together with the measured 3.58 ms mean duration of `flush()`, periodic synchronisation was not retained during the time-critical 100-Hz recording loop. Under the tested conditions,

accepting a potential loss of maximum 0.32 s of data in the exceptional case of an abrupt power interruption was considered preferable to periodically reducing the available 10-ms acquisition timing margin. The file is instead synchronised during controlled transitions out of the recording state.

6.2.4 Continuous logging integrity

A one-hour continuous acquisition test was performed to validate the complete recording path. Two electrode channels were acquired on a nominal 100-Hz time grid and stored as 16-byte binary records in a preallocated file using dedicated SPI mode. No explicit `flush()` operation was performed during acquisition.

The scheduler maintains a fixed 10-ms sampling grid. When an acquisition iteration extends beyond one or more subsequent sampling instants, the corresponding overdue slots are explicitly skipped. The sequence number stored in each record therefore remains contiguous and represents only measurements that were actually acquired and written to the SD card. This allows storage integrity to be distinguished from missed acquisition deadlines.

The main results of the one-hour experiment are summarised in Table 6.8.

Table 6.8: Summary of the one-hour continuous logging test using the final fixed-grid sampling scheduler.

Metric	Result
Nominal sampling slots	360 000
Recorded samples	359 870
Skipped sampling slots	130 (0.036%)
Missed-slot events	25
Intervals exactly 10 ms	359 818 (99.986%)
Effective stored sampling rate	99.964 Hz
Mean sampling interval	10.0036 ms
Maximum sampling interval	93 ms
Maximum SD `write()` duration	87.9 ms
Missing or duplicated stored records	0

The one-hour test produced 359 870 valid records from the 360 000 nominal sampling slots. Only 130 slots were skipped (0.036%), distributed over 25 timing events. The recovered sequence was continuous from 1 to 359 870, with no missing, duplicated or out-of-order stored records, and the final file size matched exactly the expected number of 16-byte records.

The acquisition timing remained very close to the nominal 100-Hz grid: 359 818 intervals (99.986%) were exactly 10 ms, and the effective stored sampling rate was 99.964 Hz. The largest observed interval was 93 ms.

These occasional long intervals were associated with unusually long SD-card accesses. The longest `write()` operation measured during the test was 87.9 ms. Such delays are likely related to internal SD-card operations, such as flash-page programming, block erasure or controller housekeeping, whose latency is not fully deterministic even when the file is preallocated. When this occurs, the scheduler skips sampling slots that can no longer be reached and resumes on the original 10-ms grid.

For the intended application, the test therefore shows that continuous logging is reliable over long acquisitions: no missing or duplicated stored record was detected, while only 0.036% of the nominal sampling slots were missed because of occasional SD-card latency.

6.3 Analogue front-end validation

6.3.1 Electronic offset and noise

The electronic offset and noise of the four acquisition channels were characterised with the field electrodes disconnected and the analogue inputs connected to GNDA through the DG413 switches. This configuration isolates the acquisition electronics from electrode and environmental noise.

For each channel, the `adc_noise()` routine acquired $N = 100$ consecutive samples and calculated their mean and standard deviation. The test was repeated ten times under identical conditions. The values reported below correspond to the average of these ten repetitions. With the input connected to analogue ground, the ideal transfer function predicts an ADC output of 2.5 V. The channel offset was therefore calculated as

$$V_{\text{offset},i} = \overline{V}_i - 2.5 \text{ V}. \tag{6.24}$$

The standard deviation σ_i was used as an estimate of the short-term electronic noise at the ADC input. The results are summarised in Table 6.9.

Table 6.9: Measured offset and electronic noise of the four acquisition channels with their inputs connected to analogue ground.

Channel	Mean ADC voltage	Offset from 2.5 V	Mean standard deviation
M1	2499.639 mV	-0.361 mV	0.154 mV
M2	2502.309 mV	+2.309 mV	0.158 mV
M3	2496.224 mV	-3.776 mV	0.154 mV
M4	2499.125 mV	-0.875 mV	0.156 mV

The measured mean voltages ranged from 2496.224 to 2502.309 mV, corresponding to offsets between -3.776 and $+2.309$ mV relative to the nominal 2.5-V bias point. The largest absolute offset was observed on channel M3.

The measured standard deviations were highly consistent across the four channels, ranging from 0.154 to 0.158 mV. Since the conditioning chain has a gain of 0.5 from the electrode input to the ADC, the corresponding input-referred noise is approximately

$$\sigma_{\text{electrode},i} = 2\sigma_i. \tag{6.25}$$

This corresponds to an input-referred standard deviation between 0.308 and 0.316 mV. The measurements therefore show that all four acquisition channels operate close to the expected 2.5-V bias point and exhibit comparable electronic-noise levels.

However, the resulting input-referred standard deviation of approximately 0.31 mV is substantially larger than the weakest microvolt-scale signals reported for deep-ERT applications in Section 2.1.4. Simple stacking cannot by itself be assumed to compensate for this difference. For uncor related noise, averaging N independent observations reduces the standard deviation only as

$$\sigma_N = \frac{\sigma}{\sqrt{N}}.$$

Reducing an instantaneous noise level of approximately 310 μV to 1 μV would therefore require on the order of

$$N \approx \left(\frac{310}{1}\right)^2 \approx 10^5$$

independent observations. The number of complete injection cycles available during a typical 5–10 min recording is several orders of magnitude smaller than this value. Although additional averaging can also be performed between samples within each repeated waveform, its effectiveness depends on the noise spectrum, temporal correlation and signal-processing strategy and was not quantified in the present work.

The present experiment therefore does not demonstrate that the redesigned Watcher can resolve the weakest microvolt-scale deep-ERT signals. A dedicated sensitivity experiment using representative low-amplitude periodic signals and the complete field-processing chain would be required to establish the minimum practically recoverable signal amplitude.

The fixed ADC input range used in the present implementation also represents a limitation in this respect. The ADS1115 already integrates a programmable gain amplifier, but the present single-ended acquisition scheme cannot fully exploit it because the conditioned signal is centred around 2.5 V. The ± 6.144 V range was therefore selected to accommodate the complete 0–5 V conditioned signal.

A future revision could instead exploit the differential capability of the ADS1115 by measuring each conditioned electrode signal relative to the 2.5-V bias reference. The ADC would then convert only the signal variation around this bias, allowing its smaller programmable full-scale ranges to be selected for weak signals. The range could be chosen initially from the expected signal amplitude and subsequently adjusted after a preliminary measurement to maximise use of the ADC dynamic range without saturation.

Such an architecture could substantially improve the effective resolution without necessarily requiring a different converter. However, it would require a modification of the present ADC input routing, since all four ADS1115 inputs are currently used as single-ended electrode channels. Its actual improvement in minimum recoverable signal amplitude would also have to be established experimentally.

6.3.2 DC gain and linearity

The DC transfer characteristic of the four acquisition channels was evaluated by applying known voltages between each electrode input and GNDA. The input voltage was varied from approximately -5 to $+5$ V in 1-V steps. For each operating point, the actual input voltage was measured with a digital multimeter, while the Watcher acquired ten consecutive ADC readings whose mean was used as the measured output voltage.

According to the analogue conditioning architecture described in Section 3.3, the ideal transfer function is

$$V_{\mathrm{ADC}} = 0.5V_{\mathrm{in}} + 2.5 \text{ V}. \tag{6.26}$$

Figure 6.4 compares the measured transfer characteristics with this ideal relationship.

A linear regression was performed for each channel using the central measurement range from approximately -4 to $+4$ V. The two measurements close to the -5 and $+5$ V limits were retained in Fig. 6.4 to demonstrate operation over the complete nominal input range, but were excluded from the regression because they correspond to ADC output voltages close to the 0- and 5-V limits. The resulting fitted model,

$$V_{\mathrm{ADC}} = aV_{\mathrm{in}} + b, \tag{6.27}$$

is summarised in Table 6.10.

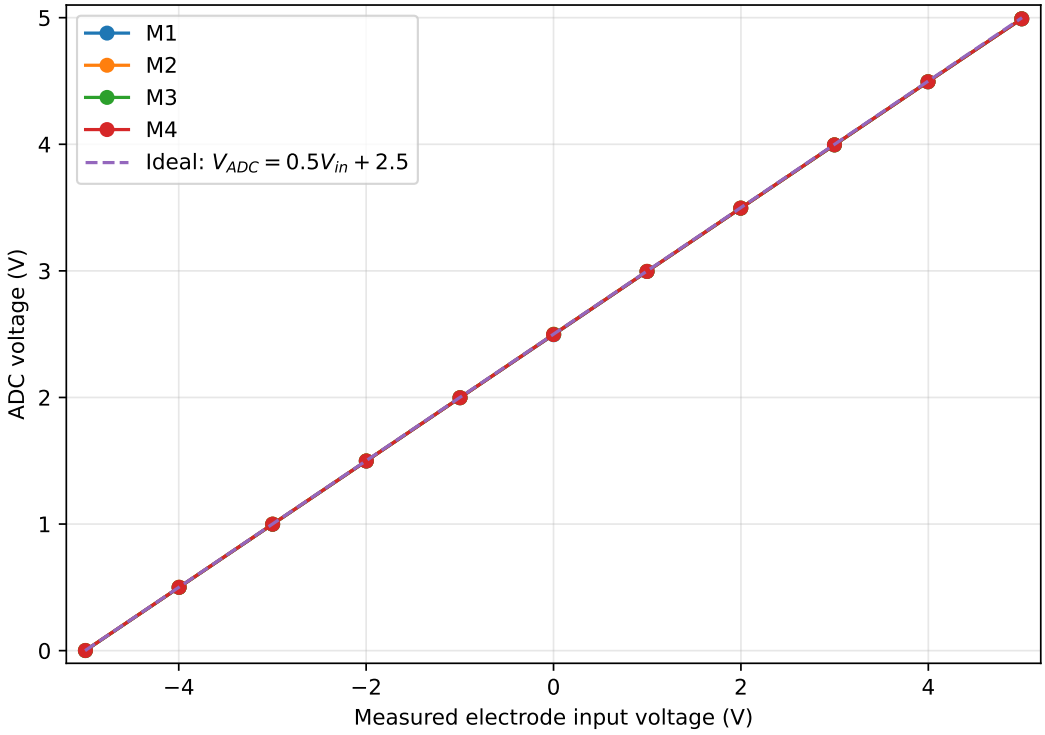

Figure 6.4: Measured DC transfer characteristics of the four analogue acquisition channels. The horizontal axis corresponds to the actual electrode voltage measured with a digital multimeter. The dashed line represents the ideal transfer function $V_{\mathrm{ADC}} = 0.5V_{\mathrm{in}} + 2.5$ V.

Table 6.10: Linear-regression results for the four analogue acquisition channels over the approximately −4 to +4 V input range.

Channel	Gain a	Offset b (V)	Deviation from ideal gain	R^2
M1	0.499747	2.496741	-0.051%	0.99999965
M2	0.499655	2.497872	-0.069%	0.99999996
M3	0.499722	2.496407	-0.056%	0.99999985
M4	0.499731	2.496682	-0.054%	0.99999994

The measured gains ranged from 0.499655 to 0.499747, compared with the theoretical value of 0.5. The corresponding deviations from the ideal gain remained below 0.07% for all four channels, while the fitted offsets ranged from 2.4964 to 2.4979 V.

The input voltage used as the reference for this characterisation was measured with a UNI-T UT61E+ digital multimeter. In the 22-V DC range, the instrument has a resolution of 1 mV and a specified accuracy of $\pm(0.05\%$ of reading $+ 5$ digits), corresponding to approximately ± 7 mV at a 4-V reading. The measured gain deviations are therefore smaller than the absolute accuracy of the reference instrument and should not be interpreted as a calibrated absolute gain accuracy. The experiment nevertheless demonstrates the high linearity and close agreement of the four acquisition channels over the tested voltage range.

All four regressions produced R^2 values greater than 0.9999996, showing an essentially linear response over the tested central input range. Together with the measurements at the approximately ± 5 V limits, these results confirm that the redesigned analogue front end provides the expected gain and level shift over the intended electrode-voltage range, with closely matched behaviour between the four acquisition channels.

6.3.3 Analogue front-end power-up settling

Since the ± 15 V analogue supply is disabled when the measurement chain is not required, the settling behaviour following its activation was characterised experimentally. For this test, the conditioning-chain inputs were connected to GNDA and the four ADC channels were sampled repeatedly immediately after enabling the TEC3 converter. Because the four channels are acquired sequentially rather than simultaneously, the time associated with each point in Fig. 6.5 was corrected to account for the finite acquisition time of each channel within the four-channel reading sequence.

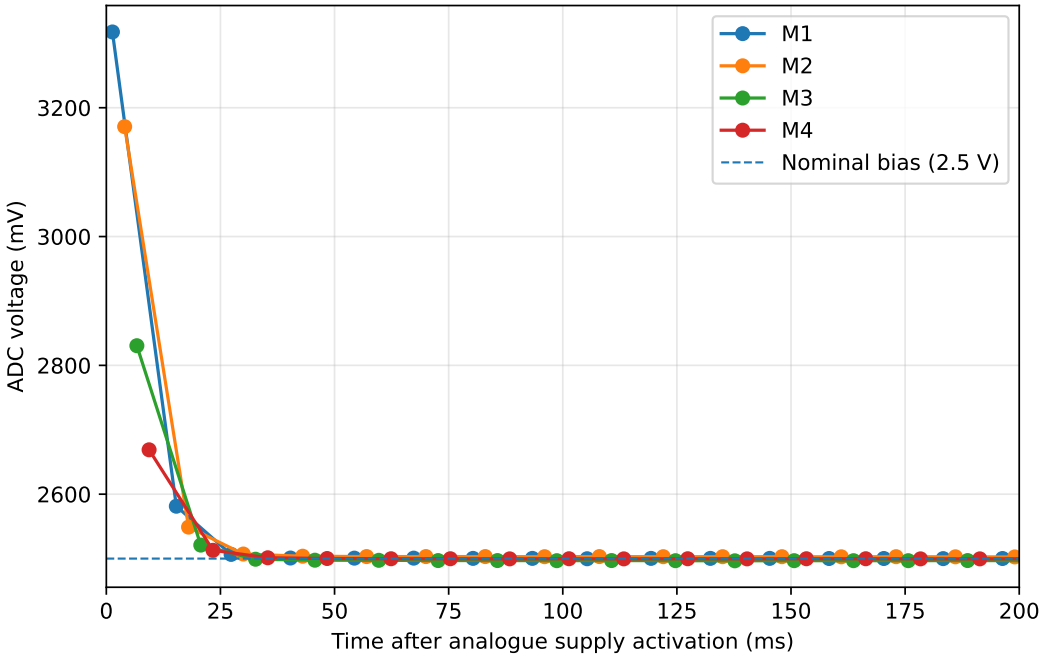

Figure 6.5: Measured settling of the four analogue acquisition channels after activation of the ± 15 V supply. The inputs were connected to GNDA during the test. The dashed line indicates the nominal 2.5-V bias point.

