



Development of a Vehicle-Agnostic Direct-Sensing Telematics System for Electric Vehicles

by

Sameer Shaboodien

25002783

Thesis presented in partial fulfilment of the requirements for the degree of
Master of Engineering (Electronic) in the Faculty of Engineering at
Stellenbosch University.

Supervisor: Professor M. J. Boysen
Department of Electrical and Electronic Engineering

October 2025



UNIVERSITEIT•STELLENBOSCH•UNIVERSITY
jou kennisvennoot • your knowledge partner

Plagiaatverklaring / *Plagiarism Declaration*

1. Plagiaat is die oorneem en gebruik van die idees, materiaal en ander intellektuele eiendom van ander persone asof dit jou eie werk is.

Plagiarism is the use of ideas, material and other intellectual property of another's work and to present it as my own.

2. Ek erken dat die pleeg van plagiaat 'n strafbare oortreding is aangesien dit 'n vorm van diefstal is.

I agree that plagiarism is a punishable offence because it constitutes theft.

3. Ek verstaan ook dat direkte vertalings plagiaat is.

I also understand that direct translations are plagiarism.

4. Dienooreenkomsdig is alle aanhalings en bydraes vanuit enige bron (ingesluit die internet) volledig verwys (erken). Ek erken dat die woordelikse aanhaal van teks sonder aanhalingstekens (selfs al word die bron volledig erken) plagiaat is.

Accordingly all quotations and contributions from any source whatsoever (including the internet) have been cited fully. I understand that the reproduction of text without quotation marks (even when the source is cited) is plagiarism

5. Ek verklaar dat die werk in hierdie skryfstuk vervat, behalwe waar anders aangedui, my eie oorspronklike werk is en dat ek dit nie vantevore in die geheel of gedeeltelik ingehandig het vir bepunting in hierdie module/werkstuk of 'n ander module/werkstuk nie.

I declare that the work contained in this assignment, except where otherwise stated, is my original work and that I have not previously (in its entirety or in part) submitted it for grading in this module/assignment or another module/assignment.

Studentenommer / 25002783	Handtekening / 
Voorletters en van / S Shaboodien	Datum / 27/10/25

Abstract

Electrifying road transport offers significant environmental and economic advantages. However, South Africa's unique transport landscape and competing developmental priorities demand that any diversion of limited public funds be justified by demonstrable economic and social viability. Achieving this requires accurate modelling of an electrified road network, supported by data linking operational behaviour and energy demand across vehicles. Yet South Africa lacks locally generated empirical datasets on electromobility, and no existing telematics platform offers universal applicability across electric vehicles (EVs) without extensive setup time.

To address this gap, this work developed and validated a retrofit, vehicle-agnostic telematics platform capable of directly and non-invasively sensing traction-battery and mobility parameters. The measurements are transmitted securely via LTE Cat 1 bis using a lightweight messaging protocol, MQTT, with backlog retransmission, and presented through an online dashboard. The system integrates HVDC voltage and bidirectional inductive current measurement, GNSS-based motion tracking, accelerometry, and temperature measurement on a compact printed-circuit assembly within an IP67-rated automotive enclosure. It safely measures DC voltages up to 880 V and bidirectional currents of ± 1.5 kA, with galvanic isolation, basic insulation, and integrated transient protection. Data aggregation, local storage, and remote transmission occur at 1 Hz, with additional backlog retransmissions, while remote interfaces enable secure processing, storage, and visualisation of fleet data.

Laboratory, EV, and on-road testing confirmed full compliance with the defined functional and non-functional requirements. Voltage and current sensing achieved sub-percent full-scale accuracy across the -5°C to 55°C temperature range, while GNSS maintained typical positional and velocity accuracies of 1.5 m and 0.05 m s^{-1} , respectively. Temperature sensing, accelerometry, and battery-backed timekeeping met their specified performance criteria for meaningful inference, and the system remained fully operational with real-time coverage under vibration and non-ideal supply conditions during high-speed automotive testing.

The validated platform provides a practical, non-invasive, and accurate means of rapidly generating synchronised electromobility datasets under South African operating conditions. These datasets provide the empirical foundation for credible modelling of an electrified transport network, enabling the state to evaluate, design, and prioritise interventions for South Africa's transition to electric mobility.

Uittreksel

Die elektrifisering van padvervoer bied beduidende omgewings- en ekonomiese voordele. Suid-Afrika se unieke vervoerlandskap en mededingende ontwikkelingsprioriteite vereis egter dat enige aanwending van beperkte openbare fondse geregverdig moet word deur bewese ekonomiese en sosiale lewensvatbaarheid. Dit vereis akkurate modellering van 'n geëlektrifiseerde padnetwerk, gesteun deur data wat operasionele gedrag en energievraag oor voertuie heen koppel. Suid-Afrika beskik egter nie oor plaaslik verkrygde empiriese elektromobiliteitsdatastelle nie, en geen bestaande telematiese platform bied universele toepasbaarheid oor elektriese voertuie (EV's) sonder aansienlike opstellingstyd nie.

Om hierdie leemte aan te spreek, is in hierdie werk 'n herbruikbare, voertuig-agnostiese telematiese platform ontwikkel en gevalideer wat in staat is om traksiebattery- en mobiliteitsparameters direk en nie-indringend te meet. Die metings word veilig via LTE Cat 1 bis oorgedra met behulp van 'n liggewigboodskapprotokol (MQTT) met agterstand-hertransmissie, en die resultate word deur 'n aanlynpaneel gevisualiseer. Die stelsel integreer HVDC-spanning- en tweerigting-induktiewe stroommeting, GNSS-gebaseerde bewegingsopsporing, versnellings- en temperatuurmeling op 'n kompakte gedrukte stroombaan (PCB) binne 'n IP67-gegradeerde motoromhulsel. Dit meet veilig DC-spannings tot 880 V en tweerigtingstrome tot $\pm 1.5\text{ kA}$, met galvaniese isolasie, basiese isolasie en geïntegreerde oorspanningsbeskerming. Data-insameling, plaaslike bering en afgeleë oordrag vind teen 1 Hz plaas, met bykomende agterstand-hertransmissies, terwyl afgeleë koppelvlakke veilige verwerking, bering en visualisering van vlootdata moontlik maak.

Laboratorium-, EV- en padtoetse het volledige voldoening aan die gedefinieerde funksionele en nie-funksionele vereistes bevestig. Spannings- en stroommeting het subpersentasie-volle-skaal-akkuraatheid bereik oor die temperatuurbereik van -5°C – 55°C , terwyl GNSS (Globale Navigasiesatellietstelsel) tipiese posisie- en snelheidsakkuraathede van onderskeidelik 1.5 m en 0.05 m s^{-1} gehandhaaf het. Temperatuurmeling, versnellingmeting en battery-ondersteunde tydhoufunksie het aan hul gespesifiseerde prestasiekriteria vir betekenisvolle afleiding voldoen, en die stelsel het volledig operasioneel gebly met regstreekse dekking onder vibrasie en nie-ideale toevoertoestande tydens hoëspoedmotortoetse.

Die gevalideerde platform bied 'n praktiese, nie-indringende en akkurate metode om vinnig gesynchroniseerde elektromobiliteitsdatastelle onder Suid-Afrikaanse bedryfstoestände te genereer. Hierdie data verskaf die empiriese grondslag vir geloofwaardige modellering van 'n geëlektrifiseerde vervoernetwerk, wat die staat in staat stel om ingrypings vir Suid-Afrika se oorgang na elektriese mobiliteit te evaluateer, te ontwerp en te prioritiseer.

Contents

Declaration	i
Abstract	ii
Uittreksel	iii
List of Figures	vii
List of Tables	x
Nomenclature	xi
1. Introduction	1
1.1. Background	1
1.2. Problem Statement	2
1.3. Aim	2
1.4. Objectives	3
1.5. Scope	3
1.6. Summary of Work	3
1.7. Thesis Structure	4
2. Literature Review	5
2.1. Electric Vehicle Fundamentals	5
2.2. Existing Data and Collection Methods	6
2.2.1. Existing Data	7
2.2.2. Existing Data Collection Methods	8
2.2.3. Research Gap	10
2.3. Hardware Architectures	11
2.3.1. Sensing Components	11
2.3.2. Processing and Control	19
2.3.3. Communication Interfaces	22
2.3.4. Power, Layout, and Mechanical Considerations	27
2.4. Cloud Architectures	31
2.4.1. Storage	32
2.4.2. Security and Identity Management	33

2.4.3. Deployment and Operation	33
2.5. Web Application Architectures	35
2.5.1. Interface Development	35
2.5.2. Deployment and Performance	36
2.5.3. Usability and Security	37
2.5.4. EV Dashboard Content	38
3. Design	40
3.1. Requirements	40
3.2. Hardware System	44
3.2.1. Sensing Components	45
3.2.2. Processing and Control	59
3.2.3. Communication Interfaces	62
3.2.4. Power, Layout, and Mechanical Integration	69
3.2.5. Hardware Versions and Modifications	78
3.2.6. Requirement Traceability	80
3.3. Firmware System	82
3.3.1. Introduction	82
3.3.2. Architecture	82
3.3.3. Sensing Interfaces	84
3.3.4. Data Handling, Storage, and Communication	85
3.3.5. Timekeeping and Synchronisation	86
3.3.6. Firmware Robustness and Diagnostic Hooks	87
3.3.7. Requirement Traceability	87
3.4. Cloud System	88
3.4.1. Architecture Overview	89
3.4.2. Messaging and Compute	89
3.4.3. Storage	90
3.4.4. Latency	91
3.4.5. Security	92
3.4.6. Cost and Scalability	92
3.4.7. Requirement Traceability	93
3.5. Web Application System	93
3.5.1. Frontend Architecture	94
3.5.2. Interface and Functional Design	95
3.5.3. Data Handling and Processing	96
3.5.4. Implementation and Configuration	97
3.5.5. Requirement Traceability	98

4. Testing and Validation	100
4.1. Approach	100
4.1.1. Structure of Chapter	100
4.1.2. Tests Conducted	100
4.1.3. Scope and Assumptions	101
4.2. Sensing	102
4.2.1. HVDC Voltage Sensing	102
4.2.2. HVDC Current Sensing	104
4.2.3. Accelerometer	108
4.2.4. GNSS	109
4.2.5. Temperature	112
4.3. Data Handling and Communication	113
4.3.1. LTE Transmission and SD Logging	113
4.3.2. System Timing	115
4.3.3. Real-Time Latency and Historical Data Access	116
4.3.4. Remote Interface Security	118
4.3.5. Frontend and Data Interface	118
4.4. System Retrofittability, Generalisability and Robustness	123
4.4.1. Calibration and Protocol Independence	123
4.4.2. PCB Sizing	124
4.4.3. Power Usage	124
4.4.4. Mechanical Robustness	124
4.4.5. Power Rails	125
5. Conclusion	129
Bibliography	131

List of Figures

2.1. Electric vehicle powertrain architecture, showing the traction battery, inverter, traction motor, and associated power conversion and control subsystems [1].	6
2.2. Hall-effect current sensing principle showing magnetic flux generated from the measured conductor interacting with the sensing element to produce a perpendicular Hall voltage. [2].	14
2.3. Fluxgate current sensor typical topologies [3].	15
2.4. Inductive clip-on CAN bus probe for non-intrusive measurement [4].	23
3.1. Overall hardware signal flow from vehicle data buses to the PCB's sensing, processing, and communication subsystems, and onward to the remote cloud and web interfaces.	45
3.2. High-voltage resistive divider chain showing staged attenuation, surge protection, and the RC filter that conditions the AMC131M01 input.	47
3.3. Isolated voltage sensing digitisation stage with local clock generation (top right), power supply filtering, and differential outputs into the isolated AMC131M01 ADC.	48
3.4. Inductive current sensors – LEM CAB-SF (left) and GMW clamp (right). .	53
3.5. Isolated ADC input chain for the GMW coreless probe, including RC anti-alias filters.	54
3.6. Isolated supply network for the GMW probe, showing the dedicated 5 V converter (bottom) and the ± 2.5 V rails (top) that power the ADS1220 AED. .	55
3.7. Galvanic digital isolation interface between the analogue sensing chain and MCU domain.	55
3.8. Acceleration sensing circuit based on the ADXL343 and its associated filtering. .	58
3.9. ESP32-S2 Mini schematic pinout showing peripheral assignments, communications buses, and sensor interfaces.	60
3.10. MCP7940N real-time clock circuit with 32.768 kHz crystal, associated load and filtration capacitors and CR2032 backup supply holder.	61
3.11. CAN bus circuits for fluxgate sensor (top) and vehicle CAN bus acquisition (bottom), showing transceivers (left) and controllers (right).	63
3.12. Clip-on CAN inductive probe and interface cable for non-intrusive signal tapping.	64

3.13. LENA module interface circuit showing level shifter (top left), module headers (top right), 3.3 V bypass network (bottom left), and GPIO-driven LED drivers (bottom right).	66
3.14. MicroSD interface circuitry showing hinged socket and pull-up network on CMD and data lines.	68
3.15. Primary power regulation chain located at the PCB input stage, showing the 12 V isolated regulation stage (top), followed by the 3.8 V regulator (bottom left) and the 3.3 V regulator (bottom right).	73
3.16. Final two-layer PCB layout with top and bottom copper overlays, highlighting the high-voltage domain (bottom left) and isolated GMW current digitisation section (centre right).	75
3.17. Connector layout showing high-voltage and vehicle inputs, GMW and fluxgate interfaces, temperature input, and DB9 port for the CSS CAN reader.	77
3.18. Comparison of first and second hardware iterations illustrating significant improvements.	79
3.19. Firmware architecture and control flow showing modular layering and second-aligned scheduler.	83
3.20. Cloud architecture linking embedded devices, AWS services, and web application.	89
3.21. Frontend architecture overview showing the Next.js dashboard entry point with analytics, fleet map, and command tabs after Cognito sign-in.	94
4.1. Laboratory configuration for HVDC measurement testing, showing the telematics PCB connected to the high-voltage source via its capacitor bank for controlled voltage sweep acquisition.	101
4.2. EV Bench experimental setup showing the current sensors and vehicle CAN logger connected to the Roam Air bench simulator, while being supplied from the 12 V Roam Air supply.	102
4.3. Measured HVDC bus voltage error relative to a precision source across the laboratory sweep range, showing both the raw error curve and the post-hoc calibrated curve demonstrating linearity.	105
4.4. Comparison of HVDC-sensed and CAN-reported voltages under stationary EV conditions, with voltage profiles plotted on the left axis and their offset (difference) on the right axis.	106
4.5. Profile comparison of the current sensors against the EV Bench.	107
4.6. Offset comparison of current sensors against the EV Bench, with Fluxgate-CAN and GMW-CAN errors shown on the left axis, and the GMW-Fluxgate offset shown on the right axis.	107

4.7. Comparison of longitudinal acceleration (<i>y</i> -axis) between prototype and reference device during road testing.	108
4.8. 2D map overlay of module and reference positions throughout the trip.	110
4.9. Vehicle speed and GNSS fix validity over time during the highway mobility test, demonstrating full coverage at high-speed.	111
4.10. Two-dimensional position error of the prototype relative to the reference receiver during the highway mobility test.	111
4.11. Three-dimensional position error of the prototype relative to the reference receiver during the highway mobility test.	111
4.12. Velocity error during highway mobility test. Maximum error: 4.23 km h^{-1}	111
4.13. Ambient temperature recorded by the enclosure-mounted sensor throughout the vehicle drive test.	112
4.14. Vehicle speed and LTE network validity during the mobility test, demonstrating full coverage at high-speed.	114
4.15. Cognito-hosted login interface presented to operators before accessing the telemetry dashboard.	119
4.16. Dashboard view of a one-month telemetry query illustrating pagination controls and downsampled trend plotting.	120
4.17. Dashboard view of a one-minute telemetry window emphasising high-resolution chart detail.	120
4.18. Frontend controls for selecting time ranges, devices, and telemetry channels before charts render.	120
4.19. Frontend GNSS map displaying latest device positions with interactive metadata pop-ups.	121
4.20. Device command console allowing remote configuration of telematics units.	122
4.21. Command help and parameter templates confirming correct interface operation.	122
4.22. Fully assembled enclosure with the mounted PCB.	125

List of Tables

2.1. Comparison of the main cellular radio access technologies (RATs), including throughput, latency, coverage, and current characteristics.	25
3.1. Stakeholder requirements (SR) capturing sensing, data handling, and retrofit needs for the telematics platform.	41
3.2. Functional requirements (FR) derived from system-level requirements, grouped by sensing, data handling, and robustness.	42
3.3. Non-functional requirements (NFR) derived from system-level requirements, grouped by sensing, data handling, and robustness.	43
3.4. Voltage measurement error budget referred to 1.1 V full-scale, highlighting component errors and temperature drift contributors.	51
3.5. ADS1220 input-chain error budget in ppm referred to ± 2.000 V.	57
3.6. Comparison of key characteristics for the two current sensing approaches. .	57
3.7. Component current requirements per rail (excluding regulators).	70
3.8. Typical and maximum currents for each supply rail (excluding regulators). .	71
3.9. Regulator specification summary listing input/output rails, peak current capability, efficiency, and topology.	72
3.10. Computed total regulator load requirements with 20% headroom and efficiency. .	72
3.11. Regulator performance metrics covering ripple, noise, accuracy, and temperature variation for each supply stage.	73
3.12. Traceability between hardware subsystems and the functional and non-functional requirements they satisfy.	80
3.13. Traceability matrix showing how firmware features satisfy the functional and non-functional requirements.	87
3.14. Summary of how the cloud architecture satisfies the key functional and non-functional requirements.	93
3.15. Frontend contributions toward requirement fulfilment.	98
4.1. Measured HVDC bus voltage error relative to applied source values across the full test range.	104
4.2. Summary of verification outcomes and traceability back to the functional and non-functional requirements.	126

Nomenclature

Key Terms and Concepts

API Gateway	Managed HTTP API service.
App Router	Next.js routing model used for pages and server routes.
Backlog	Data which is stored or buffered for later processing during temporary outages.
Bypass capacitor	A capacitor placed between a power rail and ground to shunt high-frequency noise and stabilise the local supply voltage.
Clearance distance	Shortest through-air distance between conductive parts, governing immunity to flashover under overvoltage or altitude variation.
Comparative tracking index (CTI)	Quantitative measure (in volts) of an insulating material's resistance to surface electrical breakdown when exposed to moisture or contaminants; higher CTI values allow smaller creepage distances.
Creepage distance	Shortest path along an insulating surface between conductive parts, determining resistance to surface tracking (pollution-induced carbon-path formation along the PCB surface) and partial discharge (dielectric breakdown not between conductors).
DynamoDB Keys	DynamoDB primary partition and sort keys enable efficient time-bounded queries per device.
Encryption Keys	Public and private keys with mathematical relationships that enable authentication of the other token bearer.
Full-stack Application	Integrated software solution encompassing both frontend (user-facing) and backend (server-side) components.
Insulation type	IEC 62368-1 classification defining protection level against electric shock: basic insulation provides single-fault protection; supplementary insulation adds a second, independent layer; reinforced insulation provides the equivalent of double insulation within a single system.
JWTs	JSON Web Tokens used by OIDC/OAuth 2.0; short-lived access tokens authorise API calls, while refresh tokens are used to obtain new access tokens.
Pollution degree	IEC 62368-1 classification describing expected environmental contamination and its potential conductivity: Degree 1—no conductive pollution or condensation; Degree 2—non-conductive pollution that may temporarily become conductive due to condensation; Degree 3—conductive pollution or frequent condensation.
Telemetry	Device-generated sensor data.
Telematics	Integration of telecommunications and informatics for remote measurement, monitoring, and control of vehicles.
Temporary overvoltage	Short-duration overvoltage event, such as surge, impulse, or dielectric withstand, exceeding normal working levels but typically lasting less than one minute.

TTFF	Time-to-first-fix for GNSS acquisition after startup.
URC	Unsolicited Result Code from the modem indicating network or GNSS results.
Working voltage	Maximum continuous DC or RMS voltage that can appear across an insulation barrier during normal operation.

Acronyms and Abbreviations

ADC	Analog-to-Digital Converter
API	Application Programming Interface
AWS	Amazon Web Services
CAN	Controller Area Network
CSV	Comma-Separated Values
DBC	Database CAN
FSR	Full-scale Range
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HTTP	Hypertext Transfer Protocol
HTTPS	HTTP Secure (HTTP over TLS)
ICE	Internal Combustion Engine
IMU	Inertial Measurement Unit
IoT	Internet of Things
JSON	JavaScript Object Notation
JWT	JSON Web Token
LTE	Long-Term Evolution (cellular)
MAE	Mean Absolute Error
MQTT	Message Queuing Telemetry Transport
OAuth 2.0	Open Authorization 2.0 (authorization framework for securing APIs)
OIDC	OpenID Connect (authentication layer on OAuth 2.0)
PCB	Printed Circuit Board
QoS	Quality of Service (MQTT)
RTC	Real-Time Clock
S3	Simple Storage Service (AWS)
SD	Secure Digital (memory card)
SPI	Serial Peripheral Interface
TS	TypeScript
TLS	Transport Layer Security
TTFF	Time To First Fix (GNSS)
UART	Universal Asynchronous Receiver/Transmitter

Chapter 1

Introduction

1.1. Background

Electrifying road transport is expected to substantially improve urban air quality, reduce lifecycle emissions, and lower both the cost and volatility of transport energy, especially when supported by a progressively decarbonised grid. Countries such as Norway have already realised these benefits through coordinated policy, renewable energy integration, and widespread charging infrastructure, achieving 88.9% of new vehicle sales as electric in 2024 [5].

While the transition to electric transport offers clear advantages, replicating these outcomes in South Africa entails significant fiscal implications. Without reliable quantitative modelling to identify and assess the outcomes of optimal electrification strategies, the state cannot determine the true opportunity cost of diverting limited public resources, and policy decisions risk misalignment with national objectives.

A key requirement for modelling an electrified road network is comprehensive data describing how vehicles operate in terms of their movement and energy use. These datasets require time-synchronised electrical and mobility parameters—such as voltage, current, speed, and position. When aggregated across fleets, regions, and vehicle classes, they reveal patterns of operational behaviour—such as trip lengths and duty cycles—and the influence of environmental factors including temperature, road conditions, and load on energy patterns such as power consumption, charging behaviour, and derived efficiency metrics [6]. Together, these data enable behavioural, spatial, and temporal modelling of an electrified transport network, informing strategies to assess the viability of electrification and design state interventions to ensure an economically sound and socially equitable transition [7].

Existing public data sources are insufficient for this scale of modelling. Foreign datasets cannot be directly generalised to South African conditions due to fundamental differences in factors such as grid reliability, infrastructure maturity, settlement patterns, transport modes, travel distances, and income distribution [8, 9]. Local studies have attempted to address this by leveraging internal combustion engine (ICE) vehicle data and static electric vehicle (EV) parameters to infer equivalent efficiency models; while these provide useful first-order estimates, they fail to account for real-world operating variations and therefore lack the credibility required to support large-scale investment decisions [10].

A central constraint in addressing this data gap is the time and complexity associated with high-fidelity data acquisition from EVs. Commercial telematics systems typically interface with the in-vehicle controller area network (CAN) bus; however, protocol definitions vary between vehicles, and manufacturer or third-party calibration, or inference-based decoding, requires significant time and effort. OBD-II interfaces expose little or no traction information, while existing direct-sensing telematics methods are suited only to low-power or invasive applications [11, 12]. Consequently, no existing system offers an integrated, non-invasive, and vehicle-agnostic platform capable of capturing both high-power DC and mobility data at research-grade fidelity.

To address this gap, the thesis develops a rapid-deployment, direct-sensing telematics platform capable of acquiring high-resolution electromobility data across a large range of EVs. The system measures battery pack voltage and current, mobility (position, velocity, and acceleration), temperature, and battery-backed time in a compact form factor without reliance on the CAN bus. Data are securely transmitted to a cloud-based pipeline and visualised via a modern web dashboard. The system enables the rapid generation of locally relevant datasets.

The resulting datasets are intended for public use to expedite collective research toward an economically and socially beneficial transition to cleaner transport. Any dissemination of data should exclude personal identifiers, apply appropriate aggregation, treat raw datasets as confidential, and adopt additional safeguards to minimise the risk of compromising individual privacy.

1.2. Problem Statement

South Africa lacks empirical, high-resolution datasets capturing real-world electric vehicle operation under local conditions, limiting the quantitative modelling needed to justify public investment. Obtaining such data remains challenging: existing commercial telematics systems rely on indirect sensing through proprietary CAN interfaces, involving complex and time-consuming access procedures, while direct-sensing telematics systems remain non-universal.

1.3. Aim

To design, implement, and validate a non-invasive, vehicle-agnostic EV telematics platform capable of directly capturing high-resolution electromobility data and providing secure, cloud-based access through a web dashboard.

1.4. Objectives

1. Define **stakeholder, functional, and non-functional requirements** to ensure a comprehensive, accurate, and universally deployable solution.
2. **Design and build a hardware device** that measures key electrical and vehicle-state variables with automotive-grade power protection and isolation.
3. **Develop embedded firmware** for reliable data acquisition, local storage, and remote communication.
4. **Implement a secure, scalable cloud pipeline** for data ingestion, storage, and processing.
5. **Develop a secure web dashboard** for remote data visualisation and analysis.
6. **Validate system performance** under representative operating conditions, quantifying measurement accuracy, end-to-end reliability, and data accessibility.

1.5. Scope

The system is developed as a prototype for research and pilot deployments rather than as a commercial product. Accordingly, strict safety certification, extensive environmental ingress protection, remote-control functionality, and edge-case robustness are excluded. The research prototype measures only the DC side of the drivetrain to determine power consumption and is designed for typical South African ambient temperature ranges of approximately -5°C to 55°C . Additionally, it assumes standard GNSS and mobile network coverage. Cost and power consumption are considered but not fully optimised. Validation includes both laboratory and field testing; however, available equipment does not permit exhaustive evaluation across all measurement and environmental ranges with absolute accuracy.

1.6. Summary of Work

This thesis presented the design, implementation, and validation of a telematics platform developed to address the shortage of locally derived, high-fidelity electromobility datasets in South Africa; existing collection methods are limited by lengthy setup times. The developed system offers a universal, non-invasive solution that directly and time-synchronously measures key electrical and mobility parameters, enabling large-scale analysis of real-world electric-vehicle operation.

The platform integrates isolated DC voltage and current measurement, GNSS-based motion tracking, local timekeeping, optional vehicle CAN data acquisition, and ambient temperature and acceleration sensing within a compact printed-circuit assembly. All measurements are aggregated and timestamped at 1 Hz, logged to removable storage, and transmitted securely via LTE Cat 1 bis. The accompanying cloud layer and web dashboard enable data retrieval and visualisation for validation and analysis.

Validation across laboratory, EV, and in-vehicle tests confirmed that all functional and non-functional objectives were met. Voltage and current sensing achieved sub-percent full-scale accuracy across the temperature range, up to 880 V and ± 1.5 kA, with galvanic isolation, basic insulation and transient protection. GNSS maintained a typical positional accuracy of 1.5 m and velocity accuracy of 0.05 m s^{-1} , while the accelerometer, temperature sensor, and real-time clock each met their specified tolerances. CAN acquisition operated reliably, and data logging and transmission remained continuous under all tested conditions.

The physical implementation met all retrofit, robustness, and safety requirements, with accuracy requirements maintained across -5°C to 55°C . Installation required no invasive vehicle modification, using inductive and terminal-tap connections, and the IP67-rated enclosure provided protection against vibration, shock, and moisture.

The developed platform achieved its aim of providing a non-invasive, accurate, and reliable means of rapidly collecting synchronised electromobility data across a range of electric vehicles. The system represents a practical tool for generating the locally relevant datasets necessary to inform South Africa's transition to electric road transport.

1.7. Thesis Structure

Chapter 2 (*Literature Review*) reviews related datasets and collection approaches, and surveys hardware and software building blocks. Chapter 3 (*System Design*) specifies requirements and details the hardware, firmware, cloud, and web application design. Chapter 4 (*Validation and Testing*) reports laboratory, vehicle, and field testing. Chapter 5 (*Conclusion*) summarises findings, limitations, and future work.

Chapter 2

Literature Review

This chapter identifies the gap in empirical electric-vehicle data that constrains South Africa's transition to transport electrification and establishes the technical context required for system design. It begins with an overview of electric-vehicle fundamentals to clarify the key electrical and communication principles relevant to telematics. It then reviews available datasets and collection methods, outlining their limitations and the resulting research gap. Building on this, the chapter examines existing studies and prerequisite knowledge on telematics system hardware, cloud, and web application architectures as the basis for the design choices developed in subsequent chapters.

2.1. Electric Vehicle Fundamentals

Electric vehicles (EVs) operate using electric motors powered by rechargeable battery packs, typically based on lithium-ion chemistry. The primary components of an EV include the high-voltage battery, power electronics, traction motor (the electric motor that provides the propulsive, or tractive, force), and onboard AC charger. Electrical energy stored in the battery is converted into mechanical energy to drive the wheels, while regenerative braking recovers kinetic energy during deceleration to recharge the battery.

EVs employ a DC battery, with low-voltage vehicle classes—such as electric scooters—typically operating below 100 V, while passenger and commercial EVs utilise high-voltage packs typically rated between 200 V to 800 V. Current levels vary substantially between vehicle classes, with peak currents rarely exceeding 1 kA [13]. The battery supplies an inverter that converts DC to AC for the traction motor, and the motor controller regulates torque and speed according to driver input and operating conditions. Auxiliary systems, including lighting, infotainment, and control electronics, are powered separately by an isolated low-voltage subsystem, typically 12 V.

Power flow within an EV is bidirectional, with current proportional to instantaneous power demand. Battery voltage remains largely stable under dynamic loads, where minor fluctuations arise from internal resistance, whereas longer-term variation results from changes in the battery's state of charge (SoC). The voltage follows a vehicle-specific open-circuit voltage-state of charge (OCV-SoC) curve, which models how the battery's

no-load voltage changes with its state of charge. Over time, cell degradation alters both the usable voltage range and the OCV-SoC profile, and the vehicle's battery management system (BMS) periodically recalibrates its internal model to maintain estimation accuracy.

Communication within an EV coordinates the operation of its subsystems. Commercial EVs typically employ a Controller Area Network (CAN) bus linking the BMS, inverter, and motor controller, through which variables such as speed, state of charge (SoC), voltage, current, and diagnostic information are transmitted, although message sets differ between vehicles and powertrain parameters are not mandatory. Some prototype or retrofitted EVs may instead use alternative or proprietary communication protocols [11].

As illustrated in Figure 2.1, the EV powertrain comprises the battery pack, inverter, and traction motor [1]. A clear understanding of this electrical and communication architecture is essential for identifying measurable parameters relevant to telematics and performance analysis.

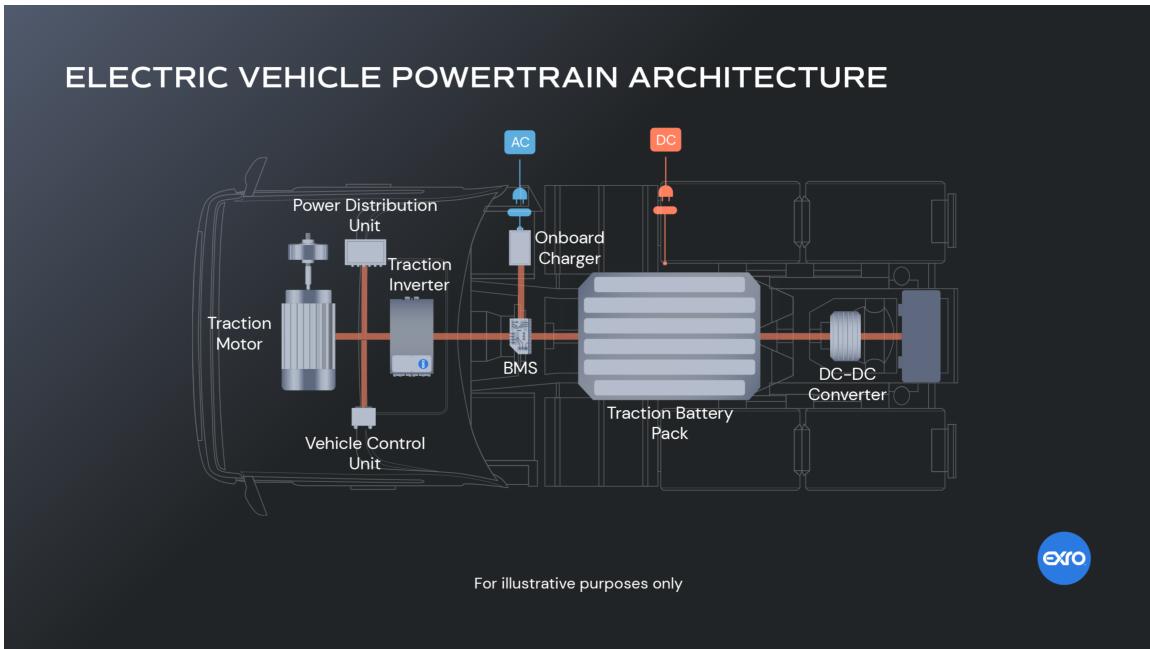


Figure 2.1: Electric vehicle powertrain architecture, showing the traction battery, inverter, traction motor, and associated power conversion and control subsystems [1].

2.2. Existing Data and Collection Methods

This section reviews existing electromobility datasets from international and South African studies to highlight the scarcity of local empirical data, then outlines the limitations of current collection methods that contribute to this shortage, together defining the research gap addressed in this study.

2.2.1. Existing Data

Research on electric-vehicle performance and infrastructure demand has expanded rapidly over the past decade, yet most existing studies remain geographically and methodologically constrained. This subsection reviews prominent international and South African studies to highlight the current limitations in available electromobility data.

Overseas, large-scale initiatives have captured empirical EV usage data. The EV Project in North America—the world’s largest plug-in vehicle infrastructure study—was conducted between 2011 and 2013 under the U.S. Department of Energy [14]. Its goal was to understand charging behaviour and infrastructure demand. Nearly 17 000 AC Level 2 and over 100 DC fast chargers were installed across 22 regions, logging about 200 million kilometres of driving and 6 million charging events from over 8 000 vehicles, including the Nissan Leaf (battery electric) and Chevrolet Volt (plug-in hybrid). Results showed that roughly 98 % of charging occurred at home or work on workdays, with the majority of public charging events concentrated around a few well-located sites.

The study provided valuable aggregated insights into charging behaviour, particularly the reliance on home and workplace charging. However, its breadth and level of output detail were insufficient for large-scale planning, constrained by participant privacy limitations. A 2025 cross-continental study further highlighted limitations in existing EV research, including restricted global representativeness, low temporal resolution, and limited data breadth. Drawing on more than 500 million charging sessions across South America, Asia, and Europe, it showed that empirical validation substantially improves the accuracy of current road-electrification modelling systems [15].

Moreover, these studies reflect regions that differ substantially from South Africa—typically safer, better-developed environments with distinct travel patterns and infrastructure. Projections based on such foreign datasets cannot reliably inform South African planning, as the country differs fundamentally in grid reliability, infrastructure maturity, settlement patterns, transport modes, travel distances, and income distribution [8]. A credible transition strategy therefore requires locally obtained electromobility data to establish realistic baselines for technical feasibility, investment scale, and policy design.

A South African study sought to address the local data gap by applying a kinetic, physics-based model to per-second GPS traces from conventional minibus taxis to estimate electric energy consumption under realistic driving conditions. The model computed aerodynamic drag, rolling resistance, and slope forces from instantaneous speed and elevation, using differenced velocity to derive acceleration and infer tractive power (the equivalent electrical power required to propel the vehicle forward). Fixed parameters for vehicle mass, drag coefficient, rolling resistance, drivetrain efficiency, auxiliary loads, and regenerative braking were then applied to integrate total energy and express efficiency in kW h km^{-1} . Temperature effects were assumed negligible. The study produced useful

context-specific estimates for paratransit electrification planning, reporting a mean energy consumption of $0.39 \text{ kW h km}^{-1}$ across various operational contexts [10].

However, the model's validity is limited by its assumptions of monotonic acceleration between 1 Hz samples and fixed vehicular and environmental parameters, which neglect real-world variability. Consequently, while ICE-to-EV conversion simulations are widely used in existing studies for preliminary approximations of fleet and grid behaviour, they remain insufficiently reliable for large-scale infrastructure or policy decisions without empirical validation [15, 16].

To the best of the author's knowledge, no public, high-resolution empirical electromobility datasets currently exist in South Africa. The generation of such data should be prioritised and made publicly accessible in aggregated, privacy-preserving form to support collaborative research and evidence-based planning for transport electrification.

2.2.2. Existing Data Collection Methods

This subsection examines why electromobility data remain scarce despite being readily available within vehicles. Although key parameters are already measured on the in-vehicle CAN bus, access to these data is restricted. Manufacturers limit distribution of Database CAN (DBC) files, third-party databases cover only selected models, OBD-II ports provide little or no traction information, while reverse-engineering CAN messages is slow and inconsistent. Existing direct-sensing EV telematics solutions are limited to low-voltage laboratory applications, lacking the required data fidelity.

Manufacturer Collaboration

Every CAN signal in a commercial vehicle is defined within a Database CAN (DBC) file that specifies message identifiers, signal scaling, offsets, and engineering units. These files form part of the vehicle's proprietary communication protocol and are therefore considered confidential intellectual property. Manufacturers rarely release them, primarily due to liability, intellectual property protection, and financial considerations [11].

In practice, access to these files is granted only through formal partnerships with established research institutions or telematics firms. The process typically involves extensive legal review and often imposes restrictions on data dissemination. Consequently, without existing partnerships, this approach demands significant time and resources, does not scale across different vehicle manufacturers, and may not allow dissemination at the required data resolution [11].

Third-Party Collaboration

Telematics providers such as Geotab [17] maintain proprietary decoded CAN databases and offer commercial hardware that interprets vehicle network data for fleet-level analytics.

However, they do not offer direct-sensing solutions, as these involve high-voltage interfacing, stringent regulatory compliance, and safety liabilities that extend beyond their software- and service-oriented business models. Moreover, such companies are not legally permitted to distribute proprietary DBC files owned by vehicle manufacturers.

These commercial systems are primarily optimised for operational and compliance metrics, operating at moderate temporal and quantitative resolutions suitable for driver behaviour, fuel efficiency, and maintenance monitoring rather than precision measurement. As a result, their capability is limited in high-fidelity or engineering-grade data acquisition scenarios. The scope of vehicle compatibility is similarly constrained, typically encompassing only mainstream commercial models within supported CAN databases. Consequently, less common, newly released, or retrofitted vehicles—particularly those lacking standard CAN access or using non-standard protocols—cannot be seamlessly integrated.

Furthermore, while modern telematics platforms emphasise data security through encryption and certifications such as ISO 27001 and SOC 2 compliance, these measures also highlight the inherent sensitivity and regulatory burden associated with telematics data handling. Overall, such systems remain valuable for fleet management and compliance, but they are not generalisable across all vehicle contexts and remain unsuitable for high-resolution electromobility data collection [17].

OBD-II Interfacing

The OBD-II port provides a convenient access point to a subset of CAN messages but often exposes little relevant data for EVs or may be absent entirely. Unlike in combustion vehicles, there is no regulatory requirement to report traction battery or motor parameters, restricting access to power-related information [12]. In addition, the breadth and temporal resolution of OBD-II data may be limited, reducing its suitability for high-fidelity telemetry.

Reverse Engineering DBC Files

In the absence of a DBC file, researchers can capture raw CAN traffic and empirically decode it by correlating message variations with observable vehicle behaviour. However, ambiguities in message structure, manufacturer-specific encoding, and the sheer volume of data—typically hundreds of unique messages—make this process complex and unlikely to yield a complete solution. It is labour-intensive, must be repeated for each vehicle type, and typically recovers only a subset of available signals [11]. Consequently, this approach is impractical for rapid, large-scale data collection across diverse vehicle platforms.

Direct Sensing Solutions

Direct-sensing systems (those not reliant on CAN network access) have been investigated in several electric vehicle prototypes, each incrementally improving sensing capability or communication performance, yet none achieving high-voltage safety, automotive resilience,

or high-fidelity data acquisition. Pearce [18] demonstrated one of the earliest GSM-linked telemetry systems, proving the feasibility of continuous remote data logging from an electric vehicle but limited to auxiliary-level sensing, a low update rate of 60 s, and a maximum input voltage of 60 V without isolation or insulation compliance.

Kaur *et al.* [19] advanced this by integrating LTE communication via the A7672E modem and combining voltage, current, temperature, and GNSS sensing—establishing four essential telemetry variables and demonstrating LTE’s superior reliability and cloud integration compared to GSM.

Recent prototypes [20–23] employed Arduino and ESP32 platforms for rapid development, maintaining similar sensing and communication functionality while remaining constrained by low-voltage operation and uncalibrated measurements, using either high-offset, low-linearity Hall-effect sensors or invasive shunt sensors. Pérez *et al.* [24] improved temporal fidelity through adaptive 1 Hz to 2 Hz sampling, demonstrating that higher refresh rates yield significantly improved accuracy in energy estimation.

However, all existing studies demonstrating retrofit, direct-sensing EV telematics systems remain unable to measure DC voltages above 100 V, lack galvanic isolation in voltage sensing, and rely on low-accuracy or invasive current sensors. Furthermore, they are not implemented in compact, vibration-resistant forms, highlighting the absence of a scalable, vehicle-agnostic, high-voltage direct-sensing telemetry solution. This gap persists primarily due to the niche requirement for rapid and safe deployment across varied high-power EV platforms, yet such a system remains essential for expanding empirical EV datasets within the South African context.

2.2.3. Research Gap

Despite the availability of international EV datasets and local ICE-to-EV inference models, empirical data capturing real-world EV operation under South African conditions remain scarce and non-transferable. Foreign datasets cannot be generalised due to South Africa’s distinct socioeconomic and infrastructural context, while local inference models rely on static EV parameters that lack empirical credibility. This constrains the development of credible, data-driven analyses and strategies needed to support an optimal, equitable, and economically viable transition to transport electrification. Existing commercial telematics systems rely on indirect sensing through proprietary CAN interfaces, while available direct-sensing approaches lack universality across vehicles due to limited high-voltage capability and insufficient compliance with automotive robustness standards.

The research gap addressed in this thesis is the absence of a rapid-deployment, non-invasive, and vehicle-agnostic method for directly measuring and collecting accurate, comprehensive electromobility data.

Summary

Existing datasets and universal data collection methods remain limited within the EV domain. Addressing this gap is a critical step toward overcoming barriers to large-scale local data acquisition and enabling the dissemination of datasets that can accelerate effective transport electrification in South Africa.

2.3. Hardware Architectures

The hardware layer forms the physical foundation of a telematics system, encompassing the sensors, processing units, data storage, and communication interfaces that enable acquisition, handling, and transmission of vehicle telemetry.

This section reviews prerequisite literature on the hardware subsystems discussed in this thesis—sensing, processing, storage, and communication. It examines representative studies and implementations to identify common architectures, performance parameters, and design considerations that inform subsequent component selection and integration.

2.3.1. Sensing Components

High-Voltage Sensing

Voltage sensing in EV telematics provides information on battery pack state, available power, and degradation trends, since voltage behaviour reflects both state of charge and internal resistance.

Existing Work

Winkelholz *et al.* [25] developed a PCB-integrated resistive high-voltage probe for precise measurement of combined DC and transient AC signals up to approximately $2.5\text{ kV}_{\text{pp}}$. The design advances prior and commercial solutions by achieving high bandwidth, high input impedance, and DC measurement capability within a compact, manufacturable PCB assembly. The circuit operates as a precision resistive divider with a nominal input impedance of $10\text{ M}\Omega$ and an attenuation ratio of 1000:1, realised through ten series-connected $1\text{ M}\Omega$ thin-film resistors forming the high-voltage arm and a single $10\text{ k}\Omega$ resistor constituting the low-voltage arm. The system directly scales the high-voltage potential relative to ground without galvanic isolation; both sections share a common reference, with safety ensured by high-impedance current limiting and PCB insulation spacing compliant with high-voltage standards. Frequency-dependent parasitic effects were compensated through a planar copper electrode implemented on the opposite PCB layer, separated by an FR-4 dielectric (dielectric strength $\approx 45\text{ kV mm}^{-1}$), thereby extending the linear bandwidth to 88 MHz. The assembly was housed within a grounded brass enclosure, providing electrostatic shielding, a stable reference potential, and a low-impedance discharge

path for leakage currents. The buffered output, limited to approximately 5 V, employed an LTC6268 operational amplifier and BUF634 high-speed driver to interface with 50Ω measurement instruments. This work demonstrates a practical PCB-based resistive divider capable of high-voltage DC measurement with excellent transient fidelity. For automotive HVDC telemetry, the study underpins the importance of distributed resistor networks to manage voltage stress and adequate creepage and clearance to prevent partial discharge. However, unlike grounded laboratory systems such as that investigated, typical EV high- and low-voltage domains float relative to each other, necessitating full signal and reference isolation to maintain operational safety under automotive conditions, as well as redundant protection in the event of resistor failure.

Sensing Topologies

Voltage can be measured by contact methods, such as aforementioned, or non-contact methods. Non-contact approaches attempt to infer voltage from surrounding electromagnetic fields. Inductive transformer techniques are limited to AC applications, while optical fibre sensors can measure DC but remain large, costly, and impractical for compact systems. Capacitive field-based techniques have demonstrated HVDC measurement capability, yet require fixed electrode spacing, auxiliary induction bars, and controlled environmental conditions, which constrain portability and retrofit use [26]. Commercial electric-field inference sensors—such as those by Advanced Energy—achieve non-contact operation but are bulky, expensive, and require grounding or manual calibration. Such systems have niche use in high-voltage instrumentation but offer limited applicability for compact EV applications.

For HVDC sensing in EV telematics, practical designs remain contact-based. The standard topology employs a resistive attenuation network, an ADC, and a galvanic isolation barrier. Integrated devices such as the LEM DVC1000-P combine scaling and isolation in a single package but remain limited by cost and size [27]. Consequently, discrete resistor chains, ADCs, and isolation components remain the preferred approach for compact and accurate implementations. Distributing the resistors across the high leg spreads voltage stress, increases dielectric distance per volt to reduce discharge risks, and limits current to safer touch levels as defined in IEC 60115-8 [28]. Using multiple elements also improves reliability by maintaining attenuation even if one resistor fails short. Resistor selection must consider tolerance, temperature coefficient, and rated voltage and current limits. The ADC must provide the required measurement accuracy (see Section 2.3.1), and the isolation interface must satisfy the insulation requirements discussed below.

High-Voltage Insulation Requirements

Insulation in high-voltage DC sensing involves selecting materials and geometries that prevent galvanically isolated electrical domains or conductive elements from forming unintended current paths. Effective insulation ensures both operator safety and long-term

reliability by maintaining adequate dielectric strength and resistance to surface tracking under sustained DC stress. IEC 62368-1 is a typical standard which links insulation strength to continuous working voltage, temporary overvoltages, and environmental factors such as pollution degree and material tracking resistance [29–32]. These and similar standards define minimum creepage and clearance distances for various applications and distinguish insulation types—basic, supplementary, and reinforced—according to protection level. In high-voltage measurement systems, adherence to such requirements ensures galvanic isolation integrity under automotive temperature, humidity, and transient conditions, thereby ensuring safety. Commercial certification requires adherence to multiple standards and independent safety audits.

High-Current Sensing

Current sensing in EV telematics varies approximately linearly with instantaneous bidirectional power flow, revealing charging states and load responses to mobility patterns.

Sensor Topologies

Current sensing topologies may be contact- or non-contact-based. Contact-based measurement employs shunt sensors, which use an inline low-impedance resistor to infer current from the voltage drop across it. These are excluded, as they require either an existing accessible shunt of precisely known resistance or invasive modification of wiring to insert one, conditions incompatible with a universal, non-invasive solution. Various non-contact topologies exist. Current transformers and Rogowski coils are unsuitable, as they can only measure AC. Fibre-optic sensors are technically applicable to HVDC but remain largely experimental and temperature dependent, with demonstrations such as [33] indicating potential rather than established commercial use. SQUID (Superconducting Quantum Interference Device) sensors can measure DC current but require cryogenic cooling and are impractical for embedded automotive applications [34]. The remaining non-contact options listed below are practically and commercially available sensors capable of measuring DC and require no permanent wiring modifications. These sensors can be used for both AC and DC applications, except single-core fluxgate types, which are limited to DC.

- **Hall-Effect Core Technologies (Open-Loop, OLH and Closed-Loop Hall, CLH):** Hall-effect sensors operate by sensing the magnetic field of a current-carrying conductor, as depicted in Figure 2.2. These sensors detect the magnetic field produced by a current-carrying conductor using a Hall plate biased by a constant current. The plate is positioned in the air gap of a ferromagnetic core, which concentrates magnetic flux orthogonal to the Hall bias current. The resulting Lorentz force induces a lateral Hall voltage that is amplified to yield a voltage output proportional to the primary current. Open-loop Hall sensors are compact and low-cost but suffer from saturation, mechanical-stress-induced and hysteresis-induced offsets, EMI, and

permeability drift, limiting accuracy to roughly 1–5 % [35,36] with appreciable offsets. In closed-loop Hall (CLH) variants, a compensation (nulling) coil drives the core flux toward zero, and the nulling current—measured across a burden resistor—forms the output. This zero-flux principle improves linearity and temperature stability, achieving accuracies of approximately 0.5–1 % [35,36].

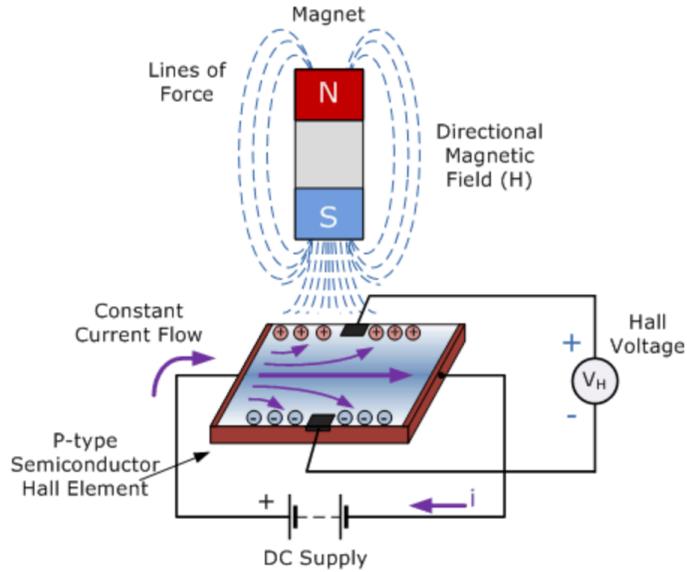


Figure 2.2: Hall-effect current sensing principle showing magnetic flux generated from the measured conductor interacting with the sensing element to produce a perpendicular Hall voltage. [2].

- **Fluxgate (FG) Sensors:** FG sensors employ an excitation coil to periodically drive a magnetic core into \pm saturation, while a fluxgate sensing element detects the primary-current-induced offset in this cycle using second-harmonic techniques. A nulling (compensation) coil is then driven to cancel the primary flux, and the resulting nulling current across a precision burden resistor represents the primary current. This topology achieves excellent linearity, with negligible core non-linearities arising from hysteresis, permeability drift, or temperature variation.

FG sensors are commonly realised in four main topologies, represented in Figure 2.3. The first is a single-core design with an air gap, similar to a CLH sensor except that the Hall probe is replaced by a saturable inductor, increasing linearity and accuracy. The second is a single-core design without an air gap, where the core itself functions as the FG sensing element; eliminating the gap improves EMC robustness and resolution. In both of these solutions, the excitation frequency is limited and the bandwidth restricted to DC applications by heating and drive losses. The third topology addresses this limitation by placing a dedicated AC current transformer next to the main core, allowing both AC and DC measurement capability without excessive loading. Finally, the balanced-core topology adds an

extra FG core positioned in opposition to the original, cancelling EMC coupling and topological offsets; this greatly improves DC accuracy and long-term stability. The first three topologies typically achieve accuracies of less than or equal to 0.5%, while balanced-core fluxgate designs can reach accuracies down to ~ 1 ppm. [3]

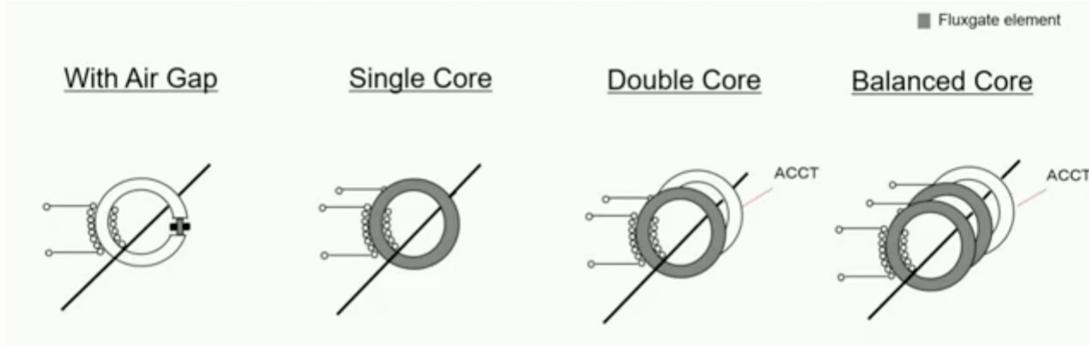


Figure 2.3: Fluxgate current sensor typical topologies [3].

- **Coreless Magnetic (Hall or magnetoresistive, MR) Sensors:** These sensors use an array of magnetic field sensors (Hall or MR) placed around the conductor to reconstruct current via Ampère's law and cancel external effects. [34]

MR sensors include anisotropic magnetoresistance (AMR), giant magnetoresistance (GMR), tunnelling magnetoresistance (TMR), and, in more specialised cases, colossal (CMR) and extraordinary magnetoresistance (EMR). All rely on the principle that an applied magnetic field alters the resistance of a thin-film structure, or quantum effects; detailed physics is omitted here to maintain the scope of this thesis. These differences determine their achievable accuracy, suitable operating environments, and the maximum field strength they can tolerate. They are typically significantly more accurate than Hall Effect sensors, but their operational conditions are uniquely limited. [34]

Without a magnetic core, Hall and MR sensors typically achieve $\sim 1\text{--}5\%$ accuracy and are sensitive to conductor angle and positioning relative to the sensor array. They exhibit low offsets and good EMC rejection, while the coreless design makes them lightweight, vibration-resistant, and clip-on, making them well suited to portable applications [34, 35].

Performance and Selection Parameters

The following requirements are relevant for choosing one of the magnetic-inferred current sensors discussed above [35]:

- **Primary Current Range:** Must measure up to $\pm 1\text{ kA}$ DC without saturation or over-range.

- **Isolation Rating:** Galvanically isolated sensors, as discussed above, do not typically require extensive additional high-voltage compliance since the buses are generally sufficiently insulated; however, providing such isolation remains beneficial as a redundancy measure or in cases where there is no insulation.
- **Accuracy:** Must account for linearity, offset, hysteresis, and temperature drift, and remain within specified limits.
- **Bandwidth and Output Data Rate:** Must support at least the Nyquist rate for the sensor bandwidth, and the output data rate must meet or exceed the main system's remote transmission rate.
- **Size and Durability:** Must withstand automotive vibration and not require disproportionate space relative to typical enclosure sizes.
- **DC Capability:** Must accurately measure DC current.

Temperature Sensing

Temperature influences EV battery efficiency, power output, and long-term health. Low temperatures increase internal resistance and reduce available range, whereas high temperatures enhance ion mobility but accelerate electrode degradation and Solid Electrolyte Interphase (SEI) growth, increasing resistance over time. Ambient and cabin temperature monitoring provides essential environmental context, with cabin placement preferred to minimise PCB self-heating bias. The following overview summarises common sensor types and Selection Criteria relevant to vehicular telemetry and embedded temperature measurement systems [37, 38].

Temperature Sensor Types

Key temperature-sensing technologies relevant to EV telemetry include:

- **RTDs:** Platinum RTDs (e.g., Pt100/Pt1000) provide high accuracy and long-term stability across wide temperature ranges (e.g., -250°C to $+750^{\circ}\text{C}$), though they require precision signal conditioning and are relatively costly.
- **Thermistors:** Inexpensive and highly sensitive within moderate temperature spans; nonlinear characteristics necessitate calibration or software linearisation, making them suitable where cost and compactness dominate.
- **Thermocouples:** Offer extremely wide temperature ranges and fast response, but offer lower sensitivity and require cold-junction compensation and signal conditioning.
- **Semiconductor (IC):** Bandgap or silicon-based devices with integrated ADCs, logic and memory, typically covering moderate temperature spans. They are compact,

low-power, durable, have good accuracy, and are digitally interfaced, allowing for long leads for remote sensing.

Selection Criteria

Selection of a temperature sensor should consider:

- **Operating Environment:** Sensors must withstand vibration, moisture, and either air or surface exposure depending on installation location.
- **Temperature Range:** Selected devices should include sufficient headroom beyond the expected environmental extremes, though this is rarely a limiting factor in automotive contexts.
- **Response Time:** Generally relatively negligible, as ambient or surface temperature changes in vehicular environments occur slowly relative to the sampling rate.
- **Accuracy:** The sensor must provide resolution sufficient to detect any appreciable variation in battery performance. At approximately 0.5 C discharge, a 5°C deviation from room temperature results in less than a 1%/°C change in available capacity, defining the threshold beyond which temperature effects become significant [39].
- **Durability:** The sensor must be robustly packaged to withstand automotive vibration and humidity, while allowing remote placement away from the PCB.

Acceleration Sensing

Accelerometers measure vehicle motion, including acceleration, braking, and turning. For telemetry applications, the focus is on low-frequency body motion rather than high-frequency vibration, as the latter depends strongly on suspension dynamics, tyre condition, and wear level, which vary between vehicles and over time, and are therefore excluded from scope. The following overview summarises the principal accelerometer types and the key selection criteria relevant to vehicular telemetry [40, 41].

Accelerometer Types

Representative accelerometer technologies for automotive telemetry include:

- **Piezoelectric (PE):** Provide wide frequency coverage—from a few Hz up to tens of kHz—making them suitable for shock and vibration measurements, but they are AC-coupled and cannot measure steady-state (DC) signals.
- **Piezoresistive (PR):** Offer DC response (down to 0 Hz), ideal for long-duration transients such as crash events, with wide bandwidth and mechanical robustness for transport applications.

- **Variable Capacitance (VC) / MEMS Capacitive:** Exhibit high sensitivity, good temperature stability, and DC response, making them well suited to low-frequency motion and steady acceleration measurements typical of vehicular dynamics.

Other accelerometer types (e.g., servo or laboratory-grade capacitive) offer higher precision but are impractical for embedded vehicular telemetry due to cost, size, and power constraints.

Selection Criteria

Accelerometer choice is guided by:

- **Dynamic range:** Full-scale range should exceed the expected peak accelerations to prevent clipping.
- **Bandwidth:** The passband must extend down to (or near) 0 Hz for automotive use, with an upper cutoff beyond the maximum expected driving accelerations. 5 Hz to 10 Hz is required to detect acceleration, braking, or harsh cornering [42].
- **Accuracy:** Overall accuracy depends on the combined effects of sensitivity, linearity, noise, bias, scale-factor error, non-linearity, and cross-axis coupling, which together determine measurement fidelity, stability, and repeatability across the operating range.

Analogue-to-Digital Converters (ADC)

Analogue-to-digital converters (ADCs) translate analogue sensor voltages into digital values for processing. The following discussion summarises the main ADC architectures and performance parameters relevant to telematics and precision measurement applications [43].

ADC Types

Common ADC architectures used in EV telemetry are:

- **SAR (Successive Approximation Register):** Provides moderate resolution (8–16 bits), very low latency (typically sub- μ s), and conversion speeds up to several MHz. Best for fast, deterministic sampling and easy input multiplexing; less effective than delta–sigma for maximum resolution in noisy, low-frequency signals.
- **Delta-Sigma:** Achieves very high resolution, low offset, and excellent linearity by oversampling (sampling in the high kHz to MHz range and outputting at much lower rates) with digital decimation filters. Decimation introduces a small latency, which is negligible for typical telematics systems.
- **Hybrid / Oversampling SAR:** Combines a SAR core with digital oversampling/averaging and optional on-chip filtering to raise effective resolution and reject noise,

providing a balance between the speed of SAR and the precision of delta-sigma architectures.

- **Other Architectures:** Dual-slope for very low rates with excellent noise rejection; flash/pipeline for very high throughput at lower resolution and higher power, used in RF or high-speed instrumentation. These are costly and provide no accuracy or performance benefits relevant to telematics applications.

Selection Criteria

ADC selection should account for:

- **Resolution:** Resolution determines the LSB size before accounting for internal limitations.
- **Input Range and Reference:** The input range must match the sensor swing, while internal references reduce layout complexity.
- **Acquisition Time:** Shorter acquisition times reduce kickback voltage and require lower input capacitance.
- **Sampling Rate:** Must exceed twice the signal bandwidth; oversampling ADCs improve SNR and transient accuracy.
- **Accuracy:** A wide range of errors affect ADCs, such as integral nonlinearity (INL), offset and gain errors, internal reference voltage errors, bias currents, input impedance and its fluctuations, differential current mismatches, and noise (which determines the maximum effective number of bits), as well as the temperature dependence of these factors. Detailed error budgets are required for ultra-precise applications.

2.3.2. Processing and Control

Microcontroller Unit (MCU)

The MCU is the central processing element of the telemetry system, responsible for acquiring, processing, and buffering sensor data, and managing communication with storage and wireless interfaces. Although FPGAs offer advantages such as high parallelism and nanosecond-level deterministic timing, their design complexity and cost make them impractical for most telematics applications. The MCU's primary role is to interface with and control peripheral devices. Compact, prebuilt MCU modules that include native USB ports and onboard regulators are typically preferred to simplify replacement in the event of failure.

The selection of a microcontroller family and its specifications involves evaluating architectures and design factors relevant to telemetry and embedded systems [44–46]. The following subsections outline common MCU families and key selection considerations.

Typical MCU Families

Representative microcontroller families suited to telemetry platforms include:

- **ARM Cortex-M Series (e.g., STM32, NXP Kinetis):** Industrial-grade MCUs offering low power consumption and extensive peripheral support (ADC, SPI, CAN, UART) with predictable real-time behaviour. They are well suited where deterministic timing, hardware reliability, and CAN integration are critical. These devices benefit from mature development ecosystems such as STM32CubeMX, CMSIS, and RTOS support, though their low-level coding model can increase firmware complexity in larger projects.
- **ESP / Arduino-Compatible MCUs:** ESP32-series and other Arduino-based devices integrate Wi-Fi/Bluetooth connectivity, ample RAM/flash, and broad open-source ecosystem support (Arduino, ESP-IDF), accelerating development. They are less rugged in extreme industrial environments and marginally less efficient than STM32-class MCUs.
- **PIC Microcontrollers (e.g., PIC18F, dsPIC):** Cost-optimised devices with moderate peripheral integration. Baseline models suit fixed-function or simple control tasks, whereas higher-end PIC32 devices provide more advanced peripherals (ADC, PWM, UART, SPI, I²C, CAN, USB) but remain limited for large-scale applications.
- **Raspberry Pi Pico (RP2040):** A dual-core Cortex-M0+ MCU featuring programmable I/O (PIO) for flexible digital interfaces; it provides high peripheral customisation at low cost.

Selection Criteria

MCU selection should consider:

- **Processing and Memory Resources:** The MCU must provide sufficient computational performance and memory to manage sensor polling, communication, and data buffering concurrently. Consider the available SRAM, optional PSRAM, Flash, and ROM capacity, as well as the clock frequency.
- **Peripherals:** Availability of interfaces such as SPI, I²C, UART, output clocks, and CAN. Ensure that the number of buses (and, for SPI, chip-select pins per bus) is adequate for all connected peripherals. A sufficient number of GPIO pins should also be available.
- **Development Toolchain:** Compatibility with established environments (e.g., Arduino, ESP-IDF, STM32Cube) to simplify firmware development and debugging.
- **Form Factor:** Compact, integrated modules with USB Type-C interfaces are preferred to minimise PCB footprint and simplify replacement during development.

Real-Time Clock (RTC)

A real-time clock provides a continuous timebase for timestamping, including when the system is powered down or when GNSS is unavailable. All RTCs exhibit drift, specified in parts-per-million (ppm) as the rate of error accumulation. Drift arises from temperature dependence, voltage dependence, long-term ageing, and mechanical stress of the resonator, and is inherently non-linear [47].

Selection Criteria

RTC selection should focus on:

- **Accuracy:** Lower ppm implies slower drift; temperature-compensated crystal oscillators (TCXO) offer superior stability compared with standard 32.768 kHz crystal designs.
- **Backup current:** Sub-microamp power usage maximises coin-cell lifespan to several years, often limited by cell self-discharge rates.

Drift Mitigation Strategies

The most effective approach is to resynchronise the RTC both periodically and upon startup following GNSS acquisition. Although retrospective correction can partially compensate for drift over time, it is not fully reliable because clock deviation is non-linear with temperature and supply variations [47]. Where precise time alignment is critical, samples should be discarded if GNSS synchronisation has not occurred recently enough to ensure the required accuracy. In less time-sensitive applications, minor timestamp offsets are typically acceptable.

Logic Level Converters

Logic level converters safely translate digital signals between devices operating at different voltage domains (e.g., 1.8 V and 3.3 V). They are essential when MCU, sensor, or bus logic levels are mismatched. The basic unidirectional level shifter is a simple resistive divider or diode clamp circuit. A more complex circuit involves MOSFET-based shifters, which can be used in various configurations and can be either unidirectional or bidirectional. However, both circuits degrade edge rates and therefore limit applicability at higher speeds, and they support only a single output type (pull-high or pull-low). MOSFETs also require additional ESD protection to avoid kV-level ESD events breaking through the oxide layer of the MOSFET. ICs overcome these limitations, offering ESD protection, high-speed operation, push-pull support, and smaller packages [48].

Selection Criteria

Suitable level shifter choice depends on:

- **Voltage range:** Must cover both high- and low-side levels (e.g., 5 V \leftrightarrow 3.3 V).

- **Directionality:** Match to the protocol – unidirectional for SPI, I²C clock, UART; bidirectional for I²C SDA or 1-Wire data.
- **Speed:** Bandwidth should exceed the target protocol's maximum expected implementation rate.
- **Signal type compatibility:** Ensure the converter matches the driver style – push-pull (e.g., SPI or UART) or open-drain (e.g., I²C or data-ready pins).

2.3.3. Communication Interfaces

CAN Bus Communication

This section outlines how the Controller Area Network (CAN) operates and how a telematics system can integrate safe access to it.

Background and Architecture

The CAN, standardised under ISO 11898, is a two-wire differential serial bus widely adopted across automotive systems. It enables distributed controllers and sensors to share data without a central host by means of priority-based arbitration, ensuring that high-priority messages are transmitted first. Communication occurs over a twisted differential pair (CAN_H, CAN_L) with 120Ω termination at each end to maintain impedance and suppress reflections. Nodes transmit only when the bus is idle; simultaneous attempts are resolved by bitwise arbitration, where the lowest identifier (highest priority) wins without collision.

Three protocol variants exist. *Classical CAN (CAN 2.0)* supports payloads up to 8 bytes at 1 Mbit s^{-1} and remains dominant across most EV subsystems. *CAN-FD* extends the payload to 64 bytes and supports data rates up to 8 Mbit s^{-1} , increasingly used for bandwidth-intensive systems such as battery management and drive control. *CAN-XL*, introduced in 2023, further expands the payload to 2048 bytes with data rates up to 20 Mbit s^{-1} , though it has not yet appeared in production EVs. Most current EVs still employ Classical CAN; newer models often integrate additional CAN-FD networks, while only two currently rely exclusively on CAN-FD [49]. Sensitive or high-speed buses may be encrypted or use proprietary formats, and there is no guarantee that such data is accessible [11]. In typical implementations, higher-layer protocols such as SAE J1939, CANopen, or OEM-specific standards define message structures, documented in a private DBC file.

Interfacing and Design Considerations

A functional CAN interface requires two components: a *transceiver* and a *controller*. The transceiver converts the small differential voltage between CAN_H and CAN_L into a digital bitstream, and in the opposite direction drives the differential lines for transmission.

The controller enforces the CAN protocol, handling arbitration, bit stuffing, CRC, and error detection, and exchanges parsed frames with the host MCU via SPI or a parallel interface.

For safe monitoring, systems must never drive the vehicle bus. This can be achieved by configuring the controller in listen-only mode, which disables transmission while still decoding messages, or by using contactless inductive probes that couple magnetically to the differential pair. The latter, shown in Figure 2.4, provides complete electrical isolation but does not yet support CAN-FD.

Key design parameters include supported bitrate (typically 125 kbit s^{-1} – 1 Mbit s^{-1} for Classical CAN and up to 8 Mbit s^{-1} for CAN-FD), buffer capacity for message queuing, availability of listen-only mode on the controller, and the transceiver's isolation rating where galvanic separation is required.

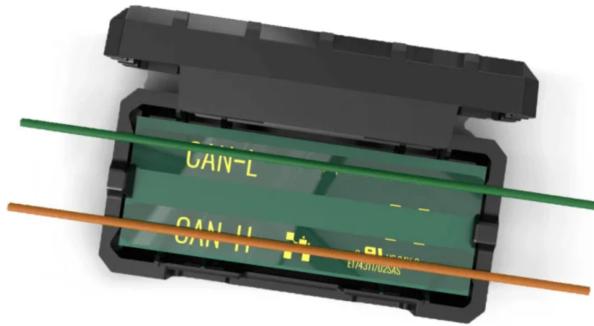


Figure 2.4: Inductive clip-on CAN bus probe for non-intrusive measurement [4].

Global Navigation Satellite Systems (GNSS)

GNSS is used to determine vehicle position, time, and speed using signals from at least four satellites. In EV applications, GNSS enables route tracking, usage profiling, and positional context for power consumption analysis. Major constellations include GPS (USA), GLONASS (Russia), Galileo (EU), and BeiDou (China). Studies show that mobility profiling performance degrades without accurate synchronisation with other onboard sensors, optimal antenna placement with an unobstructed sky view (preferably roof-mounted), multi-constellation redundancy, and update rates of 5 Hz to 10 Hz when acceleration estimation is required [42]. AT commands over UART are most commonly used for modem control (e.g. `AT+UGPS=1` to enable GNSS on some modules). While some commands are standardised, others are vendor-specific. The following considerations are informed by existing studies and guides [42, 50].

GNSS modem selection criteria

Critical GNSS modem parameters for vehicular deployments include:

- **Bands and Constellations:** Single- or dual-band; the latter (e.g. L1 + L5) enables

reception of two frequencies from the same satellite, allowing compensation for ionospheric delay and thereby achieving sub-metre accuracy. Multi-constellation capability further improves fix reliability in urban or obstructed environments.