Figure 6.5 shows the measured start-up transient of the four acquisition channels. Immediately after activation of the analogue supply, the channel outputs were significantly displaced from their nominal 2.5-V operating point. The largest initial deviation was observed on M1, whose first measured value was approximately 3.32 V, while the other channels started between about 2.67 V and 3.17 V.

All four channels then converged rapidly towards their steady-state values. By approximately 40–50 ms after activation, the four outputs were already very close to their final bias point, and no significant additional variation attributable to the start-up transient was observed thereafter. The remaining differences between channels beyond this point are consistent with the small static offsets reported in Section 6.3.1.

These results show that the analogue front end becomes effectively stable within less than 50 ms after the ±15 V supply is enabled. The 100-ms settling delay currently used in the firmware before launching diagnostic or acquisition operations is therefore conservative and provides a comfortable margin with respect to the experimentally observed start-up behaviour.

6.3.4 Low-pass filter frequency response

The frequency response of the analogue low-pass filter was characterised experimentally using the TG315 function generator and the TDS1001C-EDU oscilloscope. A sinusoidal signal was injected directly at the filter input, and the peak-to-peak voltages before and after the filter were measured over a frequency range from 5 to 500 Hz. The measured gain was calculated at each frequency as

$$G(f) = 20 \log_{10} \left(\frac{V_{\text{out,pp}}}{V_{\text{in,pp}}} \right). \tag{6.28}$$

Figure 6.6 compares the experimental measurements with the theoretical response of the nominal 4.7 kΩ, 1 μF first-order RC filter, whose cutoff frequency is

$$f_c = \frac{1}{2\pi RC} \approx 33.9 \text{ Hz}. \tag{6.29}$$

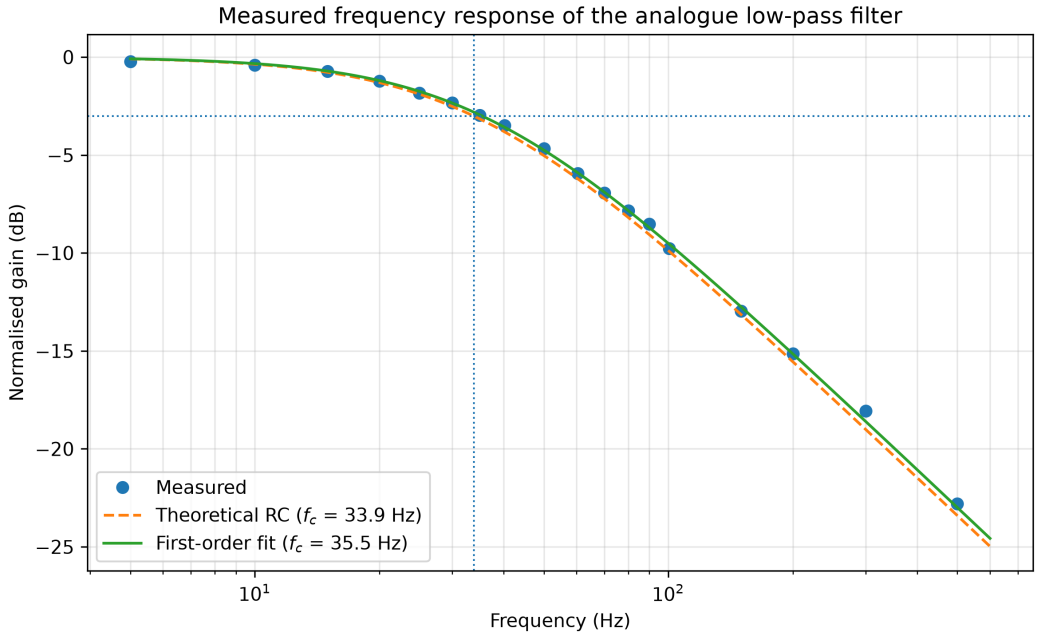

Figure 6.6: Measured frequency response of the analogue low-pass filter compared with the theoretical first-order RC response and a first-order fit to the experimental data.

The measured response closely follows the expected first-order low-pass behaviour over the investigated frequency range. A first-order fit to the experimental data gives a cutoff frequency of approximately 35.5 Hz, close to the theoretical value of 33.9 Hz. Above the cutoff frequency, the measured attenuation also follows the expected approximately −20 dB/decade trend.

Using the experimentally fitted cutoff frequency of $f_c = 35.5$ Hz, the attenuation at the 50 Hz mains frequency is -4.75 dB, corresponding to an output amplitude of about 58% of the input amplitude.

The experiment therefore confirms the expected frequency-selective behaviour of the implemented analogue low-pass filter and shows that its effective cutoff frequency remains close to the design value.

6.3.5 Contact-resistance measurement accuracy

The accuracy of the contact-resistance measurement circuit was evaluated using five discrete resistors connected successively between each electrode terminal and the Watcher ground terminal. Their actual resistances were first measured with the digital multimeter and were 988 Ω, 5.098 kΩ, 9.950 kΩ, 100.26 kΩ and 992.32 kΩ. The UT61E+ specifies a resistance-measurement accuracy of ±(0.5% of reading+10 digits) [**ut61_datasheet**]. Only one resistor and one electrode channel were connected at a time, while the three remaining electrode terminals were left open.

For each resistor value, five consecutive contact-resistance measurements were performed on each of the four channels. The standard firmware measurement principle was retained: the corresponding relay was activated for the configured injection duration, the current was estimated from the INA282 output voltage and the 2 Ω shunt, and the known 100 Ω protection resistance and shunt resistance were subsequently removed from the calculated loop resistance.

Figure 6.7 compares the mean resistance measured by each Watcher channel with the corresponding resistance measured using the digital multimeter.

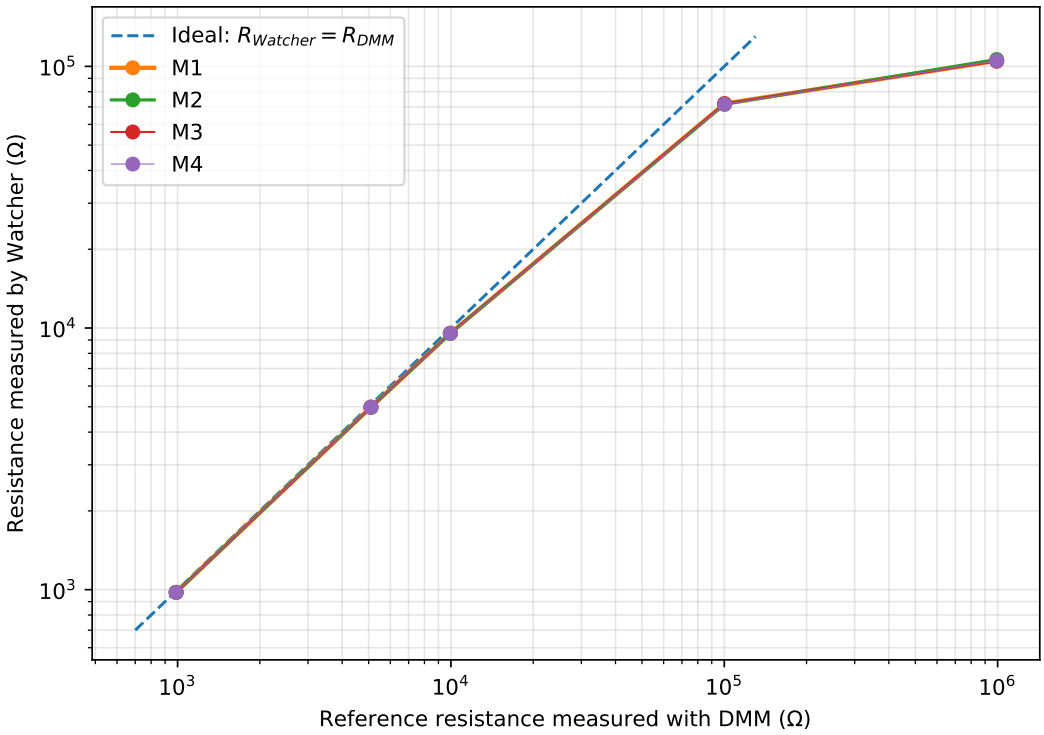

Figure 6.7: Measured contact resistance as a function of the reference resistance measured with a digital multimeter. Each point represents the mean of five measurements. The dashed line represents the ideal relationship $R_{\mathrm{Watcher}} = R_{\mathrm{DMM}}$. The four electrode channels produced very similar results, resulting in nearly overlapping curves.

The results are summarised in Table 6.11. Since the behaviour of the four channels was very similar, the table reports the mean value across the four channels for each reference resistance, together with the mean relative error and the average within-series standard deviation.

Table 6.11: Contact-resistance measurement accuracy for the four tested reference resistances.

R_{DMM}	**Mean measured** R	**Mean error**	**Mean std.**
0.988 kΩ	0.9764 kΩ	-1.17%	0.51 Ω
5.098 kΩ	4.9805 kΩ	-2.30%	4.93 Ω
9.950 kΩ	9.5404 kΩ	-4.12%	12.33 Ω
100.260 kΩ	71.823 kΩ	-28.36%	0.874 kΩ
992.320 kΩ	104.605 kΩ	-89.46%	3.48 kΩ

For the three lowest reference resistances, the circuit provided consistent measurements on all four electrode channels. The mean relative error increased from approximately -1.2% at 988 Ω to -2.3% at 5.098 kΩ and -4.1% at 9.950 kΩ. The channel-to-channel differences remained small compared with the measured resistance, indicating good reproducibility of the four replicated RS-check paths.

The deviation increased substantially at higher resistance. For the 100.260 kΩ reference, the mean measured resistance was approximately 71.8 kΩ, corresponding to an error of approximately -28%. For the 992.320 kΩ reference, the mean measured value reached only approximately 104.6 kΩ, corresponding to an error of approximately -89.5%.

The deterioration observed at high contact resistance results from the progressive reduction of the test current and, consequently, of the voltage developed across the 2 Ω shunt. As the contact resistance increases, the INA282 therefore operates with an increasingly small differential input signal. At approximately $R_{contact} = 30$ kΩ, the amplified shunt voltage approaches the specified lower linear-output limit of the INA282, around 40 mV. Beyond this point, the current can no longer be estimated with the same accuracy, which is consistent with the increasing measurement error observed experimentally.

With no external resistance connected, the Watcher reports an apparent resistance of approximately 120 kΩ. Increasingly large external resistances therefore tend towards this finite open-circuit value rather than being measured proportionally. The circuit consequently provides repeatable and reasonably accurate measurements in the approximately 1–10 kΩ range, while its accuracy deteriorates substantially at higher resistance. This remains acceptable for its intended diagnostic purpose, which is primarily to identify poor electrode contacts rather than to provide high-accuracy resistance measurements over a wide dynamic range.

6.4 Wireless communication performance

6.4.1 Antenna positioning

The influence of the Watcher antenna configuration was evaluated by comparing the original internal wire antenna with two mounting configurations of the dedicated 433-MHz antenna. The transmitting unit remained fixed throughout the experiment and transmitted one packet per second at 10 dBm using identical LoRa parameters. For each antenna configuration, 100 packets were transmitted, corresponding to a 100-s measurement period. The packet reception ratio (PRR), received signal strength indicator (RSSI) and signal-to-noise ratio (SNR) were calculated from the received packets.

The resulting performance is summarised in Table 6.12.

All three configurations achieved a PRR of 100%, indicating that the available link margin was sufficient for reliable packet reception under the tested conditions. The differences were therefore mainly observed in received signal strength and stability.

Both dedicated external antennas provided a clear improvement over the original horizontal wire antenna. The top-mounted configuration increased the mean RSSI by approximately 14.2 dB, while the side-mounted

Table 6.12: Comparison of the three Watcher antenna configurations.

Configuration	PRR (%)	Mean RSSI (dBm)	RSSI std. (dB)	Mean SNR (dB)	RSSI gain (dB)
Horizontal wire	100	-72.1	4.7	6.1	–
Side-mounted antenna	100	-59.8	1.6	6.8	+12.3
Top-mounted antenna	100	-57.9	1.2	7.1	+14.2

configuration provided an improvement of approximately 12.3 dB. The two dedicated configurations also showed much lower RSSI variability than the original wire antenna.

The difference between the two external configurations remained comparatively small, with the top-mounted antenna exceeding the side-mounted configuration by approximately 1.9 dB. Nevertheless, the top-mounted configuration combined the highest mean RSSI, the lowest RSSI variability and the highest mean SNR.

6.4.2 Radio interoperability and direct-link performance

The interoperability between the RFM96/SX1276 transceiver integrated in the Watcher and the SX1268-based radio used in the gateway was first verified through a direct bidirectional communication test. Both radios were configured with the same LoRa physical-layer parameters used in the Watcher network: a carrier frequency of 433 MHz, a bandwidth of 500 kHz, spreading factor 11, coding rate 4/5, an eight-symbol preamble and payload CRC enabled. The standard 17-byte Watcher packet format was retained.

The two devices were placed at fixed positions with their antennas vertically oriented. A sequence of 100 numbered packets was transmitted in each direction. For each received packet, the sequence number and packet contents were checked, while RSSI and SNR were recorded at the receiver. The results are summarised in Table 6.13.

Table 6.13: Direct-link interoperability between the Watcher RFM96 radio and the SX1268-based gateway radio.

Direction	Sent	Received	PRR (%)	Mean RSSI (dBm)	RSSI std. (dB)	Mean SNR (dB)
Gateway → Watcher	100	100	100	-58.4	1.3	7.2
Watcher → Gateway	100	100	100	-59.7	1.6	6.8

All transmitted packets were successfully decoded in both directions, resulting in a packet reception ratio of 100%. No malformed packets or sequence-number inconsistencies were observed. The mean RSSI remained close to -60 dBm in both directions, with a standard deviation below 2 dB, indicating a stable direct radio link under the tested geometry.

A small difference was observed between the two transmission directions, with the gateway-to-Watcher link exhibiting an approximately 1.3 dB higher mean RSSI. Such a difference is compatible with the use of different radio front ends, antenna implementations and receiver architectures on the two devices, and does not affect the practical interoperability of the link.

This test therefore confirms that the previous-generation SX1276-based Watcher radio can communicate directly with the newer SX1268-based gateway while retaining the same LoRa configuration and Watcher packet format. The SX1268-based unit can consequently be inserted into the existing communication architecture without requiring a modification of the Watcher radio hardware or physical-layer settings.

6.4.3 Transmission and forwarding time

The communication delay introduced by packet transmission and gateway forwarding was characterised using the standard 17-byte Watcher packet format.

The Watcher transmission duration was first measured by timing the blocking LoRa transmission function with the microcontroller `micros()` timer. One hundred consecutive packets were transmitted using the normal radio configuration. To independently verify that this software timing corresponded to the physical transmission interval, the supply-current transient was also observed with a digital oscilloscope. As shown in Fig. 6.8, the high-current transmission plateau lasted approximately 168.0 ms, in close agreement with the duration obtained from the microcontroller timing.

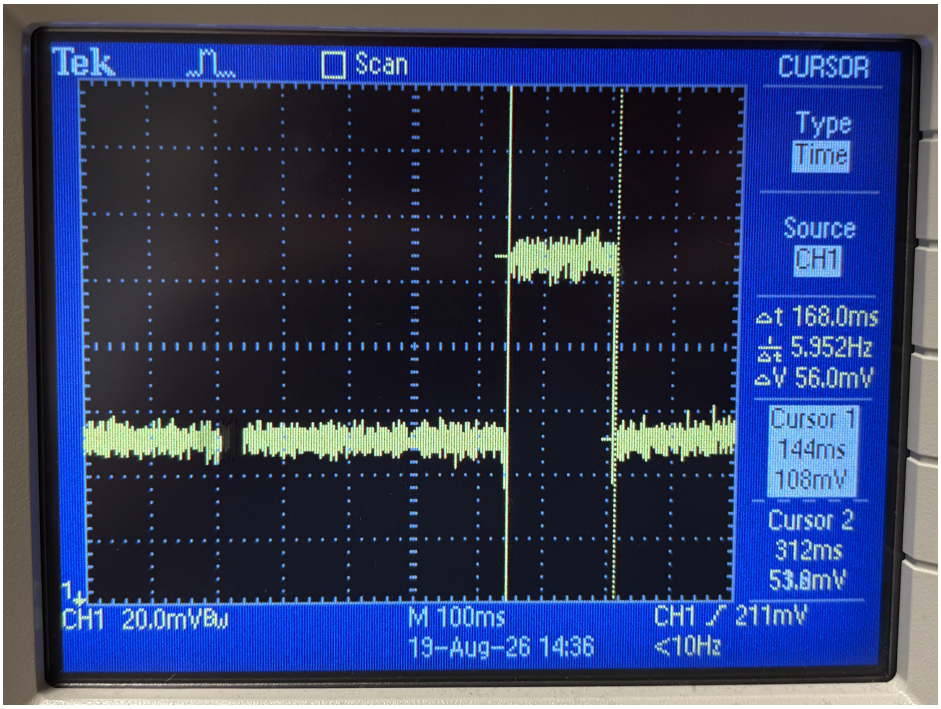

Figure 6.8: Oscilloscope measurement of the Watcher LoRa transmission interval. The increased supply-current plateau associated with the transmission of a 17-byte packet lasted approximately 168.0 ms.

The gateway forwarding time was then characterised separately. A dedicated test packet was transmitted towards the gateway and immediately retransmitted after reception. High-resolution timestamps were recorded at the SX1268 `RX_DONE` and `TX_DONE` events. Additional timestamps were used to determine the time required to process the received packet, switch from RX to TX, and finally restore continuous reception after retransmission.

The resulting timing measurements are summarised in Table 6.14.

Table 6.14: Measured packet-transmission and gateway-forwarding times.

Measured interval	Mean (ms)	Std. (ms)	Range (ms)
Watcher transmission	167.579	0.020	167.555–167.758
Gateway processing before TX	56.118	0.522	55.275–57.063
Gateway TX stage	209.722	0.001	209.721–209.725
Gateway RX done → TX done	265.840	0.522	264.997–266.785
Gateway RX done → RX rearmed	321.108	0.522	320.265–322.053

A 17-byte Watcher transmission required on average 167.579 ms, with very little variation between successive packets. This result is consistent with the oscilloscope measurement.

At the gateway, approximately 56.1 ms elapsed between complete packet reception and the start of the measured transmission stage. The complete interval from `RX_DONE` to `TX_DONE` was therefore approximately 265.8 ms.

After retransmission, the gateway required approximately 55 ms more before continuous reception was restored. The radio was consequently unavailable for reception for approximately 321.1 ms after completion of an incoming packet. This interval represents an important limitation if several packets arrive within a short time window.

The measured Watcher transmission duration can also be used to estimate the minimum propagation time through a long Watcher chain. For a network of 100 Watchers with the remote unit located immediately before the first Watcher, a command addressed to the final Watcher requires 100 successive radio transmissions. Neglecting processing and retransmission overhead at the intermediate Watchers, the theoretical minimum one-way delay is therefore

$$t_{\text{chain,min}} = 100\, t_{\text{TX,Watcher}} = 100 \times 168 \text{ ms} = 16.8 \text{ s}. \tag{6.30}$$

A complete request–response transaction with the final Watcher would require the packet to travel through the same 100 hops in both directions, corresponding to a minimum cumulative transmission time of approximately 33.6 s.

The impact becomes more significant if all Watchers must be queried individually, for example during a complete network status check. A request–response transaction with W_k requires $2k$ packet transmissions. Sequentially querying all N Watchers therefore requires

$$N_{\text{TX,all}} = 2\sum_{k=1}^{N} k = N(N+1). \tag{6.31}$$

For a 100-Watcher network, this corresponds to 10100 packet transmissions. Using the measured Watcher transmission duration, the minimum cumulative radio transmission time is therefore

$$t_{\text{poll,all}} = 10100 \times 168 \text{ ms} \approx 1696 \text{ s} \approx 28.2 \text{ min}. \tag{6.32}$$

These values represent theoretical lower bounds, since they include only the packet transmission time and neglect processing delays, packet losses, retries and any additional gateway forwarding. The actual end-to-end delay would therefore be larger. While the latency associated with reaching a single distant Watcher may remain acceptable for occasional commands, sequentially polling the complete network becomes prohibitively slow at the intended deployment scale. This further motivates the investigation of more efficient routing and network-supervision strategies.

The LoRa configuration used throughout this work was retained as a functional baseline rather than being systematically optimised. Parameters such as spreading factor, bandwidth, coding rate and transmit power directly affect the trade-off between link robustness, packet duration and energy consumption. The optimum configuration is strongly dependent on the actual node spacing, terrain, vegetation and antenna installation. A meaningful optimisation should therefore be performed during a representative deep-ERT field campaign. Such a campaign was not available during the present work, so the current settings should be regarded as experimentally validated for interoperability and relaying, but not as an optimised configuration for a full-scale deployment.

6.4.4 Gateway-assisted link recovery

The ability of the intermediate gateway to restore communication across an obstructed radio link was evaluated using a deliberately unfavourable deployment geometry. The transmitting E22-based remote unit was placed at the far end of a garden, while the receiving Watcher was installed at the opposite side of the building, in the basement. The radio path therefore crossed both the exterior distance and several building structures before reaching the Watcher.

To make the test representative of a degraded link rather than relying on the maximum available transmitter power, the E22 transmitter output power was reduced to 5 dBm. All other LoRa parameters were kept identical to those used during the interoperability tests.

The direct link was first evaluated without an intermediate gateway. A sequence of 100 numbered packets was transmitted by the remote unit. None of these packets was successfully decoded by the Watcher, corresponding to a packet reception ratio of 0%. During the preliminary positioning of the remote unit, the last intermittently received packets exhibited RSSI values close to −121 dBm and SNR values around −12 dB, indicating operation close to the receiver sensitivity limit. Once the remote unit was moved to the final test position, the link was completely lost and no valid packet-level RSSI or SNR could therefore be recorded.

An intermediate gateway was then positioned on top of the basement stairwell, a location providing a usable radio path towards both devices. The gateway received packets transmitted by the remote unit and immediately forwarded them towards the Watcher using the same 17-byte packet format. The end-to-end results are summarised in Table 6.15.

Table 6.15: Gateway-assisted recovery of an obstructed radio link.

Link	**Sent**	**Received**	**PRR** (%)	**Mean RSSI** (dBm)	**Mean SNR** (dB)
Remote → Watcher	100	0	0	–	–
Remote → Gateway	100	100	100	-88.6	4.1
Gateway → Watcher	100	100	100	-82.9	5.6

The introduction of the gateway restored the previously unavailable communication path. All 100 packets transmitted by the remote unit were successfully received by the gateway, with a mean RSSI of −88.6 dBm and a mean SNR of 4.1 dB. All 100 forwarded packets were successfully received by the Watcher, with a mean RSSI of −82.9 dBm and a mean SNR of 5.6 dB.

These results illustrate the main purpose of the proposed gateway architecture. Although the direct remote-to-Watcher link was outside the usable reception range in the selected geometry, dividing the path into two shorter radio links provided sufficient link margin for reliable end-to-end communication. The gateway therefore allows a communication path to be recovered without modifying the radio hardware of the existing Watchers.

6.4.5 Duplicate suppression

The duplicate-suppression mechanism was validated using a dedicated radio test in which the E22-based transmitter deliberately generated repeated copies of the same Watcher packet. The purpose of this test was to reproduce the situation in which the same transaction reaches a Watcher through more than one radio path, for example directly and through an intermediate gateway.

Five distinct transactions were generated. For each transaction, the transmitter sent five byte-identical copies of the same 17-byte packet. The Watcher identifies a transaction using the combination of sender

identifier, sequence number and timestamp (`fromId`, `seq`, and `utcms`). The sequence used for the test is summarised in Table 6.16.

Table 6.16: Packet identifiers used for the duplicate-suppression test.

Group	seq	utcms	Copies transmitted
1	42	100000	5
2	43	100000	5
3	43	100001	5
4	255	100002	5
5	0	100003	5

This pattern was selected to verify several cases. Between the first two groups, only the sequence number was modified. Between Groups 2 and 3, the sequence number was deliberately kept unchanged while `utcms` was modified. The final two groups additionally verified that the sequence-number wrap from 255 to 0 does not cause a new packet to be incorrectly rejected.

For each group, the first packet was expected to be accepted as a new transaction, while the four following identical copies were expected to be identified and discarded as duplicates. The resulting terminal output is illustrated in Fig. 6.9.

```
------------------------------------------
RX 01 | seq= 42 | utcms=100000 | RSSI= -41.0 dBm | SNR=  6.8 dB | ACCEPTED – NEW TRANSACTION
RX 02 | seq= 42 | utcms=100000 | RSSI= -28.0 dBm | SNR=  6.2 dB | DUPLICATE SUPPRESSED
RX 03 | seq= 42 | utcms=100000 | RSSI= -29.0 dBm | SNR=  5.8 dB | DUPLICATE SUPPRESSED
RX 04 | seq= 42 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.0 dB | DUPLICATE SUPPRESSED
RX 05 | seq= 42 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.8 dB | DUPLICATE SUPPRESSED
RX 06 | seq= 43 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.5 dB | ACCEPTED – NEW TRANSACTION
RX 07 | seq= 43 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 08 | seq= 43 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.2 dB | DUPLICATE SUPPRESSED
RX 09 | seq= 43 | utcms=100000 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 10 | seq= 43 | utcms=100000 | RSSI= -30.0 dBm | SNR=  4.2 dB | DUPLICATE SUPPRESSED
RX 11 | seq= 43 | utcms=100001 | RSSI= -30.0 dBm | SNR=  5.5 dB | ACCEPTED – NEW TRANSACTION
RX 12 | seq= 43 | utcms=100001 | RSSI= -30.0 dBm | SNR=  6.0 dB | DUPLICATE SUPPRESSED
RX 13 | seq= 43 | utcms=100001 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 14 | seq= 43 | utcms=100001 | RSSI= -31.0 dBm | SNR=  5.8 dB | DUPLICATE SUPPRESSED
RX 15 | seq= 43 | utcms=100001 | RSSI= -31.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 16 | seq=255 | utcms=100002 | RSSI= -30.0 dBm | SNR=  5.2 dB | ACCEPTED – NEW TRANSACTION
RX 17 | seq=255 | utcms=100002 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 18 | seq=255 | utcms=100002 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 19 | seq=255 | utcms=100002 | RSSI= -31.0 dBm | SNR=  5.2 dB | DUPLICATE SUPPRESSED
RX 20 | seq=255 | utcms=100002 | RSSI= -31.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 21 | seq=  0 | utcms=100003 | RSSI= -31.0 dBm | SNR=  5.8 dB | ACCEPTED – NEW TRANSACTION
RX 22 | seq=  0 | utcms=100003 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 23 | seq=  0 | utcms=100003 | RSSI= -30.0 dBm | SNR=  5.5 dB | DUPLICATE SUPPRESSED
RX 24 | seq=  0 | utcms=100003 | RSSI= -31.0 dBm | SNR=  5.2 dB | DUPLICATE SUPPRESSED
RX 25 | seq=  0 | utcms=100003 | RSSI= -31.0 dBm | SNR=  5.2 dB | DUPLICATE SUPPRESSED

==========================================
                RESULTS
==========================================
Physical packets received : 25 / 25
Unique transactions       : 5 / 5
Duplicates suppressed     : 20 / 20
```

Figure 6.9: Serial output recorded during the duplicate-suppression validation. The first packet of each transaction is accepted, whereas subsequent identical copies are identified and suppressed.

A total of 25 physical packets were received by the Watcher. Of these, five were correctly identified as unique transactions and the remaining 20 were suppressed as duplicates.

This experiment validates the recent-history mechanism itself: once a transaction has been accepted and stored, subsequent copies carrying the same transaction information are rejected before being processed

or forwarded. The test also showed that changing `utcms` while keeping `seq` unchanged allows a new transaction to be distinguished from the preceding one.

The experiment was, however, performed with the prototype `fromId–seq–utcms` comparison and does not constitute a direct experimental validation of the simplified `fromId–utcms` identifier proposed afterwards. In the revised protocol, each new transaction generated by a given source must be assigned a new `utcms` value. Under this condition, `seq` becomes redundant, while the 15-s recent-history lifetime remains sufficient to reject additional copies arriving through the direct and gateway-assisted paths.

Removing `seq` therefore does not invalidate the demonstrated ability of the history mechanism to suppress repeated copies, but the simplified transaction identifier itself should be regarded as a protocol revision rather than as an independently validated result.

6.5 Implementation cost and scalability

The implementation cost of the redesigned Watcher was estimated from the bill of materials used for the final prototype. The purpose of this estimate is to assess the approximate hardware cost per field unit and to show how this cost is distributed among the main components.

Except for the OLED display, all listed components are required for each Watcher. The OLED is only needed for configuration and diagnostics and could therefore be shared between several units.

Table 6.17 summarises the resulting estimate. All values are expressed excluding VAT, for August 2026. Prices initially available including VAT were converted using a 21% VAT rate. Shipping and assembly labour are not included.