- **Augmentation Systems:** Recreational-grade systems provides 1 m to 15 m accuracy; RTK and DGPS achieve cm–dm accuracy using external correction data, often with subscription fees.
- **Antenna Interface:** Modules support passive or active antennas; active types require a 3 V to 5 VDC bias via bias-tee to power the LNA. RF traces must maintain 50Ω controlled impedance (e.g. CPWG or microstrip) to minimise mismatch and loss.
- **Filters and LNAs:** SAW filters provide narrow band-pass filtering, while front-end LNAs improve SNR by amplifying weak GNSS signals before downconversion.
- **Sensitivity:** GNSS signals at the Earth’s surface are typically at worst -130 dBmW during clear conditions, but may degrade to as low as -150 dBmW under forest covered or highly developed urban conditions [51]. Ideally, acquisition sensitivity should be around -150 dBmW and tracking sensitivity ≤ -150 dBmW for robust operation under weak-signal and multipath conditions.
- **Update Rate:** Commonly 1 Hz to 25 Hz; higher rates are typically used for acceleration sensing.

Cellular Radio Access Technologies

In EV telemetry, cloud connectivity is typically provided by cellular networks due to their wide-area coverage. Radio Access Technologies (RATs) vary most significantly in throughput, antenna requirements, latency, coverage, and power usage. Modems are generally controlled by AT commands over UART (e.g. AT+CSQ for signal strength), with a mix of standardised and vendor-specific commands. This study does not consider ultra-low-power communication systems, such as LoRa, which are typically suited to low-bandwidth static applications in rural contexts. The trade-off of reduced data rate for energy efficiency is unnecessary in this case, as the system operates from the vehicle’s supply (several kilowatt-hours, typically recharged weekly).

Legacy Networks: 2G and 3G

2G introduced digital voice and basic data, while 3G added mobile broadband capability and improved peak efficiency. Both networks are informally scheduled for retirement in South Africa by 2027 [52], which limits their long-term relevance despite current availability. 2G is a useful fallback due to its low-frequency coverage but it is comparatively slow and power-inefficient.

4G LTE and Variants

LTE transitioned cellular systems to an all-Internet-Protocol (IP) architecture, where all communication including voice, data, and control is carried over IP packets. Cat 1 and Cat 1 bis remain highly practical for telemetry applications, supporting single-antenna, low-power designs with only marginally reduced edge coverage compared to dual-antenna implementations (2.5 dB–3 dB lower receiver sensitivity [53]). Cat 1 bis formalises single-antenna operation to reduce modem cost and power consumption while maintaining compatibility with Cat 1 networks. Higher LTE categories (Cat 4 and above) target broadband consumer use and are overprovisioned for telemetry.

5G NR

5G New Radio offers ultra-low latency and extremely high data rates, but the hardware remains costly and power-intensive, making it overprovisioned for telemetry. In South Africa, deployments are limited, and IoT coverage is not yet widespread.

Low-Power Wide-Area (LPWA) Variants: LTE-M and NB-IoT

Defined alongside 5G, but deployed on LTE networks, these low-power wide-area technologies reduce throughput in exchange for deep coverage and long battery life [54]. LTE-M is not deployed locally nor planned to be [55], while NB-IoT is offered by Vodacom on bands B3 and B8 but constrained by limited coverage and mobility support. Although NB2 (the enhanced version introduced in NB-IoT Rel-14/15) offers improved handover capability over NB1, it remains restricted at higher vehicular speeds and is not suited to fast-moving objects [56]. Its latency is also between 1.6 and 10 s, which limits real-time application [57].

RAT comparison

Table 2.1 summarises the main cellular RATs relevant to EV telemetry, with performance figures for peak latency and downlink/uplink (DL/UL) speeds drawn from [56–63]. Coverage information for South Africa is taken from [52]. NB-IoT corresponds to LTE Cat-NB2, while LTE-M refers to LTE Cat-M1.

Table 2.1: Comparison of the main cellular radio access technologies (RATs), including throughput, latency, coverage, and current characteristics.

RAT	DL/UL	Latency	SA Coverage	Current
2G	100 kbit s ⁻¹ to 400 kbit s ⁻¹	300 ms to 1000 ms	Significant	2.0 A
3G	0.5 Mbit s ⁻¹ to 5 Mbit s ⁻¹	100 ms to 500 ms	Significant	0.8 A
LTE Cat 1	10 / 5 Mbit s ⁻¹	< 100 ms	Significant	0.8 A
LTE Cat 1 bis	10 / 5 Mbit s ⁻¹	< 100 ms	Significant	0.5 A
LTE Cat 4+	≥150 / 50 Mbit s ⁻¹	< 100 ms	Significant	0.8 A
5G NR	Up to Gbps	< 1 ms	Major cities	1.1 A
LTE-M	1200 / 375 kbit s ⁻¹	50 ms to 100 ms	None	0.4 A
NB-IoT	140 / 125 kbit s ⁻¹	1.6 s to 10 s	Significant	0.3 A

Relevance in telematics systems

Low latency and seamless handover are important in telematics to ensure real-time performance and uninterrupted connectivity at vehicular speeds. While multi-antenna configurations can marginally improve edge coverage, they add cost and complexity, with 5G requiring multiple antennas. LTE-M, NB-IoT, LTE Cat 1, LTE Cat 1 bis, as well as legacy 2G and 3G, operate with single-antenna designs.

Although 2G and 3G remain useful fallback options due to their longer range per unit transmit power [53], their high latencies and imminent decommissioning make them unsuitable as primary solutions [52]. NB-IoT, with latency between 1.6 and 10 s and poor handover performance at high speeds, is likewise unsuited to real-time automotive applications [52, 56]. LTE-M is neither commercially deployed nor planned for deployment in South Africa [55]. By contrast, LTE Cat 1 (when operated in single-antenna mode) and Cat 1 bis offer a practical balance of latency, throughput, and coverage while retaining the simplicity of a single-antenna design [53].

Antennas

Both the GNSS and cellular modem subsystems rely on external antennas to ensure efficient signal transmission and reception. Antenna performance depends on several measurable properties that determine signal quality and reliability. Frequency coverage defines the range of bands the antenna can effectively transmit or receive, ensuring compatibility with the intended radio systems. The gain pattern describes directional sensitivity, with peak gain indicating the direction of strongest radiation or reception; the ideal isotropic antenna has 0 dB gain, representing perfect omnidirectionality. Radiation efficiency and return loss quantify how effectively input power is radiated versus reflected, and the antenna should be impedance matched (typically to a 50Ω system). These parameters are often band-dependent, varying with frequency and shielding or ground-plane conditions. Polarisation alignment determines how efficiently signals are transferred between transmitting and receiving antennas. For instance, GNSS satellite transmitters typically implement right-hand circular polarisation (RHCP), while cellular base stations employ linear polarisation; antennas should match this to minimise coupling loss. The most critical parameters in practical selection are frequency bands, radiation efficiency (overall power conversion), return loss (matching quality), and gain pattern (directional coverage), all evaluated relative to expected inbound signal strength and receiver sensitivity [64].

SD Card Interface

SD cards are widely employed in telemetry systems due to their endurance, compact form factor, removability, compatibility with standard computers, low cost, high storage density, simplicity, and fast transfer rates. Alternative storage options (e.g., SSD, EEPROM, NOR flash, FRAM, or soldered eMMC) lack this combination of advantages – particularly

flexibility – and are therefore not considered in detail.

When selecting an SD card for telemetry, several factors must be evaluated to ensure long-term reliability. Most cards provide sufficient bandwidth for logging, support a widely compatible filesystem (typically FAT32), and offer adequate endurance for the expected service life. The main considerations are selecting appropriate capacity for long-term data retention and ensuring reliability through industrial-grade cards with built-in power-failure protection to maintain data integrity during shutdowns.

2.3.4. Power, Layout, and Mechanical Considerations

Power Regulation

Voltage regulation keeps each rail within limits despite input and load changes. In automotive telematics the supply must tolerate wide input swings and transients while preventing switching noise from corrupting analogue, RF, and ADC rails.

Regulator Classes

The principal voltage-regulator categories are:

- **Switch-mode regulators.** A high-speed switch and a control loop modulate duty cycle so that energy alternately flows into an LC circuit, producing the desired average output with high efficiency. A tight current-loop layout, shielding and distancing components, and low-impedance grounding prevent this ripple and its harmonics from coupling into sensitive domains. [65]
 - *Non-isolated topologies.* The main DC-DC converter types are buck, boost, inverting buck-boost, and non-inverting buck-boost. Each uses an inductor and switching element to step voltage down, up, invert it, or a combination thereof, depending on duty cycle. Modern designs typically use synchronous MOSFET switching over diodes for higher efficiency [65].
 - *Isolated topologies.* These use a transformer to provide galvanic isolation between input and output, improving safety and noise immunity. Common variants include flyback, forward, and bridge converters, which differ mainly in power level and efficiency trade-offs [65].
- **Linear regulators (LDO).** A series transistor, controlled by an error amplifier, maintains a constant output voltage. LDOs offer low noise and high power-supply rejection, making them suitable for sensitive analogue and RF rails; however, their efficiency decreases sharply with voltage drop and load current, resulting in significant heat dissipation that requires careful management [66].

Selection Criteria

Regulator selection should consider:

- **Input Voltage Range:** Must tolerate nominal and transient voltages without loss of regulation.
- **Efficiency:** Critical for battery-powered and thermally constrained enclosures.
- **Minimum Load:** Ensure stable regulation under lowest expected load; important for some isolated topologies.
- **Accuracy and Stability:** Maintain target voltage over load and line variations; assess tolerance, ripple, load/line regulation.
- **Load Transient Response:** Must sustain output during bursts (e.g., LTE transmit peaks) without excessive droop or overshoot.
- **Noise and Ripple:** Low-noise outputs are essential for GNSS, LTE, and precision ADC rails; filtering and post-regulation with LDOs are common.
- **Size and Integration:** Minimise footprint and external component count where possible.
- **Thermal Performance:** Low thermal coefficients and exposed-pad packages requirements help limit thermal buildup in sealed enclosures.

PCB Layout Guidelines

This section outlines key layout considerations for designing a reliable and functional PCB for an EV telematics system. It focuses on power integrity, signal integrity, grounding, thermal management, and EMI mitigation strategies relevant to typical EV telematics applications.

Layer Stack-up

In multilayer PCB designs, dedicated power and ground planes provide low-impedance return paths and excellent noise isolation between subsystems. Two-layer boards, however, are more cost-effective and commonly used for compact or low-power designs. In such layouts, a continuous ground plane is typically placed on the bottom layer, while the top layer is used for signal and power routing. Orthogonal routing between adjacent layers (e.g., vertical on the top, horizontal on the bottom) reduces crosstalk and simplifies layout. Limiting components on ground pours minimises loop areas, though small elements such as inline SPI resistors may be used sparingly.

RC Filtering

RC filters attenuate unwanted high-frequency components according to

$$f_c = \frac{1}{2\pi RC}$$

where R limits current flow and C limits the rate of voltage change. The selected component values define the cutoff frequency, which must balance noise suppression against signal bandwidth. Cascaded RC stages or active filters are used where steeper roll-off is required to attenuate out-of-band interference under low signal-to-noise conditions.

Bypass Capacitors

Power traces to integrated circuits (ICs) exhibit parasitic resistance and inductance, leading to voltage dips during switching events. Bypass capacitors should therefore be placed as close as possible to IC supply pins to provide local charge storage and maintain rail stability. Typical practice is:

- 10 μF for bulk energy storage against low-frequency fluctuations.
- 100 nF for high-frequency transients and switching noise.
- Additional capacitors for highly sensitive ICs.

Ceramic capacitors with low ESR are preferred for fast transient response. Voltage ratings should exceed the maximum rail voltage, with a typical design margin of 1.5 \times .

Grounding Strategy

A continuous, unbroken ground plane on the bottom layer minimises return impedance. Sensitive analogue and digital domains may require local separation zones in mixed-signal designs. When a component includes both analogue and digital sections, each portion should be placed above its corresponding ground region so that return currents remain confined. Except for specific cases, star grounding (where domains connect at a single central point) is not recommended, as it increases loop area and susceptibility to interference if signals cross the gap [67].

Signal Integrity

To maintain signal integrity:

- Series resistors on fast digital lines (e.g., SPI) reduce edge rates and damp reflections.
- Differential pair routing with controlled impedance is used to minimise reflections and noise arising from impedance or length mismatches (e.g., CAN, USB, high-speed ADC links).
- Avoid stubs and minimise vias on sensitive or high-speed traces.

Thermal and Current Management

High-current paths are routed with wide traces or copper pours to reduce resistive heating and voltage drop. Thermal vias beneath heat-generating components (e.g., regulators) spread heat into internal or rear copper areas. Trace widths and via sizes are determined using IPC-2152-compliant calculators, based on current-carrying requirements, PCB material parameters, and permissible temperature rise [68].

EMI and EMC Practices

To further minimise electromagnetic interference:

- Use compact high-current loops, shielded inductors, low-ESR capacitors, and input/output filtering for regulator design.
- Ferrite beads can filter high-frequency noise, but excessive series inductance limits transient current delivery, requiring larger supply capacitance to maintain voltage stability.
- Physically separate noisy elements (e.g., inductors, switching regulators) from sensitive analogue circuits.
- Keep high-frequency loop areas as small as possible.

Isolation

Isolation in signal interfaces preserves integrity by preventing ground loops, common-mode noise coupling, and cross-domain interference. It is typically implemented using digital or analogue isolation components, which enable accurate data transfer between domains while maintaining independent ground references to maximise the signal-to-noise ratio of sensitive analogue-to-digital conversions.

Power Usage Considerations

Power consumption is an important consideration in EV telemetry systems, as all circuitry imposes a parasitic load on the vehicle battery. Sealed enclosures, required to prevent condensation and contamination, further increase the risk of thermal buildup. HVDC battery packs typically range from 30 kWh to 100 kWh [69], so even a continuous 30 W drain (approximately 1000 h to full discharge) is negligible unless the vehicle remains unused for several weeks. Parasitic load can be minimised by wiring the system supply to an ignition-controlled line and ensuring that passive subsystems draw negligible current. The primary motivation for limiting power consumption is therefore thermal management, achieved through efficient component selection and careful PCB layout.

Enclosure

Enclosures for automotive electronics must provide ingress protection while accommodating connectors, cable glands, and mounting features. The choice of enclosure depends on the required IP rating, mechanical robustness, and insulation properties, with the IP rating being most critical to prevent condensation. Penetrations such as cable entries or ventilation slots represent potential weak points for dust or moisture ingress; therefore, glands and seals are employed to maintain enclosure integrity while allowing reliable cable pass-through. For high-voltage circuits, standards such as IEC 62368-1 require that openings and clearances preserve safety barriers against user contact. Enclosure geometry and mounting arrangements also influence vibration resistance and long-term reliability in automotive environments.

Summary

The review of sensing, processing, storage, and communication hardware components shows how each contributes distinct constraints and opportunities in an electric vehicle telematics system design. Understanding their parameters and PCB-level interactions enables informed selection of solutions, while ensuring accuracy and robustness in demanding automotive environments. This broad foundation supports the more detailed design choices that follow in subsequent sections.

2.4. Cloud Architectures

Cloud computing forms the foundation of modern telematics infrastructures by enabling the ingestion, processing, storage, and retrieval of vehicular data without the need for locally managed servers [70]. Its elasticity allows computational and storage resources to scale dynamically with data volume, while global accessibility and managed security frameworks make it well-suited to distributed fleets.

This section reviews prerequisite literature on cloud-based infrastructures supporting EV and IoT telemetry. It examines representative studies and implementations describing messaging and data transmission, computation and analytics, storage management, and security and identity control, highlighting operational considerations such as latency, monitoring, provider selection, and cost that inform the architecture adopted in this work.

Messaging and Data Transmission

Messaging services enable bidirectional communication between devices, cloud components, and web application clients. Communication between cloud components is handled internally by the cloud provider, while design choices must be made for the device-to-cloud and cloud-to-application links.

For **device-to-cloud transmission**, lightweight publish-subscribe protocols such as **MQTT** are widely adopted because they provide low message overhead, configurable Quality of Service (QoS) levels, and resilience on unreliable mobile links [71]. An alternative is **CoAP**, which runs over UDP and achieves even lower overhead but offers weaker delivery guarantees, does not buffer until success, and is not typically supported by IoT cloud services [72]. Transmission strategies can be either *streaming*, where each data record is forwarded with minimal delay, or *batching*, where multiple records are buffered before upload. Hybrid approaches are often used to ensure streaming of real-time data while batching historical logs to reduce system load and bandwidth cost.

For **cloud-to-application communication**, the dominant mechanisms are **HTTPS request-response APIs** and **WebSockets**. HTTPS APIs are well-suited for asynchronous or low-rate requests, since each connection is created and torn down per request. By contrast, WebSockets upgrade an HTTP connection into a persistent, always-open channel that supports low-latency delivery of continuous updates [73]. Many cloud services offer MQTT bridging to WebSockets, allowing telemetry to be pushed in real time directly from ingestion to frontend dashboards with negligible delay. While WebSockets are advantageous for dashboards requiring this update rate, they increase server resource load significantly and may not operate correctly on networks with strict firewall policies. The choice between HTTPS and WebSockets is primarily governed by the application's need for sub-second updates.

Compute and Processing Layer

Cloud compute resources replace fixed backend servers with on-demand execution environments. Event-driven *serverless functions* execute stateless processing in response to triggers, enabling payload validation, transformation, enrichment, and conditional routing without the cost of continuous hosting [74, 75]. Container-based services allow stateful or dependency-specific workloads while retaining elastic scaling [76]. For most telematics systems, serverless functions are preferred for their simplicity and cost-effectiveness. These processing layers are key to data processing and communication between cloud services, and can also initiate analytics jobs, generate downsampled datasets, or trigger alerting workflows [74].

2.4.1. Storage

Storage systems provide the persistent retention of data and define how it can be organised, retrieved, and queried. Block and object storage form the fundamental models, while time-series, NoSQL, and relational databases are higher-level systems typically built on top of these. The main systems used are:

- **Object storage:** Stores data as discrete objects (e.g., CSV logs) within a flat

namespace. It is ideal for large, immutable (not editable) archives and is typically used for retrieving extensive raw historical datasets.

- **Block storage:** Divides data into fixed-size blocks and supports file systems, making it efficient for operating systems and small random reads. However, it is not typically relevant for telemetry applications.
- **Time-series databases:** Specialised for timestamped data such as sensor readings, enabling efficient queries of both recent values and long-term trends. They support advanced time-based operations (e.g., aggregation, transformation) and are widely used for real-time dashboards and vehicle tracking.
- **NoSQL stores:** Schema-flexible databases suited to metadata or telemetry with variable structures across devices. They support real-time operations but offer limited scalability and efficiency for complex, time-based queries compared to dedicated time-series databases.
- **Relational databases:** Structured, table-based systems optimised for complex queries, multi-table joins, and analytical workloads. These are typically used for offline data analytics and management rather than real-time telemetry.

2.4.2. Security and Identity Management

Security mechanisms ensure that only authorised devices and users can send or access data, and that information remains protected both in transit and at rest. Device authentication commonly relies on unique X.509 certificates with mutual TLS [77, 78], while user authentication is typically token-based over TLS. Least-privilege access control further restricts each role or service to the minimum required permissions, reducing the risk that a compromise can escalate beyond limited access [79]. Cloud providers implement these mechanisms as managed services, making them straightforward to integrate across devices and applications.

2.4.3. Deployment and Operation

Provider choice, latency, monitoring, and autoscaling govern system performance, reliability, and cost.

Cloud Provider Landscape

The main commercial cloud providers—Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)—offer broadly similar categories of services that are sufficient to enable various telematics use cases. They differ in niche applications, beginner-friendliness, number of regions, and ecosystem maturity. AWS is the market leader,

offering the largest community, widest range of services, and most extensive server coverage; however, its pricing models are often unpredictable. Microsoft Azure provides comparable capabilities but excels in seamless integration with other Microsoft applications. Finally, GCP offers fewer services overall but stands out in the machine learning domain for both cost and performance [80]. All three providers are viable for telematics deployments, though the most appropriate choice depends on factors such as required services, latency across regions, system integration effort, and long-term cost predictability.

Latency and Quality of Service

End-to-end delay in cloud telemetry arises from wireless uplinks, broker processing, compute execution, and database access. Although cloud providers replicate services across regions to minimise downtime, latency remains strongly dependent on the geographic location of data ingestion and querying [80]. For real-time telemetry, workloads should be deployed in regions geographically close to the fleet to reduce path delay, which is fundamentally limited by signal propagation time. To achieve sub-second responsiveness, MQTT over WebSockets can be used, whereas routing data only to storage via HTTP requests offers lower cost but higher latency.

QoS levels are defined for MQTT transmitters and receivers. QoS 0 provides best-effort delivery without acknowledgement; QoS 1 ensures at-least-once delivery through retries and local buffering but may introduce duplicates; and QoS 2 guarantees exactly-once delivery by verifying that a message has not already been processed before transmission, at the expense of additional overhead. The choice of QoS level represents a trade-off between latency, tolerance for duplication, and the requirement for assured delivery. Telematics systems typically employ QoS 1, as they require confirmation of delivery but can tolerate and filter duplicates using timestamps.

Scalability, Observability, and Cost

Scalability and observability are central to maintaining reliable cloud-based telemetry systems. Observability involves configuring dashboards and alerts to track message throughput, ingestion failures, and system latency. Cloud monitoring platforms provide visibility into metrics such as dropped packets, queue depth, API error rates, and database throttling, while logging pipelines capture device identifiers, timestamps, and error codes for post-event analysis [81]. This observability framework becomes critical when operating at scale or serving large client bases.

Cloud pricing models typically scale with the number of transmitted messages and the volume of compute and storage consumed [82]. In telematics systems, costs are primarily driven by message frequency, payload size, and data query volume. Efficient service selection should therefore minimise recurring costs while maintaining reliability. Techniques such as downsampling reduce data volume and storage costs, while autoscaling

dynamically allocates resources to match workload demand. However, autoscaling can slightly reduce performance and may result in uncontrolled cost growth if not properly configured.

Summary

Cloud computing forms the foundation of scalable and secure EV telematics, enabling efficient data ingestion, processing, storage, and access without reliance on local infrastructure. Effective cloud architecture must address elasticity, managed security, provider selection, latency, monitoring, and cost to ensure reliable and economical fleet telemetry operations.

2.5. Web Application Architectures

Web applications serve as the interface through which users engage with cloud-based telemetry data, transforming raw logs and APIs into accessible dashboards and analytical visualisations. They comprise two functional layers: the frontend, responsible for rendering visual elements and user input, and the backend, which mediates cloud communication, authentication, and data access.

This section reviews prerequisite literature on web application architectures relevant to telematics systems. It examines representative studies and frameworks addressing interface development, hosting and rendering strategies, performance optimisation, accessibility, and security, identifying approaches most applicable to scalable and responsive EV dashboard design.

2.5.1. Interface Development

Low-code and drag-and-drop platforms such as WordPress are often adopted for accessibility to development, efficiency, and maintainability; however, they remain constrained in areas of customization, integration, and performance when handling complex or high-load scenarios [83]. At the other extreme, directly writing in the lowest level HTML, CSS, and JavaScript provides maximum flexibility but becomes difficult to maintain in large-scale, high-feature systems, since developers must manually handle state, events, and DOM updates.

Most production dashboards are therefore built using JavaScript frameworks such as React, Angular, or Vue. These frameworks provide abstractions for building modular components that update efficiently when underlying data changes. React in particular has become the dominant choice for data-heavy dashboards, as confirmed by academic surveys of framework adoption and performance [84, 85]. Its component-based structure and virtual DOM allow developers to build reusable elements and render frequent updates

with minimal performance penalties. As noted by [85], selecting the most widely adopted framework—currently React—can be advantageous for most typical applications, as this ensures availability of open-source components and extensive support communities.

Meta-frameworks such as Next.js (based on React) and Nuxt.js (based on Vue) extend these libraries by integrating server routing, build pipelines, and rendering strategies. This reduces developer backend integration effort while improving maintainability and scalability, which is especially valuable for complex telemetry platforms [86].

2.5.2. Deployment and Performance

This subsection covers the deployment, hosting, rendering models, and performance considerations essential for delivering scalable and efficient dashboards.

Hosting and Domain Management

Once a web application is developed, it must be deployed for user access. Traditional file servers or virtual machines can serve static assets but provide limited resilience and scalability. In contrast, managed hosting platforms such as Vercel, Netlify, or AWS Amplify automate scaling, redundancy, and global delivery through Content Delivery Networks (CDNs). A CDN distributes cached content across geographically dispersed nodes, reducing latency for users regardless of location [87].

These platforms also streamline domain and security configuration by automatically provisioning and renewing TLS certificates, ensuring that all dashboard communication remains encrypted without manual intervention [88, 89]. Authentication tokens such as JSON Web Tokens (JWTs) are issued by the identity provider (e.g., AWS Cognito) and stored client-side—typically in secure, HTTP-only cookies or local session storage—while the hosting platform merely forwards them during authenticated requests. This ensures that sensitive credentials are never persisted on the hosting server itself, and that API calls to the cloud backend originate from verified users. For telematics systems requiring secure and reliable remote access, this approach reduces operational overhead while maintaining consistent user experience. The hosting region should be selected to minimise the additional network path introduced by routing requests through the hosting server between the user’s browser and the cloud backend.

Rendering Models

Different rendering strategies are suited to different dashboard requirements. Client-Side Rendering (CSR) constructs the interface in the browser after the initial page load, supporting interactive real-time charts but increasing time-to-first-paint and depending on client device performance [86, 90]. Server-Side Rendering (SSR) sends pre-rendered HTML from the server by rendering components and fetching data prior to transmission, thereby

reducing initial latency (since the request does not originate from a browser, minimising path and handshake delay) and improving accessibility across devices. Hydration logic is then added to make pages interactive [86, 90].

Static Site Generation (SSG) creates pages at build time and serves them as static files. It is suitable for non-dynamic elements such as buttons or layout components that do not require re-rendering per request (with CSR adding interactivity where required). Incremental Static Regeneration (ISR), a hybrid approach, regenerates static content on-demand or at defined intervals [91]; for example, it may be used for daily updates of aggregated telemetry metrics such as vehicle trip lengths. In practice, these rendering models are often combined—CSR for live charts and user input, SSR for faster initial loading, and ISR for periodic data refresh—to optimise both performance and interactivity in telemetry dashboards.

Performance and Data Efficiency

Performance in telematics dashboards depends on update frequency and delivery efficiency. For sub-second updates, WebSockets provide a persistent, push-capable channel with significantly lower overhead than repeated HTTP requests [73, 92, 93]. In practice, these are often linked directly to MQTT streams, minimising latency. For applications with lower refresh demands, HTTP polling offers a simpler and more cost-effective alternative.

On the frontend, frameworks such as React enhance efficiency by re-rendering only components whose data has changed rather than rebuilding the full page [84]. For instance, when only a vehicle's location marker changes, the affected component is re-rendered rather than the entire map. Page transitions (e.g., map view to time-series chart) also occur without a full reload, reducing reflows and supporting smooth visualisation even under frequent updates.

2.5.3. Usability and Security

This subsection examines how dashboards remain accessible and secure across devices and user groups, addressing both inclusivity and protection of user data.

Accessibility Across Devices

Telemetry dashboards are accessed in varied contexts, from control-centre desktops to tablets in garages and mobile devices in the field. Responsive Web Design principles ensure that layouts adapt seamlessly to different screen sizes, resolutions, and orientations [94]. Applications should also adjust data density and resolution according to device sizes to prevent overloading smaller screens.

Accessibility guidelines such as WCAG specify requirements like minimum text-to-background contrast to ensure dashboards remain usable for people with reduced vision

[95]. Using proper semantic HTML elements enables screen readers to interpret the structure of the interface, and ARIA attributes are used to make non-standard components understandable to users who rely on assistive technologies (e.g., audio narration of the component) [96]. Following these practices is particularly useful as dashboards scale, helping to maintain inclusivity and broad usability when the user base becomes large.

Web Application Security

In the client layer, the primary security requirement is to ensure that only authorised users gain access and that their sessions cannot be hijacked or misused. In practice, this is achieved through session tokens. When a user authenticates via Amazon Cognito, the service issues tokens that are stored as cookies in the browser. These cookies are marked with the `HttpOnly`, `Secure`, and `SameSite` flags: the `HttpOnly` attribute prevents JavaScript from reading the token, the `Secure` attribute ensures transmission only over HTTPS, and `SameSite` restricts cross-site request inclusion—together mitigating XSS, interception, and CSRF risks [97–101]. The browser automatically attaches these cookies to subsequent requests, allowing the cloud service to validate the session without exposing sensitive credentials. Three token types are employed: the access token, which authorises API requests; the ID token, which conveys user identity information; and the refresh token, which allows new tokens to be issued once the others expire. Because access and ID tokens are short-lived, their theft yields limited value, while logout actions revoke the refresh token to terminate future access [98].

Beyond authentication, additional mechanisms harden the interaction between client and server. All traffic is served over HTTPS, enforced by the hosting platform (such as Vercel or CloudFront), ensuring confidentiality and integrity of communication [87–89]. Cross-Origin Resource Sharing (CORS) is configured at the API layer to restrict which domains may issue requests, mitigating cross-site request forgery [102]. A Content Security Policy (CSP), delivered by the application server, defines which resources the browser may execute or load, blocking unauthorised scripts and reducing the attack surface for code injection [103]. Finally, the application or its hosting configuration can include the `frame-ancestors` directive, which instructs the browser not to display the site inside untrusted iframes—thus preventing clickjacking (an attack that tricks users into unintended clicks)—and is recommended in modern practice, with fallback `X-Frame-Options` headers for broader compatibility [104–106]. Together these measures establish a layered defence that secures both authentication and all subsequent interactions between client and server.

2.5.4. EV Dashboard Content

The data displayed on the dashboard must be directly aligned with user requirements. For example, fleet management platforms commonly track state-of-charge (SoC) in real time,

issuing alerts when vehicles are not charging or when the remaining range is insufficient for upcoming trips [107,108]. Researchers, by contrast, may require access to time-series charts, real-time vehicle locations overlaid on maps with associated metrics, and the ability to download logs [107]. It is also advantageous to support remote command execution without requiring physical access to the hardware or reflashing of firmware. Such commands may include adjusting programmable gain amplifier (PGA) settings to accommodate vehicles with varying operating ranges, or updating system certificates to maintain authentication validity.

Summary

Web applications provide the user-facing layer for EV telematics, translating raw data into accessible, secure, and responsive dashboards. Effective application design must prioritise accessibility, robust security, and real-time data visualisation using appropriate frameworks and rendering strategies, while considering user context, update frequency, and device diversity to optimise usability and performance.

Chapter 3

Design

This chapter outlines the requirements and subsequent design of the telematics system to achieve the system goal. It outlines the stakeholder and functional requirements, followed by the design of the hardware, firmware, cloud system, and frontend system.

3.1. Requirements

To ensure the product aligns with its intended use-case, the project must derive its requirements from key stakeholders and translate them into corresponding functional and non-functional specifications.

Stakeholder Requirements

Stakeholder requirements define what the system must achieve to satisfy the practical needs of its end-users and the broader research context. In this project, these requirements capture the needs of researchers studying the operation of electric vehicles where internal CAN data access is unavailable.

The primary representative stakeholder is the Stellenbosch University Electric Mobility Laboratory, which operates an electrified minibus taxi but lacks access to its internal CAN communications. The sensing tool must facilitate data acquisition from this electric vehicle as well as others used in various research contexts.

The key stakeholder requirements included power usage and mobility parameters—specifically position and speed profiles. Together, these enable empirical derivation of key insights such as duty cycles, charging demand, and fuel efficiency, which are critical for evaluating vehicle viability and infrastructure planning. Additional sensing parameters, such as ambient temperature and acceleration, were included to capture secondary influences on performance. Additionally, a CAN acquisition interface was required to support contexts where partial or complete access to vehicle CAN data is available. Cellular connectivity and battery-backed timekeeping were required to support reliable data synchronisation and remote analysis through cloud infrastructure and a secure web interface. These requirements are summarised in Table 3.1.

Table 3.1: Stakeholder requirements (SR) capturing sensing, data handling, and retrofit needs for the telematics platform.

ID	Requirement
Sensing	
SR1.1	The system shall directly sense a comprehensive range of parameters sufficient to profile electromobility and environmental characteristics.
SR1.2	The system shall achieve measurement accuracy sufficient for credible electromobility analysis.
SR1.3	The system shall sense and log data at a time resolution sufficient for fleet-level electromobility profiling.
SR1.4	The system shall support CAN bus reading for redundancy.
Data Handling	
SR2.1	The system shall store all measurement data both locally and remotely.
SR2.2	The system shall report data remotely in near real-time.
SR2.3	The system shall retransmit data not transmitted due to connectivity outage.
SR2.4	The system shall allow real-time and historical data to be visualised remotely via a secure web interface.
SR2.5	The system shall maintain nationwide connectivity and support operation at expected vehicular speeds.
SR2.6	The system shall maintain time during power-loss.
Retrofitability, Generalisability and Robustness	
SR3.1	The system shall be powered from the vehicle's 12 V supply.
SR3.2	The system shall not require appreciable per-device or per-vehicle calibration.
SR3.3	The system shall not require permanent alterations to the host vehicle's electronic systems.
SR3.4	The system shall not affect the operation of the vehicle's internal communications (CAN bus non-intrusiveness).
SR3.5	The system shall be electrically safe when interfacing with any power systems on the vehicle.
SR3.6	The system shall be robust to vibration experienced by vehicles.
SR3.7	The system shall minimise parasitic load to avoid significant battery drain during prolonged vehicle inactivity.
SR3.8	The system shall ensure secure handling of all remote interfaces, including gateways, cloud services, and frontend applications.

Functional Requirements

The functional requirements address sensing, logging, communication, and safety needs for EV telematics in the South African context. They are grouped into **Sensing**, **Data Handling**, and **Retrofit, Generalisability and Robustness**. Table 3.2 summarises the functional requirements, which are derived from the stakeholder requirements in Table 3.1.

Table 3.2: Functional requirements (FR) derived from system-level requirements, grouped by sensing, data handling, and robustness.

ID	Requirement	Derived from
Sensing		
FR1.1	Measure HVDC pack voltage up to 880 V with FSR error $\leq 1\%$.	SR1.2
FR1.2	Measure bidirectional HVDC pack current up to $\pm 1 \text{ kA}$ with FSR error $\pm 1.5\%$.	SR1.2
FR1.3	Acquire GNSS position with 3D position accuracy $\leq 10 \text{ m}$ and velocity accuracy $\pm 5 \text{ km h}^{-1}$.	SR1.1
FR1.4	Acquire 3-axis drivetrain acceleration up to $\pm 20 \text{ m s}^{-2}$ (approximately $2g$) with FSR error $\pm 10\%$.	SR1.1
FR1.5	Acquire ambient temperature from -5 to $+55^\circ\text{C}$ with resolution $\pm 1^\circ\text{C}$.	SR1.1
FR1.6	Acquire vehicle CAN frames.	SR1.4
Data Handling and Communication		
FR2.1	Log all outputs locally to removable, PC-readable storage at 1 Hz.	SR2.1
FR2.2	Transmit all measurement channels to a cloud endpoint at 1 Hz when available.	SR2.2
FR2.3	Ensure all missed historical data is retransmitted to the cloud after connectivity outages.	SR2.3
FR2.4	Maintain timekeeping locally using GNSS with RTC fallback.	SR2.6
FR2.5	Allow for complete raw dataset downloading, and downsampled live data viewing via the frontend, including GNSS maps and time-series plots.	SR2.4, SR2.5
FR2.6	Ensure the startup time, excluding GNSS TTFF, is $\leq 10 \text{ s}$.	SR2.1, SR2.2
FR2.7	Ensure the GNSS TTFF is $\leq 45 \text{ s}$ from first line-of-sight occurrence, and warm reacquisition time is $\leq 3 \text{ s}$.	SR2.1, SR2.2

Functional Requirements Defence

Voltage and current ranges are selected to cover up to 800 V and/or 1 kA battery packs, including charging headroom, and the associated accuracy targets are set to enable end-to-end profiling. GNSS performance is aligned with typical commercial multi-band receivers (e.g., u-blox M8) and is sufficient given vehicle speeds and trip lengths [109]. A 1 Hz reporting rate is sufficient, as electromobility profiles vary meaningfully over seconds rather than milliseconds. The selected sensors provide configurable bandwidths above 1 Hz, allowing future extensions to support higher-rate sampling for transient analysis. The RTC time drift is acceptable given periodic GNSS resynchronisation. Accelerometer accuracy is chosen for typical vehicle speeds and drivetrain accelerations rather than vehicle-specific vibration detection. Ambient temperature accuracy is adequate, as smaller variations than this resolution have negligible impact on battery capacity, as discussed in Section 2.3.1. CAN acquisition is included to capture additional state information that may be partially available to support the existing telemetry acquisition (e.g., SoC estimation, or speed

redundancy when GNSS is unavailable). The startup time ensures that negligible trip time is lost while allowing headroom for startup sequences and MQTT setup. The GNSS TTFF is consistent with typical commercial modules and assumes that line of sight to satellites is available during initialisation.