Table 6.17: Estimated hardware cost of one redesigned Watcher. The 100-unit estimate remains conservative because quantity discounts on electronic components are not included.

Item	**Prototype basis (€/Watcher, excl. VAT)**	**100-unit basis (€/Watcher, excl. VAT)**
Main-PCB Mouser components	53.52	53.52[a]
Adafruit Feather M0 RFM96	47.04	47.04[a]
GPS FeatherWing	29.52	29.52[a]
Adalogger FeatherWing	21.46	21.46[a]
Two ADS1115 modules	4.96	4.96[a]
SD card	8.26	8.26
Bare PCB	2.02	1.13
OLED display	12.86	0.13[b]
Electronics excluding battery	**179.64**	**166.02**
12.8-V, 6-Ah $LiFePO_4$ battery	24.79–165.29	24.79–165.29
Estimated hardware total	**204.43–344.93**	**190.81–331.31**

[a] Based on the prototype purchase price; quantity discounts available for a larger order are not included.
[b] One OLED shared between 100 Watchers.
Headers, enclosure, antenna, external wiring, shipping and assembly labour are not included.

The battery is responsible for most of the uncertainty in the final hardware cost because equivalent-capacity $LiFePO_4$ batteries are available from suppliers covering a very wide price range. A 12.8 V, 6 Ah battery stores a nominal

$$E_{\text{bat}} = 12.8\ \text{V} \times 6\ \text{Ah} = 76.8\ \text{Wh}, \tag{6.33}$$

which exceeds the approximately 50.3 Wh estimated for the five-day reference deployment in Section 6.1.4. This provides a useful nominal energy margin, although the effective available capacity will depend on temperature, ageing and discharge conditions.

At the 100-unit scale, the present conservative estimate gives a total hardware cost of approximately 191–331 € per Watcher, excluding VAT, shipping, assembly labour and the remaining mechanical accessories. For the remainder of this work, a representative cost of approximately 250 € per Watcher is used, corresponding to a configuration equipped with a mid-range battery.

Assembly time represents another important scalability consideration. The prototype was assembled manually, which remains practical for a small number of units but becomes increasingly restrictive for a network containing approximately 100 Watchers. Partial or fully automated assembly would therefore become attractive for larger production volumes.

Finally, the current modular architecture relies on several daughter boards and removable modules. This approach was useful during prototype development because individual subsystems could be replaced or tested independently. A future production-oriented revision could integrate some of these functions directly onto the main PCB, reducing the number of headers, connectors and manual assembly operations. Such an integration could further reduce both cost and assembly time, but would require an additional hardware design and validation cycle.

6.6 Validation summary and remaining limitations

The experimental results can be compared directly with the design requirements established in Section 2.8. Table 6.18 summarises the degree to which each requirement was validated during the present work.

The laboratory validation therefore confirms the main functional improvements introduced in the redesigned Watcher, particularly the reduction in power consumption, reliable local storage, analogue-channel behaviour and gateway-assisted communication. The principal remaining limitations concern four-channel acquisition at 100 Hz and the absence of a representative large-scale field campaign. Consequently, long-range radio performance, optimum LoRa parameters and complete multi-day operation with a large number of deployed receivers remain to be validated under actual deep-ERT field conditions.

Table 6.18: Validation of the main Watcher design requirements defined in Table 2.2.

Requirement	**Experimental outcome**	**Status**
Continuous acquisition and flexibility	Two-channel acquisition was sustained on a nominal 100-Hz grid during the one-hour logging test, with 99.964 Hz effectively stored. Four-channel sequential acquisition required 10.709 ms and therefore cannot presently sustain 100 Hz.	Partially validated
Measurement integrity	The four analogue channels exhibited gain errors below 0.07% and input-referred electronic noise of approximately 0.31 mV. The low-pass filter exhibited the expected first-order response, with a fitted cutoff frequency of approximately 35.5 Hz. Contact-resistance measurements were repeatable in the approximately 1–10 kΩ range.	Partially validated
Time synchronisation	GNSS/PPS synchronisation and RTC-based local timekeeping were reused as-is in the redesigned firmware, but no dedicated end-to-end timing-accuracy experiment was performed in the present validation campaign.	Not experimentally validated
Reliable local storage	One hour of continuous logging produced no missing, duplicated or out-of-order stored records. Abrupt power-loss tests showed a maximum confirmed loss of 26 records (0.26 s), without exceeding the expected 32-record (0.32 s) maximum loss associated with one partially buffered 512-byte sector.	Validated
Low average consumption	Two-channel recording current decreased from 174.6 mA to 77.6 mA, while state-based power management allows the analogue domain to be disabled between measurements. The five-day requirement was evaluated from the measured currents rather than through a complete five-day battery test.	Validated / estimated
Robust communication	Direct Watcher–gateway interoperability reached 100% PRR in both directions. The gateway restored a selected obstructed link from 0% to 100% end-to-end PRR, and duplicate suppression was successfully validated. Kilometre-scale field operation was not tested.	Partially validated
Scalable architecture	The redesigned PCB, four-channel architecture, state-based firmware and intermediate gateway provide a reproducible basis for a larger network. The estimated hardware cost is approximately 250 € per Watcher including the battery, but excluding wiring, enclosure, shipping and assembly labour. A complete 100-node deployment was not performed.	Partially validated

Chapter 7

Conclusion and future work

This work investigated the redesign of the OhmPi Watcher as an open-source autonomous receiver for large-scale deep electrical resistivity tomography. The initial prototype had already demonstrated the feasibility of distributed full-wave acquisition, but its power consumption, hardware integration and wireless communication architecture limited its suitability for large-scale deployment.

The objective was therefore to improve the existing platform while preserving its main acquisition functions. The developments focused on four complementary aspects: reducing power consumption and improving hardware integration, extending and protecting the analogue acquisition electronics, reorganising the embedded software, and improving wireless communication through multi-hop forwarding and intermediate portable gateways.

7.1 Main outcomes

At the hardware level, the redesigned Watcher integrates the revised power, acquisition, switching and protection circuitry on a single 100×160 mm PCB. The acquisition electronics were extended from two to four independently conditioned electrode channels, while remaining compatible with the original two-channel configuration. The analogue front end was also redesigned to improve compactness and energy efficiency, while adding analogue filtering, input protection and revised electrode-switching circuitry.

The most significant improvement concerned power consumption. The original oversized converters were replaced by devices better matched to the actual loads, while the bipolar analogue supply was redesigned with a high-side switch, reducing its off-state consumption to negligible leakage currents. During two-channel recording, the measured current decreased from approximately 174.6 mA to 77.6 mA, corresponding to a reduction of approximately 56%. For the reference daily cycle of 6 h recording, 6 h idle operation and 12 h sleep, the resulting five-day consumption is estimated at approximately 4.16 Ah, excluding the smaller contributions of transmissions and state transitions.

The embedded firmware was reorganised around an explicit finite-state machine separating idle, recording, transition and sleep operations. This structure also manages the SPI bus shared by the SD card and LoRa radio and allows the RTC to wake the Watcher directly from deep sleep. The existing storage strategy based on fixed-size binary records, SD-card preallocation, sector buffering and deferred file synchronisation was retained. During a one-hour two-channel test, an effective sampling rate of 99.964 Hz was obtained without missing, duplicated or out-of-order stored records. Abrupt power-loss tests showed a maximum confirmed loss of 26 records, or 0.26 s, without exceeding the expected 32-record interval associated with one partially buffered SD-card sector.

The four analogue channels were also experimentally characterised. The input-referred short-term electronic-noise standard deviation was approximately 0.31 mV. The added low-pass filter exhibited the expected

first-order response with a fitted cut-off frequency of approximately 35.5 Hz, while the revised contact-resistance circuit provided repeatable measurements over the tested range. Four-channel sequential acquisition nevertheless required approximately 10.709 ms, meaning that the present implementation cannot yet sustain 100 Hz when all four channels are used simultaneously.

Wireless communication was improved through the use of a new antenna and the development of an autonomous STM32/SX1268-based gateway intended to bypass locally obstructed links. Interoperability between the Watcher and gateway radios was demonstrated experimentally. In the obstructed-link experiment, insertion of the gateway restored end-to-end packet reception from 0% to 100% PRR for the 100 transmitted packets. Duplicate suppression was also successfully validated, preventing repeated copies of the same transaction from being processed multiple times.

Overall, the work resulted in a substantially revised Watcher platform combining lower power consumption, a more integrated four-channel hardware architecture, structured autonomous firmware, controlled local storage and gateway-assisted communication. The estimated hardware cost is approximately 250 € per Watcher including the battery, supporting the objective of replication across larger networks. Several limitations nevertheless remain before the system can be considered fully validated for large-scale deep-ERT field operation.

7.2 Limitations and future improvements

Although the redesigned Watcher improves several aspects of the original prototype, the system has not yet been validated during a complete large-scale deep-ERT field campaign. Most experiments in this work were performed in a laboratory or under controlled radio conditions. The next major step should therefore be to test the complete system under representative field conditions, including actual deep-ERT signals, environmental noise, kilometre-scale radio links and multi-day operation. A comparison with a commercial receiver under the same injection sequence would also provide a useful reference for determining the minimum signal amplitude that can be recovered reliably.

Direct participation in a deep-ERT field campaign before the main design phase would also have provided useful information for several design choices. In particular, the required analogue dynamic range and bandwidth, as well as the radio conditions along a real profile, could have been assessed more directly. The present design was therefore based mainly on literature, expected operating conditions and laboratory characterisation. Future analogue and radio optimisation should be based as much as possible on representative field measurements.

The analogue acquisition chain is one area where this information is particularly important. The present fixed input range provides sufficient margin for electrode offsets and spontaneous potentials, but makes less efficient use of the ADC range when the useful deep-ERT signal is much smaller. In addition, the electronic-noise measurements presented in Section 6.3.1 do not demonstrate that the weakest expected signals can be recovered. A future revision could therefore use a programmable measurement range, for example by exploiting the differential inputs and programmable gain of the ADS1115 or by introducing programmable analogue gain. If another ADC is considered, its input-referred noise and effective resolution at the required sampling rate should be prioritised over nominal bit depth. The AD4111 discussed in Section 3.3.4 is one possible candidate. Absolute gain calibration should also be repeated with more accurate reference instrumentation.

The analogue filtering could also be improved. The present first-order low-pass filter provides only moderate attenuation near and above the Nyquist frequency. As discussed in Section 3.3.2, the unused fourth amplifier of each OPA4188 could be used to implement a second-order active low-pass filter. Its cut-off frequency should be selected from representative field recordings so that interference is reduced without removing useful transient or induced-polarisation information.

Four-channel acquisition also requires further development. Although the four analogue channels are operational, the current sequential acquisition does not sustain 100 Hz when all four channels are used. The software and communication overhead of the ADC sequence should first be reduced, since the measured acquisition time is significantly longer than the nominal conversion time alone (Section 6.2.1). If this is insufficient, the channels could be distributed across multiple ADCs or acquired using parallel or simultaneous conversion. The binary record format described in Section 4.2.3 would also have to be extended for four-channel storage. The relative sampling time of each channel should then be represented explicitly or corrected during post-processing.

Time synchronisation is another point that still requires dedicated validation. The GNSS/PPS mechanism was inherited from the original Watcher and was not modified in this work, while the RTC was mainly extended for low-power scheduling and wake-up. A future multi-Watcher test should therefore compare several receivers against a common timing reference and the current-injection sequence. This would allow absolute timestamp accuracy, receiver-to-receiver synchronisation and long-term drift to be quantified (Section 4.3).

Further reductions in power consumption are also possible. In the Watcher, the residual sleep current mainly comes from circuitry that remains powered rather than from the sleeping microcontroller. A deeper component-level power analysis could identify which remaining loads can be switched off (Section 6.1.2). The estimated five-day autonomy should also be confirmed with a complete battery-powered test under a representative operating schedule. For the gateway, the main limitation is the continuous consumption of the E22 evaluation board. A dedicated gateway PCB and low-power STM32 operation while keeping the SX1268 available for reception could significantly reduce the required battery capacity (Section 6.1.5).

Finally, the wireless architecture should be tested at larger scale. The gateway successfully restored an obstructed link during controlled experiments, but kilometre-scale propagation and networks containing tens or around one hundred Watchers remain to be evaluated. These tests should also be used to select a regulatory-compliant LoRa configuration from measured field link margins. Once the basic communication architecture has been validated, additional protocol functions could be introduced. These include the `PING`/`PONG` link diagnostic described in Section 5.5.4, support for a mobile remote controller (Section 5.7.1), and, in the longer term, topology-aware routing and network supervision (Sections 5.7.2 and 5.7.3).

Taken together, these improvements provide a clear path from the present prototype towards a field-ready receiver. Table 7.1 summarises the main development directions and points to the related discussions in this report. For now, the priority is to validate the current architecture under representative deep-ERT conditions and use the resulting data to decide which analogue, timing, power and communication improvements are most useful.

Table 7.1: Main directions identified for future development of the OhmPi Watcher platform.

Development	**Main objective**	**Detailed discussion**
Deep-ERT field validation	Validate measurement sensitivity, timing, autonomy and communication under representative operating conditions	Sections 6.3, 6.1 and 6.4
Analogue dynamic range	Improve sensitivity to weak signals through adaptive gain or ADC range selection and a lower-noise conversion architecture	Sections 6.3.1 and 3.3.4
Anti-alias filtering	Replace or complement the present first-order filter with a higher-order response selected from representative field data	Section 3.3.2
Four-channel acquisition	Sustain 100 Hz on four channels and explicitly account for inter-channel acquisition timing	Sections 6.2.1 and 4.2.3
Power optimisation	Reduce residual Watcher sleep consumption and gateway continuous-RX overhead	Sections 6.1.2 and 6.1.5
Network diagnostics	Automatically identify weak or interrupted communication links	Section 5.5.4
Network routing and supervision	Support a mobile remote controller and eventually introduce topology-aware routing and global network supervision	Sections 5.7.1–5.7.3

Bibliography

[1] *ohmpi / ohmpi-watcher · GitLab*. en. Jan. 2026. URL: https://gitlab.com/ohmpi/ohmpi-watcher (visited on 13/08/2026).

[2] Marianna Balasco et al. 'Deep Electrical Resistivity Tomography for Geophysical Investigations: The State of the Art and Future Directions'. en. In: *Geosciences* 12.12 (Nov. 2022), p. 438. ISSN: 2076-3263. DOI: 10.3390/geosciences12120438. URL: https://www.mdpi.com/2076-3263/12/12/438 (visited on 25/07/2026).

[3] A Carrier et al. 'Deep electrical resistivity tomography for the prospection of low- to medium-enthalpy geothermal resources'. en. In: *Geophysical Journal International* 219.3 (Dec. 2019), pp. 2056–2072. ISSN: 0956-540X, 1365-246X. DOI: 10.1093/gji/ggz411. URL: https://academic.oup.com/gji/article/219/3/2056/5570585 (visited on 26/07/2026).

[4] *Einstein Telescope*. en. URL: https://einsteintelescope.eu/ (visited on 31/07/2026).

[5] *V-FullWaver*. URL: https://www.iris-instruments.com/v-fullwaver.html (visited on 29/07/2026).

[6] Rémi Clement et al. 'OhmPi: An open source data logger for dedicated applications of electrical resistivity imaging at the small and laboratory scale'. en. In: *HardwareX* 8 (Oct. 2020), e00122. ISSN: 24680672. DOI: 10.1016/j.ohx.2020.e00122. URL: https://linkinghub.elsevier.com/retrieve/pii/S2468067220300316 (visited on 25/07/2026).

[7] Andrew Binley and Andreas Kemna. 'DC Resistivity and Induced Polarization Methods'. en. In: *Hydrogeophysics*. Ed. by Yoram Rubin and Susan S. Hubbard. Vol. 50. Series Title: Water Science and Technology Library. Dordrecht: Springer Netherlands, 2005, pp. 129–156. ISBN: 978-1-4020-3101-4. DOI: 10.1007/1-4020-3102-5_5. URL: http://link.springer.com/10.1007/1-4020-3102-5_5 (visited on 25/07/2026).

[8] William Daily et al. 'Electrical resistance tomography'. en. In: *The Leading Edge* 23.5 (May 2004), pp. 438–442. ISSN: 1070-485X. DOI: 10.1190/1.1729225. URL: https://pubs.geoscienceworld.org/tle/article/23/5/438/59774/Electrical-resistance-tomography (visited on 03/08/2026).

[9] Stanley H. Ward, ed. *Geotechnical and Environmental Geophysics: Volume I: Review and Tutorial*. en. Society of Exploration Geophysicists, Jan. 1990. ISBN: 978-1-56080-278-5. DOI: 10.1190/1.9781560802785. URL: https://library.seg.org/doi/book/10.1190/1.9781560802785 (visited on 29/07/2026).

[10] *ABEM Terrameter LS 2*. en-US. URL: https://guidelinegeo.com/products/abem/abem-terrameter-ls-2/ (visited on 27/07/2026).

[11] *Surveys — GPG 0.0.1 documentation*. URL: https://gpg.geosci.xyz/content/DC_resistivity/DC_surveys.html (visited on 26/07/2026).

[12] Torleif Dahlin and Bing Zhou. 'A numerical comparison of 2D resistivity imaging with 10 electrode arrays'. en. In: *Geophysical Prospecting* 52.5 (Sept. 2004), pp. 379–398. ISSN: 0016-8025, 1365-2478. DOI: 10.1111/j.1365-2478.2004.00423.x. URL: https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2478.2004.00423.x (visited on 27/07/2026).

[13] Mariane Peter-Borie et al. 'Electrical resistivity monitoring with buried electrodes and cables: noise estimation with repeatability tests'. en. In: *Near Surface Geophysics* 9.4 (Aug. 2011), pp. 369–380.

ISSN: 1569-4445, 1873-0604. DOI: 10.3997/1873-0604.2011013. URL: https://onlinelibrary.wiley.com/doi/10.3997/1873-0604.2011013 (visited on 27/07/2026).

[14] M.H. Loke and R.D. Barker. 'Rapid least-squares inversion of apparent resistivity pseudosections by a quasi-Newton method[1]'. en. In: *Geophysical Prospecting* 44.1 (Jan. 1996), pp. 131–152. ISSN: 0016-8025, 1365-2478. DOI: 10.1111/j.1365-2478.1996.tb00142.x. URL: https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2478.1996.tb00142.x (visited on 27/07/2026).

[15] F. Nguyen et al. 'Characterization of seawater intrusion using 2D electrical imaging'. en. In: *Near Surface Geophysics* 7.5-6 (Oct. 2009), pp. 377–390. ISSN: 1569-4445, 1873-0604. DOI: 10.3997/1873-0604.2009025. URL: https://onlinelibrary.wiley.com/doi/10.3997/1873-0604.2009025 (visited on 29/07/2026).

[16] Van Wijk Kasper et al. 'Data and model uncertainty estimation for linear inversion'. en. In: *Geophysical Journal International* 149.3 (June 2002), pp. 625–632. ISSN: 1365-246X, 0956-540X. DOI: 10.1046/j.1365-246X.2002.01660.x. URL: https://academic.oup.com/gji/article/149/3/625/592733 (visited on 27/07/2026).

[17] Zhengyong Ren and Thomas Kalscheuer. 'Uncertainty and Resolution Analysis of 2D and 3D Inversion Models Computed from Geophysical Electromagnetic Data'. en. In: *Surveys in Geophysics* 41.1 (Jan. 2020), pp. 47–112. ISSN: 0169-3298, 1573-0956. DOI: 10.1007/s10712-019-09567-3. URL: http://link.springer.com/10.1007/s10712-019-09567-3 (visited on 27/07/2026).

[18] Adrian White et al. 'Assessing the effect of offline topography on electrical resistivity measurements: insights from flood embankments'. en. In: *Geophysical Journal International* 239.2 (Sept. 2024), pp. 1117–1132. ISSN: 0956-540X, 1365-246X. DOI: 10.1093/gji/ggae313. URL: https://academic.oup.com/gji/article/239/2/1117/7750635 (visited on 27/07/2026).

[19] Carsten Rücker, Thomas Günther and Florian M. Wagner. 'pyGIMLi: An open-source library for modelling and inversion in geophysics'. In: *Computers & Geosciences* 109 (Dec. 2017), pp. 106–123. ISSN: 0098-3004. DOI: 10.1016/j.cageo.2017.07.011. URL: https://www.sciencedirect.com/science/article/pii/S0098300417300584 (visited on 27/07/2026).

[20] Enzo Rizzo and Valeria Giampaolo. 'New deep electrical resistivity tomography in the High Agri Valley basin (Basilicata, Southern Italy)'. en. In: *Geomatics, Natural Hazards and Risk* 10.1 (Jan. 2019), pp. 197–218. ISSN: 1947-5705, 1947-5713. DOI: 10.1080/19475705.2018.1520150. URL: https://www.tandfonline.com/doi/full/10.1080/19475705.2018.1520150 (visited on 25/07/2026).

[21] University of Liège. *Géophysique appliquée*. fr. URL: https://www.geologie.uliege.be/cms/c_5911077/fr/geologie-geophysique-appliquee (visited on 26/07/2026).

[22] Douglas J. LaBrecque et al. 'Autonomous Monitoring of Fluid Movement Using 3-D Electrical Resistivity Tomography'. en. In: *Journal of Environmental & Engineering Geophysics* 9.3 (Sept. 2004), pp. 167–176. ISSN: 1083-1363, 1943-2658. DOI: 10.4133/JEEG9.3.167. URL: https://pubs.geoscienceworld.org/jeeg/article/9/3/167/655978/Autonomous-Monitoring-of-Fluid-Movement-Using-3-D (visited on 29/07/2026).

[23] Oliver Kuras et al. 'Monitoring hydraulic processes with automated time-lapse electrical resistivity tomography (ALERT)'. en. In: *Comptes Rendus. Géoscience* 341.10-11 (Oct. 2009), pp. 868–885. ISSN: 1631-0713, 1778-7025. DOI: 10.1016/j.crte.2009.07.010. URL: https://comptes-rendus.academie-sciences.fr/geoscience/articles/10.1016/j.crte.2009.07.010/ (visited on 29/07/2026).

[24] Adrien Dimech et al. 'A Review on Applications of Time-Lapse Electrical Resistivity Tomography Over the Last 30 Years : Perspectives for Mining Waste Monitoring'. en. In: *Surveys in Geophysics* 43.6 (Dec. 2022), pp. 1699–1759. ISSN: 1573-0956. DOI: 10.1007/s10712-022-09731-2. URL: https://doi.org/10.1007/s10712-022-09731-2 (visited on 29/07/2026).

[25] M. Farzamian et al. 'Advancing Permafrost Monitoring With Autonomous Electrical Resistivity Tomography (A-ERT): Low-Cost Instrumentation and Open-Source Data Processing Tool'. en. In: *Geophysical Research Letters* 51.6 (Mar. 2024), e2023GL105770. ISSN: 0094-8276, 1944-8007.

DOI: 10.1029/2023GL105770. URL: https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023GL105770 (visited on 25/07/2026).

[26] Rémi Clement et al. 'Design and advances in OhmPi: an open hardware resistivity meter'. en. In: *HardwareX* 27 (Sept. 2026), e00811. ISSN: 24680672. DOI: 10.1016/j.ohx.2026.e00811. URL: https://linkinghub.elsevier.com/retrieve/pii/S2468067226000714 (visited on 31/07/2026).

[27] I.F. Akyildiz et al. 'Wireless sensor networks: a survey'. en. In: *Computer Networks* 38.4 (Mar. 2002), pp. 393–422. ISSN: 13891286. DOI: 10.1016/S1389-1286(01)00302-4. URL: https://linkinghub.elsevier.com/retrieve/pii/S1389128601003024 (visited on 07/08/2026).

[28] V. Raghunathan et al. 'Energy-aware wireless microsensor networks'. en. In: *IEEE Signal Processing Magazine* 19.2 (Mar. 2002), pp. 40–50. ISSN: 10535888. DOI: 10.1109/79.985679. URL: http://ieeexplore.ieee.org/document/985679/ (visited on 02/08/2026).

[29] Tifenn Rault, Abdelmadjid Bouabdallah and Yacine Challal. 'Energy efficiency in wireless sensor networks: A top-down survey'. en. In: *Computer Networks* 67 (July 2014), pp. 104–122. ISSN: 13891286. DOI: 10.1016/j.comnet.2014.03.027. URL: https://linkinghub.elsevier.com/retrieve/pii/S1389128614001418 (visited on 02/08/2026).

[30] Giuseppe Anastasi et al. 'Energy conservation in wireless sensor networks: A survey'. en. In: *Ad Hoc Networks* 7.3 (May 2009), pp. 537–568. ISSN: 15708705. DOI: 10.1016/j.adhoc.2008.06.003. URL: https://linkinghub.elsevier.com/retrieve/pii/S1570870508000954 (visited on 02/08/2026).

[31] Ricardo C. Carrano et al. 'Survey and Taxonomy of Duty Cycling Mechanisms in Wireless Sensor Networks'. en. In: *IEEE Communications Surveys & Tutorials* 16.1 (2014), pp. 181–194. ISSN: 1553-877X. DOI: 10.1109/SURV.2013.052213.00116. URL: http://ieeexplore.ieee.org/document/6524461/ (visited on 02/08/2026).

[32] Traco Power. *Traco THD15N Datasheet*. URL: https://www.tracopower.com/sites/default/files/products/datasheets/thd15n_datasheet.pdf?t=1785711601 (visited on 03/08/2026).

[33] Traco Power. *Traco THD12 Datasheet*. URL: https://www.tracopower.com/sites/default/files/products/datasheets/thd12_datasheet.pdf?t=1785711601 (visited on 03/08/2026).

[34] V. Raghunathan et al. 'Design considerations for solar energy harvesting wireless embedded systems'. en. In: *IPSN 2005. Fourth International Symposium on Information Processing in Sensor Networks, 2005.* Los Angeles, CA, USA: IEEE, 2005, pp. 457–462. ISBN: 978-0-7803-9201-4. DOI: 10.1109/IPSN.2005.1440973. URL: http://ieeexplore.ieee.org/document/1440973/ (visited on 03/08/2026).

[35] R Torah et al. 'Self-powered autonomous wireless sensor node using vibration energy harvesting'. en. In: *Measurement Science and Technology* 19.12 (Dec. 2008), p. 125202. ISSN: 0957-0233, 1361-6501. DOI: 10.1088/0957-0233/19/12/125202. URL: https://iopscience.iop.org/article/10.1088/0957-0233/19/12/125202 (visited on 02/08/2026).

[36] Jennifer Yick, Biswanath Mukherjee and Dipak Ghosal. 'Wireless sensor network survey'. en. In: *Computer Networks* 52.12 (Aug. 2008), pp. 2292–2330. ISSN: 13891286. DOI: 10.1016/j.comnet.2008.04.002. URL: https://linkinghub.elsevier.com/retrieve/pii/S1389128608001254 (visited on 02/08/2026).

[37] Bharath Sundararaman, Ugo Buy and Ajay D. Kshemkalyani. 'Clock synchronization for wireless sensor networks: a survey'. en. In: *Ad Hoc Networks* 3.3 (May 2005), pp. 281–323. ISSN: 15708705. DOI: 10.1016/j.adhoc.2005.01.002. URL: https://linkinghub.elsevier.com/retrieve/pii/S1570870505000144 (visited on 02/08/2026).

[38] Ki Young Koo, David Hester and Sehoon Kim. 'Time Synchronization for Wireless Sensors Using Low-Cost GPS Module and Arduino'. en. In: *Frontiers in Built Environment* 4 (Jan. 2019), p. 82. ISSN: 2297-3362. DOI: 10.3389/fbuil.2018.00082. URL: https://www.frontiersin.org/article/10.3389/fbuil.2018.00082/full (visited on 02/08/2026).

[39] Danniel R. Fuhrmann et al. 'Node Synchronization in a Wireless Sensor Network Using Unreliable GPS Signals'. en. In: *2014 IEEE Military Communications Conference*. Baltimore, MD, USA: IEEE,

Oct. 2014, pp. 630–636. ISBN: 978-1-4799-6770-4. DOI: 10.1109/MILCOM.2014.111. URL: http://ieeexplore.ieee.org/document/6956832/ (visited on 02/08/2026).

[40] NXP Semiconductors. *PCF8523 Real-Time Clock (RTC) and calendar*. en. Datasheet. 2015.

[41] Adafruit Industries. *Adalogger FeatherWing - RTC + SD Add-on For All Feather Boards*. en-US. URL: https://www.adafruit.com/product/2922 (visited on 10/08/2026).

[42] Analog Devices. *OP07 - Datasheet*. URL: https://www.analog.com/media/en/technical-documentation/data-sheets/op07.pdf (visited on 04/08/2026).

[43] Texas Instruments. *ADS1115 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/ads1115.pdf (visited on 04/08/2026).

[44] Veepsa Bhatia, Neeta Pandey and Asok Bhattacharyya. 'APPLICATION BASED COMPARISON OF DIFFERENT ANALOG TO DIGITAL CONVERTER ARCHITECTURES'. en. In: *International Journal of Engineering Science and Technology* 2 (2010).