Non-Functional Requirements

The non-functional requirements constrain usability, robustness, endurance, and efficiency. Table 3.3 summarises the non-functional requirements, which are derived from the stakeholder requirements in Table 3.1.

Table 3.3: Non-functional requirements (NFR) derived from system-level requirements, grouped by sensing, data handling, and robustness.

ID	Requirement	Derived from
Sensing		
NFR1.1	Vehicle CAN bus data acquisition shall be non-intrusive.	SR3.4
NFR1.2	All sensing channels and subsystems shall operate correctly from -5 to $+55^{\circ}\text{C}$; sensing accuracy shall be maintained across this range.	SR1.2
NFR1.3	Sampling bandwidth shall meet or exceed the Nyquist rate for each channel.	SR1.3
NFR1.4	Inter-channel and timestamp misalignment from true sample centres shall be ± 1 s.	SR1.3
Data Handling and Communication		
NFR2.1	Local storage shall support continuous logging for at least six months at 1 Hz without media replacement.	SR2.1
NFR2.2	The system shall maintain LTE connectivity in covered regions under automotive conditions up to 120 km h^{-1} .	SR2.5
NFR2.3	RTC time drift shall not exceed 1 s/hour.	SR2.6
NFR2.4	End-to-end latency from MCU transmission to frontend rendering of real-time data shall be ≤ 2 s under typical conditions.	SR2.2
NFR2.5	All remote interfaces (gateways, cloud services, frontend applications) shall enforce secure access.	SR3.8
Retrofittability, Generalisability and Robustness		
NFR3.1	PCB assembly shall be below $15 \text{ cm} \times 15 \text{ cm} \times 5 \text{ cm}$.	SR3.1
NFR3.2	The system shall not require permanent alterations to the host vehicle's electronic systems.	SR3.3
NFR3.3	The HVDC power measurements shall remain electrically safe under automotive conditions up to 880 V.	SR3.5
NFR3.4	The system shall operate correctly when powered from a non-ideal vehicle 12 V supply.	SR3.1
NFR3.5	The enclosure shall provide a high degree of protection against dust, water, and mechanical impact suitable for in-vehicle installation.	SR3.1

Continued on next page

ID	Requirement	Derived from
NFR3.6	Inductive probes shall accommodate typical EV power and typical CAN bus diameters.	SR3.3
NFR3.7	The system shall operate within required accuracy without the need for per-device calibration or per-vehicle CAN protocol knowledge.	SR3.2
NFR3.8	The PCB assembly and internal electronics shall operate reliably under vibration and mechanical shock expected within the case and automotive contexts.	SR3.6
NFR3.9	The total power draw shall not exceed an average of 5 W and a peak of 50 W.	SR3.7

Non-Functional Requirements Defence

CAN acquisition is required to be non-intrusive so that no active load or signalling from erroneous firmware is introduced onto the vehicle bus. The temperature range aligns with typical ambient conditions within South Africa with some headroom for case and cabin heating. Meeting Nyquist ensures reliability when measuring, and inter-channel latency is based on the assumption that significant changes will not occur within that period, in line with the original 1 Hz output data rate choice.

The more-than-one-year endurance of local storage eliminates routine field maintenance, while connectivity at highway speeds ensures uninterrupted reporting. A 2 s end-to-end latency is sufficient for real-time fleet analysis, as profiling occurs on trip-scale timescales. Time drift is constrained so that unrecoverable gaps from time-source switching or inter-sample offsets remain negligible relative to the dataset. Security is enforced at all remote interfaces (MCU-cloud, cloud-frontend), and within the cloud and frontend themselves.

Retrofitability and robustness requirements guarantee ease of installation and safe operation across vehicles. The PCB footprint ensures fitment in typical automotive enclosures; no permanent modifications to wiring or busbars are required; and probes accommodate common conductor and bus geometries. Accuracy is held without calibration per device or per vehicle, and the system must remain robust to vibration, shock, and parasitic load from the vehicle battery. Safety is ensured by maintaining basic insulation, single-fault tolerance, and surge absorption appropriate for 880 V EV applications. However, this prototype does not aim to exhaustively ensure safety-standards compliance [13].

3.2. Hardware System

The hardware design targets a retrofittable, accurate, compact, safe, and vehicle-agnostic architecture for EV telematics. This section defines the physical platform that enables reliable acquisition, processing, and communication of vehicle telemetry signals, as illustrated in Figure 3.1. The system integrates high-voltage and current sensing, auxiliary sensors,

timing and processing subsystems, communication interfaces, and power regulation within a robust enclosure, ensuring safe operation under automotive environmental conditions.

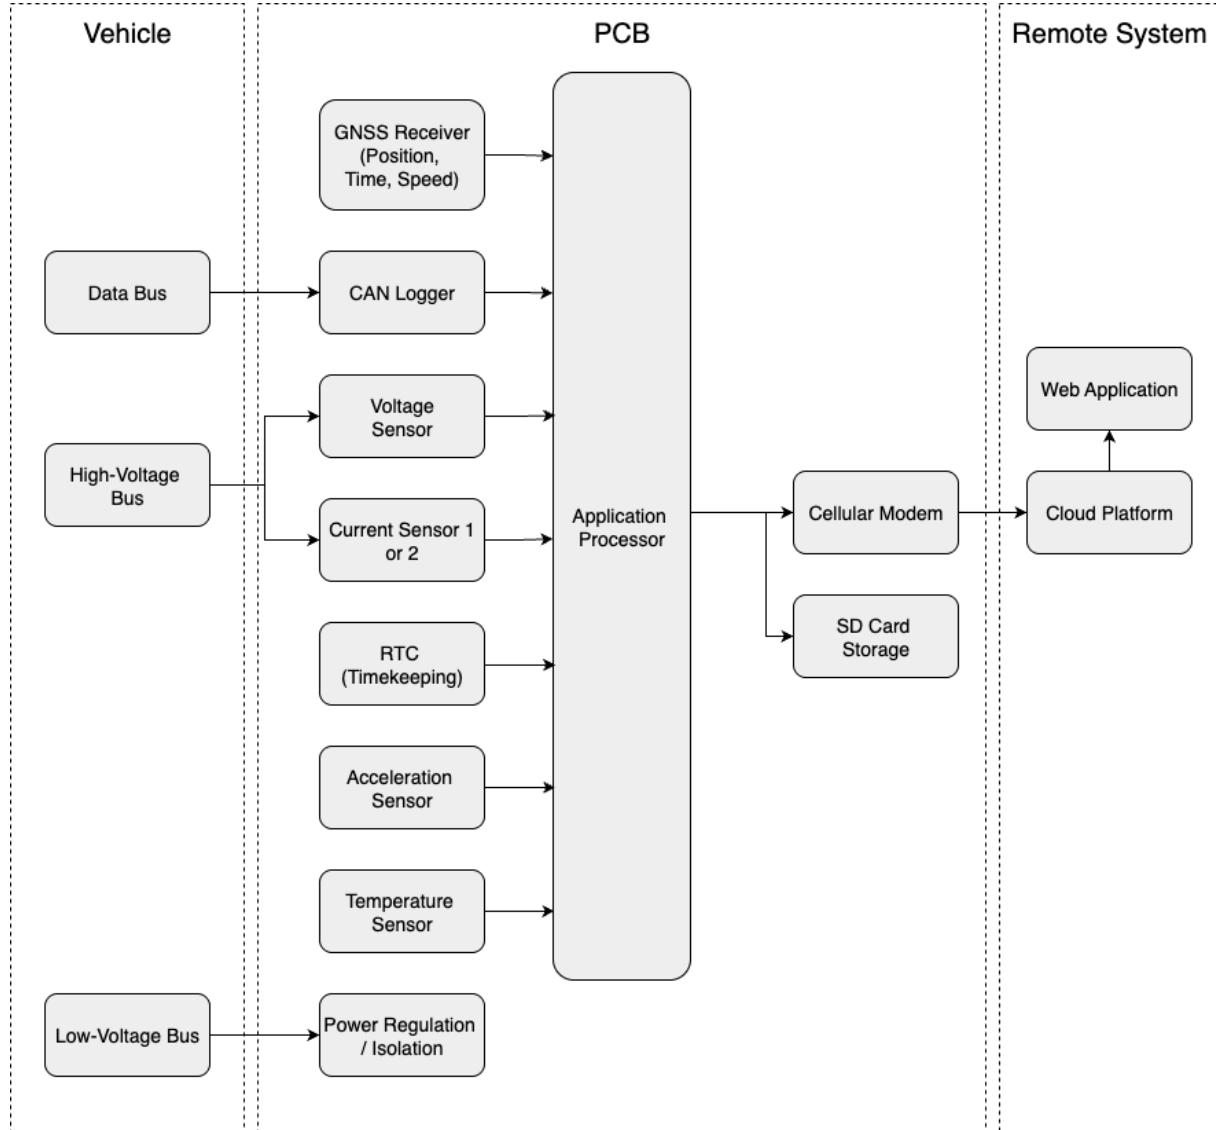


Figure 3.1: Overall hardware signal flow from vehicle data buses to the PCB's sensing, processing, and communication subsystems, and onward to the remote cloud and web interfaces.

3.2.1. Sensing Components

This subsection details the implementation of the sensing components, including high-voltage sensing, high-current sensing, GNSS, accelerometry, temperature sensing, and CAN bus data acquisition. These components form the core of the telematics system, enabling comprehensive data collection for electromobility profiling. The detailed register configuration of the sensors is provided in the Firmware Section 3.3.

High-Voltage Sensing Implementation

Accurate traction pack voltage measurement is critical for electromobility profiling. The chosen architecture is a contact-based resistive divider with an isolated ADC, which requires non-invasive terminal connections, meeting NFR3.2. The subsystem is designed to maintain precision across operational temperatures, reject vehicle electrical noise, and provide galvanic isolation between the high-voltage (HV) and low-voltage (LV) domains. The system meets FR1.1 (HVDC pack voltage error $\pm 1\%$ up to 880 V), maintains this accuracy across the operational temperature range (NFR1.2), and remains safe under automotive conditions up to 880 V (NFR3.3).

Measurement Architecture

The HV pack voltage is first attenuated by a precision resistive divider with a gain of $\frac{1}{800}$, scaling the input from a maximum of 880 V to 1.1 V. The distributed high-leg network spreads electric stress across multiple precision resistors, sustaining creepage distances while limiting divider current. An impedance-matched RC low-pass filter then rejects high-frequency interference before the isolated delta-sigma ($\Delta\Sigma$) ADC captures the signal in differential mode and transfers it across the isolation barrier. Protective components are placed at both the HV and LV sides to ensure transient immunity; basic insulation is achieved up to a working voltage of 1 kV. The AMC131M01DFMR ADC operates over a temperature range of -40 to $+125^\circ\text{C}$ [110], and all other components comply with the -5°C to 55°C temperature range requirement.

Figure 3.2 shows the resistive divider chain implementation, the inline fuse (F2), the 1.2 V clamping Zener diode (U18), the scaled-voltage test points (TP8), as well as the impedance-matched RC low-pass filter. Figure 3.3 shows the low-noise design of the AMC131M01DFMR ADC (IC1) and 8 MHz clock generation (Y3). The low-noise filtering of the integrated regulator output follows Texas Instruments' equivalent multi-channel AMC131M03 ADC reference design [111]. Post-PCB manufacturing, the assembler could not supply the fuse, so it was bypassed for this prototype with negligible effect. The crossed out jumper may be bridged to add additional capacitance.

Safety Compliance

This section describes the components used to ensure compliance with safety standards for both continuous and transient overvoltage conditions.

Protective Devices and Interfaces

Protective elements incorporated into the high-voltage interface include:

- **MOV protection:** An 895 V DC-rated MOV (B72220S0681K101) at the high-voltage input clamps surges to a maximum of 1.815 kV. Secondary clamping is designed to suppress this clamping voltage down to clearance and creepage requirements of 1 kV and the AMC131M01 input requirements.

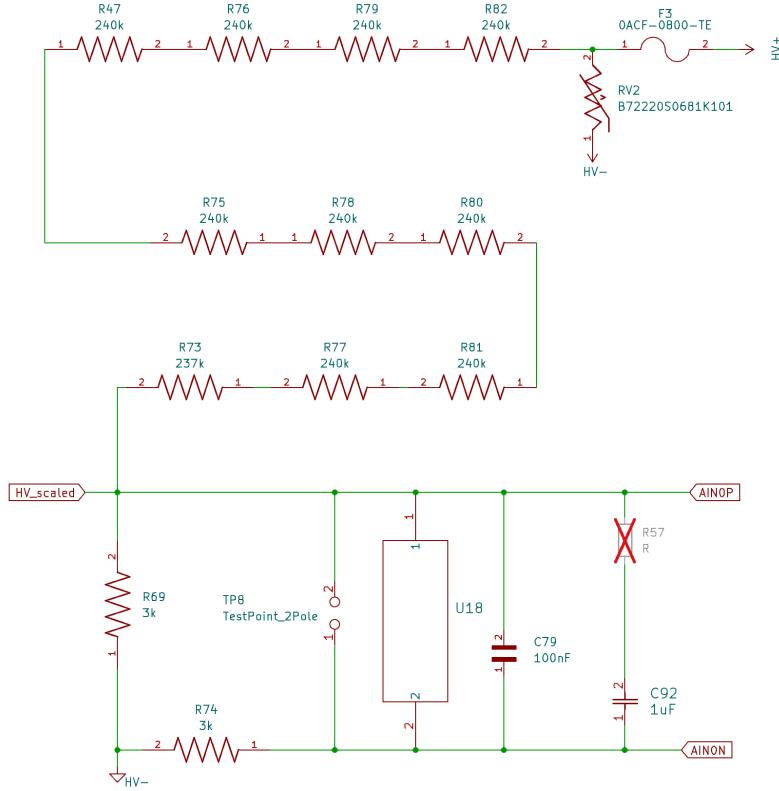


Figure 3.2: High-voltage resistive divider chain showing staged attenuation, surge protection, and the RC filter that conditions the AMC131M01 input.

- **Secondary clamping:** A Zener clamp at the ADC side suppresses fast transient spikes not attenuated by the MOV. For future AC measurement capability, this would be replaced with a bidirectional TVS diode (not used due to availability at the required voltage), and an additional MOV added at the HV- conductor to provide symmetrical protection.
- **Isolation barrier:** The AMC131M01 is rated for 1.2 kV RMS working isolation, 1.7 kV DC withstand, and 5 kV RMS transient isolation. Positive ADC inputs can tolerate up to $V_{HGND} + 2.7$ V, with clamping ensuring all transients remain below this limit.
- **Resistive divider:** The high-leg uses ten series precision resistors, each experiencing ≤ 100 V under a 1 kV system voltage, providing margin against the 200 V rating per element and distributing electrical stress evenly.
- **Connectors:** High-voltage connectors are 1 kV-rated spring-cage types meeting PD2 creepage and clearance requirements.

Insulation and Protection

The high-voltage sensing subsystem is designed for continuous operation at up to 895 V DC, as limited by the MOV continuous working voltage. It tolerates short-duration surges

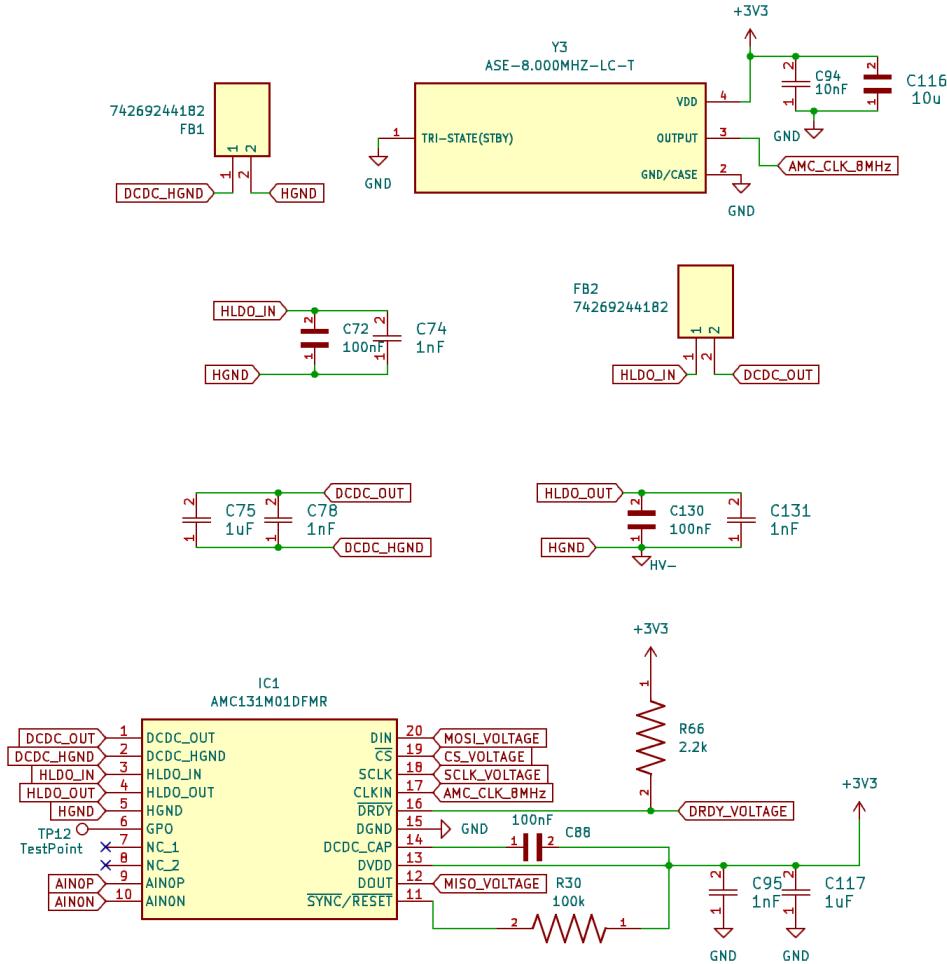


Figure 3.3: Isolated voltage sensing digitisation stage with local clock generation (top right), power supply filtering, and differential outputs into the isolated AMC131M01 ADC.

up to 1.815 kV (8/20 μ s, 250 J) and, independent of the MOV, the PCB creepage/clearance and AMC131M01 barrier provide basic insulation for working voltages up to 1 kV DC under IEC 62368-1. In practice, this means the system can operate continuously up to 895 V, withstand transient surges up to 1.815 kV, and achieve basic insulation to ensure safety up to a continuous 1 kV DC working voltage.

As this is not meant for commercial application, the PCB high-voltage domain employs basic insulation as opposed to reinforced insulation; higher-voltage protection is still ensured by multiple redundancy safety measures and the sealed Allbro enclosure. This section assumes operation under an altitude of 2 km and positive polarity DC operation.

Creepage and Clearance

Electric vehicle high-voltage systems are typically classified under Pollution Degree 2 (PD2) due to condensation and non-conductive pollution likelihood, even in sealed cases. However, certain PCB creepage paths can be evaluated as Pollution Degree 1 (PD1) in accordance with IEC 60664-1 where conductive surfaces are completely covered by

a durable protective coating. In this design, HV copper pours that do not meet PD2 creepage distances are fully covered by solder mask with no exposed copper edges or vias, allowing them to be treated as and meet requirements for PD1. All exposed pads, leads, traces, and connector terminations meet PD2 spacing requirements.

For FR-4 (Material Group IIIa), IEC 62368-1 Table 18 specifies a minimum working creepage of 3.2 mm at 1 kV_RMS for PD1 and 10.0 mm for PD2. Table 11 specifies a minimum working clearance of 0.26 mm at 1 kV_RMS for both PD1 and PD2. All PD2-exposed conductors in this design meet or exceed these requirements with a minimum creepage exceeding 10.5 mm and clearance exceeding 5 mm. For PD1-protected areas, the HV-LV creepage and clearance also exceed these limits with a minimum creepage exceeding 3.5 mm and clearance also exceeding 3.5 mm. The clearance between the high-voltage domain and the enclosure exceeds the Table 9 requirement of 0.2 mm with around 4 mm of clearance, and internal component spacings are maintained above 1 mm per 100 V-well above the breakdown threshold for air with typical humidity. All clearances and creepage requirements therefore exceed with sufficient tolerance the criteria for basic insulation.

Markings and Diagnostic Features

All high-voltage nets are colour-coded in orange and marked with the high-voltage symbol. Low-leg test points enable continuity checks before energisation to detect open-circuit faults in the low-leg of the divider.

Future Certification Enhancements

In part with achieving reinforced insulation and touch compliance, potential additions include:

- Adding jagged slots or increasing spacing to meet IEC 62368-1 Table 18 reinforced creepage requirements for PD2 at 1 kV.
- Using dual off-board high-voltage resistor chains in each lead, combined with two on-board sub-SELV TVS clamps and an upstream fuse for energy limitation. This arrangement ensures the circuit remains within SELV limits during normal operation, and maintains compliance under single-fault conditions through redundancy.
- Selecting PCB materials with higher CTI to reduce required creepage distance for a given working voltage. This may include conformal coating which would protect against moisture, tampering, vibration, accidental touch, and decrease creepage and clearance requirements.

Safety Compliance

The system satisfies creepage and clearance requirements for basic insulation, incorporates protection against transients and faults, and includes adequate marking, ensuring safety for automotive applications up to 880 V. The combination of the precision divider, isolated

ADC, and MOV and clamp protection results in a robust high-voltage sensing system that is safe for controlled research use, with identified enhancements enabling reinforced insulation for broader deployment.

Accuracy Analysis

The accuracy analysis considers an 880 V FSR application, analysed per the AMC131M01DFMR (AMC13) ADC datasheet [110]. Errors concerning the AMC13 that are already specified across the FSR are included directly in Table 3.4, whereas the following sections provide contextual derivations of FSR and offset errors.

Resistive Divider Errors

The resistive divider consists of a high leg of ten series 0.1% resistors (totalling $2.397\text{ M}\Omega$, implemented as nine $240\text{ k}\Omega$ resistors and one $237\text{ k}\Omega$ resistor) and a low leg jumper-selectable between $3\text{ k}\Omega$, $6\text{ k}\Omega$, or $12\text{ k}\Omega$ using 0.05% resistors. The temperature coefficients are $25\text{ ppm/}^\circ\text{C}$ for the high leg and $10\text{ ppm/}^\circ\text{C}$ for the low leg. Worst-case tolerance and ppm values are applied to both typical and maximum cases, as no typical values are specified within the given tolerance. The worst-case tolerance error occurs when the low leg is at its lowest value and the high leg at its highest, resulting in an error of -0.15015% at 25°C . For temperature drift, the worst case is taken as the maximum differential in ppm values within tolerance, assuming both increase monotonically with temperature. In this case, the high leg has $25\text{ ppm/}^\circ\text{C}$ and the low leg $0\text{ ppm/}^\circ\text{C}$, resulting in an error of -0.0834028% .

AMC13 ADC Errors

The AMC13 ADC operates in high-resolution global-chop mode with a unity gain, yielding a typical input-referred (scaled down voltage) offset of $6\text{ }\mu\text{V}$ (maximum $100\text{ }\mu\text{V}$) with a drift of approximately $\pm 0.1\text{ }\mu\text{V/}^\circ\text{C}$ (maximum $\pm 0.3\text{ }\mu\text{V/}^\circ\text{C}$). The drift results in an additional input-referred offset typical error of $3\mu\text{V}$ and maximum error of $9\mu\text{V}$. Table 3.4 refers these values 1.1 V FSR on the AMC13 input side.

The datasheet only specifies the typical offset between bias currents as 15 nA ; this is assumed to be appreciably stable at this value. Given approximately $3\text{ k}\Omega$ source impedance, this results in a input-referred voltage error of $45\text{ }\mu\text{V}$ at 25°C ; a 10% headroom is added, resulting in $49.5\text{ }\mu\text{V}$, to represent the maximum error at 25°C . The offset bias current drift is typically $\pm 0.35\text{ nA/}^\circ\text{C}$ and maximally $\pm 1\text{ nA/}^\circ\text{C}$. This results in an additional input-referred voltage error of $31.5\text{ }\mu\text{V}$ and $90\text{ }\mu\text{V}$ respectively. Table 3.4 refers these values 1.1 V FSR on the AMC13 input side.

The differential input impedance is typically $275\text{ k}\Omega$. No peak error or detailed graphical model is specified; it is therefore reasonable to believe it is at least within 1% typically and 10% maximally of this value under all operational conditions. Assuming that firmware provides a constant gain correction factor (static per device) the remaining non-constant error, given the true parallel resistance of the lower-leg, can be assumed to be a gain error

of 0.011% typically and 0.11% maximally. This error is considered under the static section, but includes expected drift due to temperature changes.

The bias current is typically $0.65\ \mu\text{A}$ with a maximum of $0.9\ \mu\text{A}$. Since the input impedance is chosen to be matched within approximately 0.1001% (estimated under worst-case scenario given the individual tolerances of 0.05%), this error becomes negligible ($\sim 2\ \text{mV}$ primary referred or around 2.3 ppm at 880 V full-scale). Input-referred noise is in the low single-digit μV range, integral non-linearity is approximately 6 ppm, and the common-mode rejection ratio (CMRR) is 110 dB (316,228 \times attenuation of equivalent signals on both inputs); these are sufficiently low to also be neglected in the error budget. Kickback voltage is not explicitly specified but is minimised by ensuring the input capacitance is sufficiently high to supply current given the acquisition time and source impedance. This error is considered negligible, as the default capacitor of 100 nF is assumed to mitigate the effect. Ferrite beads and multiple capacitors limit the influence of switching noise, and this effect is also considered negligible. This RC filter results in a cutoff frequency of approximately 600 Hz, and thereby removes any appreciable expected higher-frequency components.

Voltage Sensing Accuracy Table

The complete voltage measurement error budget is shown in Table 3.4, demonstrating compliance across the full operational range up to 880 V from $-5\ ^\circ\text{C}$ to $55\ ^\circ\text{C}$. The input-referred offset error (IROE) values are also provided to ascertain the 0 V offset error in μV .

Table 3.4: Voltage measurement error budget referred to 1.1 V full-scale, highlighting component errors and temperature drift contributors.

Contribution	Typ [%FS]	Max [%FS]
Static error at 25 °C		
Divider accuracy	0.15015	0.15015
ADC gain error	0.02500	0.20000
ADC offset error	0.00055	0.00909
Bias current offset	0.00409	0.00450
Input impedance offset	0.01100	0.11000
Subtotal static (%)	0.19079	0.47374
Subtotal IROE (μV)	51.0	149.5
Additional drift from $-5\ ^\circ\text{C}$ to $55\ ^\circ\text{C}$		
Divider drift	0.08340	0.08340
ADC gain drift	0.00240	0.00750
ADC offset drift	0.00027	0.00082
Bias current drift	0.00286	0.00818
Subtotal drift (%)	0.08893	0.09990
Subtotal IROE drift (μV)	34.5	99.0
Total (%)	0.27972	0.57364
Total IROE (μV)	85.5	248.5

The typical error remains below 0.3% at standard temperature, while the maximum

error remains below 0.6% across the temperature range of -5°C to 55°C . The input-referred offset error remains below 51.0 μV typically and 248.5 μV at maximum, or approximately 0.041 V typically or 0.22 V maximally primary-referred at the given 880 V full-scale, which is acceptable for low-voltage systems. The system therefore meets FR1.1 by maintaining a maximum error below 1% across the full operational range (NFR1.2).

Outcome

This implementation meets FR1.1 and NFR1.2 with a maximal error of 0.57%FS across the full measurement and temperature range, below the 1%FS requirement. Additionally, it meets NFR3.2 with non-invasive tapping installations. It further meets NFR3.3 with basic insulation compliance up to 1 kV and additional protective measures. It is also noteworthy that the system can, with minor modifications, support AC voltages, and scale to higher voltage requirements.

High-Current Sensing

Accurate traction pack current measurement is critical for electromobility profiling. This section details the achievement of FR1.2 (HVDC pack current error $\pm 1.5\%$ up to 1 kA), NFR1.2 (accuracy across the operational temperature range), NFR3.2 (non-invasive installation), NFR3.3 (safety under automotive conditions up to 880 V), and NFR3.6 (accommodation of typical EV conductor or busbar sizes).

Fluxgate sensors cannot achieve their extreme accuracy ($\approx 0.5\%$ or better) with split-core designs and therefore require temporarily disconnecting the HVDC conductor, slightly increasing installation complexity. Coreless clip-on designs connect easily in confined spaces with minimal effort but typically do not achieve the same accuracy ($\approx 1\%$ at best). As different operators may have varying priorities and access constraints, both requirement-compliant sensor types are implemented to provide a choice between the most accurate and the most easily retrofittable solution. Figure 3.17 shows the current-sensor selection switch provided on the PCB.

The two sensors chosen were: a high-accuracy, mountable fluxgate transducer by LEM and a portable, light, vibration-resistant, clip-on coreless probe by GMW. The sensors may be seen in Figure 3.4.

LEM Fluxgate Sensor

The LEM fluxgate sensor (CAB-SF 1500-006) provides high-accuracy DC current measurement with a bandwidth of 20 Hz. It uses a single-core, air-gap-free fluxgate design that has excellent accuracy, temperature stability and EMC robustness. It employs internal digitisation, outputting 24-bit current data over CAN, which is received by a dedicated CAN transceiver-controller pair detailed in Section 3.2.3. The sensor aperture accommodates circular conductors up to 16.2(3) mm in diameter, or busbars up to 24.4(3) mm \times 4.0(1) mm. This ensures compliance with NFR3.2 and NFR3.6.



Figure 3.4: Inductive current sensors – LEM CAB-SF (left) and GMW clamp (right).

The sensor supports an operational range of ± 1500 A and provides a CAN bus output at 100 Hz, meeting the Nyquist criterion, with internal $120\ \Omega$ termination (a second $120\ \Omega$ terminator is placed on the CAN transceiver). It has a bandwidth of 20 Hz and an error $\leq 0.5\%$ over its operational temperature range of $-40\text{ }^{\circ}\text{C}$ to $85\text{ }^{\circ}\text{C}$, improving to below 0.3% at $25\text{ }^{\circ}\text{C}$. The supply requirement is a unipolar 8–16 V (automotive compliance) with current draw proportional to the primary current (400–1000 mA at 16–8 V, respectively, for a 1000 A primary current). Additionally, its automotive compliance ensures vibration and moisture resistance, while its compact form factor minimises physical impact [112]. This meets FR1.2 ($\leq 1.5\%$ error) up to 1000 A.

GMW Coreless Probe

The GMW coreless current probe (CPC-1000-27-BP2) provides non-intrusive measurement with portability and simple repositioning, making it well-suited for constrained or temporary installations such as electric bikes. It offers an intrinsic accuracy of 1%, and requires additional digitisation circuitry. It features a larger 27 mm diameter aperture, and is compliant with moisture standards while being applicable in high vibration environments. Moreover, its lightweight, slim form factor enables use in space-constrained environments without the need for dedicated mounting. Additionally, it supports a bandwidth of 75 kHz, allowing use for both DC and AC or transient analysis [113]. This ensures compliance with NFR3.2 and NFR3.6, with the larger aperture accommodating thicker cables or busbars.

The probe supports an operational range of ± 1000 A, with a ± 2 V differential output corresponding to 2 mV A^{-1} sensitivity. It achieves an error of $\pm 1\%$ of full scale, with a bandwidth of 75 kHz. The device operates from a 3.5–5.5 V supply, drawing approximately 85 mA. It operates over a temperature range of -40°C to 100°C .

Conditioning Architecture

The coreless probe output is routed through an impedance-matched RC anti-alias filter directly into a Texas Instruments ADS1220 24-bit $\Delta\Sigma$ ADC [114], operated with bipolar ± 2.5 V analogue rails derived from an isolated 5 V supply. The ADS1220 provides high-precision conversion with sub-millivolt input offset, low drift, and an integrated 2.048 V reference and precision clock, enabling direct digitisation of the probe's bidirectional ± 2 V differential span without level shifting. The design positions digital and analogue circuitry on their respective sides of the ADC's analogue-digital boundary to minimise high-frequency noise coupling into the analogue front end. The entire analogue and digital power rails of the architecture are galvanically isolated from the rest of the circuitry.

Figure 3.5 illustrates the isolated ADC input chain for the GMW probe, including the RC filtering which is routed differentially (that is, through a fixed adjacent trace gap and equivalent trace lengths). Bridging the jumper (JP1) allows for additional capacitance to be added. Figure 3.6 shows the generation of the bipolar analogue rails required for the ADC front end, and Figure 3.7 highlights the isolation barrier separating the isolated sensor and main system domains.

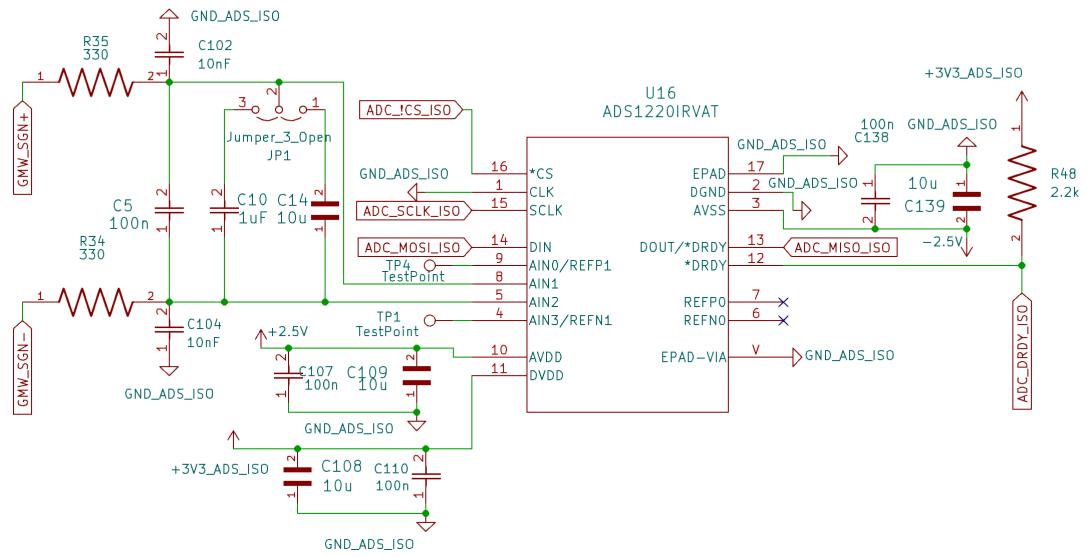


Figure 3.5: Isolated ADC input chain for the GMW coreless probe, including RC anti-alias filters.