[45] R.H. Walden. 'Analog-to-digital converter survey and analysis'. en. In: *IEEE Journal on Selected Areas in Communications* 17.4 (Apr. 1999), pp. 539–550. ISSN: 07338716. DOI: 10.1109/49.761034. URL: http://ieeexplore.ieee.org/document/761034/ (visited on 04/08/2026).

[46] Texas Instruments. *ADS1219 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/ads1219.pdf?ts=1785864170339&ref_url=https%253A%252F%252Fwww.google.com%252F (visited on 04/08/2026).

[47] R.W. Brodersen, P.R. Gray and D.A. Hodges. 'MOS switched-capacitor filters'. In: *Proceedings of the IEEE* 67.1 (Jan. 1979), pp. 61–75. ISSN: 1558-2256. DOI: 10.1109/PROC.1979.11203. URL: https://ieeexplore.ieee.org/document/1455472/ (visited on 07/08/2026).

[48] Geoff Werner-Allen et al. 'Fidelity and yield in a volcano monitoring sensor network'. In: *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. OSDI '06. USA: USENIX Association, 2006, pp. 381–396. ISBN: 1-931971-47-1.

[49] R. Bez et al. 'Introduction to flash memory'. en. In: *Proceedings of the IEEE* 91.4 (Apr. 2003), pp. 489–502. ISSN: 0018-9219. DOI: 10.1109/JPROC.2003.811702. URL: http://ieeexplore.ieee.org/document/1199079/ (visited on 04/08/2026).

[50] Gaurav Mathur et al. 'Ultra-low power data storage for sensor networks'. en. In: *ACM Transactions on Sensor Networks* 5.4 (Nov. 2009), pp. 1–34. ISSN: 1550-4859, 1550-4867. DOI: 10.1145/1614379.1614385. URL: https://dl.acm.org/doi/10.1145/1614379.1614385 (visited on 04/08/2026).

[51] Hung-Wei Tseng, Laura Grupp and Steven Swanson. 'Understanding the impact of power loss on flash memory'. en. In: *Proceedings of the 48th Design Automation Conference*. San Diego California: ACM, June 2011, pp. 35–40. ISBN: 978-1-4503-0636-2. DOI: 10.1145/2024724.2024733. URL: https://dl.acm.org/doi/10.1145/2024724.2024733 (visited on 04/08/2026).

[52] Rashmi Sharan Sinha, Yiqiao Wei and Seung-Hoon Hwang. 'A survey on LPWA technology: LoRa and NB-IoT'. en. In: *ICT Express* 3.1 (Mar. 2017), pp. 14–21. ISSN: 24059595. DOI: 10.1016/j.icte.2017.03.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S2405959517300061 (visited on 25/07/2026).

[53] Martin Bor, John Vidler and Utz Roedig. 'LoRa for the Internet of Things'. In: *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks*. EWSN '16. International Conference on Embedded Wireless Systems and Networks (EWSN), 2016, pp. 361–366. ISBN: 978-0-9949886-0-7.

[54] Siddhartha S. Borkotoky, Udo Schilcher and Christian Bettstetter. 'Cooperative Relaying in LoRa Sensor Networks'. en. In: *2019 IEEE Global Communications Conference (GLOBECOM)*. Waikoloa, HI, USA: IEEE, Dec. 2019, pp. 1–5. ISBN: 978-1-7281-0962-6. DOI: 10.1109/GLOBECOM38437.2019.9014071. URL: https://ieeexplore.ieee.org/document/9014071/ (visited on 25/07/2026).

[55] Semtech. *Semtech an1200.22 - LoRa Modulation Basics*. URL: https://www.semtech.com/uploads/documents/an1200.22.pdf.

[56] Dong-Hoon Kim, Eun-Kyu Lee and Jibum Kim. 'Experiencing LoRa Network Establishment on a Smart Energy Campus Testbed'. en. In: *Sustainability* 11.7 (Mar. 2019), p. 1917. ISSN: 2071-1050.

DOI: 10.3390/su11071917. URL: https://www.mdpi.com/2071-1050/11/7/1917 (visited on 05/08/2026).

[57] Denis Stanescu et al. 'Spread spectrum modulation recognition based on phase diagram entropy'. en. In: *Frontiers in Signal Processing* 3 (July 2023), p. 1197619. ISSN: 2673-8198. DOI: 10.3389/frsip.2023.1197619. URL: https://www.frontiersin.org/articles/10.3389/frsip.2023.1197619/full (visited on 05/08/2026).

[58] Lorenzo Vangelista. 'Frequency Shift Chirp Modulation: The LoRa Modulation'. en. In: *IEEE Signal Processing Letters* 24.12 (Dec. 2017), pp. 1818–1821. ISSN: 1070-9908, 1558-2361. DOI: 10.1109/LSP.2017.2762960. URL: http://ieeexplore.ieee.org/document/8067462/ (visited on 25/07/2026).

[59] Marco Chiani and Ahmed Elzanaty. 'On the LoRa Modulation for IoT: Waveform Properties and Spectral Analysis'. en. In: *IEEE Internet of Things Journal* 6.5 (Oct. 2019), pp. 8463–8470. ISSN: 2327-4662, 2372-2541. DOI: 10.1109/JIOT.2019.2919151. URL: https://ieeexplore.ieee.org/document/8723130/ (visited on 05/08/2026).

[60] *LoRa Technology Guide: CSS, Spreading Factor & Meshtastic*. en. URL: https://www.haraldkreuzer.net/en/news/lora-radio-technology-internet-things (visited on 07/08/2026).

[61] Adafruit Industries. *Adafruit Feather M0 RFM96 LoRa Radio - 433MHz*. en-US. URL: https://www.adafruit.com/product/3179 (visited on 10/08/2026).

[62] Adafruit Industries. *Adafruit Ultimate GPS FeatherWing*. en-US. URL: https://www.adafruit.com/product/3133 (visited on 10/08/2026).

[63] Microchip. *MCP23017 - Datasheet*. URL: https://ww1.microchip.com/downloads/aemDocuments/documents/APID/ProductDocuments/DataSheets/MCP23017-MCP23S17-16-Bit-IO-Expander-with-Serial-Interface-DS20001952.pdf (visited on 10/08/2026).

[64] Omron. *Omron G6L - Datasheet*. URL: https://components.omron.com/system/files/datasheet_pdf/K119-E1.pdf (visited on 10/08/2026).

[65] Traco Power. *Traco TSR1 - Datasheet*. URL: https://www.tracopower.com/sites/default/files/products/datasheets/tsr1_datasheet.pdf?t=1786230001 (visited on 09/08/2026).

[66] Texas Instruments. *OPA4188 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/opa4188.pdf?ts=1786282913875 (visited on 10/08/2026).

[67] Analog Devices. *DG413 - Datasheet*. URL: https://www.analog.com/media/en/technical-documentation/data-sheets/dg411-dg413.pdf (visited on 10/08/2026).

[68] Analog Devices. *REF03 - Datasheet*. URL: https://www.analog.com/media/en/technical-documentation/data-sheets/ref01_02_03.pdf (visited on 10/08/2026).

[69] Traco Power. *Traco TEC3 - Datasheet*. URL: https://www.tracopower.com/sites/default/files/products/datasheets/tec3_datasheet.pdf?t=1786230001 (visited on 09/08/2026).

[70] STMicroelectronics. *STN3P6F6 - Datasheet*. URL: https://www.st.com/resource/en/datasheet/stn3p6f6.pdf (visited on 10/08/2026).

[71] Onsemi. *2N7002 - Datasheet*. URL: https://www.onsemi.com/pdf/datasheet/nds7002a-d.pdf (visited on 10/08/2026).

[72] Texas Instruments. *ULN2003 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/uln2003a.pdf (visited on 10/08/2026).

[73] Texas Instruments. *INA282 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/ina282.pdf?ts=1786362259804&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FINA282 (visited on 10/08/2026).

[74] Eric Bogatin. *Bogatin's practical guide to prototype breadboard and PCB design*. en. Norwood, MA: Artech House, 2021. ISBN: 978-1-63081-848-7.

[75] B. P. Welford. 'Note on a Method for Calculating Corrected Sums of Squares and Products'. en. In: *Technometrics* 4.3 (Aug. 1962), pp. 419–420. ISSN: 0040-1706, 1537-2723. DOI: 10.1080/00401706.1962.10490022. URL: http://www.tandfonline.com/doi/abs/10.1080/00401706.1962.10490022 (visited on 12/08/2026).

[76] Bill Greiman. *SdFat: Arduino FAT16/FAT32 and exFAT Library*. original-date: 2014-08-06T11:56:59Z. July 2026. URL: https://github.com/greiman/SdFat (visited on 04/08/2026).

[77] *IBPT - Spécifications d'interface radio*. URL: https://www.ibpt.be/file/cc73d96153bbd5448a56f19d925d05b1379c7f21/3a32bbf04a46bed4de560b2292819f29e685b70a/25-11-b-fr.pdf (visited on 20/08/2026).

[78] *ETSI EN 300 220*. URL: https://www.etsi.org/deliver/etsi_en/300200_300299/30022001/01.03.01_20/en_30022001v010301c.pdf (visited on 20/08/2026).

[79] EBYTE. *E22-M+Series+Module_UserManual_FR*. URL: https://fr-ebyte.com/pdf-down.aspx?id=5920.

[80] EBYTE. *TX433-JKS-IPX20 - Datasheet*. URL: https://cdn.webshopapp.com/shops/304271/files/479031547/tx433-jks-ipx20-en-v332.pdf (visited on 13/08/2026).

[81] Semtech. *SX1268 - Datasheet*. URL: https://semtech.my.salesforce.com/sfc/p/#E00000000JelG/a/RQ000008nGa9/v.ea8jp.cso4PrL7koonfOfzpF4krk_NglUXEPVQRzg (visited on 14/08/2026).

[82] EBYTE. *E22-400MBL-SC - Datasheet*. URL: https://www.cdebyte.com/pdf-down.aspx?id=3314 (visited on 14/08/2026).

[83] STMicroelectronics. *STM32F103C8 - Datasheet*. URL: https://www.st.com/resource/en/datasheet/stm32f103c8.pdf (visited on 14/08/2026).

[84] Texas Instruments. *TPS63802 - Datasheet*. URL: https://www.ti.com/lit/ds/symlink/tps63802.pdf?ts=1786765834538&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FTPS63802%252Fpart-details%252FTPS63802DLAR (visited on 15/08/2026).

[85] *Meshtastic - Introduction*. en. URL: https://meshtastic.org/docs/introduction/ (visited on 19/08/2026).

[86] Uni-Trend. *UT61E+ - Datasheet*. URL: https://meters.uni-trend.com/download/ut61b-d-e-flyer/?wpdmdl=7139&refresh=6a7d98ca552771786616010 (visited on 15/08/2026).

Appendix A

Complete electrical schematic

The following pages present the complete electrical schematic of the redesigned OhmPi Watcher developed in this work.

H3
MountingHole

H1
MountingHole

H2
MountingHole

H4
MountingHole

A5
Adafruit_Feather_Generic

A1
Adafruit_Feather_Generic

A3
Adafruit_Feather_Generic

A4
Adafruit_Feather_Generic

U23
ADS1115
I²C address
0x48

U21
ADS1115
I²C address
0x49

ADS1115

U22
ADS1219
I²C address
0x40

ADS1219

U18
MCP23017x-x-SO
address= 0x27

Analog chain
J1
Screw_Terminal_01x05
File: conditioner.kicad_sch

Testpads
J3
Connector_I2C_for_MC

J2
Screw_Terminal_01x02
Power Supply
File: power_supply.kicad_sch
Current consumption :
12V / 300kOhms = 0.04mA

PWR_FLAG
PWR_FLAG

U9
PCA9306D
I2C translator 5V <-> 3.3V

Sheet: /	
File: WatcherV2.kicad_sch	
Title:	
Size: A3	Date:
KiCad E.D.A. 10.0.3	Id: 1/3

Input attendu :
de -5V à +5V

Divise le signal par 2

AO Suiveur
Pour éviter que le non-inverting
summing OpAmp influence la sortie
du diviseur par 2

Multiplie par 2

RC LP filter
Fc = 34Hz, Tau = 4,7ms
période du courant appliqué : 500ms
à modif si trop long

Clippe entre 0 et 5V

Output :
de 0V à 5V

Expected entry
range = ± 5V

Follower OpAmp + Voltage divider

Range = ± 2.5V

Follower OpAmp + Non-Inverting summing OpAmp

Range = 0-5V

Low Pass Filter

Clipping Zenner diodes

Expected output
range = 0-5V

Darlington transistor array

R_Protection_3W
I = 12 / (100 + 2) ≈ 118 mA
Pshunt = 0.118² × 2 ≈ 0.028 W
Pserie = 0.118² × 100 ≈ 1.4 W

Alimentations AO

AO inutilisés des OPA4188

Sheet: /Analog chain/	
File: conditioner.kicad_sch	
Title:	
Size: A2	Date:
KiCad E.D.A. 10.0.3	**Rev:** Id: 2/3

OPA4188 : 570µA de quiescent current / AO
soit 16*570µA = 9.12mA au total (par rail)

Les OP07 consommaient 10mA / AO at no load (MAX)
soit 12*10 = 120mA total (par rail)
Sur le graphique (TYP), c'est 5.33mA par AO, soit 12*5.33 = 64mA par rail

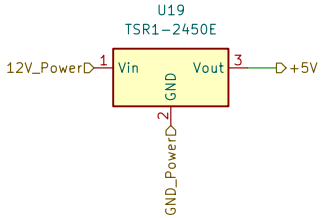

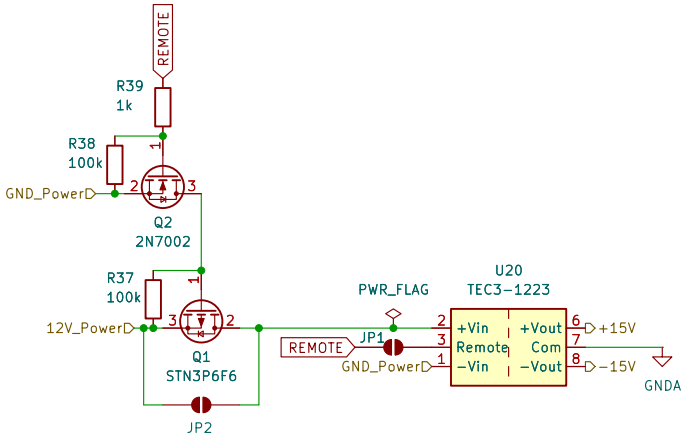

Sheet: /Power Supply/		
File: power_supply.kicad_sch		
Title:		
Size: A4	Date:	**Rev:**
KiCad E.D.A. 10.0.3		Id: 3/3

Appendix B

E22-400MBL-SC evaluation board schematic

The following pages present the complete manufacturer electrical schematic of the Ebyte E22-400MBL-SC evaluation board used as the hardware platform for the gateway developed in this work.

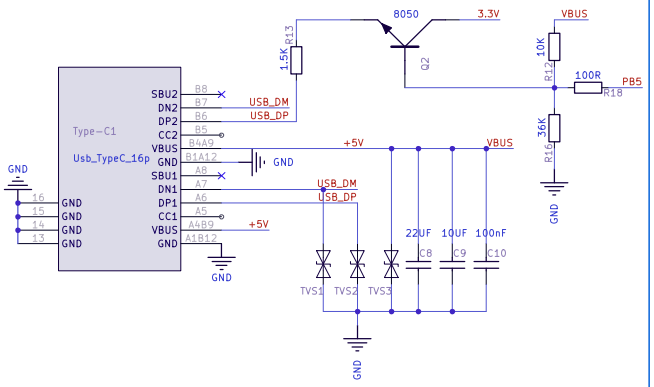

typc接口

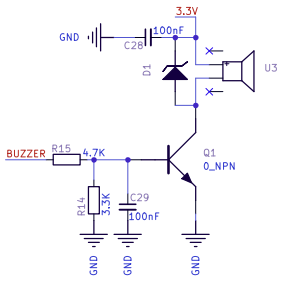

蜂鸣器

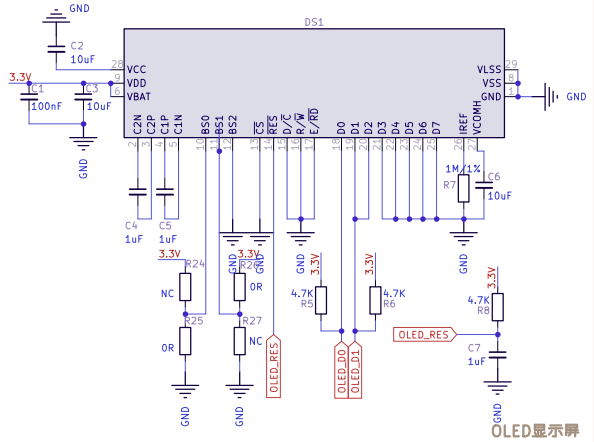

OLED显示屏

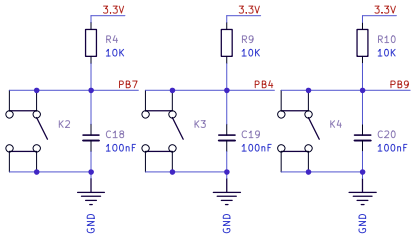

确认按键　向上按键　向下按键

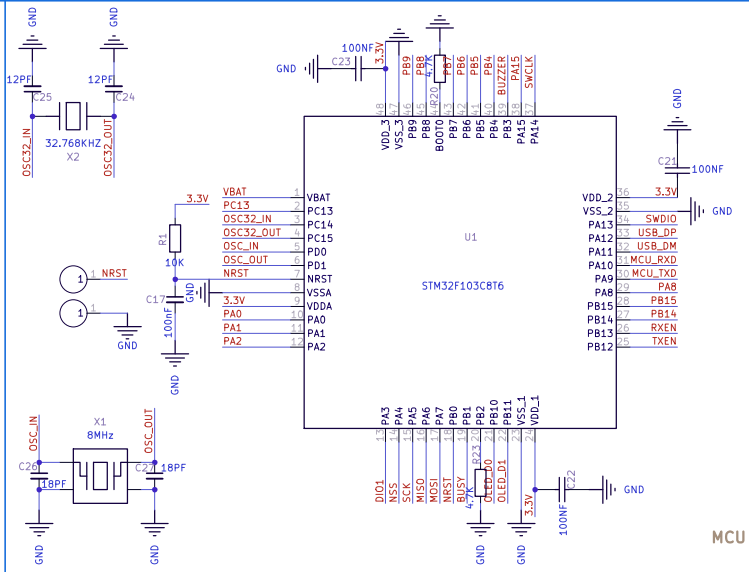

MCU

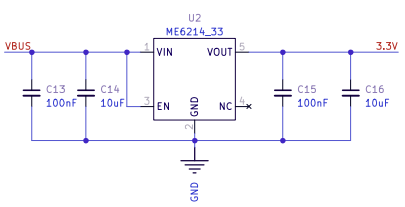

LDO

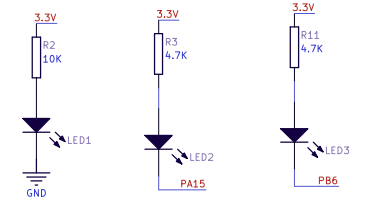

PWR指示灯　TXD显示灯　RXD显示灯

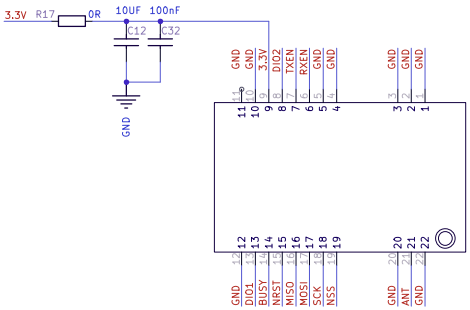

E22/E220

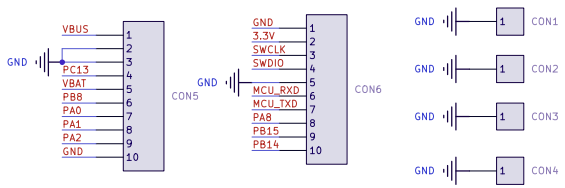

排针

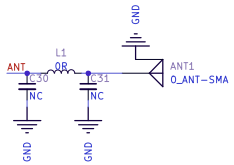

SMA头

Sheet: /
File: E22-400MBL-OLED - 副本.kicad_sch

Title:

Size: A3	Date:	Rev:
KiCad E.D.A. 8.0.4		Id: 1/1

Appendix C

Adafruit Feather M0 RFM9x pinout

The following page presents the pinout of the Adafruit Feather M0 RFM9x board used as the main microcontroller and LoRa communication platform of the OhmPi Watcher.

Adafruit Feather M0 RFM9x LoRa

https://www.adafruit.com/products/3178

The Microchip (nee Atmel) SAMD21 is an ARM Cortex-M0+ running at 48 MHz with 32kB on-chip SRAM, 256KB Flash memory and built in USB. All GPIO is 3.3V in/out max unless otherwise stated. SERCOMs can be used as UART (TX on SERCOM pad 0 or 2, RX on any pad), I2C (SDA on pad 0, SCL on pad 1), or SPI (SCK on pad 1 or 3, MOSI on pad 0 or 3 (SCK on pad 1) or pad 0 or 2 (SCK on pad 3), MISO on any pad remaining)

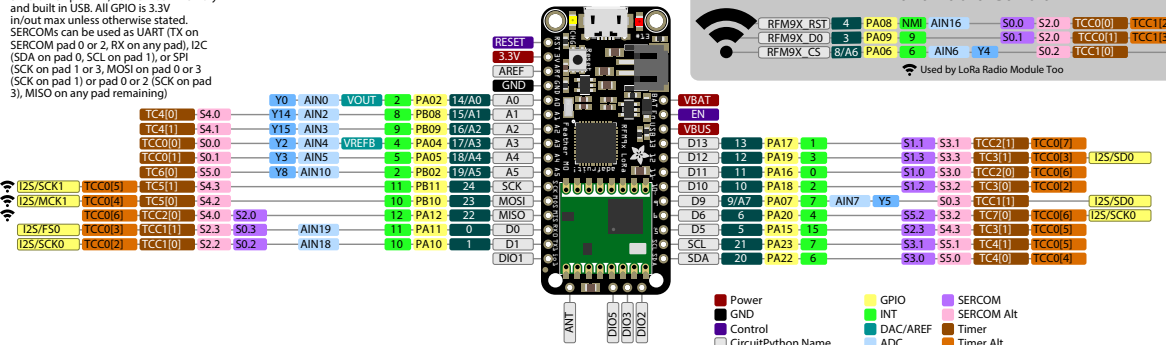

Appendix D

Detailed firmware execution diagrams

The following diagrams provide the detailed execution flows underlying the simplified representations presented in Fig. 4.1.

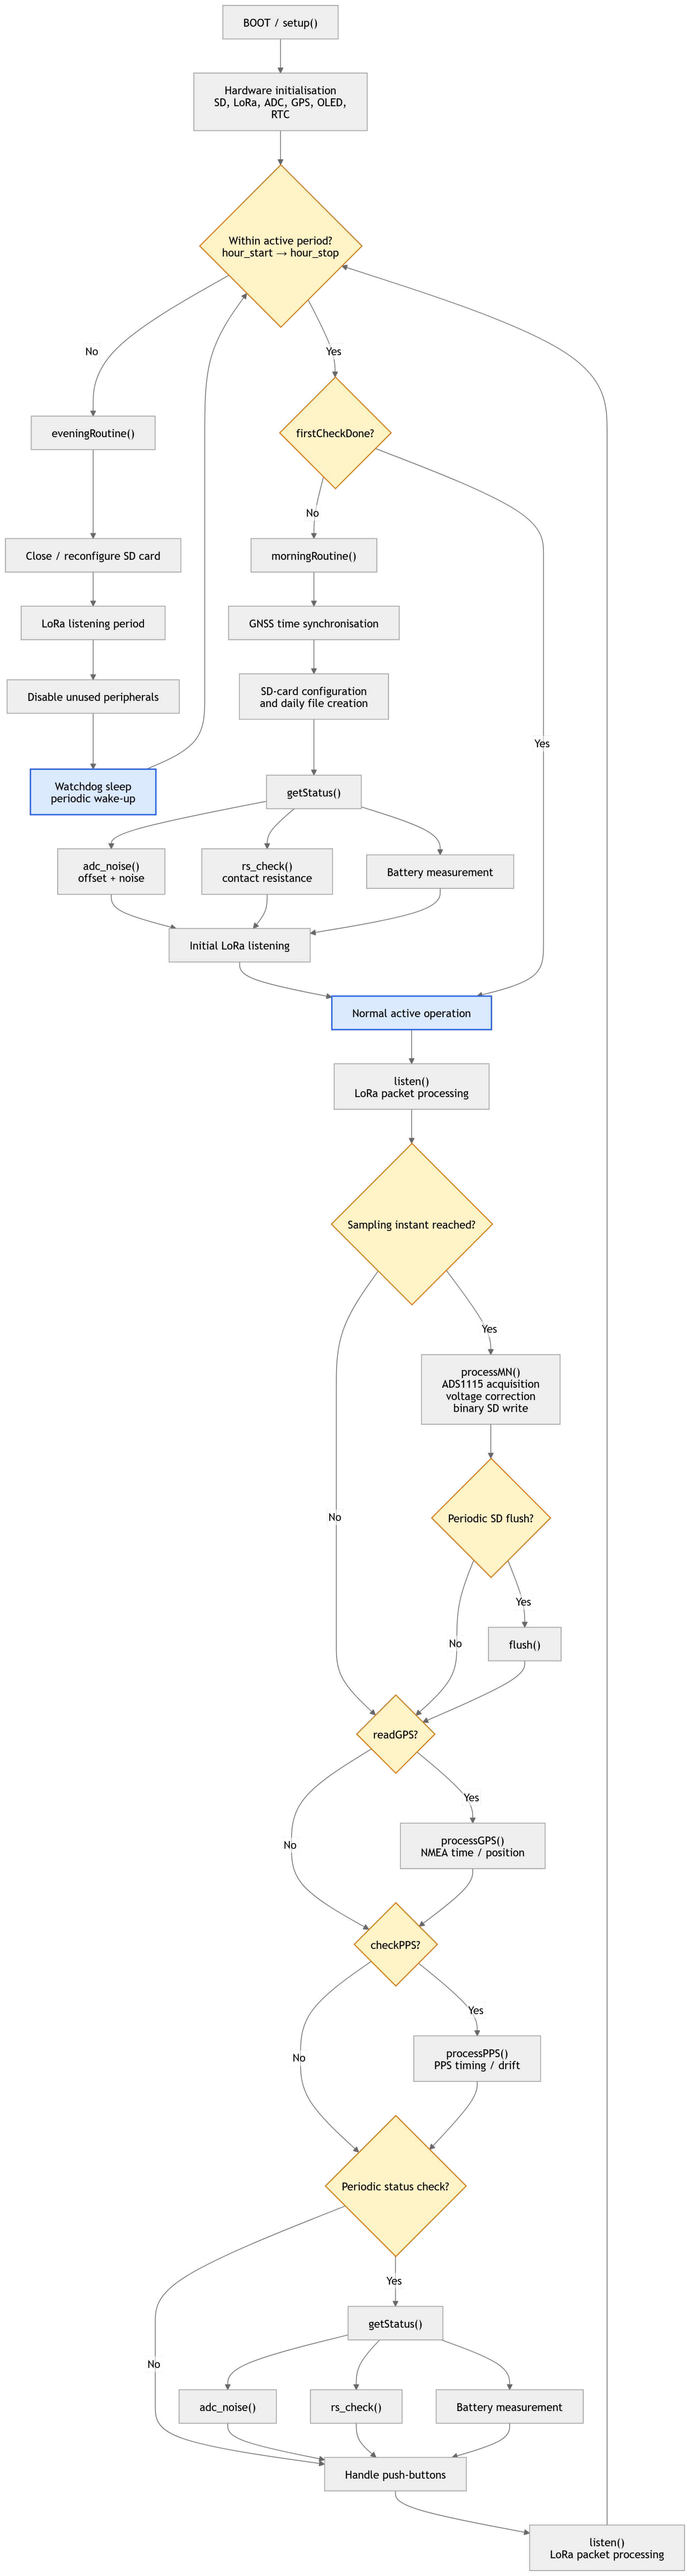
BOOT / setup()
Hardware initialisation
SD, LoRa, ADC, GPS, OLED,
RTC
Within active period?
hour_start → hour_stop
No
Yes
eveningRoutine()
firstCheckDone?
No
Yes
Close / reconfigure SD card
morningRoutine()
LoRa listening period
GNSS time synchronisation
Disable unused peripherals
SD-card configuration
and daily file creation
Watchdog sleep
periodic wake-up
getStatus()
adc_noise()
offset + noise
rs_check()
contact resistance
Battery measurement
Initial LoRa listening
Normal active operation
listen()
LoRa packet processing
Sampling instant reached?
Yes
processMN()
ADS1115 acquisition
voltage correction
binary SD write
No
Periodic SD flush?
Yes
No
flush()
readGPS?
Yes
No
processGPS()
NMEA time / position
checkPPS?
Yes
No
processPPS()
PPS timing / drift
Periodic status check?
Yes
getStatus()
No
adc_noise()
rs_check()
Battery measurement
Handle push-buttons
listen()
LoRa packet processing