Isolation Interface

Complete galvanic isolation is implemented to suppress noise coupling from other circuitry, using an isolated 5 V regulator that supplies the ADC and its downstream regulators. Two

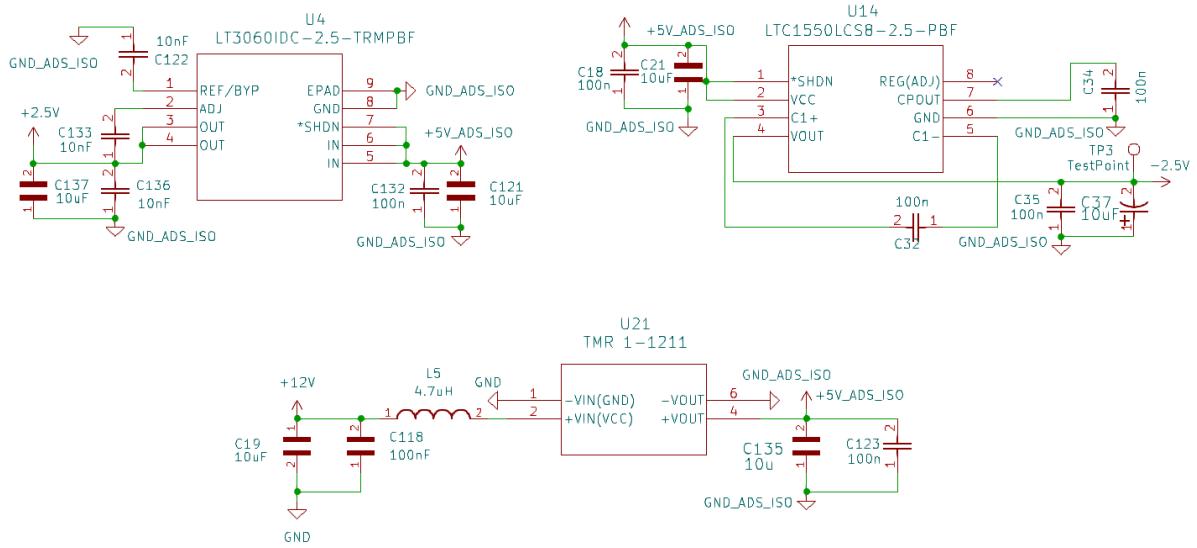


Figure 3.6: Isolated supply network for the GMW probe, showing the dedicated 5 V converter (bottom) and the ± 2.5 V rails (top) that power the ADS1220 AED.

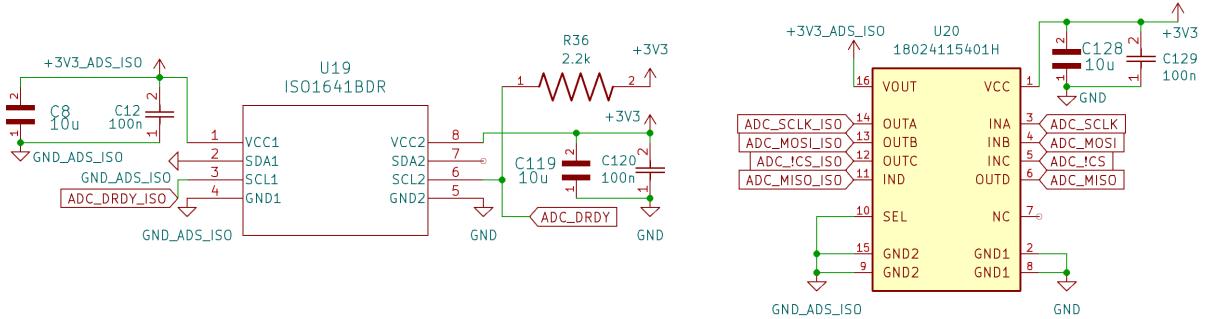


Figure 3.7: Galvanic digital isolation interface between the analogue sensing chain and MCU domain.

series linear regulators derive the ± 2.5 V analogue rails required for bipolar operation, with stability ensured through appropriate decoupling at each stage to prevent oscillation given the cascaded regulator design. The digital interface employs a bidirectional SPI isolator that transfers communication signals and generates an isolated 3.3 V rail for the ADC's digital domain, while a separate digital isolator, also using this isolated supply, handles the nDRDY signal. All three isolation barriers share a common isolated ground on the sensor side, which does not introduce significant coupling due to the physical separation about the ADC's analogue-digital boundary. The combination of external isolation and internal analogue-digital segregation provides strong noise immunity, supporting sub-millivolt measurement accuracy given the 2mV A^{-1} sensor resolution.

GMW ADC Error Budget

The ADS1220 [114] is configured for differential operation (PGA disabled, turbo mode, 20 SPS) with the internal 2.048 V reference. The sensor drives a ± 2.000 V full-scale range (FSR), well within the $\pm V_{\text{REF}}$ limits.

INL is specified as 6 ppm typical and 15 ppm maximum. The RMS 1 Hz downsampled noise is estimated at 0.83 μ V (typical), while the peak worst-case noise is 13.5 μ V across the temperature range. Offset error is typically 4 μ V and maximally 30 μ V, while offset drift is typically $0.25 \mu\text{V}^{\circ}\text{C}^{-1}$; a maximum offset drift is not specified, and it is assumed to remain stable, estimated conservatively as no worse than $0.275 \mu\text{V}^{\circ}\text{C}^{-1}$ (10% tolerance). Assuming worst-case differential drift (twice the individual channel drift), the resulting offset error is 0.8 μ V typical and 60 μ V maximum, while the offset drift is 15 μ V typical and 16.5 μ V maximum over the full temperature range.

Gain error is typically 150 ppm, and no maximum is specified; it is therefore assumed to be appreciably similar in the worst case, conservatively estimated as no worse than 165 ppm (10% tolerance). Gain drift is typically 1 ppm/ $^{\circ}\text{C}$ and maximally 4 ppm/ $^{\circ}\text{C}$, corresponding to 30 ppm and 120 ppm respectively over the full temperature range. The reference exhibits a typical error of 244.1 ppm and a maximum of 488.3 ppm, with a typical drift of 5 ppm/ $^{\circ}\text{C}$ and a maximum of 30 ppm/ $^{\circ}\text{C}$, equivalent to 150 ppm and 900 ppm respectively over the full temperature range. Long-term drift is specified as 110 ppm per 1000 h and is added to the maximum static reference error.

Errors due to input currents and differential impedances are considered negligible, as the primary-referred induced errors are within the sub-milliampere range (less than 1 ppm of FSR). Furthermore, the common-mode, normal-mode, and power-supply rejection ratios (PSRR) are consistently above 80 dB ($10\,000\times$ attenuation), ensuring that any common-mode or supply noise is negligible. The RC filter cutoff frequency of approximately 8 Hz effectively suppresses high-frequency noise, and the low source impedance with 100 nF input capacitance ensures that kickback voltage is negligible.

Table 3.5 summarises the error budget for the ADS1220. The input-referred offset errors (IROE) are also shown in μ V for context.

The error budget shows typical errors below 0.06%FS at standard temperature and below 0.19%FS across the full temperature range. Offsets remain below a maximum of 76.5 μ V, corresponding to 38.3 mA primary-referred given a $2\,\text{mV}\,\text{A}^{-1}$ sensitivity, which is relatively negligible for typical EV applications. Including the sensor accuracy of 1%, the total error of 1.2% remains within the FR1.2 requirement of 1.5%FS across the entire temperature range.

Safety Considerations

The primary safety consideration is the potential exposure of the user to hazardous primary-side voltages through the applied system. As typical EV installations use insulated conductors compliant with relevant insulation standards, direct arcing is not a significant concern. Both sensors safely tolerate over-current-induced magnetic fields. Capacitive coupling risks are only relevant for high-frequency AC conditions, which are beyond the scope of this design. As an added redundancy, sensors provide reinforced isolation: the

Table 3.5: ADS1220 input-chain error budget in ppm referred to ± 2.000 V.

Contribution	Typ [ppm]	Max [ppm]
Static error at 25 °C		
Noise	0.415	6.75
Offset	0.400	30.00
Gain	150.00	165.00
Reference	244.10	598.30
INL	6.00	15.00
Subtotal static error (ppm)	400.90	815.10
Subtotal static IROE (μV)	0.80	60.00
Additional drift from -5 °C to 55 °C		
Gain drift	30.00	120.00
Reference drift	150.00	900.00
Offset drift	7.50	8.25
Subtotal drift error (ppm)	187.50	1028.25
Subtotal IROE drift (μV)	15.00	16.50
Total error (ppm)	588.40	1843.30
Total IROE (μV)	15.80	76.50
Sensor intrinsic (ppm)	—	± 10000

LEM sensor is rated for a 5 kV_{AC} withstand voltage for 1 min at 60 Hz, and the GMW sensor for 50 Hz. Additionally, the LEM sensor includes enhanced moisture and dust protection and is certified for 800 V applications per IEC 60664-1. The system is therefore considered electrically safe for its intended use.

Outcome

Table 3.6 summarises the key characteristics of the two current-sensing approaches. The LEM fluxgate offers higher accuracy ($\leq 0.5\%$), whereas the GMW coreless probe with its digitisation circuitry provides approximately 1.2% accuracy and significantly simpler installation. The system ensures compliance with FR1.2, NFR1.2, NFR3.2, NFR3.3, and NFR3.6, delivering accurate, non-invasive and safe current measurement for electric vehicle telemetry.

Parameter	Fluxgate Sensor	Coreless Probe
Measurement range	± 1500 A	± 1000 A
Output type	CAN bus	± 2 V differential
Reading accuracy	$\pm 0.5\%$	$\pm 1.2\%$
Supply voltage	8-16 V	3.5-5.5 V
Ease of installation	High	Extreme
Aperture diameter	15 mm	27 mm

Table 3.6: Comparison of key characteristics for the two current sensing approaches.

Temperature Sensing

A digital, semiconductor-based one-wire ambient sensor monitors cabin air temperature to provide environmental context. The DS18B20 [115] offers a 93.75 ms conversion time,

$\pm 0.5^\circ\text{C}$ accuracy across the -55 to $+125^\circ\text{C}$ operating range, and vibration-tolerant packaging, ensuring durability. Placement away from the PCB or air-conditioning vents ensures it captures ambient conditions that correlate to the battery pack. The circuit is simply implemented with a pull-up resistor, as can be seen in Figure 3.17. This implementation meets FR1.5 (ambient temperature accuracy $\leq \pm 1^\circ\text{C}$) and satisfies NFR1.2 (operation across -5 to $+55^\circ\text{C}$).

Acceleration Sensing

The ADXL343 3-axis MEMS capacitive accelerometer records vehicle motion and vibration. Configured for a $\pm 2\text{ G}$ full-scale range with 25 Hz ODR (12.5 Hz bandwidth), it provides DC response and low-noise performance suited to drivetrain events and is downsampled to 1 Hz for transmission. Its operating temperature range covers -40 to $+85^\circ\text{C}$ [116]. Figure 3.8 shows the acceleration sensing circuit and filtering.

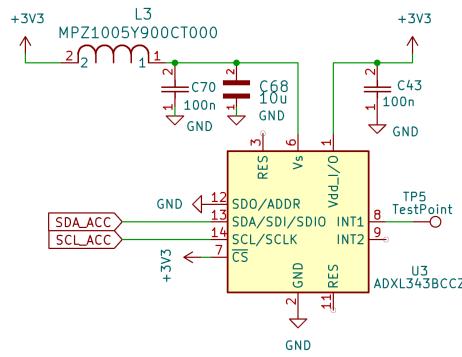


Figure 3.8: Acceleration sensing circuit based on the ADXL343 and its associated filtering.

The ADXL343 error at the $\pm 2\text{G}$ range is contributed by several factors: a zero-G offset of about 35 mG per axis, a sensitivity error of 1% of reading, a nonlinearity of 0.5%FSR, and a cross-axis sensitivity of 1%FSR, with inter-axis misalignment negligible. The axial 0G offset error (per axis) is 35 mG at 25°C and 83 mG over the operational temperature range defined by scope (-5°C to 55°C). The axial full-scale-range error at 25°C is 2.88% FSR, and over the full operating range it is 4.38% FSR [116].

Further, these errors must be adjusted if the device is not at a multiple of 90° with respect to the vehicle; using the 2D planar resultant gives a worst-case of $\sqrt{2} \times 4.38\% = 6.19\%$ over the full temperature range. This, with modest headroom for cases beyond the datasheet's 1σ typicals, remains under the $\leq 10\%$ requirement. However, installation non-levelness can dominate if not controlled, and the installer must level the device or accept the added error. This study excludes accident events, which are not material to assessing economic feasibility with respect to mobility and power profiles. This implementation meets FR1.4 (3-axis acceleration accuracy $\leq 10\%$ FSR) and satisfies NFR1.2 (operation across -5 to $+55^\circ\text{C}$).

3.2.2. Processing and Control

This subsection describes the processing and control subsystems, including the microcontroller unit (MCU), real-time clock (RTC), and power management. These subsystems aim to ensure reliable operation, accurate timekeeping, and efficient power usage.

Microcontroller Unit (MCU)

The central controller is a Wemos ESP32-S2 Mini, based on the single-core 240 MHz Xtensa LX7 ESP32-S2. Its 320 kB SRAM, 4 MB flash, and auxiliary 2 MB PSRAM cover concurrent sensing, buffering, and network tasks without stressing headroom, while the 4×8 header format keeps the footprint compact. Integrated full-speed USB Type-C, a native bootloader, and mature Arduino/ESP-IDF tooling shorten firmware iteration and align with the rapid-development criteria established in the hardware survey. The module exposes 23 GPIOs with two SPI, two I²C, and two UART controllers, satisfying the peripheral mix required for the sensor suite. It is header-mounted for easy field replacement and is rated for -40 to $+105^\circ\text{C}$ operation along with its passives [117].

Programming and serial communication are performed directly over USB. The board enters programming mode by holding `Boot` for three seconds, then pressing `RST`. In the Arduino IDE, during application development, upload mode is set to Internal USB with USB-CDC enabled on boot.

The MCU runs from the regulated 3.3 V rail; Wi-Fi remains disabled in normal service to reserve power, but the radio stack provides an OTA update path if later required. The GPIO matrix permits flexible reassignment of peripherals, and the dual SPI/I²C buses avoid contention between the high-rate ADC chain, SD interface, and auxiliary sensors. Two UARTs remain available even with USB-CDC active, keeping the LENA modem on a dedicated channel while preserving a service console.

SPI1 and SPI2 split between the 5 peripherals to limit bus loading. The LTE modem occupies one UART, with the second used for debugging or reprogramming. External pull-ups and small series resistors on the digital buses maintain signal integrity. Interrupt lines are assigned to both CAN transceivers, the temperature sensor, and DRDY pins from SPI peripherals. A switch input allows firmware to detect which current sensor is connected, and another controls 3V3 power allocation when switching between USB and board power.

This MCU is capable of facilitating the achievement all functional and non-functional requirements related to all sensing, communication, and control tasks.

Figure 3.9 illustrates the MCU's pin mapping and peripheral interconnections. The LENA-R8 power control line is routed to Test Point 6 after PCB assembly.

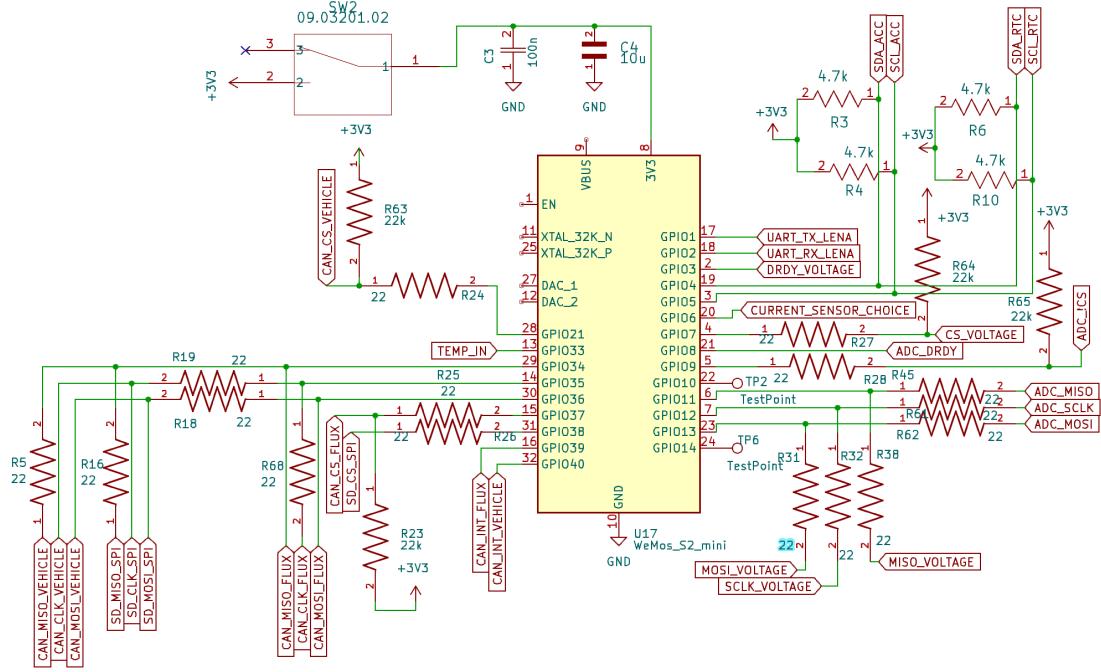


Figure 3.9: ESP32-S2 Mini schematic pinout showing peripheral assignments, communications buses, and sensor interfaces.

Real-Time Clock (RTC)

The MCP7940N was selected for its I²C interface, integrated battery-backup switching, and frequency trimming capability—though trimming is not used, as the system follows the zero-calibration design requirement. Pairing it with a 32.768 kHz tuning-fork crystal and discrete load capacitors ensures meeting of the accuracy expectations for vehicular timestamping.

Architecture

Both the crystal and RTC operate over the -5 to $+55$ °C range [118, 119]. Firmware disciplines the clock to the GNSS time every 15 min, coupling the low-power standby characteristics of a crystal RTC with the long-term stability needed for accurate dataset alignment.

The VBAT pin draws approximately 700 nA [118], so the vertical CR2032 (220 mA h) comfortably delivers the multi-year autonomy anticipated in the component survey; in practice the cell's

approx 1%/year self-discharge limits service life to roughly a decade before replacement is advisable.

The schematic implementation is shown in Figure 3.10. The crystal is placed adjacent to the RTC package with short, symmetrical traces and a solid ground plane to minimise parasitic capacitance and EMI susceptibility.

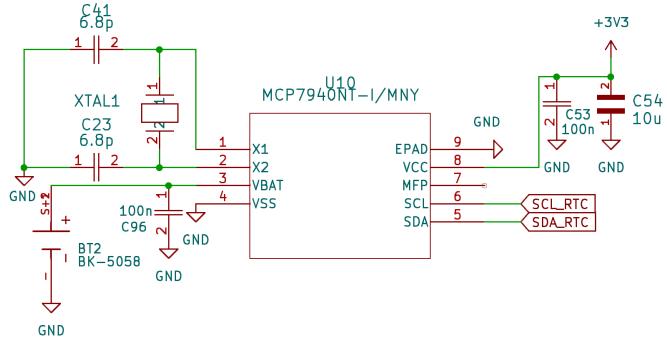


Figure 3.10: MCP7940N real-time clock circuit with 32.768 kHz crystal, associated load and filtration capacitors and CR2032 backup supply holder.

Accuracy Analysis

The RTC employs an SC32S-7PF20PPM tuning-fork crystal with a ± 20 ppm initial frequency tolerance and ± 3 ppm/year ageing. This corresponds to approximately 1.73 s/day initial drift, increasing to 2.16 s/day after one year, and up to 4.32 s/day (50 ppm) after ten years.

Temperature-induced deviation follows the parabolic relation

$$\frac{\Delta f}{f} = B(T - T_i)^2,$$

where $B = -3.0 \times 10^{-8}/^{\circ}\text{C}^2$ and $T_i = 25^{\circ}\text{C}$. Across the -5 to $+55^{\circ}\text{C}$ range ($\Delta T = 30^{\circ}\text{C}$), the frequency decreases by approximately 27 ppm on either side of T_i . The total worst-case frequency deviation accounting for initial tolerance and ten-year drift is approximately -47 ppm (-4.06 s/day) at room temperature, and -77 ppm (-6.65 s/day) across the full temperature range.

Further frequency deviation arises from capacitive-loading mismatches within the oscillator network. The oscillation frequency depends on the reactive impedance formed by the external load capacitors, the PCB parasitics, and the internal oscillator pin capacitances. The effective load capacitance for the oscillator experienced between the crystal terminals is expressed as

$$C_L = \frac{C_{X1} \times C_{X2}}{C_{X1} + C_{X2}} + C_{\text{STRAY}},$$

per [120]. For this configuration, $C_{X1} = C_{X2} = 6.8$ pF and $C_{\text{STRAY}} = 3.5$ pF, the latter representing the mean of the typical 2-5 pF range [121]. The resulting effective load capacitance of approximately 6.9 pF aligns closely with the 7 pF requirement of the SC32S crystal and lies within the MCP7940N oscillator's recommended driveable range (6-9 pF), thereby minimising frequency pulling.

To accommodate tolerance variations in both the external capacitors and stray capacitance, a conservative 50% margin is included in the error budget to ensure stable

oscillation under all conditions.

Combining all effects, the total expected drift remains within 70.5 ppm (0.306 s/hour) at room temperature and 115.5 ppm (0.4158 s/hour) across the full temperature range, satisfying NFR2.3 (RTC drift ≤ 1 s/hour or < 277.78 ppm) and NFR1.2 (accuracy maintained from -5 to $+55$ °C). Additionally, this fulfils the latter component of FR2.4 (GNSS timekeeping with RTC fallback).

3.2.3. Communication Interfaces

This subsection describes the communication interfaces, including CAN bus, GNSS, and cellular (RAT) connectivity. These interfaces ensure reliable data acquisition from vehicle systems and sensors, as well as robust data transmission to cloud services.

CAN Bus Interfaces

The system employs two independent CAN networks. One interfaces with the vehicle's existing CAN bus, while the other connects exclusively to the fluxgate current sensor. Keeping these buses separate prevents cross-interference, simplifies integration, and ensures compliance with system requirements that avoid permanent vehicle modifications.

Both CAN buses employ a separate transceiver and controller pair, although the models are the same. The transceiver is the TCAN3413DRBR and the controller is the MCP2515-I/ML. Both devices support 5–1000 kbps operation. Both are chosen as they operate entirely on 3V3, and the TCAN3413 operates from -40 to $+150$ °C, while the MCP2515 operates from -60 to $+125$ °C [122, 123]. The circuit is shown below in Figure 3.11.

Vehicle CAN

The aim of the vehicle CAN interface is to provide optional supplementary or redundant information where message definitions are partially or fully known, for example maintaining speed data during GNSS outages.

Connection is made via a CSS Electronics contactless CAN reader, shown in Figure 3.12 which inductively couples to the bus to provide full galvanic isolation and eliminate any risk of interfering with vehicle systems. The device supports Classical CAN communication over the full range of 5–1000 kbps and interfaces with typical communication on conductors around (≈ 20 AWG) [4]. Exact minimum and maximum conductor diameters were not provided; the device worked with 30 AWG conductors. A conservative range is estimated as 18 to 30 AWG, which is consistent with normal CAN wiring systems which are low current. The output is then handled by the aforementioned dedicated transceiver-controller pair, which translates the bus signalling into CAN frames accessible to the MCU.

Since commercial inductive probes do not yet support CAN-FD, the hardware is limited to Classical CAN, which remains sufficient for reading low-speed signals in most vehicles in operation today, as per Section 2.3.3. A removable 120Ω termination jumper is included

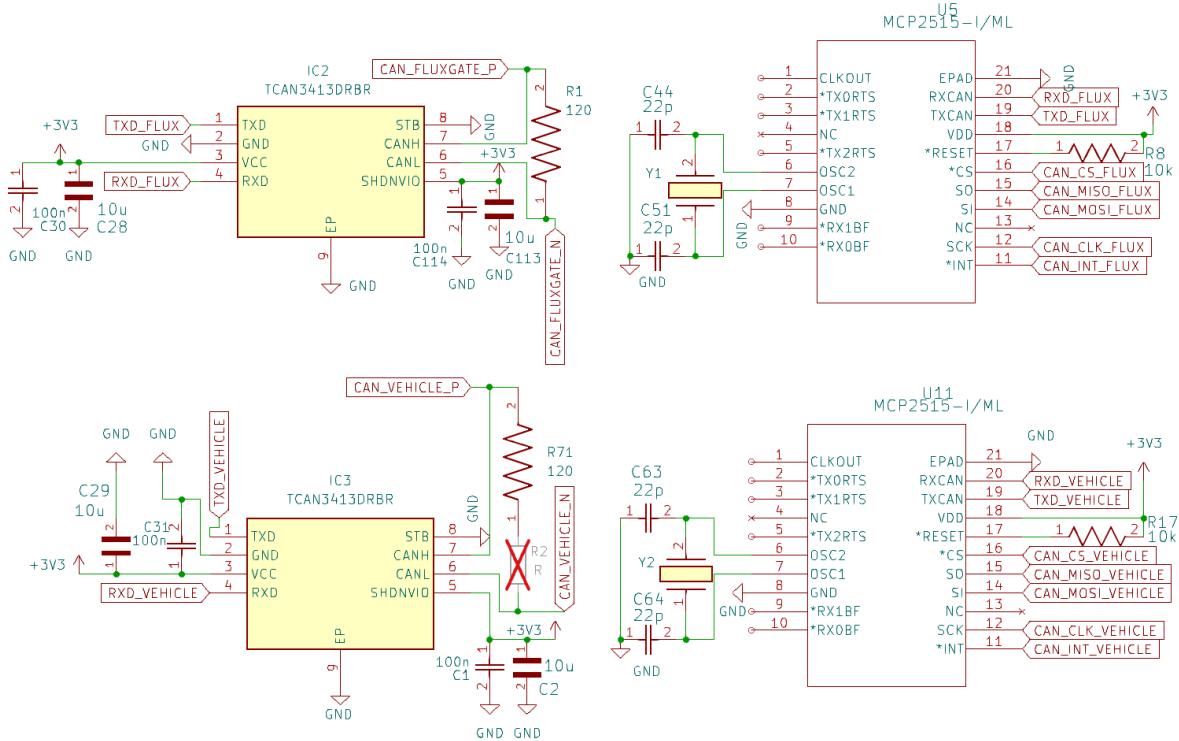


Figure 3.11: CAN bus circuits for fluxgate sensor (top) and vehicle CAN bus acquisition (bottom), showing transceivers (left) and controllers (right).

for test purposes. Future versions may add CAN-FD-capable controllers and transceivers with required additional safety measures.

Sensor CAN

The LEM fluxgate current sensor is interfaced over a dedicated CAN bus with its own transceiver-controller pair. This bus is not connected to the vehicle CAN system, preventing any possible influence on the vehicle bus or requirement for permanent bus alterations.

The sensor includes an internal 120Ω termination resistor, with a second 120Ω at the transceiver side to provide the correct bus loading (60Ω total) and maintain signal integrity by minimising reflections. The sensor transmits current readings at 100 Hz (10 ms intervals) with a measurement bandwidth of 20 Hz (DC-only), providing oversampled and accurate output. Both sensor and vehicle CAN buses operate at 500 kbps.

Outcome

This system meets FR1.6 (vehicle CAN frame acquisition) and satisfies NFR1.1 (non-intrusive vehicle CAN access) through read-only galvanic isolation. It further fulfils NFR3.2 (non-invasive installation) via inductive coupling and NFR3.6 (typical vehicle conductor compatibility) through its broad supported conductor range, while also providing the dedicated CAN interface required for FR1.2 (fluxgate current sensing).



Figure 3.12: Clip-on CAN inductive probe and interface cable for non-intrusive signal tapping.

GNSS and Cellular (RAT)

Module Selection

This design uses the u-blox LENA-R8001M10-00C module, which combines LTE Category 1 bis, 2G fallback, and a multiband GNSS receiver (supporting GPS, GLONASS, Galileo, and BeiDou) in a single compact package with a single UART interface. This integrated approach reduces PCB space, simplifies antenna layout, and ensures proper coexistence management between GNSS and cellular functions internally. The embedded UBX-M10050 engine tracks up to four constellations concurrently with a default 1 Hz navigation rate and -167 dBm tracking sensitivity, providing resilient fixes in weak-signal conditions. LTE Cat 1 bis offers sufficient throughput (up to 10.3 Mb/s downlink and 5.2 Mb/s uplink) for the targeted telemetry application, including periodic GNSS and sensor uploads. The inclusion of 2G fallback ensures continued operation in areas with limited LTE coverage—still common in certain regions. Vodacom is utilised since it has the highest LTE coverage in South Africa as of May 2024 [124]. The module operates within specification from -20 to $+65$ °C [60]. Together, these features enable remote connectivity, accurate timekeeping, and mobility and speed profiling. The LTE aspect of this module ensures compliance with FR2.2, FR2.3 (real-time and backlog transmission), NFR2.2 (reliable connectivity during high-speed automotive use), NFR2.5 (remote interface security), and plays a significant role in achieving NFR2.4 (latency below 2 s, covered in Section 3.4). The GNSS interface

supports FR1.3 (position error $\leq 10\text{ m}$ and velocity error $\leq 5\text{ km/h}$), the first part of FR2.4 (GNSS timekeeping with RTC fallback), FR2.7 (GNSS TTFF $\leq 45\text{ s}$), and NFR1.4 (timestamp alignment $\leq \pm 1\text{ s}$).

Breakout Board Justification and Features

For the second and final version, a ready-made breakout board was used primarily for cost reasons, as it was provided at no cost. It includes breakout headers, GNSS and LTE U.FL connectors via a $50\ \Omega$ CPWG, required I/O connections, a bias-tee for an active GNSS antenna, a USB port with the latest firmware preloaded, and a nano-SIM slot.

External Interface Circuitry

The module communicates via a 1.8 V UART port, level-shifted to 3.3 V using an SN74AXC1T45 bidirectional translator. The low-side supply of the level shifter is driven by the 1.8 V V_{INT} pin on the LENA module. The circuit shown in Figure 3.13 integrates the LENA module interface, including voltage-level translation, GPIO-controlled LED drivers, and local decoupling for stable operation. The SN74AVC4T774 (U24) enables push-pull unidirectional level shifting and adapts UART logic levels using the module's internal V_{INT} supply. It supports up to 200 Mbit s^{-1} operation and provides four independent channels for translation between 1.8 V and 3.3 V , with two channels utilised, each controlled via dedicated direction pins. The two headers (J15 and J13) on the top right provide mechanical and electrical connections for the LENA module. The lower section includes 3.3 V bypass capacitors for the level shifter and LED circuitry, and two transistor-based LED drivers (bottom right) controlled by LENA GPIO outputs for status indication.

GNSS Performance Summary

The integrated GNSS system utilises the u-blox UBX-M10050-KB-A0100A receiver [60]. It tracks up to four of the five global navigation satellite constellations (GPS, Galileo, GLONASS, BeiDou, and QZSS) across their primary L1/L5 bands, allowing the receiver to draw on whichever signals are available to maximise positioning reliability.

The receiver is capable of operating under high-dynamic conditions – up to 4 g acceleration, altitudes up to 800 m , and speeds up to 500 m s^{-1} – without loss of tracking. It achieves a velocity error of 0.05 m s^{-1} (0.18 km h^{-1}), which is sufficient for land vehicle navigation. The time-to-first-fix (TTFF), which measures how long the module takes to determine its position from startup, is $23\text{--}28\text{ s}$ under cold-start conditions (no prior satellite or location data), and 1 s when warm-started or assisted. Position updates are available at a rate of $1\text{--}18\text{ Hz}$ in default mode, and up to 25 Hz in high-performance configurations. RF sensitivity refers to the lowest satellite signal strength that can be used for tracking, with more negative values indicating better performance. The module achieves -167 dBm for tracking/navigation, -160 dBm for reacquisition (after short signal loss), -159 dBm for hot start, and -148 dBm for cold start. These values are considered excellent, allowing reliable operation even in partially obstructed environments such as urban canyons or

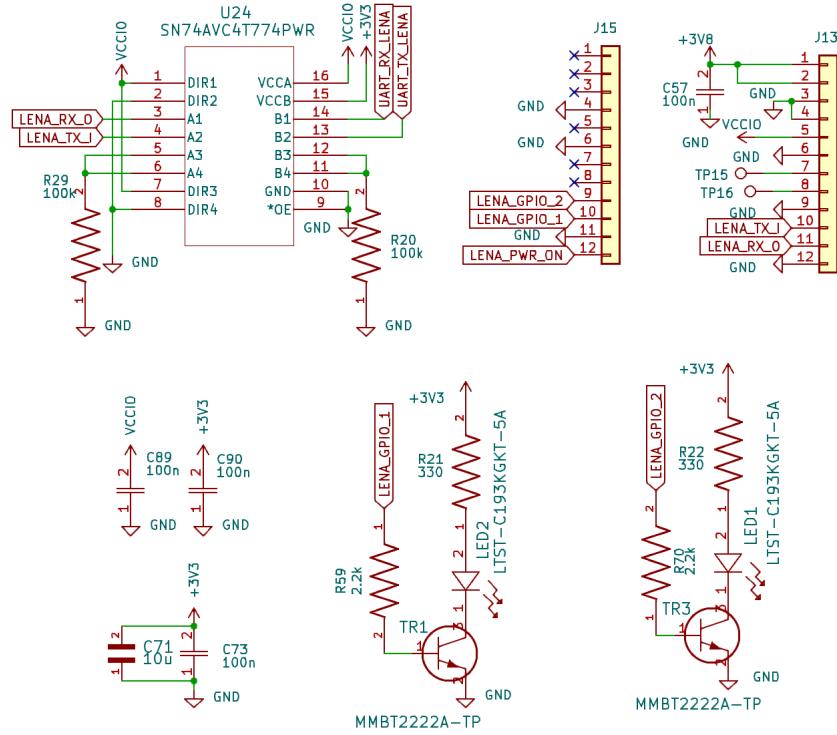


Figure 3.13: LENA module interface circuit showing level shifter (top left), module headers (top right), 3.3 V bypass network (bottom left), and GPIO-driven LED drivers (bottom right).

under light foliage. The typical horizontal error (measured as Circular Error Probable, CEP) is approximately 1.5 m under open-sky conditions when multiple constellations are used. In worst-case conditions (e.g., GLONASS-only), errors may increase to around 4 m. For complete parameter details, refer to the u-blox UBX-M10050-KB-A0100A datasheet.

Therefore, this configuration meets FR1.3 (GNSS position error \leq 10 m; velocity error \leq 5 km/h) and FR2.7 (TTFF \leq 45 s; warm reacquisition \leq 3 s).

Cellular Performance Summary

The LENA-R8 supports all major FDD LTE bands, including Vodacom's primary bands: B1 (2100 MHz), B3 (1800 MHz), and B8 (900 MHz). It transmits as a UE Power Class 3 device, which corresponds to a maximum output of 23 dBm (approximately 200 mW). This is typical for mobile user equipment and ensures adequate range for cellular communication without requiring excessive power.

The module implements LTE Category 1 bis, a classification that defines its peak data rates – 10.3 Mb/s downlink and 5.2 Mb/s uplink – which are more than sufficient for periodic telemetry capture and transmission at the required 1 Hz, or even at significantly higher rates in future applications. As per literature Section 2.3.3, LTE latency is below 100 ms when coverage is available.

Receiver sensitivity refers to the minimum inbound signal strength that the modem may decode accurately 95% of the time. On LTE B1, the sensitivity is -101 dBm at 5

MHz channel bandwidth (-95 dBm at 20 MHz); on B3, -106 dBm at 1.4 MHz (-101 dBm at 5 MHz); and on B8, -101 dBm at 5 MHz (-98 dBm at 10 MHz). In all cases, lower (more negative) dBm values indicate the ability to decode weaker signals. Typical inbound signals are almost always above -100 dBm in covered areas, so the device has excellent coverage.