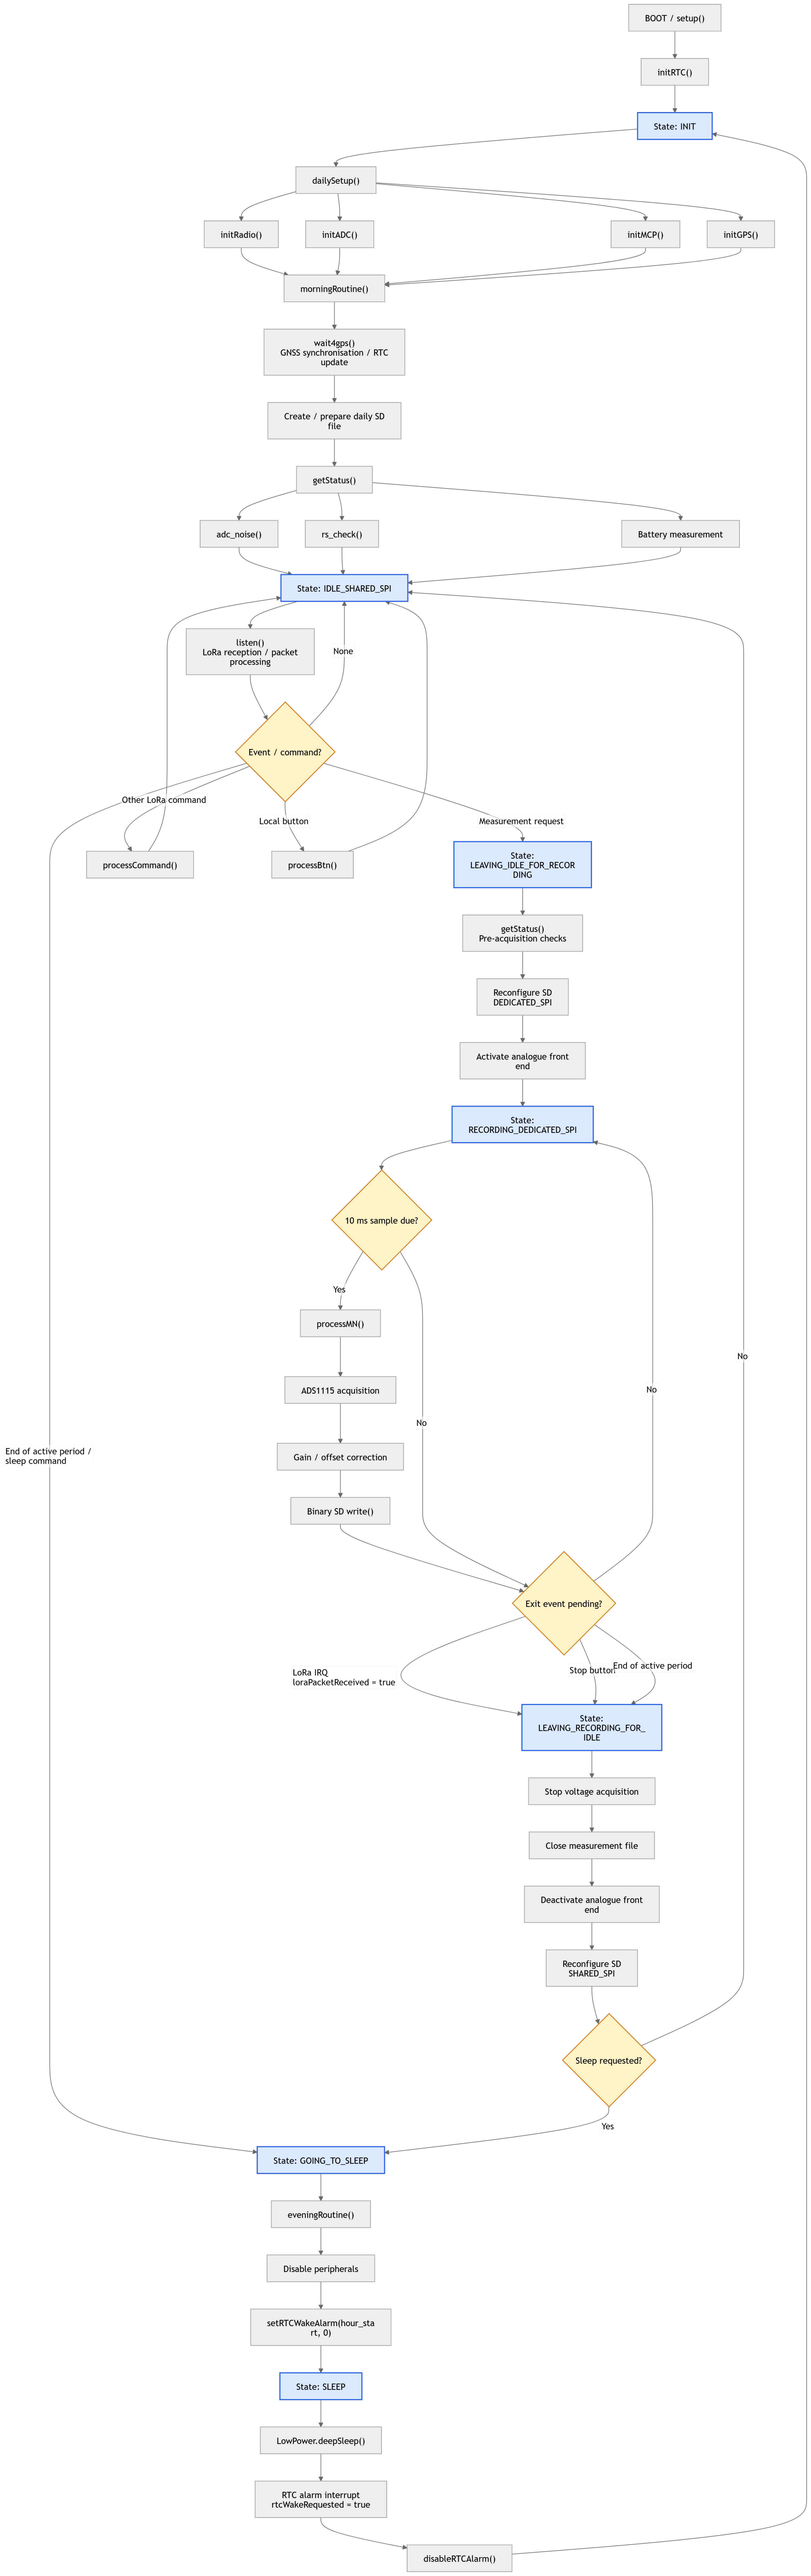
BOOT / setup()
initRTC()
State: INIT
dailySetup()
initRadio()
initADC()
initMCP()
initGPS()
morningRoutine()
wait4gps()
GNSS synchronisation / RTC update
Create / prepare daily SD file
getStatus()
adc_noise()
rs_check()
Battery measurement
State: IDLE_SHARED_SPI
listen()
LoRa reception / packet processing
None
Event / command?
Other LoRa command
Local button
Measurement request
processCommand()
processBtn()
State: LEAVING_IDLE_FOR_RECORDING
getStatus()
Pre-acquisition checks
Reconfigure SD DEDICATED_SPI
Activate analogue front end
State: RECORDING_DEDICATED_SPI
10 ms sample due?
Yes
No
processMN()
ADS1115 acquisition
Gain / offset correction
Binary SD write()
End of active period / sleep command
No
No
Exit event pending?
LoRa IRQ
loraPacketReceived = true
Stop button
End of active period
State: LEAVING_RECORDING_FOR_IDLE
Stop voltage acquisition
Close measurement file
Deactivate analogue front end
Reconfigure SD SHARED_SPI
Sleep requested?
No
Yes
State: GOING_TO_SLEEP
eveningRoutine()
Disable peripherals
setRTCWakeAlarm(hour_start, 0)
State: SLEEP
LowPower.deepSleep()
RTC alarm interrupt
rtcWakeRequested = true
disableRTCAlarm()

Design Improvements and Optimisation of the OhmPi Watcher System for Large-Scale Autonomous Deep ERT Measurements

Améliorations de conception et optimisation du système OhmPi Watcher pour des mesures ERT profondes autonomes à grande échelle

Auteur	EYEN Guillaume
Section	Master ingénieur civil électricien — Systèmes et dispositifs électroniques
Promoteurs	Prof. Jean-Michel Redouté, Prof. Frédéric Nguyen
Conseiller scientifique	Dr Guillaume Blanchy
Année académique	2025–2026

La tomographie de résistivité électrique profonde (*deep-ERT*) étend l'imagerie électrique du sous-sol à des profils kilométriques et à de grandes profondeurs d'investigation, avec des applications en exploration minière et géothermique, en cartographie de zones de faille et en caractérisation de sites destinés à de grandes infrastructures souterraines. Ces campagnes exigent de nombreuses stations réceptrices autonomes, capables d'enregistrer de faibles potentiels de sol, de rester synchronisées et accessibles à distance pendant plusieurs jours. Les récepteurs commerciaux répondant à ce besoin coûtent plusieurs milliers d'euros par unité, ce qui limite directement la couverture spatiale atteignable. Ce travail porte sur la refonte de l'OhmPi Watcher, un récepteur autonome et open-source dérivé de la plateforme OhmPi, pour un déploiement de référence de 100 stations réparties sur plus de 10 km pendant au moins cinq jours.

La refonte s'articule autour de quatre axes complémentaires. Sur le plan matériel, l'alimentation, l'acquisition, la commutation et les protections ont été intégrées sur une carte unique de 100 $\times$ 160 mm ; les convertisseurs surdimensionnés ont été remplacés, l'alimentation analogique bipolaire est coupée par un interrupteur *high-side* lorsqu'elle n'est pas nécessaire, et la chaîne d'acquisition est passée de deux à quatre voies conditionnées indépendamment, avec filtrage passe-bas, protection d'entrée et circuit de mesure de résistance de contact revu. Le micrologiciel a été réorganisé autour d'une machine à états finis explicite gérant l'enregistrement, la veille et le sommeil profond, le réveil par alarme RTC et l'arbitrage du bus SPI partagé entre la carte SD et la radio. La communication sans fil a été améliorée par une nouvelle intégration d'antenne 433 MHz, un relayage multi-saut avec suppression des doublons et le développement d'une passerelle portable autonome STM32/SX1268 permettant de contourner un lien radio localement obstrué.

La validation en laboratoire montre une réduction du courant d'enregistrement à deux voies de 174,6 mA à 77,6 mA (soit -56 %), correspondant à une consommation estimée de 4,16 Ah sur cinq jours pour un cycle journalier de 6 h d'enregistrement, 6 h de veille et 12 h de sommeil. Les quatre voies analogiques présentent un bruit électronique ramené à l'entrée de 0,31 mV et des erreurs de gain inférieures à 0,07 %, avec un filtre du premier ordre de fréquence de coupure ajustée à 35,5 Hz. Une heure d'enregistrement continu a été réalisée à une fréquence effective de 99,964 Hz sans enregistrement manquant, dupliqué ou désordonné, les coupures d'alimentation brutales entraînant au plus 26 enregistrements perdus (0,26 s). L'insertion de la passerelle a rétabli un lien obstrué de 0 % à 100 % de paquets reçus de bout en bout. Le coût matériel estimé est d'environ 250 € par unité, batterie comprise.

Le travail aboutit ainsi à une plateforme sensiblement révisée, plus économe, mieux intégrée, structurée logiciellement et reproductible, adaptée à une réplication sur un grand nombre de stations. Les principales limites restantes concernent l'acquisition simultanée des quatre voies à 100 Hz, la validation expérimentale de la synchronisation temporelle et l'absence de campagne de terrain représentative, qui constituent les priorités des développements futurs.

Design Improvements and Optimisation of the OhmPi Watcher System for Large-Scale Autonomous Deep ERT Measurements

Author	EYEN Guillaume
Section	MSc in Electrical Engineering — Electronic Systems and Devices
Academic supervisors	Prof. Jean-Michel Redouté, Prof. Frédéric Nguyen
Scientific advisor	Dr Guillaume Blanchy
Academic year	2025–2026

Deep electrical resistivity tomography (deep-ERT) extends electrical subsurface imaging to kilometre-scale profiles and greater investigation depths, with applications in mineral and geothermal exploration, fault-zone mapping and site characterisation for major underground infrastructure. Such surveys require large numbers of autonomous receiving stations that must record weak ground potentials, remain time-synchronised and stay remotely accessible throughout multi-day campaigns. Commercial receivers meeting these needs cost several thousand euros per unit, which directly limits the achievable spatial coverage. This work addresses that gap by redesigning the OhmPi Watcher, an open-source autonomous receiver derived from the OhmPi platform, for a reference deployment of 100 stations spread over more than 10 km for at least five days.

The redesign follows four complementary directions. At hardware level, the power, acquisition, switching and protection circuitry has been integrated on a single 100 × 160 mm board: oversized converters were replaced by devices matched to the actual loads, the bipolar analogue supply is disconnected through a high-side switch when it is not required, and the acquisition chain was extended from two to four independently conditioned electrode channels, with low-pass filtering, input protection and a revised electrode contact-resistance measurement circuit. The embedded firmware was reorganised around an explicit finite-state machine handling recording, idle and deep-sleep operation, RTC-alarm wake-up and arbitration of the SPI bus shared by the SD card and the LoRa radio. Wireless communication was improved through a new 433 MHz antenna integration, multi-hop packet forwarding with duplicate suppression, and the development of an autonomous STM32/SX1268 portable gateway able to bypass locally obstructed radio links.

Laboratory validation shows a two-channel recording current reduced from 174.6 mA to 77.6 mA (−56 %), corresponding to an estimated five-day consumption of 4.16 Ah for a daily cycle of 6 h recording, 6 h idle operation and 12 h sleep. The four analogue channels exhibit an input-referred electronic noise of 0.31 mV and gain errors below 0.07 %, with a first-order filter whose fitted cut-off frequency is 35.5 Hz. One hour of continuous logging was achieved at an effective sampling rate of 99.964 Hz without missing, duplicated or out-of-order records, while abrupt power-loss tests caused at most 26 lost records (0.26 s). Inserting the gateway restored an obstructed link from 0 % to 100 % end-to-end packet reception. The estimated hardware cost is approximately 250 € per unit including the battery.

The work therefore results in a substantially revised platform that is more energy efficient, better integrated, equipped with a better-structured software architecture and reproducible, supporting replication across a large number of stations. The main remaining limitations concern simultaneous four-channel acquisition at 100 Hz, the experimental validation of time synchronisation and the absence of a representative field campaign, which define the priorities for future development.