2G GSM/GPRS supports up to 85.6 kbps downlink (DL) and uplink (UL), with typical latency in the 500–1000 ms range. Receiver sensitivity is approximately –104 dBm on GSM 900 and –102 dBm on GSM 1800, comparable to LTE coverage thresholds but at much lower throughput.

The LTE and GSM/2G transmit power and receiver sensitivity follow the mandatory 3GPP conformance specifications (3GPP TS 36.521-1 for LTE and 3GPP TS 51.010-1 for GSM). Since the LENA-R8 is a certified module, it fully complies with these requirements, ensuring adequate uplink power and minimum guaranteed sensitivity without further custom validation.

Due to low relative performance of 2G, it is suited only as a short-duration fallback in rural areas, maintaining basic connectivity and limiting backlog growth during LTE outages. Regulator design enables the usage of 2G, but it is not implemented.

For full performance details across all supported bands, refer to Table 17 in the LENA-R8 datasheet [60].

Antenna Choices

The chosen Mobile Network Operator, Vodacom, utilises band 1 (2100 MHz), 3 (1800 MHz), and 8 (900 MHz) for LTE connections. The chosen antenna (X9000984-4GDRMW [125]) operates between 790–960 MHz and 1710–2690 MHz, which meets these requirements. The input impedance is $50\ \Omega$, consistent with the LTE module and coaxial feed.

In free space, the LTE antenna covers 790–960 MHz and 1710–2690 MHz, aligning with national LTE bands. The return loss ranges from –4 to –6 dB, corresponding to $\approx 60\text{--}75\%$ of RF power delivered. The radiation efficiency is $>40\%$ in the low band and $>60\%$ in the high band, and the peak gain is 2.2–3.8 dBi, indicating mild directionality relative to an isotropic radiator. Overall, the antenna represents mid-level commercial performance, adequate for vehicular telemetry in urban areas, though in rural contexts a higher-gain and more efficient antenna would improve link robustness; for cost reasons, this was not purchased.

There are a significant number of reliable GNSS bands available in South Africa. This includes GPS, Galileo, BeiDou, QZSS, and GLONASS, all owned by separate entities. The selected GNSS antenna (ANT-GNRM-L1A [126]) supports all L1 bands, enabling multi-constellation operation. L2 and L5 are also widely used in GPS systems, with dual-frequency combinations (e.g. L1+L5) enabling direct correction of ionospheric delay and achieving sub-metre accuracy. This level of precision is not required for the present

application, where L1-only reception suffices for electromobility profiling.

Across the supported bands (1561-1602 MHz), the return loss lies between -14.5 and -19.9 dB (around 96-99% of power delivered), the peak gain including the integrated LNA is 32.8-33.5 dBi, and the axial ratio ranges from 3.5 to 11.4 dB. The LNA provides 28 dB gain with a noise figure of 0.3 dB, ensuring high sensitivity for weak signals. Overall, the antenna provides strong multi-constellation L1 support, suitable for vehicular GNSS positioning in South Africa.

Outcome

The selected LTE-GNSS module, antennas, and associated circuitry provide robust cellular connectivity and accurate positioning capabilities. These subsystems jointly address FR1.3, FR2.2, FR2.3, FR2.4, FR2.7, NFR1.4, NFR2.2, NFR2.3, NFR2.5 and NFR2.4.

SD-Card Interface

A hinged microSD socket is employed to minimise horizontal insertion space while ensuring reliable field accessibility, aligning with the serviceability and retention considerations identified for removable logging media. Pull-up resistors of $10\text{ k}\Omega$ are placed on the CMD and data lines to guarantee defined logic levels during card absence or power-up, thereby preventing unintended bus activity, as shown in Figure 3.14.

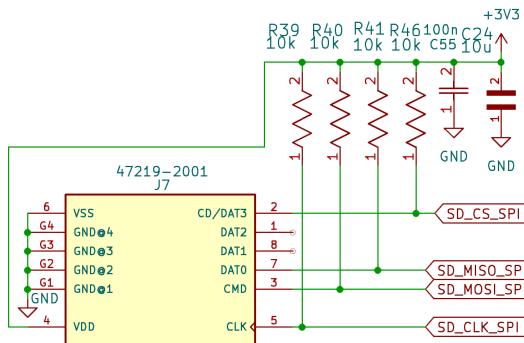


Figure 3.14: MicroSD interface circuitry showing hinged socket and pull-up network on CMD and data lines.

The preferred storage medium is Kingston's industrial microSD card, specified for -40°C to 85°C operation, dynamic wear levelling, and health telemetry suited to preventative maintenance. With a capacity of 16 GB, formatted as FAT32 according to the SDHC specification, the device significantly exceeds the two-year continuous logging requirement: even at a sustained 500 B s^{-1} , capacity would support more than 13 months of logging. Therefore, the SD subsystem meets FR2.1 (removable, PC-readable logging at 1 Hz) and NFR2.1 (at least 6 months of capacity at 1 Hz). For the prototype, a standard 16 GB commercial microSD card was used due to cost constraints, which has the same interface as the industrial version and meets the operating temperature requirement.

3.2.4. Power, Layout, and Mechanical Integration

This subsection describes the power and thermal design, PCB layout considerations, and mechanical integration. These aspects ensure efficient power delivery, signal integrity, and reliable operation under various environmental conditions.

Power Regulation

This section describes the power regulation system, which ensures all components receive stable and appropriate voltage levels. It ensures compliance with NFR3.4 (correct operation on a non-ideal vehicle 12 V supply) and NFR3.9 (the PCB power draw shall not exceed a typical 5 W and maximum 50 W). The section details the load requirements, regulator sizing, specifications, and regulator circuitry design.

The regulators were selected to ensure stable, efficient, and low-noise operation across all subsystems. The main 12 V isolated vehicle supply regulator and the 3.8 V regulator supplying the LENA modem are buck-boost converters, chosen to maintain strict output stability under input variations. The chosen wide-input buck-boost, buck, and isolated converters—together with a single LDO supplying the positive rail of the low-power ADS1220 ADC—support the wide voltage ranges, accuracy, and output stability required for typical vehicular environments while ensuring efficiency to minimise thermal buildup in the sealed case. All converters were chosen for extremely compact packages, small footprint requirements, no minimum load conditions, and good thermal coefficients, alongside design decisions encompassing careful layout, adequate copper area, and proper bypassing to minimise noise, ripple, and thermal buildup.

Load Assumptions and Conditions

The following assumptions clarify the operating modes of loads, applied simplifications, and notation.

To ensure sufficient regulator sizing, conservative worst-case scenarios are assumed wherever ambiguity in load demands arises. The MCU Wi-Fi is not utilised in this system, while SD-card current—unspecified by the manufacturer—is conservatively estimated from a sample-set study [127]. The fluxgate CAN bus, using two $60\ \Omega$ terminations, is conservatively assumed to draw one-fifth of its peak current at a 2 V differential. The fluxgate sensor is modelled for a mean of 100 A primary current and 1 kA at the system maximum. Both current sensors are budgeted to operate simultaneously to allow for mutual referencing during testing. Loads below 100 μA (e.g., SPI leakage) are excluded unless dominant for that device.

LTE and 2G operation are listed separately to allow future removal if 2G is no longer required. These rows are not summed; LTE defines the typical radio-access-technology (RAT) load, while 2G defines the maximum. The pull-up network, $(5 \times 22\ \text{k}\Omega + 4 \times 4.7\ \text{k}\Omega + 1 \times 2.2\ \text{k}\Omega + 4 \times 10\ \text{k}\Omega + 1 \times 4.7\ \text{k}\Omega)$, is evaluated assuming all lines are logic-low to capture

the worst-case draw.

The isolated current-sensing (“GMW”) domain is included in the analysis, whereas the isolated domain of the HVDC measurement circuit is excluded due to its passive nature. The 12 V input is divided into a “dirty” vehicle side and a “clean” isolated side; All subsystems are powered from the cleaned 12 V rail; the qualifier “clean” is omitted elsewhere for brevity.

Device-Level Load Summary

Table 3.7 lists power usage by individual device, while Table 3.8 groups loads by supply rail.

Table 3.7: Component current requirements per rail (excluding regulators).

Part	V_{rail} (V)	I_{typ} (mA)	I_{max} (mA)
<i>MCU</i>			
Wemos-S2-Mini	3.300	23.000	35.000
Pull-ups (all 3.3 V)	3.300	6.430	6.430
<i>GMW current sensing</i>			
ADS1220	2.500 _{GMW}	0.240	0.490
ADS1220	-2.500 _{GMW}	0.240	0.490
ADS1220	3.300 _{GMW}	0.075	0.110
ISO1641B	3.300	4.600	6.100
ISO1641B	3.300 _{GMW}	3.800	5.200
Main GMW IC	5.000 _{GMW}	86.000	86.000
<i>HVDC sensing</i>			
AMC131M01DFMR	3.300	17.000	22.000
ASE-8.000 MHz-LC-T	3.300	2.500	7.000
<i>LEM fluxgate sensor</i>			
Fluxgate sensor	12.000	100.000	1000.000
<i>RTC</i>			
MCP7940NT-I/MNY	3.300	0.300	0.400
<i>Acceleration sensor</i>			
ADXL343	3.300	0.030	0.140
<i>CAN bus</i>			
CAN termination	3.300	6.666	33.300
MCP2515 ($\times 2$)	3.300	10.000	20.000
TCAN3413 ($\times 2$)	3.300	60.000	120.000
CSS inductive reader	12.000	6.900	6.900
<i>MicroSD card interface</i>			
Hiksemi Neo 16 GB	3.300	50.000	200.000
<i>LTE/GNSS interface</i>			
LENA-R8 (LTE+GNSS)	3.800	170.000	690.000
LENA-R8 (2G+GNSS)	3.800	205.000	2000.000
Level shifter	3.300	0.005	0.008
Level shifter	1.800	0.005	0.008
LED bias ($\times 2$)	1.800	0.410	0.545
LED drive ($\times 2$)	3.300	3.330	6.500

Rail-Level Load Summary

Table 3.8 aggregates the load data by supply rail. These values represent only direct device power usage and exclude regulator headroom, efficiency losses, and downstream regulator dependencies.

Table 3.8: Typical and maximum currents for each supply rail (excluding regulators).

Rail	I_{typ} (mA)	I_{max} (mA)
1.8	0.008	0.011
2.5 _{GMW}	0.240	0.490
-2.5 _{GMW}	0.240	0.490
3.3	185.068	458.226
3.3 _{GMW}	3.875	5.310
5.0 _{GMW}	86.000	86.000
3.8	170.000	2000.000
12.0	106.900	1006.900

Regulator Requirements and Sizing

Regulator sizing accounts for voltage ratio, efficiency, and a 20% headroom margin for long-term reliability. The output current requirement of each regulator includes both direct loads and any downstream regulators, with the corresponding input current derived by applying efficiency and voltage ratio adjustments. Except for the LT3060-2.5, implemented as an LDO since it supplies less than 102 mW, all others are switching types to minimise thermal buildup within the sealed enclosure. All regulators operate within the required -5°C to $+55^{\circ}\text{C}$ range and are specified for wider limits. Efficiency figures are taken as worst-case for conservatism. Output capacitances are small enough that inrush-induced thermal stress is brief and non-destructive.

Table 3.9 summarises each chosen regulator's key specifications, namely their input/output rails, current rating, efficiency, and type for verification against load and thermal requirements. Additionally, the inline 12 V isolated regulator is automotive-grade and capable of compensating for voltage ripple amplitudes up to 2 V_{pp} (peak-to-peak) while maintaining high output stability. As an added safeguard, all other 12 V-connected regulators operate reliably over a wider input range than 9 V to 18 V, encompassing the full span of expected automotive supply conditions with high output stability. This ensures compliance with NFR3.4. Input ranges and detailed specifications are provided in the respective datasheets cited within the table.

Table 3.10 shows the computed input/output typical and maximum requirements per regulator. The vehicle 12 V supply draws a typical load of 3.41 W and a maximum of 32.18 W under simultaneous 1 kA primary current and 2G transmission bursts per Table 3.10. These values confirm compliance with NFR3.9 (typical 5 W, maximum 50 W) while maintaining sufficient thermal margin.

Table 3.9: Regulator specification summary listing input/output rails, peak current capability, efficiency, and topology.

Part	V_{in} (V)	V_{out} (V)	I_{out} (mA)	η (%)	Type
LMZM23601V3SILR [128]	12.0_clean	3.3	1000	81.0	Buck
LT3972IMSE [129]	12.0_clean	3.8	3500	82.0	Buck
CCG30-24-12S [130]	12.0_dirty	12.0_clean	2500	89.0	Iso DC-DC
TMR 1-1211 [131]	12.0_clean	5.0 _{GMW}	400	83.0	Iso DC-DC
LT3060-2.5 [132]	5.0 _{GMW}	2.5 _{GMW}	100	41.0	LDO
LTC1550-2.5 [133]	5.0 _{GMW}	-2.5 _{GMW}	100	1.7	Pump
18024115401H [134]	3.3	3.3 _{GMW}	50	54.0	Iso DC-DC

Table 3.10: Computed total regulator load requirements with 20% headroom and efficiency.

Part	$I_{out,typ}$ (mA)	$I_{out,max}$ (mA)	$I_{in,typ}$ (mA)	$I_{in,max}$ (mA)
LMZM23	230.693	561.671	78.322	190.691
LT397	204.000	2400.000	78.780	926.829
CCG30	341.618	2386.649	383.840	2681.628
TMR 1	112.022	121.211	56.236	60.849
LT306	0.288	0.588	0.351	0.717
LTC15	0.288	0.588	8.471	17.294
18024	4.650	6.372	8.611	11.800

Regulator Circuit Design and Performance

Figure 3.15 shows the main power regulation section positioned near the vehicle input connector. The top subsection stabilises, isolates, and filters the 12 V supply before distributing it to the downstream converters. The bottom right circuit generates the regulated 3.3 V rail supplying digital and logic domains, while the bottom left regulator provides the 3.8 V rail dedicated to the LTE modem. Figure 3.6 illustrates the isolated analogue regulation stage and subsequent dual-rail supply for the GMW current-sensing subsystem, producing independent 5 V and bipolar ± 2.5 V rails for the ADC and amplifier domains.

Table 3.11 compares secondary performance parameters such as noise, line/load regulation, accuracy, and temperature drift. All specifications are worst-case scenarios and confirm that regulator performance meets the needs of sensitive analogue and RF subsystems. A temperature variation of 0 is indicated where other specifications are guaranteed across the full rated temperature range. The automotive-grade input supply, wide input ranges of all subsequent regulators, and their stable outputs collectively ensure reliable operation under non-ideal vehicle supply conditions.

Outcome

The regulators meet all load and margin requirements with conservative assumptions. This ensures long-term power stability, with performance metrics confirming suitability for sensitive analogue and RF subsystems. The total typical PCB power consumption is

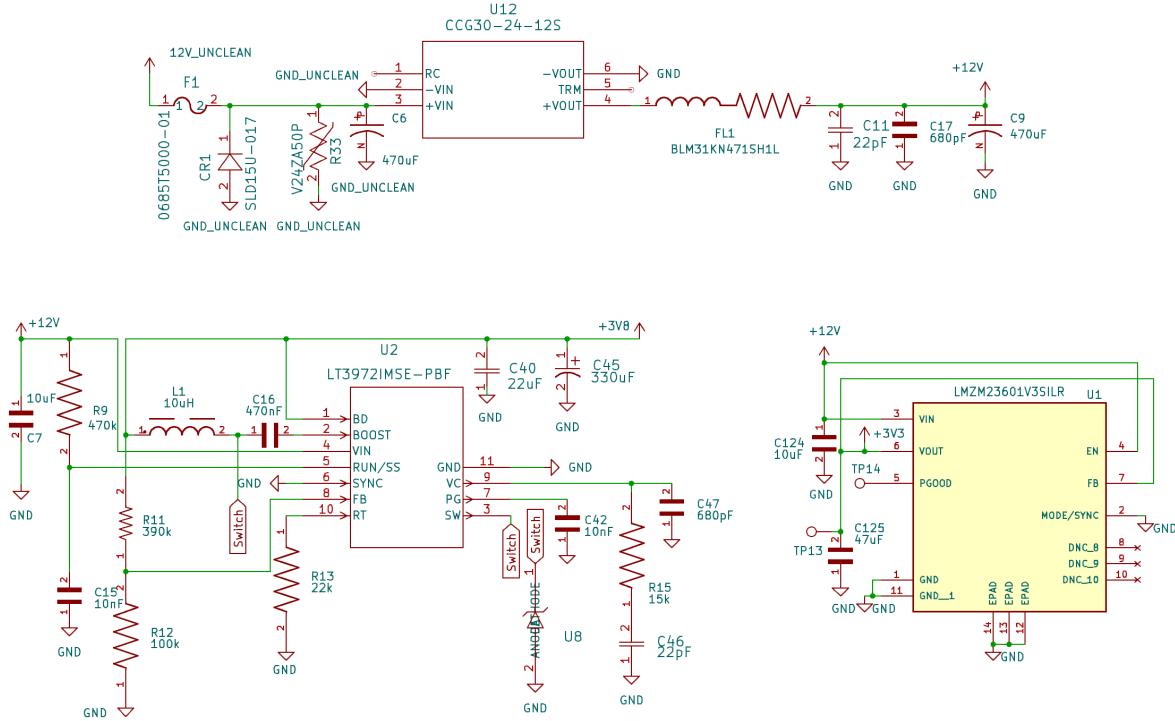


Figure 3.15: Primary power regulation chain located at the PCB input stage, showing the 12 V isolated regulation stage (top), followed by the 3.8 V regulator (bottom left) and the 3.3 V regulator (bottom right).

Table 3.11: Regulator performance metrics covering ripple, noise, accuracy, and temperature variation for each supply stage.

Part Unit	Noise [mV _{pp}]	Line reg [%]	Load reg [%]	Accuracy [%FS]	Temp var. [%FS]
LMZM23	3.30	2.0	2.0	1.50	0.00
LT397	15.00	0.29	3.0	2.00	0.00
CCG30	95.00	0.4	0.4	2.00	0.60
TMR 1	110.00	0.2	0.5	1.00	0.60
LT306	0.04	2.0	2.0	1.00	2.00
LTC15	1.00	2.5	2.5	2.50	2.50
18024	50.00	0.2	1.0	5.20	0.00

approximately 3.41 W, with a maximum of 32.18 W under rare peak conditions, ensuring compliance with NFR3.9. The regulation system maintains high stability under non-ideal vehicle supply conditions, thereby complying with NFR3.4.

PCB Layout

This subsection details the PCB layout design, which in part meets specified requirements, as well as general expected practices for EMC, thermal management, and mechanical integration. The PCB layout is shown below in Figure 3.16.

PCB Size

A central stakeholder requirement was that the PCB remain as compact as possible without compromising performance or safety, with a maximum allowed form factor of $\leq 150 \text{ mm} \times 150 \text{ mm} \times 50 \text{ mm}$ (NFR3.1). The final design measures 97.4 mm \times 127.4 mm \times 37 mm, meeting this requirement while providing sufficient spacing for high-voltage isolation, thermal management, and connector placement. The reduction in PCB size is illustrated in Section 3.2.5.

High-Speed Interfaces

SPI traces are kept short with controlled impedance to minimise reflections and crosstalk. Series 22 Ω resistors are used on data lines longer than 2 cm and on all clock lines, whose higher edge rates are more susceptible to high-frequency errors. Pull-ups are added on all SPI chip-select (CS) lines to prevent floating during power-up, avoiding unintended peripheral activation and spurious data transmission to the MCU.

Decoupling and Thermal Management

Decoupling and bulk capacitors are placed within 2 mm of each IC supply pin to stabilise the local supply and suppress high-frequency noise. Thermal vias are provided under voltage regulators and other heat-dissipating devices to transfer heat into internal planes. Exposed pads (EPADs) are employed on all regulators: the high-power 3.3 V and 3.8 V regulators additionally include adjacent ground and power pours to minimise loop area and reduce impedance in the return paths. Adequate via stitching connects top and bottom ground pours. Lower-power regulators rely on EPADs alone. The main 12 V filter regulator, which lacks an EPAD, instead uses its large package body and multiple pins to spread heat into the PCB and dissipate it to air. Wide traces are used on long power runs to minimise voltage drop, as well as for handling RF burst currents during LTE and 2G transmission.

Analogue-Digital Separation and Isolation

The GMW analogue current-sensing circuitry is physically separated into digital and analogue sections, without star-grounding or cutouts that could increase loop area. This reduces the risk of noise coupling from digital switching activity. Local decoupling

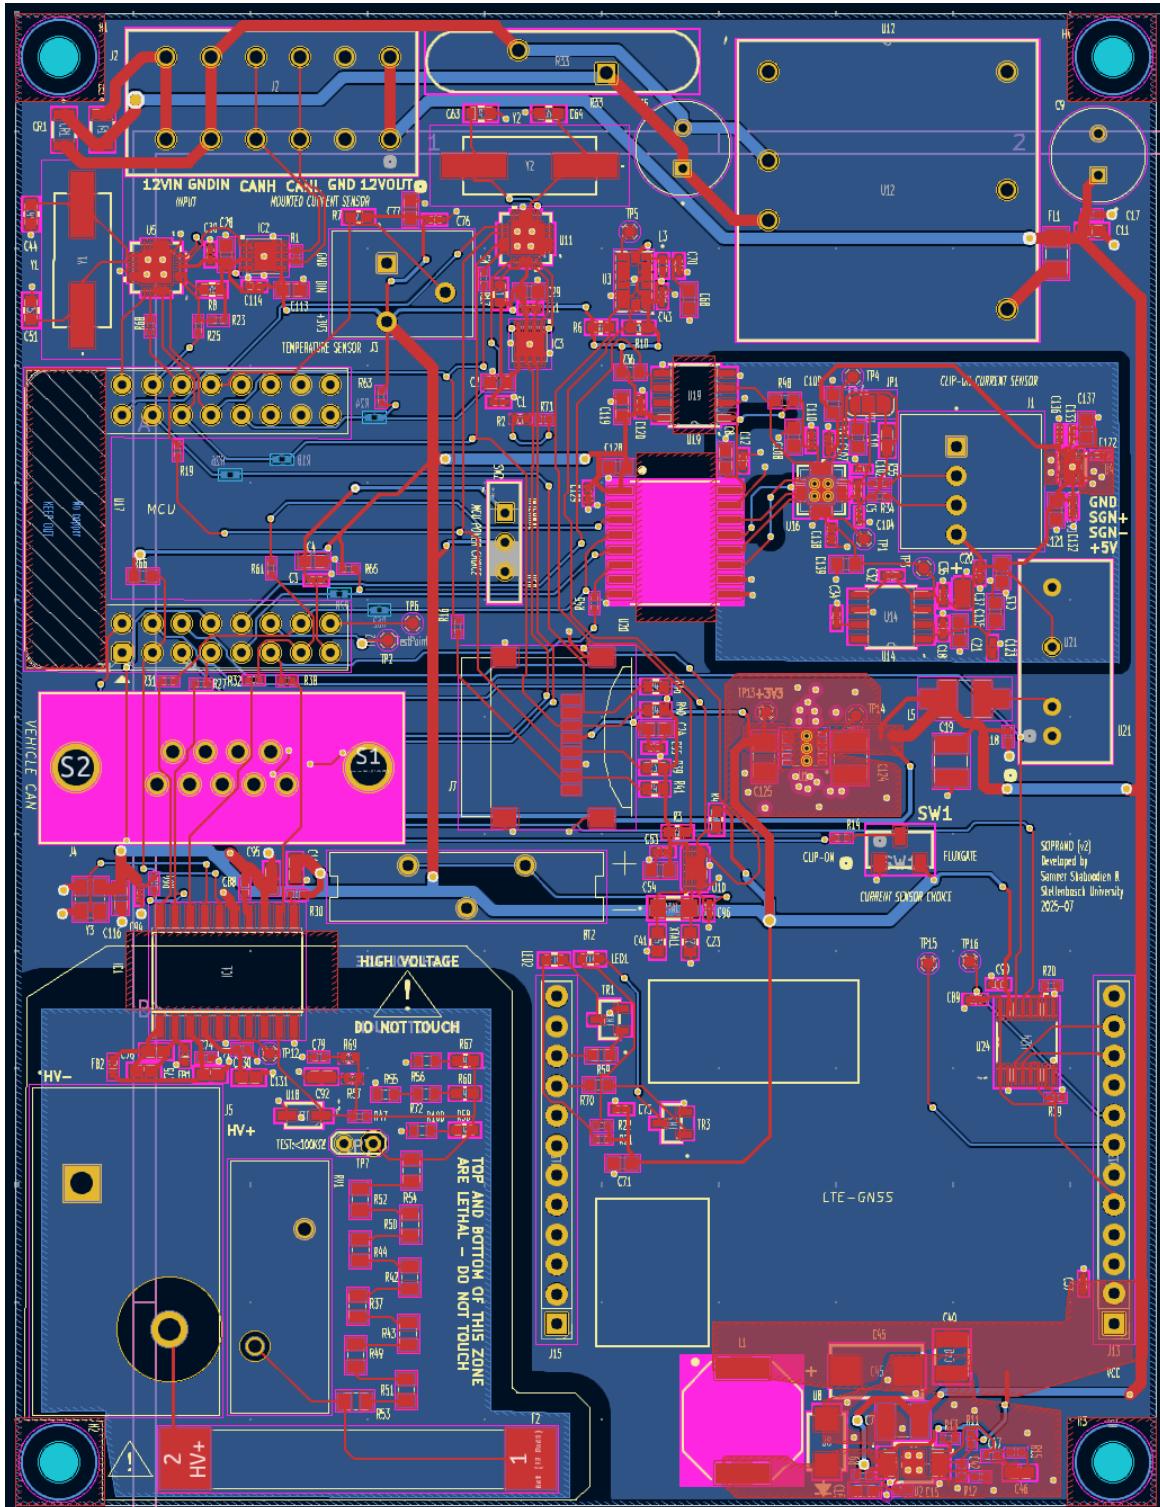


Figure 3.16: Final two-layer PCB layout with top and bottom copper overlays, highlighting the high-voltage domain (bottom left) and isolated GMW current digitisation section (centre right).

capacitors are placed within 2 mm of each analogue IC supply pin, and regulator filtering with loop minimisation further limits EMI on the ADC. The high-voltage sensing section is also physically isolated from the rest of the PCB, both to improve accuracy and, as explained in the voltage-sensing section, to form part of the basic insulation barrier required for safety.

Layer Stack-up and Routing

The PCB uses a two-layer stack-up with a continuous ground plane on the bottom layer, except for five 0402 resistors placed in series with SPI lines. Top-layer traces are predominantly routed vertically and bottom-layer traces horizontally to simplify routing and reduce crosstalk. Ferrite beads are used at the output of the 12 V input filter to suppress automotive conducted high-frequency noise that makes it through the main filter, travelling along the supply rails into sensitive circuitry.

Test Points and Debugging

Test points are included on unused MCU pins, the unused LENA-R8 module pins, the 3.3 V regulator output and its power-good pin, and the unused channel pins of the ADS1220 ADC used for GMW current sensing. In addition, a double test point is provided to measure the low-leg resistance to HV-, confirming that the sensing circuit is correctly closed and preventing an open-circuit condition that could damage the functional integrity of the component.

Connectors and Mechanical Integration

All I/O connectors are mounted vertically, or near-vertically, to simplify cable access and minimise horizontal space. All connectors are spring-cage based, providing secure retention and vibration resistance under automotive conditions without the loosening risk associated with screw terminals, thereby complying with NFR3.8 (shock and vibration resistance). Silkscreen markings identify pins (e.g., SGN+) and label subsystems (e.g., CAN, HVDC, GNSS/LTE) to prevent wiring errors. Mounting holes are positioned at the PCB corners to align with the standoffs of the selected Allbro enclosure.

The following subsystems employ the specified connectors or switches:

- **HV+ Input:** 2-pin spring-cage connector (part: 1098175) rated for 1 kV and 41 A, exceeding the requirements of the application.
- **12 V Input, Fluxgate Sensing:** 6-pin spring-cage terminal (part: 1751516), with two pins for the vehicle 12 V nominal input, two pins for the CAN output, and two pins for CANH and CANL. This connector is rated for 160 V and 17.5 A.
- **Temperature Sensor:** 3-pin compact spring-cage connector for 5 V, GND, and one-wire digital signal (part: 1990012).

- **GMW Sensor:** 4-pin spring-cage connector for 5 V, GND, and differential signal pair (part: 1989764).
- **Vehicle CAN Bus:** DB9 male connector interfacing with the inductive CAN bus sniffer (part: 61800929221).
- **Current Sensor Selector:** CAS-120TA mechanical switch, 3-pin, 2-position (part: CAS-120TA). The common pin connects to a GPIO via a series resistor, toggling between 3.3 V and GND.

Figure 3.17 shows the connector layout and pin assignments for all external interfaces on the PCB. The top right region contains the high-voltage spring-cage input from the vehicle's HVDC bus, while the top left houses the current-selection switch used to toggle between the GMW coreless probe and the LEM fluxgate sensor. The middle right section includes the vehicle 12 V input and the isolated 12 V output that powers the Fluxgate CAN interface. The middle left connector accepts the temperature-sensor input with its pull-up resistor network. The bottom left features the DB9 male connector for the CSS contactless CAN bus reader, and the bottom right provides the GMW current-sensor input.

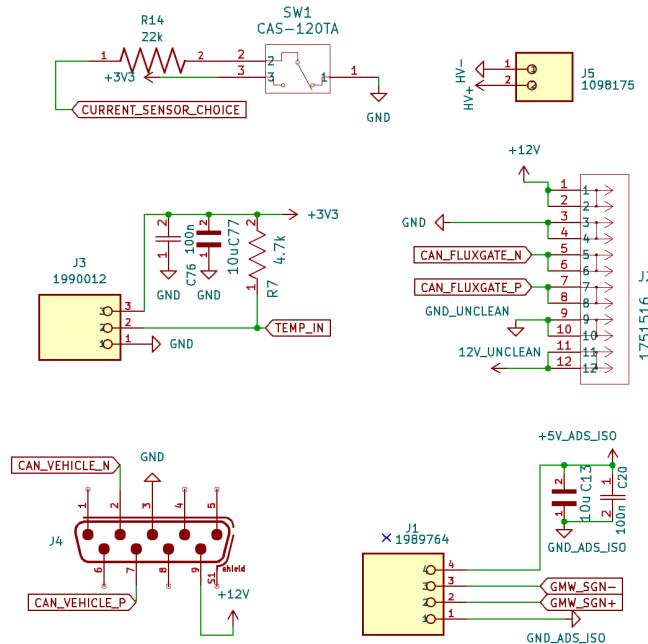


Figure 3.17: Connector layout showing high-voltage and vehicle inputs, GMW and fluxgate interfaces, temperature input, and DB9 port for the CSS CAN reader.

Enclosure

The PCB is housed in an Enlec® polycarbonate enclosure (part ENL161209P), rated to IP66 and IK10. The enclosure provides dust- and water-tight protection, mechanical robustness, and resistance to UV exposure and flammability (UL94 V0). The selected size

($160 \times 120 \times 90$ mm external, $146 \times 120 \times 86$ mm internal) accommodates the PCB with sufficient clearance for connectors and wiring. Integrated mounting posts (90×120 mm pattern) are used to secure the PCB, into which M5 mounting holes have been placed on the board. The self-tapping and lid screws are included with the product.

To ensure ingress protection, cables must enter the case through dedicated sealed glands with at least the ingress rating required (IP66, if matched to the enclosure). Sealants are typically added to fill any excess volume if the glands are oversized for the cable bundle. For development and testing, glands and sealants were not installed.

The enclosure selection meets NFR3.5 (ingress protection and mechanical robustness).

3.2.5. Hardware Versions and Modifications

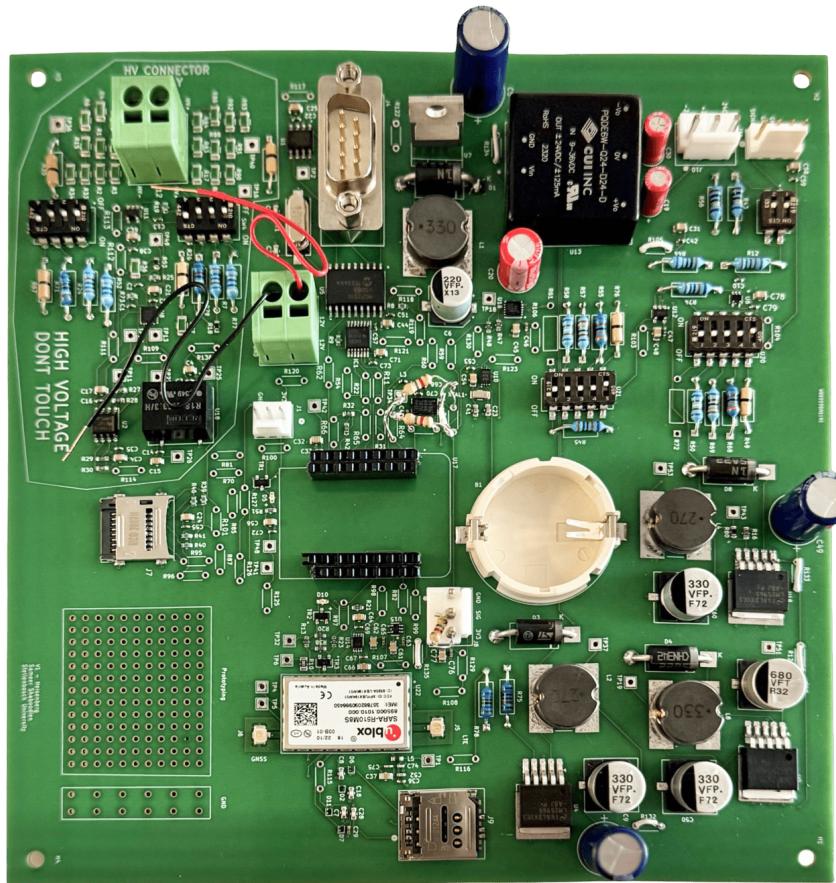
Two hardware iterations were developed during this thesis, with minor modifications applied to the final version.

Hardware Versions

The first iteration served as a functional prototype, allowing rapid design flexibility and expedited development. The second version achieved substantial improvements in accuracy, compactness, and safety.

Version 1, shown in Figure 3.18a at 1:2 scale, employed NB-IoT (NB2) via the SARA-R5 modem, incorporating RF design using a CPWG layout and coexistence techniques. However, subsequent research revealed the inapplicability of NB-IoT for vehicular contexts due to poor mobility handover performance and high communication latencies, both of which limited real-time use. Additionally, the Version 1 high-voltage section failed to meet basic insulation requirements. The accuracy of both the HVDC and GMW current-sensing stages were constrained by cascaded analogue stages, rail limitations, inaccurate bias voltages, and inadequate isolation between analogue and digital domains. Several regulators and passives used oversized packages, while cabling vibration resistance, regulator thermal dissipation, and EMC compliance were not sufficiently addressed. Hardware Version 1 also involved several modifications to test different configurations while realising these improvements.

Version 2, shown in Figure 3.18b at 1:2 scale, was designed to resolve these shortcomings. It adopted a single-antenna LTE Cat 1 bis architecture offering lower latency and higher throughput. The PCB area was reduced by 65.26% from $165.5 \text{ mm} \times 171.5 \text{ mm}$ to $77.4 \text{ mm} \times 127.4 \text{ mm}$, representing a major advancement in compactness. Simultaneously, sensing accuracy was substantially enhanced through rigorous research and selection of state-of-the-art ICs, optimised analogue design, and ultra-compact component packaging. The accuracy of the voltage and current subsystems was improved through the usage of integrated high-accuracy, dual-rail ADCs with matched-impedance filtering and physical isolation between analogue and digital domains. This reduced cascaded errors, rail



(a) First hardware iteration (Version 1) of the telematics system, shown at 1:2 scale.



(b) Final hardware iteration (Version 2) of the telematics system, shown at 1:2 scale with a 65% size reduction while achieving significant improvements in sensing accuracy and safety.

Figure 3.18: Comparison of first and second hardware iterations illustrating significant improvements.

limitations, and noise coupling while maintaining compactness. The high-voltage subsystem was also redesigned to meet basic insulation requirements, with further incremental improvements enhancing robustness and functionality.

Modifications

During Version 2 development, minor hardware alterations were made. The LENA modem's power-on pin, omitted in the initial schematic, was rerouted directly to the ESP32 microcontroller. A short-circuit incident damaged the voltage-sensing ADC clock line, the current-sensing SPI isolator, the onboard 3.3 V regulator, and the inline 12 V regulator. The ADC clock was reassigned to an equivalent 8 MHz clock source, and the isolator was removed with the ADC interface connected directly to the ESP32, showing no measurable performance degradation. The 3.3 V regulator was replaced with an equivalent component, while the 12 V regulator was bypassed and additional bulk capacitance added to the supply rail, as all downstream subsystems were tolerant of ripple and voltage variation. The HVDC fuse was not delivered and was shorted during development.

3.2.6. Requirement Traceability

The hardware design contributes to the satisfaction of several functional and non-functional requirements defined in Section 3.2 and Section 3.3. These contributions are summarised in Table 3.12.

Table 3.12: Traceability between hardware subsystems and the functional and non-functional requirements they satisfy.

ID	Requirement	Hardware contribution
Sensing		
FR1.1	Measure HVDC pack voltage up to 880 V with FSR error $\leq 1\%$.	Section 3.2.1: precision divider and isolated $\Delta\Sigma$ ADC yield a quantified 0.57%FS worst-case error.
FR1.2	Measure bidirectional HVDC pack current up to ± 1 kA with FSR error $\leq 1.5\%$.	Section 3.2.1: LEM fluxgate (± 1500 A, $\leq 0.5\%$) and GMW coreless probe subsystem (± 1000 A, 1.2%) provide complementary current coverage.
FR1.3	Acquire GNSS position with 3D position accuracy ≤ 10 m and velocity ≤ 5 km/h.	Section 3.2.3: multi-constellation LENA-R8 GNSS yields 1.5 m CEP and 0.18 km h^{-1} velocity accuracy.
FR1.4	Acquire 3-axis drivetrain acceleration up to ± 2 G with FSR error $\leq 10\%$.	Section 3.2.1: ADXL343 configured for ± 2 G and documented error budget $< 6.2\%$ FS.
FR1.5	Acquire ambient temperature from -5°C to 55°C with $\pm 1^\circ\text{C}$ resolution.	Section 3.2.1: DS18B20 sensor ($\pm 0.5^\circ\text{C}$) on one-wire bus.
FR1.6	Acquire vehicle CAN frames.	Section 3.2.3: inductive CAN coupler and TCAN3413/MCP2515 pair provide reliable access.
NFR1.1	Vehicle CAN bus data acquisition shall be non-intrusive.	Section 3.2.3: inductive coupler provides galvanic isolation and prevents drive onto the vehicle bus.

Continued on next page

ID	Requirement	Hardware contribution
NFR1.2	All sensing channels shall operate correctly from -5°C to 55°C while maintaining accuracy.	All cited error budgets are evaluated across at least -5°C to 55°C .
NFR1.3	Sampling bandwidth shall meet or exceed the Nyquist rate for each channel.	All sensors have bandwidths well above 1 Hz or provide internally downsampled outputs.
Data Handling		
FR2.1	Log all outputs locally to removable, PC-readable storage at 1 Hz.	Section 3.2.3: hinged microSD socket and 16 GB capacity.
FR2.2	Transmit all measurement channels to a cloud endpoint at 1 Hz when available.	Section 3.2.3: LTE Cat 1 bis modem with matched antenna provides sufficient uplink throughput for continuous telemetry.
FR2.3	Ensure missed historical data is retransmitted after connectivity outages.	Sections 3.2.3 and 3.2.3: modem plus buffered SD storage enable backlog upload once coverage returns.
FR2.4	Maintain timekeeping locally using GNSS with RTC fallback.	Sections 3.2.3 and 3.2.2: GNSS timekeeping plus MCP7940N RTC with battery backup and periodic resynchronisation.
FR2.7	Ensure GNSS TTFF $\leq 45\text{ s}$ (cold) and warm reacquisition $\leq 3\text{ s}$.	Section 3.2.3: LENA-R8 specifications report 23–28 s cold start and 1 s warm start.
NFR2.1	Local storage shall support continuous logging for ≥ 6 months at 1 Hz.	Section 3.2.3: 16 GB SD capacity covers more than 13 months at 500 B s^{-1} .
NFR2.2	Maintain LTE connectivity in covered regions up to 120 km h^{-1} .	Section 3.2.3: certified Cat 1 bis modem meeting 3GPP Tx power/sensitivity limits and vehicular-rated antennas ensure highway operation.
NFR2.3	RTC time drift shall not exceed 1 s h^{-1} .	Section 3.2.2: RTC drift less than 0.4158 s h^{-1} .
Retrofittability, Generalisability and Robustness		
NFR3.1	PCB assembly shall be below $15\text{ cm} \times 15\text{ cm} \times 5\text{ cm}$.	Section 3.2.4: final PCB measures $97.4\text{ mm} \times 127.4\text{ mm} \times 37\text{ mm}$.
NFR3.2	The system shall not require permanent alterations to the host vehicle's electronic systems.	Sections 3.2.1, 3.2.1, and 3.2.3: clamp-on voltage taps, fluxgate/coreless probes, and inductive CAN reader avoid cutting conductors.
NFR3.3	HVDC power measurements shall remain electrically safe up to 880 V.	Section 3.2.1: galvanic isolation meeting 1 kV basic insulation, distributed resistances, and transient protection components.
NFR3.4	The system shall operate correctly from a non-ideal vehicle 12 V supply.	Section 3.2.4: automotive-grade, wide-input-range regulators with high stability.
NFR3.5	The enclosure shall protect against dust, water, and mechanical impact.	Section 3.2.4: polycarbonate IP66/IK10 enclosure with captive screws secures the PCB.
NFR3.6	Sensing probes shall accommodate typical EV power bus and CAN wire sizes.	Sections 3.2.1 and 3.2.3: fluxgate aperture 16.2 mm, coreless probe 27 mm, inductive CAN clamp supports approximately 18–30 AWG conductors.
NFR3.7	Maintain accuracy without per-device calibration or vehicle-specific CAN knowledge.	Accuracy budgets in this table assume zero calibration, and required signals remain measurable without CAN access.
NFR3.8	Operate reliably under vibration and mechanical shock.	Sections 3.2.4 and 3.2.4: enclosure standoffs and spring-cage terminations.

Continued on next page

ID	Requirement	Hardware contribution
NFR3.9	Total power draw shall not exceed 5 W average and 50 W peak.	Section 3.2.4: load analysis shows typical 3.4 W draw and < 33 W peaks during LTE bursts.

Summary

The hardware design combines precision sensing, reliable processing, robust communication, and efficient power management within a protected enclosure, thereby fulfilling the functional and non-functional requirements necessary for a safe, accurate, and deployable EV telematics system.

3.3. Firmware System

This section describes the firmware for the telematics node: how it satisfies the data-handling and communication requirements, how sensing is orchestrated at 1 Hz, and the design decisions that prioritise reliability, time alignment, and non-intrusiveness on the vehicle. The firmware Git repository is available at <https://github.com/Samshabz/Masters-Firmware>.

3.3.1. Introduction

The firmware acquires all sensing channels, produces one data record every second, stores it locally, and transmits it to the cloud when connected. It keeps time using GNSS with RTC fallback (no LTE clock use) and guarantees eventual delivery by retransmitting missed data after outages. The code is organised into clear modules (drivers, data handling, communications, time) with a second-aligned schedule driven by a 100 ms time poll and a 500 ms post-rollover commit.

3.3.2. Architecture

The firmware follows a modular architecture, shown in Figure 3.19, which prioritises timing determinism, non-blocking execution, and robust communication. Each subsystem operates independently but is synchronised through a sample-aligning scheduler to ensure consistent sampling, logging, and transmission at 1 Hz.

At startup, peripheral buses, sensors, storage, and LTE/GNSS modules are initialised sequentially, followed by MQTT session establishment. The runtime loop then continuously cycles through non-blocking sensor polling, LTE event handling, GNSS or RTC UTC second rollover detection, and both real-time and backlog record publishing. If GNSS is unavailable or the device has been running for more than 15 minutes without resynchronisation, the RTC is updated from the next available GNSS fix to preserve time accuracy.

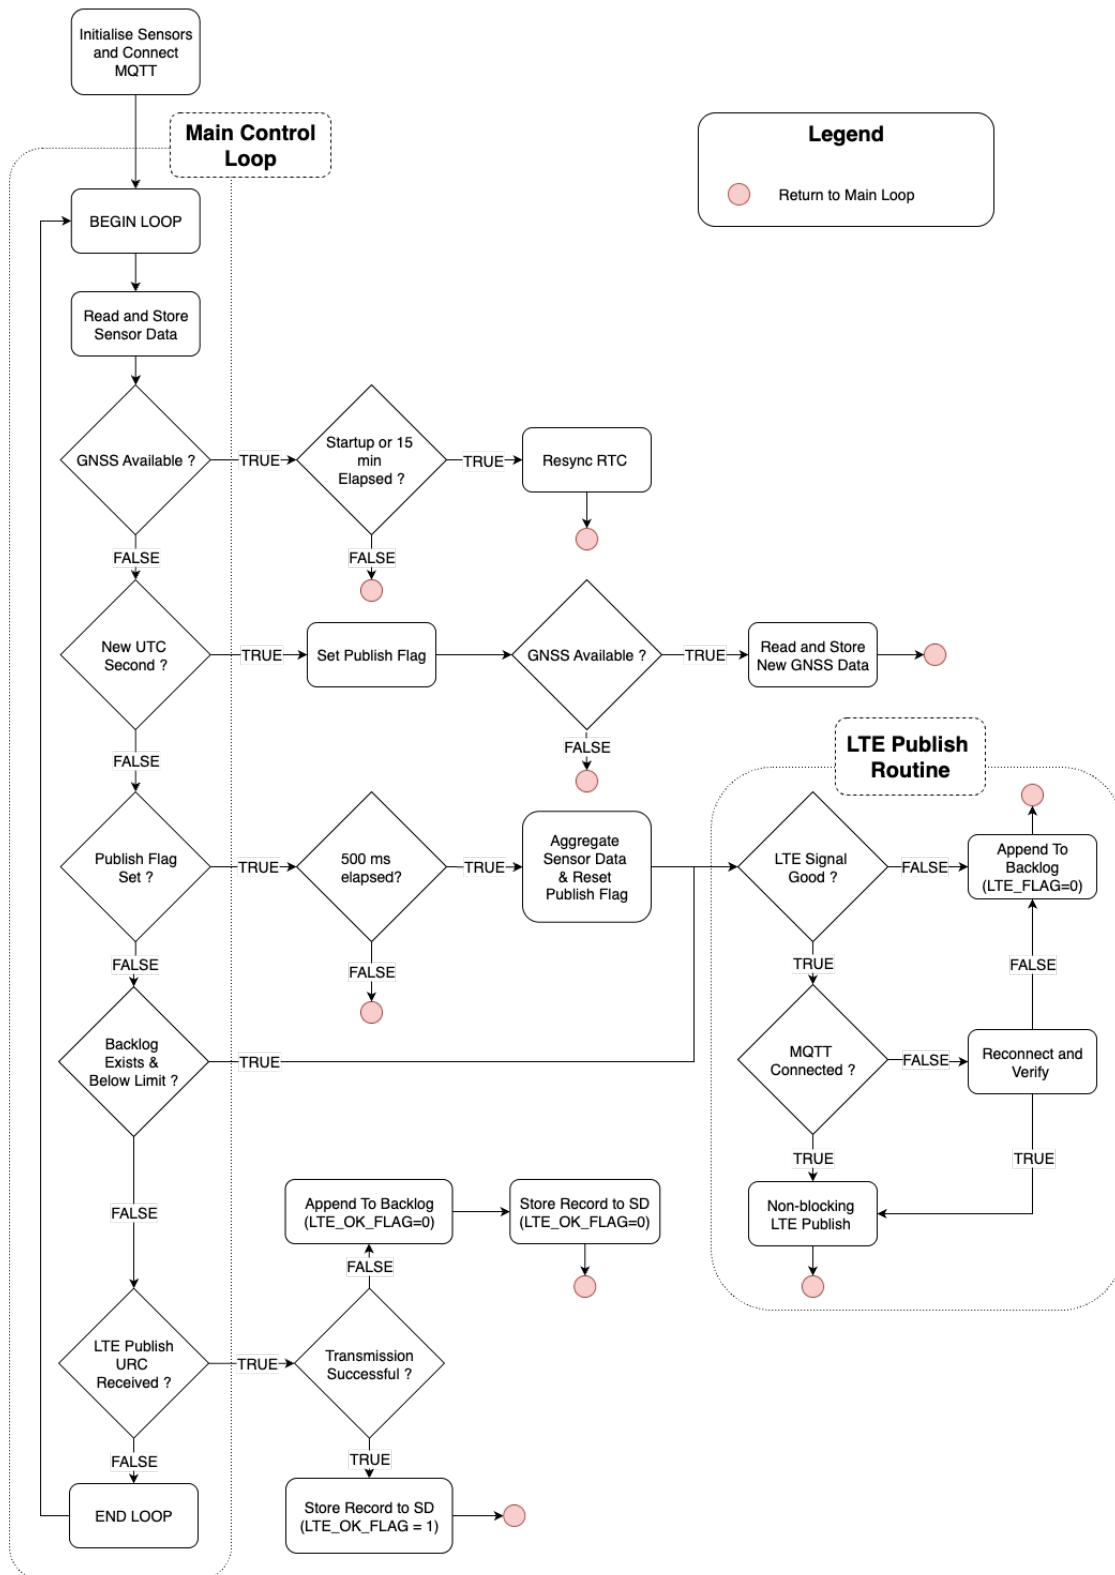


Figure 3.19: Firmware architecture and control flow showing modular layering and second-aligned scheduler.

Sensor drivers push data into shared buffers, and the aggregated samples represent the complete preceding one-second window. Each record is published 0.5 s after the UTC second rollover, ensuring that averaged sensor values are temporally centred and accurately aligned to real time. When LTE transmission fails or signal quality is poor, the record is appended to a backlog for later retransmission once connectivity resumes. MQTT connection checks and LTE publish attempts occur non-blocking within each loop, with retries only triggered on failed transmissions or disconnected states.

The firmware architecture is therefore designed to maintain deterministic one-second cycles with non-blocking handling of LTE unsolicited result codes (URCs), GNSS updates, and sensor reads, ensuring that sensing, logging, and communication remain responsive and temporally aligned under all operating conditions.

3.3.3. Sensing Interfaces

Driver-level configurations and sampling characteristics are as follows:

- **High-voltage pack voltage:** Measured via an isolated AMC131M01 series ADC over SPI. The driver operates in 24-bit mode with an oversampling ratio (OSR) of 2048 and global-chop enabled, yielding an effective rate of approximately 2000 SPS. Firmware reads on DRDY interrupts and applies hardware scaling to convert ADC input volts to pack volts.
- **Pack current:** Acquired from either the ADS1220 (via the GMW probe) or a fluxgate sensor over dedicated CAN. Firmware supports ADS-only, flux-only, or autoselect modes based on the external current-sensor selection switch. The ADS1220 samples every 22 ms (\approx 45 SPS), while the fluxgate broadcasts 24-bit current frames at 100 Hz (10 ms period) with a DC bandwidth of approximately 20 Hz.
- **Vehicle CAN:** Receives vehicle-transmitted frames from the MCP2515 RX buffers via interrupt-driven polling. This optional, vehicle-specific interface is not part of the main firmware and is implemented only for the electric vehicle test, where power frames were received at 10 Hz and downsampled to 1 Hz as test references. The hardware interface is inductively coupled and thereby receive-only.
- **Accelerometer (ADXL343):** Receives 3D acceleration data via SPI at a 100 Hz output rate and $\pm 2G$ range.
- **Temperature (DS18B20):** Measures ambient temperature over a One-Wire interface with 9-bit resolution and a 93.75 ms conversion time (\approx 10 Hz).
- **GNSS (LENA module):** Enables potential concurrent tracking of GPS, Galileo, GLONASS, and BeiDou B1C constellations, configured to output position, speed,

time, heading, and diagnostic fields (GGA, GSA, VTG, and RMC messages) upon poll after each UTC second rollover. The modem internally downsamples to its configured output rate of 1 Hz and synchronises updates to UTC time. Querying the module every 100 ms is sufficient to detect the rollover and align both the record timestamp and the mean sensor data within ± 100 ms of UTC.

Drivers operate from 10 Hz to several hundred hertz and aggregate data within a one-second window. The firmware polls GNSS or RTC time every 100 ms to detect the UTC rollover, then commits each record at 500 ms past the second. This ensures that both timestamps and aggregated sensor means are centred on UTC, achieving internal synchronisation and UTC alignment within ± 100 ms (NFR1.4). All drivers are internally band-limited or sample at rates sufficient relative to their signal bandwidths to satisfy Nyquist sampling requirements (NFR1.3).

3.3.4. Data Handling, Storage, and Communication

This section describes the firmware's complete data pipeline: how sensor records are stored locally, transmitted to the cloud, and recovered during connectivity loss. The architecture is designed to ensure integrity, continuity, and compliance with FR2.1–FR2.5.

Local Storage and Logging Rate

Records are appended to date-rotated CSV files (e.g., `/log_YYYYMMDD.csv`) on the SD card mounted over HSPI. Each file contains one record per second (1 Hz, FR2.1). Field definitions and schema alignment with the MQTT payload are described in Section 3.3.4. On card absence or write failure, the logger retries on the next loop without blocking real-time acquisition.

MQTT Uplink and Logging Rate

When LTE connectivity is available, data is published at 1 Hz (FR2.2) via MQTT over TLS using the modem's AT command stack. Each valid `Sample` is serialised to JSON, wrapped in the modem's required hex mode, and sent on the device's topic. Publish acknowledgements are captured through URCs within a non-blocking window. TLS is enforced end-to-end (NFR2.5), with credentials pre-provisioned via a separate certificate tool.

Backlog and Recovery Policy

If a publish fails or times out, the payload is appended to a date-rotated backlog file on the SD card. Upon reconnection, the oldest backlog entries are flushed in small batches alongside real-time data (FR2.3). This ensures deterministic loop timing and prevents runaway processing. LTE attach and MQTT reconnects follow an exponential backoff (5 s–60 s) gated by the measured RSSI (`AT+CSQ`). When signal quality improves, retries resume immediately.

Certificate Provisioning

Certificates and keys are uploaded once by a standalone Arduino provisioning tool. The tool transfers the CA, client certificate, and private key (PEM format) to the modem via `AT+UDWNFILE`, registers them with `AT+USECMNG`, and links them to a TLS profile using `AT+USECPRF`. The runtime firmware only references this stored profile and does not handle certificate updates.

Payload Format (FR2.5)

Each record is emitted once per second as a JSON object with nulls for unavailable fields. The parameter `vehCAN` is application-specific, while `src` indicates the active current-sensing source (ADS, fluxgate, or both). The modem hex-encodes the JSON payload before publish; IoT Core automatically decodes it back to JSON. The SD-card CSV logs use the same schema without hex encoding. Example:

Listing 3.1: Example MQTT payload and CSV-equivalent JSON structure.

```
{
  "deviceId": "BMW_XIII",
  "timestamp": 1727808005123,
  "src": 1,
  "lte_ok": 1,
  "voltage_V": 387.125,
  "current_A": -54.213,
  "power_kW": -20.989,
  "accel_x": 0.14, "accel_y": -0.02, "accel_z": 9.72,
  "temp_C": 26.8,
  "lat": -33.9321, "lon": 18.8605, "alt": 104.2,
  "speed_m_s": 12.3, "heading_deg": 190.1, "gnss_quality": 0.9,
  "gnss_ok": 1,
  "vehCAN": 0
}
```

3.3.5. Timekeeping and Synchronisation

Timekeeping and timestamping: GNSS is preferred; RTC is a fallback. The firmware polls GNSS/RTC time every 100 ms, detects the UTC second rollover, then commits the record at 500 ms past the second. On rollover, fresh GNSS fields are read and included so all channels are centred on UTC. RTC is periodically resynchronised to GNSS when drift exceeds 500 ms. This achieves a maximum timing offset of ± 100 ms relative to UTC (NFR1.4).

Startup and connectivity (FR2.6): peripheral initialisation and SD mount are kept lean. LTE attach and MQTT connect are attempted in `setupSensors()` via `initLteGnss()/initLTE()`; subsequent reconnect attempts run near the start of each loop

using `ensureLTEConnected()` with an exponential backoff (5 s, 10 s, 20 s, 40 s; capped at 60 s). GNSS is enabled at boot; timekeeping follows the second-centred strategy described above.

3.3.6. Firmware Robustness and Diagnostic Hooks

The firmware is designed for robust operation under partial subsystem failure and degraded connectivity. Failed sensors do not disrupt execution or extend loop times, as they are automatically flagged as inactive when initialisation fails. Absence of GNSS or LTE coverage does not cause blocking delays due to pre-checks of signal validity, non-blocking transmission routines, and controlled timeout handling. Operation continues even without an SD card, while separating the backlog file from the main log reduces the risk of data corruption and preserves recoverability after outages.

Concise runtime diagnostics support verification of correct operation and compliance with functional and non-functional requirements. A 1 Hz console summary reports the epoch, SD and MQTT success flags, as well as all sensor values and states. GNSS diagnostics output the fix dimension, HDOP, and satellite count, while timekeeping hooks record RTC resynchronisation and slow-loop events. The backlog flusher reports transmitted counts when debug mode is active, assisting evaluation of transmission behaviour.

3.3.7. Requirement Traceability

Table 3.13: Traceability matrix showing how firmware features satisfy the functional and non-functional requirements.

ID	Requirement	Firmware contribution
Data Handling and Communication		
FR2.1	Log all outputs locally to removable, PC-readable storage at 1 Hz.	Sections 3.3.2 and 3.3.4 describe the second-aligned scheduler that calls <code>processSampleAndPublish()</code> each rollover, appending CSV rows to SD storage for PC-readable archives.
FR2.2	Transmit all measurement channels to a cloud endpoint at 1 Hz when available.	Section 3.3.2 details the LTE MQTT uplink executed after each commit; <code>flushBacklog()</code> streams the JSON payloads over the modem whenever coverage is present, sustaining the 1 Hz telemetry rate.
FR2.3	Ensure missed historical data is retransmitted after connectivity outages.	Section 3.3.2 buffers unsent records and budgets retransmission inside <code>flushBacklog()</code> , while Section 3.3.4 explains backlog persistence and oldest-first replay until the SD queue is cleared.

Continued on next page

ID	Requirement	Firmware contribution
FR2.4	Maintain timekeeping locally using GNSS with RTC fallback.	Section 3.3.5 outlines the GNSS-first epoching, periodic RTC resynchronisation, and second-centred commit that preserve absolute timing when GNSS is unavailable.
FR2.5	Provide secure, cloud-accessible telemetry for remote analysis.	Section 3.3.2 shapes the 1 Hz JSON payload for the cloud pipeline, and Section 3.3.6 keeps the schema aligned with downstream processing, enabling remote visualisation and downloads.
FR2.6	Achieve reliable startup and reconnection behaviour.	Section 3.3.2 performs LTE attach during setup and enforces bounded waits, while Section 3.3.5 applies exponential backoff retries (5 s–60 s) to re-establish connectivity without blocking the 1 Hz loop.
FR2.7	Limit GNSS time-to-first-fix (TTFF) and reconnection delays.	Section 3.3.5 leverages the modem's GNSS subsystem for cold and warm starts, polling at 100 ms to capture the second rollover and maintain sub-second reacquisition.
Sensing and Timing		
NFR1.3	Sampling bandwidth shall meet or exceed the Nyquist rate for each channel.	Section 3.3.3 configures ADC, CAN, and sensor drivers at ≥ 10 Hz up to 2 kSPS, aggregating within the one-second window to avoid aliasing before downsampling to 1 Hz.
NFR1.4	Inter-channel and timestamp misalignment shall be ± 1 s.	Section 3.3.5 keeps UTC-centred commits at 500 ms past rollover, constraining timestamp skew to ± 100 ms.
Security		
NFR2.5	Remote interfaces shall enforce secure access.	Section 3.3.2 uses TLS-authenticated MQTT sessions with device certificates, and Section 3.3.4 maintains credential handling within the firmware to protect the uplink path.

The firmware therefore meets the functional and non-functional requirements through a modular architecture that sustains reliable sensing, logging, synchronisation, and secure communication.

3.4. Cloud System

The cloud platform links embedded telematics hardware to the operator-facing web application, providing a secure and scalable environment for data ingestion, processing, storage, and retrieval. This section is divided into four parts, namely: messaging and compute; storage; security and identity management; and scalability, observability, and cost.

3.4.1. Architecture Overview

Figure 3.20 illustrates the overall data and control flow from embedded telematics hardware to cloud databases and web dashboards. Embedded devices communicate with the cloud through lightweight, event-driven protocols, where serverless compute validates and stores payloads in managed databases. Authenticated users retrieve data or issue commands asynchronously via secure APIs, forming a cost-effective architecture since sub-second latency is unnecessary for telematics applications that do not require control-system-grade response times.

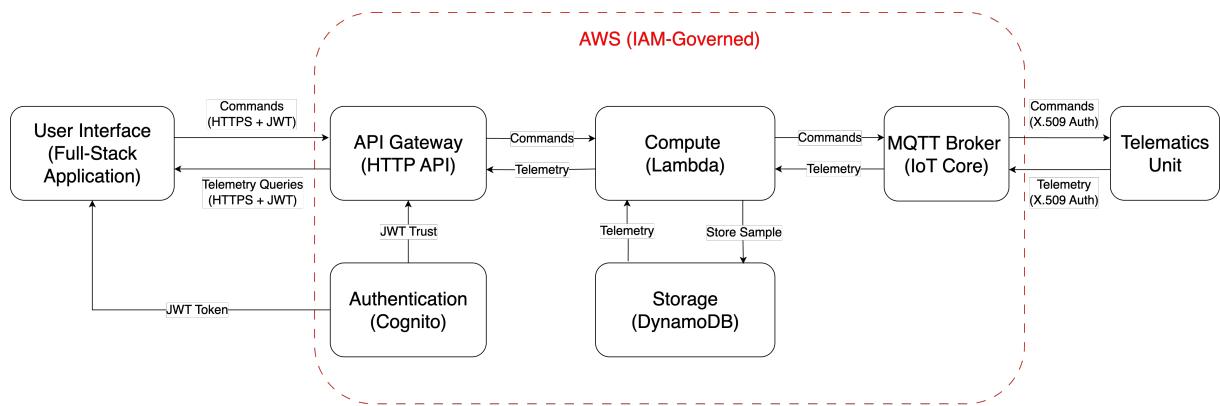


Figure 3.20: Cloud architecture linking embedded devices, AWS services, and web application.

3.4.2. Messaging and Compute

This section describes how telemetry and commands are transmitted across the network and cloud layers. It combines the messaging layer, which manages communication between devices, cloud components, and web applications, with the compute layer, which validates, transforms, and distributes data. Three dedicated data flows are implemented: telemetry ingestion from embedded devices, telemetry querying from the web application, and command publishing from the web application to embedded devices. Each flow is handled by its own AWS Lambda compute function that formats and routes the message, and involves a secure, limited-access connection to the cloud system.

Device-to-Cloud Communication

Telematics devices communicate through AWS IoT Core, which provides an MQTT broker with built-in authentication, asynchronous ingestion, and message routing. Each device uses a unique X.509 certificate for mutual TLS authentication, ensuring trusted communication between the embedded controller and cloud. Messages are published to dedicated telemetry topics using QoS 1, guaranteeing at-least-once delivery under variable mobile network conditions. This platform choice ensures that messaging is efficient and lightweight, enabling reliable transmission over embedded devices with limited bandwidth.

Incoming data is processed through IoT Core rules, which trigger an ingestion Lambda function. This function validates and normalises payloads, converts numerical fields into DynamoDB-compatible types, and writes the results to the database. Items are stored using a composite key consisting of deviceId (partition key) and timestamp (sort key), enabling efficient range queries per device and inherently avoiding duplicate entries with identical keys. This design provides horizontal scalability and schema consistency across heterogeneous vehicles while maintaining predictable query performance.

Cloud-to-Device Communication

Commands and configuration updates are sent from authenticated users through API Gateway, which validates each request and triggers a control Lambda compute function. The function republishes the command as an MQTT message to the corresponding device topic using QoS 1, ensuring at-least-once delivery under variable network conditions. In a full implementation, embedded devices would subscribe to these topics and execute the received commands; however, this functionality is reserved for future work, and only the cloud-side command path is currently implemented and verified.

Cloud-to-Application Communication

The web application dashboards access telemetry exclusively through HTTPS via API Gateway (HTTP API). Each request invokes a read Lambda that queries DynamoDB using the specified deviceId and time window (timestamp between start_ms and end_ms). Results are paginated using LastEvaluatedKey (returned as nextKey). Large responses are gzip-compressed by the Lambda to reduce transfer size; Next.js forwards the compressed data, and modern browsers automatically decompress it based on the Content-Encoding: gzip header.

The web application performs periodic HTTP polling at the desired refresh interval instead of using persistent WebSockets. This approach simplifies deployment, avoids firewall traversal issues, and remains cost-effective, as sub-second latency is unnecessary for telematics applications.

3.4.3. Storage

Telemetry was initially stored in a time-series database, AWS Timestream, which provided native downsampling, aggregation, and time querying techniques. However, due to large minimum costs it was decided that a cheaper option should instead be used for the prototype. As such, an almost-free but less scalable NoSQL storage system was used, AWS DynamoDB.

DynamoDB data is keyed by deviceId (partition) and timestamp (sort), supporting time windowing. Since DynamoDB does not support native aggregation, downsampling of large

datasets is handled in the web application. AWS Timestream provides native aggregation features offering marginally higher efficiency. However, performing downsampling in the web application enables flexible time-window adjustments and direct access to raw data without re-querying the database, thereby substantially improving analysis speed at the expense of slightly higher initial latency. When querying, DynamoDB limits the return size (1 MB per query), so pagination is implemented. Pagination bookmarks the last key to allow multiple queries to retrieve a complete dataset.

For larger queries of hundreds of thousands of data points, object storage (particularly S3) would be required, with Athena, Parquet or Redshift used for analytics over large periods. Further, for the time-series chart, AWS Timestream must be reimplemented to ensure downsampling prior to web application data processing to avoid extreme rendering delays. This forms part of a more efficient commercial product, excluded per scope.

3.4.4. Latency

NFR2.4 specifies that the real-time latency between MCU transmission and web application visualisation must remain below 2 s. This latency encompasses the LTE uplink delay, message propagation through successive cloud services (in order: IoT Core, Lambda, DynamoDB, Lambda, API Gateway), transmission to the web application, and final rendering on the dashboard. It excludes the intentional averaging delay introduced by 0.5 s downsampling, as well as the offset between the web application request time and the record's placement in DynamoDB. The real-time case refers to single-sample, negligibly sized telemetry transmission that is requested using an individual HTTPS request while the embedded system is active.

Frontend rendering contributes minimally to total latency due to lightweight React processing. Intra-cloud delays are small given the small payload size, limited compute, and regional colocation of AWS services. Since Vercel servers (Washington DC) are geographically close to AWS US-East1 (Northern Virginia), the intermediary hop adds only a small overhead. Assuming warm Lambda starts, summing these cloud and rendering delays are expected to be well below 200 ms of latency.

The dominant latency arises from LTE transmission and physical propagation between the embedded device in Cape Town, South Africa and the AWS servers in the United States, as well as between AWS and the user of the web application, also in Cape Town. The distance from Cape Town to Northern Virginia is approximately 13 000 km. With fibre propagation at about 5 ns m^{-1} and a conservatively assumed effective path $1.5 \times$ the great-circle distance, the one-way delay is roughly 100 ms. Therefore, summing the ESP-to-AWS fibre-optic transmission delay and the return-path fibre-optic delay for the bidirectional HTTPS exchange yields an estimated total of approximately 300 ms. Typical Cape Town LTE uplink delay contributes about 100 ms [63]. The total delay is therefore

conservatively assumed to be around 600 ms. Additional delays arising from modelling assumptions may cause minor variations, but the total latency is not expected to approach the 2 s end-to-end limit specified in NFR2.4.

Future deployments could further minimise latency by selecting regions with closer geographic proximity between the operational area and the cloud servers, either within AWS or through alternative providers offering localised infrastructure.

3.4.5. Security

Security in this architecture is enforced at multiple trust boundaries, ensuring that devices, users, and backend services are each authenticated and authorised through distinct mechanisms.

At the **device layer**, every MCU authenticates to AWS IoT Core using an X.509 certificate issued per device. This ensures that only authorised hardware can publish telemetry or subscribe to command topics, preventing spoofed devices or credential reuse. IoT Core therefore secures the device–cloud boundary and provides the foundation for trusted ingestion of telemetry.

At the **frontend layer**, user authentication is managed by Amazon Cognito, which provides a hosted login interface, handles identity management, and issues JSON Web Tokens (JWTs) to authenticated users. These tokens are securely stored in the browser and attached to API Gateway requests, where they are validated before invoking backend Lambdas or accessing DynamoDB. By offloading identity management to Cognito, authentication flows remain consistent across clients, and security controls are separated from the application logic.

Within the **AWS zone**, fine-grained IAM policies govern permissions between services, enforcing least privilege so that each Lambda or subsystem has access only to the resources it requires. For example, a Lambda function that is allowed to read from DynamoDB cannot change table values, or interact with IoT Core.

Encryption is applied consistently across all layers: TLS secures MCU–IoT Core MQTT sessions, HTTPS (HTTP with TLS) protects Cognito authentication flows and API Gateway requests, and inter-service traffic within AWS is encrypted by default. This ensures confidentiality and integrity for all data in transit, regardless of which boundary it crosses.

3.4.6. Cost and Scalability

All services are serverless, autoscaling, and event-driven, incurring charges only when invoked and avoiding any fixed minimum costs. This architecture scales without upfront provisioning: additional devices, higher data rates, or larger storage requirements seamlessly expand to meet demand. As a result, the system retains a low entry cost during

prototyping; however, as Section 3.4.3 notes, capacity increases require additional scaling effort. Runaway costs should be monitored, since AWS does not enforce strict budget caps.

3.4.7. Requirement Traceability

The cloud design contributes to the satisfaction of several functional and non-functional requirements defined. These contributions are summarised in Table 3.14.

Table 3.14: Summary of how the cloud architecture satisfies the key functional and non-functional requirements.

ID	Requirement	Cloud contribution
FR2.2	Transmit all measurement channels to a cloud endpoint at 1 Hz when available.	AWS IoT Core enables telemetry transmission to the cloud securely via MQTT.
FR2.5	Allow for complete raw dataset downloading, and downsampled data viewing via the web application.	API calls and storage through DynamoDB with pagination allows control over requesting complete datasets to the web application.
NFR2.2	System shall maintain reliable cellular connectivity at vehicular speeds up to 120 km h^{-1} .	Lightweight MQTT sessions with IoT Core support efficient reconnects and session persistence suitable for high mobility.
NFR2.4	Real-time latency from MCU to web application rendering shall be $\leq 2 \text{ s}$ typical.	Total expected latency calculated to be below 600 ms.
NFR2.5	All remote interfaces (MCU–cloud, cloud–web application, etc.) shall enforce secure access.	TLS for MQTT/HTTPS, Cognito-issued JWTs via secure hosted login for the web application, and IAM least-privilege between AWS services.

Summary

The cloud architecture integrates AWS IoT Core, Lambda, DynamoDB, API Gateway, and Cognito into a cohesive environment that bridges embedded hardware and the operator-facing web application. By separating telemetry, command, and query flows, the architecture ensures reliability, decoupled operation between devices and users, and scalability under variable conditions, while serverless execution maintains low entry cost. Strong authentication, encryption, and fine-grained access control preserve data integrity and confidentiality across all layers. Together, these design choices meet the system requirements for secure ingestion, storage, and retrieval of EV telemetry in real time and at scale.

3.5. Web Application System

This chapter presents the frontend design of the Telematics Hub. Design choices prioritise minimalism, responsiveness, secure credential handling, and maintainability.

3.5.1. Frontend Architecture

This subsection describes the structural foundations of the Telematics Hub frontend, including rendering choices, routing patterns, and session management. The main system page may be seen in Figure 3.21. The hosted link is available at: <https://telematicshub.vercel.app>; access credentials are required. The Git repository is available at: <https://github.com/Samshabz/Masters-Site>.

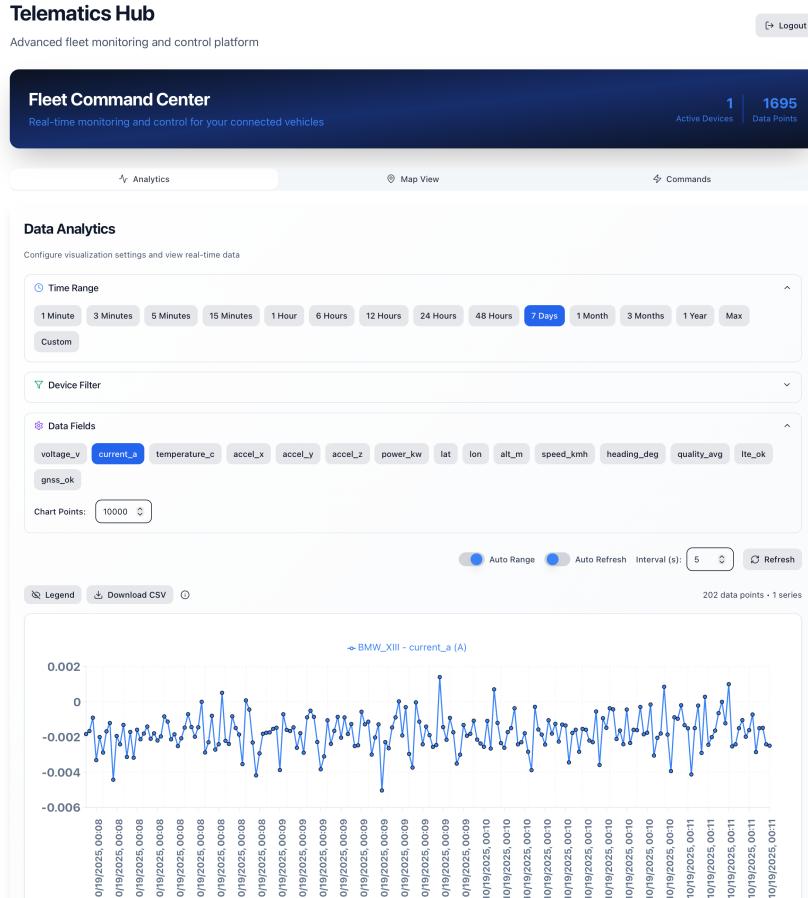


Figure 3.21: Frontend architecture overview showing the Next.js dashboard entry point with analytics, fleet map, and command tabs after Cognito sign-in.

System Overview and Rendering Model

The frontend is a Next.js application that exposes three capabilities: time-series analytics, a fleet map of latest locations, and device command submission. Authenticated calls to backend services are channelled through route handlers in `app/api/*`. View logic resides in pages and modular components, while integration logic and token handling live in server routes and shared utilities under `lib/*`. Rendering is predominantly client-side for interactive dashboards: the main page `(/)` and its tabs (Analytics, Fleet Map, Device Commands) hydrate in the browser, enabling responsive updates to charts and maps. The sign-in page initiates authentication on the client before control passes to a server callback,

which redirects to the main page. Logout flows run entirely on the server and use the server callback URL to redirect to the sign-in page. Static Site Generation and Incremental Static Regeneration are not applied to these dynamic views; deployment therefore serves static assets from the CDN while executing route handlers as serverless functions while preserving the rendering model defined in code.

Routing and Data Exchange

Navigation relies on the Next.js App Router. Client components in `app/` orchestrate data fetching, while server routes mediate data exchange with the cloud backend. Each request to `/api/*` verifies the refresh token cookie, trades it for a short-lived access token, and forwards the call upstream. Returned records are normalised before presentation, and polling intervals can be tuned per tab; WebSockets remain an option for future subsecond updates. The primary endpoints are:

- `/api/raw-data` (GET) – returns paginated telemetry for a selected device and time window, returning `{items, nextKey}`.
- `/api/command` (POST) – forwards a device command after exchanging credentials.
- `/api/save-refresh-token` and `/api/revoke-token` (POST) – set and clear the refresh token cookie during login and logout.

Authentication Lifecycle

Session management follows OpenID Connect (OIDC) with Amazon Cognito. After a successful sign-in, the refresh token is persisted as an *HttpOnly*, *Secure*, *SameSite* cookie set by a server route, preventing access from client-side code. Whenever data is required, the relevant route handler exchanges that refresh token for a short-lived access token and then invokes the backend API. Logout clears client state, clears the cookie on the server, and completes the provider's sign-out sequence before redirecting to `/signin`.

3.5.2. Interface and Functional Design

The user experience is organised into a small number of pages and tabs, each supported by focused React modules that deliver the required telemetry views.

Dashboard Overview

The application exposes four routes: `/signin`, `/` (Dashboard), `/logout`, and `/logout-callback`. Within the dashboard, three tabs mirror the literature synthesis in Section 2.5: Analytics for time-series inspection, Fleet Map for spatial awareness, and Device Commands for remote actions. `AuthWrapper` protects these views, redirecting unauthenticated users to the sign-in page.

Analytics (Charts)

The Analytics tab allows users to select time ranges (one minute to one year, plus *Max* and *Custom*), devices, and telemetry fields such as voltage, current, temperature, and GNSS signals. The chart supports legend toggling, a brush selector positioned beneath the x-axis, manual and automatic refresh intervals, CSV export of the unfiltered dataset, vertical time labels for readability, and theme-aware styling that adjusts markers and tooltips.

Fleet Map

The Fleet Map tab displays the latest valid position per device, presenting markers with pop-ups that summarise identifiers and metadata. Users can trigger manual refreshes or enable periodic updates, and the view automatically fits bounds when multiple devices are visible.

Device Commands

The Device Commands tab provides a structured form with templates for common operations (Wi-Fi, LTE, IoT configuration). Submissions are routed through `/api/command`, which validates the session, exchanges credentials, and forwards the payload to the backend service.

Performance and Accessibility

Responsive operation is achieved through frontend downsampling, pagination of large queries, user-tunable refresh intervals, and incremental map updates that reuse a single Leaflet instance. Accessibility considerations include keyboard navigation, adherence to system themes, adequate contrast in both light and dark modes, and Radix-based components that preserve ARIA roles, focus order, and screen-reader semantics.

3.5.3. Data Handling and Processing

Telemetry is transformed through a dedicated pipeline that normalises records, constrains data volume, and prepares outputs for visualisation and export.

Normalisation and Downsampling

`lib/data-processor.ts` harmonises heterogeneous inputs by flattening nested GNSS structures, reconciling scalar and statistical current measurements, and coercing boolean-like values to `{0,1}`. When series exceed a target point budget, the utility partitions them into buckets of size $\lceil N/\text{targetPoints} \rceil$ and applies the requested reducer—average, minimum, maximum, or latest. Boolean channels use the minimum operator so that any fault condition persists. Derived power is computed as $(V \cdot I)/1000$ when voltage and current are available. The result is sorted chronologically and emitted with complete unit

metadata for reuse in CSV exports.

Telemetry Transformation Pipeline

`app/page.tsx` coordinates dashboard state: active tab, device selection (including an *All devices* option), chosen metrics, time windows, auto-refresh interval, and command forms. It calls `/api/raw-data` with pagination support, merges and sorts records, and hands chart-ready structures to `DataChart1`. The chart component renders multi-series outputs via Recharts, reports transport and UI latency, and exposes CSV downloads. `VehicleMap.tsx` maintains a single Leaflet instance, updates markers when new locations arrive, and recalculates bounds only when the device set changes, thereby reducing redraw overhead.

3.5.4. Implementation and Configuration

Delivering the dashboard reliably requires disciplined environment management, a clear repository layout, consistent technology choices, and explicit hosting assumptions.

Environment and Secrets Management

Configuration is centralised in `lib/config.ts` with supporting environment files. During development, settings reside in `.env.local`; production pipelines inject the same keys via the hosting platform so that secrets remain outside version control and can vary across development, preview, and production environments. Server-only values (for example `COGNITO_CLIENT_SECRET`) are never prefixed, whereas browser-visible settings require the `NEXT_PUBLIC_...` prefix. Shared constants—including host URLs, Cognito identifiers, default device IDs, and internal API paths—are exported from `lib/config.ts` for reuse across modules.

Repository Layout

Paths are relative to the project root; inline comments summarise responsibilities.

```
app/
  layout.tsx          # App shell, global styles, ThemeProvider
  page.tsx            # Dashboard (charts, map, commands)
  AuthWrapper.tsx     # Route protection using OIDC (Cognito)
  signin/page.tsx     # Initiates OIDC redirect
  logout/page.tsx     # Revokes cookie, triggers provider logout
  logout-callback/page.tsx # Completes logout and returns to /signin

api/
  (each folder contains route.ts)
```

```

raw-data/          # GET /api/raw-data
command/          # POST /api/command
save-refresh-token/ # POST: store refresh token cookie
revoke-token/     # POST: remove refresh token cookie

app/components/
  DataChart1.tsx      # Recharts line chart
  VehicleMap.tsx      # Leaflet markers and fit-to-bounds logic

components/ui/      # Radix-based UI primitives (buttons, etc.)

lib/
  config.ts           # Central configuration exports
  data-processor.ts   # Normalisation and downsampling

```

Libraries and Frameworks Used

The interface employs Tailwind CSS for layout and spacing, Radix primitives for accessible controls, Recharts for time-series visualisation, Leaflet for geospatial rendering, and MUI pickers for date and time selection. These libraries integrate with React's component model, enabling modular reuse and theme-aware styling.

Hosting and Deployment Considerations

Deployments target managed platforms such as Vercel, which serve static assets through a CDN while executing Next.js route handlers as serverless functions in the chosen region. Selecting a hosting region close to the backend minimises additional network latency. HTTPS is enforced end-to-end, with the platform provisioning TLS certificates and honouring session cookies marked *Secure*. Because rendering behaviour is dictated by the codebase, deployments preserve the client-rendered dashboard while keeping authentication and API logic on server routes.

3.5.5. Requirement Traceability

Table 3.15: Frontend contributions toward requirement fulfilment.

Requirement	Frontend contribution
FR2.5 (raw download and downsampled views)	Charts with selectable fields and brush; CSV download of underlying (non-downsampled) data; map view of latest positions.
NFR2.5 (secure remote interfaces)	OIDC authentication; refresh token stored via HTTP-only cookie; server routes exchange for short-lived access tokens; access checks on server.
NFR2.4 (typical latency \leq 2 s)	Efficient rendering paths (downsampling before draw and pagination).

Summary

The frontend presents a modern interface and separate tabs for different viewing goals. Client pages manage interaction, server routes handle authentication and forwarding, and shared utilities prepare data for efficient visualisation. The result is a secure, responsive, and maintainable design aligned with the requirements of a production-ready telematics dashboard.

Chapter 4

Testing and Validation

This chapter verifies the system against the requirements defined in Chapter 3, Subsection 3.1, with detailed targets in Tables 3.2 and 3.3. Results are grouped by subsystem; each subsection posits rational methods of testing to achieve the requirements, specifies the results of these tests, and presents a statement of compliance.

4.1. Approach

4.1.1. Structure of Chapter

This chapter is structured according to the three fields of requirements: Sensing, data handling and communication, and system retrofittability, generalisability, and robustness. Certain requirements are implicitly satisfied across multiple sections; however, the final summary (Table 4.2) explicitly outlines the subsection(s) in which each criterion is demonstrated.

4.1.2. Tests Conducted

The system was evaluated across three primary contexts: laboratory characterisation, in-vehicle testing, and stationary EV testing. Laboratory characterisation encompassed all controlled, non-automotive tests, with the HVDC laboratory setup shown in Figure 4.1. The in-vehicle test was conducted within the passenger compartment of a conventional ICE vehicle using a straightforward setup, powered from the 12 V accessory outlet to verify system and communication stability under alternator-driven voltage fluctuations and high-speed real driving conditions. EV testing employed the stationary EV Bench Simulator—a wall-mounted replica of the Roam Air electric motorbike rated at 70 A and 80 V—which transmits CAN messages at approximately 300 ms intervals. The developed telematics system was powered from the simulator’s 12 V supply, as shown in Figure 4.2.

Together, these experiments validate the system’s measurement accuracy, connectivity, generalisability, retrofit compatibility, and operational robustness. The findings confirm its suitability for deployment within the South African retrofit context, enabling rapid acquisition of critical electromobility datasets to accelerate national EV integration.

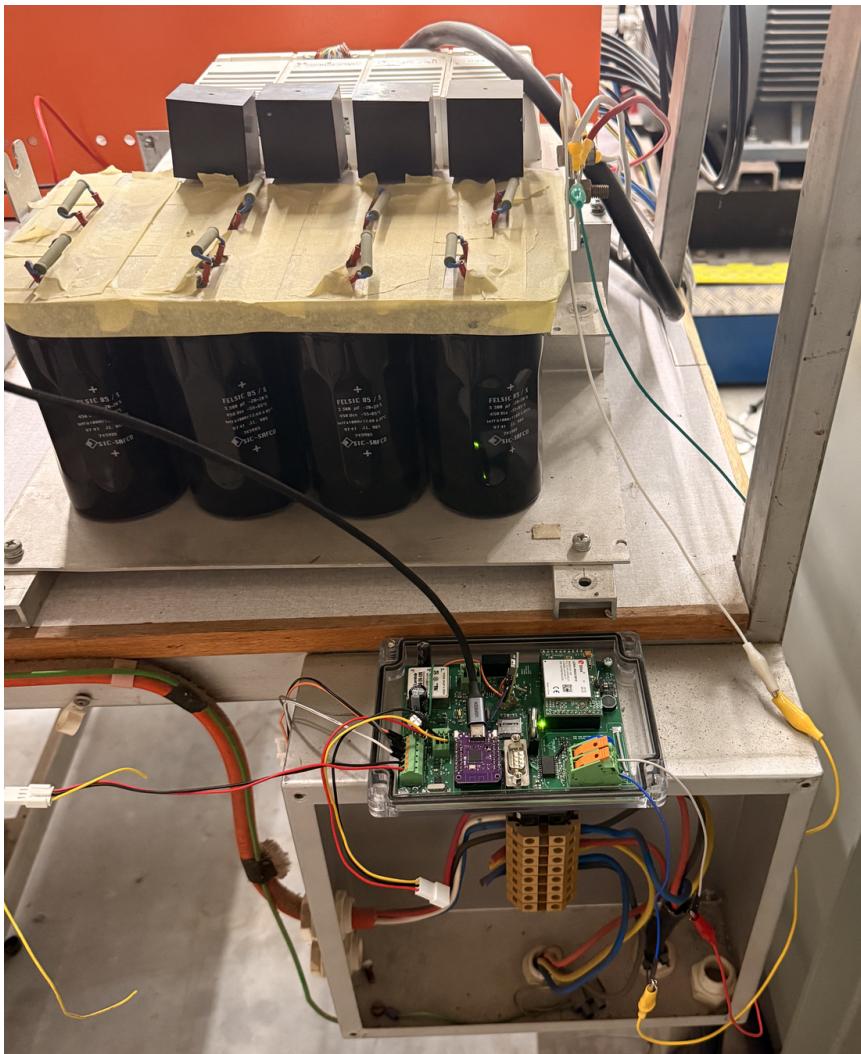


Figure 4.1: Laboratory configuration for HVDC measurement testing, showing the telematics PCB connected to the high-voltage source via its capacitor bank for controlled voltage sweep acquisition.

4.1.3. Scope and Assumptions

Due to the limited availability of high-precision references and controlled test facilities—such as high-power, acceleration-controlled, and temperature-regulated environments—full verification of all datasheet specifications was not feasible. Accordingly, datasheet performance parameters are assumed valid, and testing focuses on verifying correct operation, reasonable accuracy, contextual validity, and safety. In-vehicle, EV, and/or laboratory results, together with design data referencing manufacturer specifications, are used to infer compliance with requirements that could not be directly tested. This forms **Assumption 1**. In the following tests, unless otherwise specified (such as in the HVDC laboratory test), sensing errors are interpreted under the assumption that the reference instrument is sufficiently accurate to demonstrate operational validity, though not considered an absolute reference.

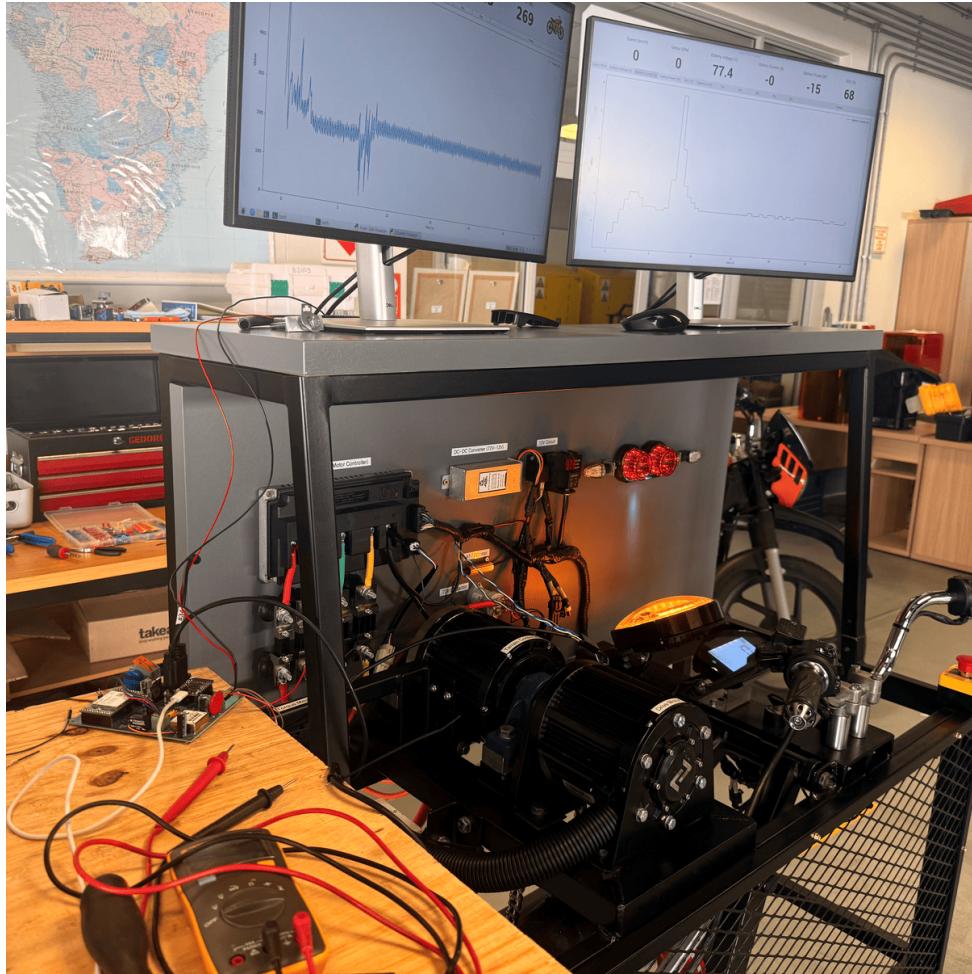


Figure 4.2: EV Bench experimental setup showing the current sensors and vehicle CAN logger connected to the Roam Air bench simulator, while being supplied from the 12 V Roam Air supply.

4.2. Sensing

This section addresses how the system captures the key physical telemetry required for EV profiling. Accurate and robust measurement of high-voltage pack current and voltage, acceleration, GNSS metrics, and ambient temperature is critical, as these sensors underpin the validity of all subsequent analysis. Each sensing channel is therefore tested to confirm accuracy, stability, and compliance with the specified functional and non-functional requirements.

4.2.1. HVDC Voltage Sensing

This subsection validates the HVDC voltage sensing chain, aiming to confirm accuracy, robustness, and safety per FR1.1 (error below 1% up to 880 V).

Laboratory testing establishes quantitative accuracy, while EV bench tests verify operational behaviour under realistic conditions.

Lab Testing

Method

A DC sweep from 0 V to 700 V was applied using a source with an error below 20 mV, negligible relative to the applied voltage magnitudes. The sweep was limited to 700 V at room temperature due to laboratory constraints, below the 880 V operational requirement and operational temperature range. Multiple redundancy mechanisms, as detailed in Section 3.2.1, make system failure above 700 V improbable if no failure occurred within this range.

Results

The measured full-scale range (FSR) error was 0.044%, and all non-zero measurement points remained below 1% error. Additionally, the measured gain error was approximately 0.35 mV V⁻¹ (0.035%), while an offset of approximately 67 mV (primary-referred) or 83.75 μ V (ADC-input-referred) was observed. Both the offset and FSR errors are below the worst-case room-temperature design accuracy limits of 149.5 μ V and 0.47374%FS, respectively, indicating that the hardware exceeded its design performance limits and can therefore be regarded as conservatively specified. Consequently, based on empirical validation, the worst-case design accuracy across the full temperature range of -5 to $+55$ °C (0.57364% FSR), which remains below the 1% limit, is considered a conservative reliability margin, ensuring compliance with FR1.1 (error below 1% up to 880 V) and NFR1.2 (accuracy maintained across temperature).

A post-hoc linear regression fit produced a calibrated curve with a mean absolute deviation of only 9 ppm. This demonstrates potential for ultra-high accuracy improvements in future work if calibration is implemented.

Stationary-EV Testing

Method

The voltage sensing chain was tested on the EV bench setup. The MCU measurement was compared with the vehicle's CAN-reported pack voltage under stationary operating conditions. As with current sensing, the CAN reference accuracy and bandwidth are unspecified; results are therefore interpreted as operational validation and relative agreement.

Results

Figure 4.4 shows that the sensed HVDC voltage closely matches the CAN-reported value, with a mean offset of 0.088 V (0.121%) and instantaneous deviations up to 300 mV, attributed mainly to sampling-time differences and an unknown EV Bench accuracy. This indicates consistent and reliable operation.

Table 4.1: Measured HVDC bus voltage error relative to applied source values across the full test range.

Applied Voltage (V)	Measured Voltage (V)	Error (%)
0.00	0.067	—
50.00	50.078	+0.156
100.00	100.095	+0.095
150.00	150.116	+0.077
200.00	200.134	+0.067
250.00	250.151	+0.060
300.00	300.164	+0.055
350.00	350.180	+0.051
400.00	400.196	+0.049
450.00	450.219	+0.049
500.00	500.233	+0.047
550.00	550.259	+0.047
600.00	600.265	+0.044
650.00	650.289	+0.045
700.00	700.309	+0.044

Outcome

In laboratory testing, the system achieved a total full-scale error of 0.044% and a worst-case accuracy of 0.4738% FSR at room temperature, remaining within the 1% limit across the full operational temperature range. Safety was verified up to 700 V, with protection further ensured by the MOV, Zener clamp, resistor chain, basic insulation compliance, and the isolated ADC's reinforced insulation rating of 1.2 kV. Together, these provide tolerance for transient events up to 1.85 kV, single-fault protection, and continuous operation within the 880 V range (FR3.3).

In stationary EV testing, the measured voltage closely matched the CAN-reported value, with a mean offset of 0.088 V (0.121%) and instantaneous deviations up to 300 mV, attributed primarily to sampling-time differences and reference uncertainty. The sensing chain remained stable under noisy supply conditions, validating robust performance in realistic environments.

Overall, the HVDC voltage sensing chain met the required accuracy and insulation margins, confirming compliance with FR1.1 (HVDC voltage accuracy below 1% FSR), NFR1.2 (accuracy maintained across temperature), and FR3.3 (insulation and robustness). Full reinforced insulation verification is recommended prior to deployment.

4.2.2. HVDC Current Sensing

This subsection validates the performance of the two current sensing channels: the GMW coreless probe and the fluxgate sensor. The aim is to confirm accuracy, responsiveness, and robustness across the measuring range. The GMW digitisation chain is first tested in

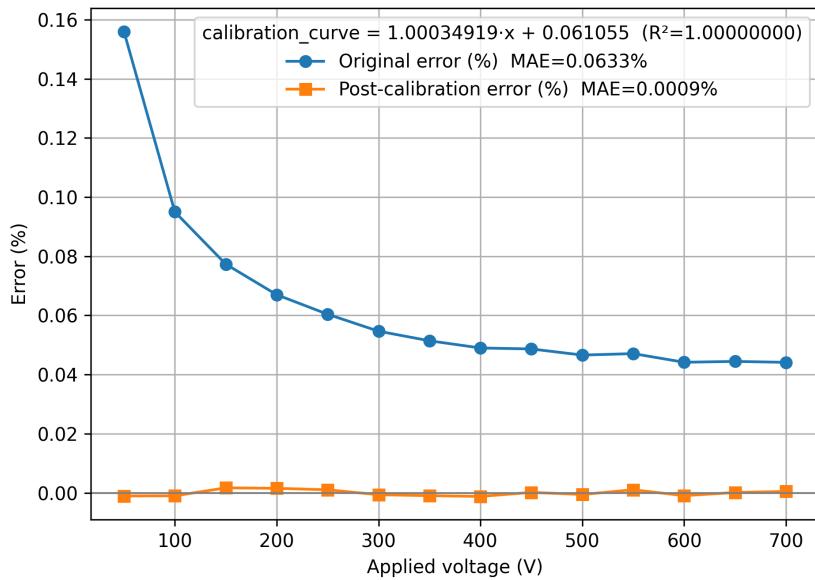


Figure 4.3: Measured HVDC bus voltage error relative to a precision source across the laboratory sweep range, showing both the raw error curve and the post-hoc calibrated curve demonstrating linearity.

the laboratory, while both sensors are subsequently validated in the EV context.

Lab Testing

This section tests the only analogue section for either current sensor in isolation, namely the GMW probe digitisation section.

Method

The GMW analogue output (FSR ± 2 V \leftrightarrow ± 1000 A primary) was swept at 0, 0.5, 1.0, 1.5, and 2.0 V in both polarities. A reference Fluke Digital Multimeter (FDMM) with 1 mV resolution was used to measure the ADC/frontend error.

Results

The source could be finely adjusted such that the ADC output matched without any discernible deviation on the 1 mV-resolution DMM, implying that the digitisation error lies below the meter's resolution, i.e. $\leq 0.05\%$ FS. Further discrimination is not possible, as the error is already dominated by the sensor itself (1 %). Accounting for ageing, worst-case design, and temperature drift (0.35217 percentage points), the total remains below 0.5%.

Outcome

The GMW digitisation chain contributes a negligible error at room temperature ($< 0.05\%$ FS) and remains $< 0.5\%$ across drift, ageing, and design variation. With specified sensor performance, total error is $\leq 1.5\%$ FS under all conditions, satisfying FR1.2 for the GMW probe.

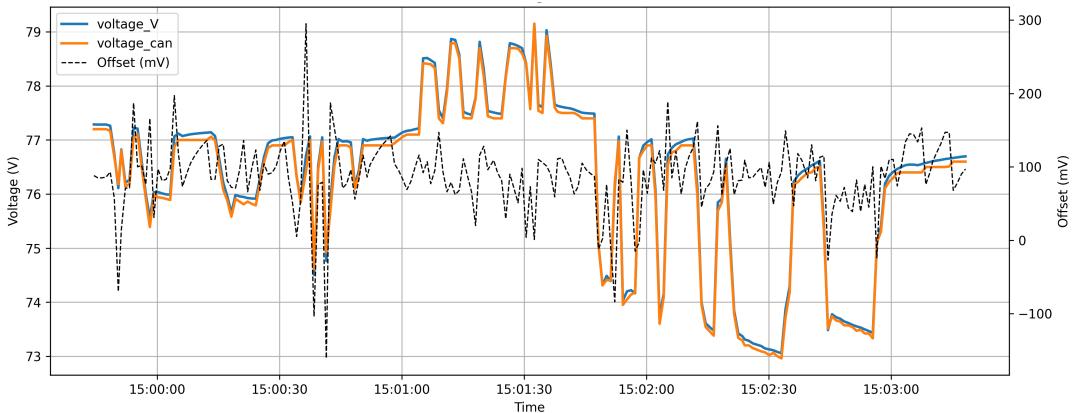


Figure 4.4: Comparison of HVDC-sensed and CAN-reported voltages under stationary EV conditions, with voltage profiles plotted on the left axis and their offset (difference) on the right axis.

Stationary-EV Testing

This test evaluates the profiling capability of the current sensing subsystem in relation to an EV’s internal CAN-reported current. The objective is to assess operational behaviour, response time, and resilience to switching noise under realistic conditions rather than absolute accuracy. Absolute accuracy cannot be guaranteed because the CAN reference has unspecified accuracy and bandwidth, and asynchronous sampling means the sensor and CAN timestamps are not aligned.

Method

The test was conducted on the Roam Air electric bike (70 A, 80 V peak) under stationary conditions. Current measurements from the fluxgate and GMW sensors were compared against the vehicle’s CAN-reported signal while all subsystems operated concurrently, including dual SPI buses, SD card logging, and both CAN interfaces.

Results

Figure 4.5 shows the current profiles from all three sensors, while Figure 4.6 illustrates their relative offsets. The left axis indicates the offsets between the CAN reference and each current sensor, and the right axis shows the offset between the GMW and fluxgate sensors. The fluxgate and GMW measurements align closely with the CAN signal during steady-state and slow transitions; however, rapid throttle transitions produce amplitude offsets of up to 40 A, primarily due to sampling-time differences between the CAN reporting rate and the faster analogue sensors. The two sensors produced nearly identical waveforms, with instantaneous deviations up to 2 A during peaks, arising from minor sampling-time differences and the GMW sensor’s higher bandwidth combined with its digital decimation, which yielded a slightly sharper transient response.

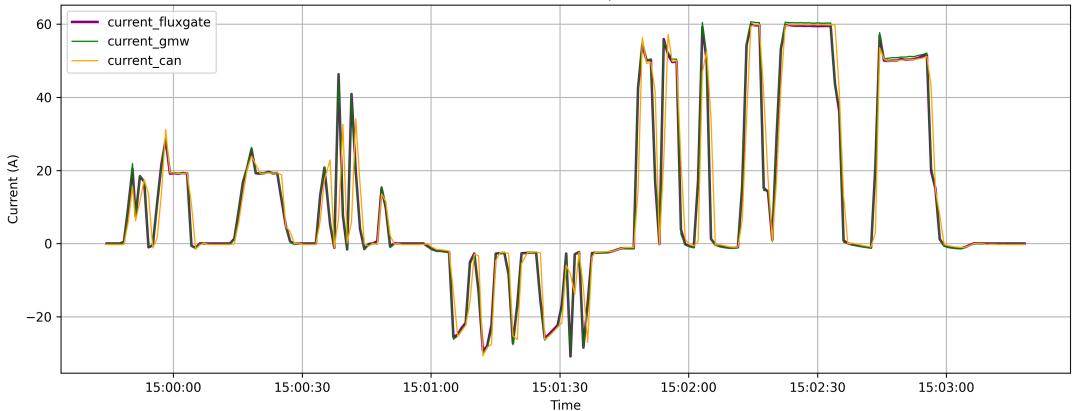


Figure 4.5: Profile comparison of the current sensors against the EV Bench.

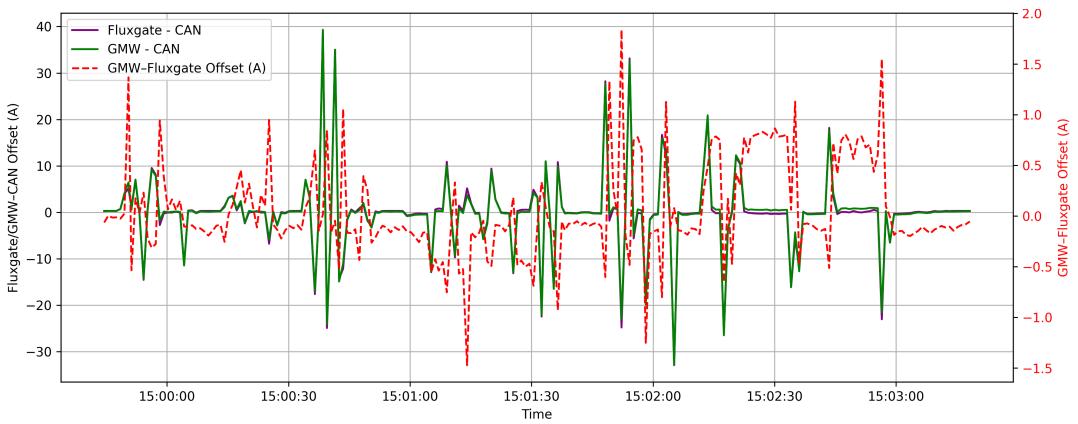


Figure 4.6: Offset comparison of current sensors against the EV Bench, with Fluxgate-CAN and GMW-CAN errors shown on the left axis, and the GMW-Fluxgate offset shown on the right axis.

Furthermore, the sensors accommodated insulated conductors up to 27 mm in diameter, and the system operated reliably from the vehicle's non-ideal 12 V supply.

Outcome

Both sensors reproduced the EV current profile with comparable accuracy and response times, confirming operational robustness within an EV context. Combined with the GMW sensor's digitisation accuracy demonstrated in laboratory testing and manufacturer-guaranteed specifications, these results verify correct operation and compliance with FR1.2 (current sensing range and accuracy). The GMW probe achieved an accuracy below 1.5%, while the LEM fluxgate sensor remained below 0.5% across the full temperature range. The GMW probe exhibited a marginally superior transient response owing to its significantly higher sampling rate and digital decimation, though this difference is unlikely to have any appreciable impact on operational analyses. Overall, the test confirms the achievement of non-invasive, ultra-high current-sensing accuracy of 0.5% up to the required 1 kA through the LEM fluxgate sensor, while the GMW probe remains a strong candidate where its clip-on form factor is advantageous. The test also verifies compliance with NFR3.6 (probe

fit for voltage, current, and CAN), FR1.6 (vehicle CAN bus data acquisition), and NFR1.1 (non-intrusive vehicle CAN acquisition). SD logging operated correctly under the vehicle's non-ideal 12 V supply, satisfying FR2.1 (local SD logging) and NFR3.4 (automotive supply robustness), with all subsystems and buses functioning concurrently without conflict.

4.2.3. Accelerometer

This subsection evaluates the ability of the digital accelerometer module to provide a sufficiently accurate profile of vehicle acceleration, and assesses its applicability within an automotive context.

Method

Validation was performed through a controlled in-vehicle test representative of typical urban driving. Stop–start road conditions were selected to induce clear longitudinal accelerations. The ADXL343 accelerometer was compared against a reference measurement from an iPhone 2025 model. As the iPhone accelerometer has no published accuracy specification, the test evaluated contextual reliability and reasonable accuracy. Both devices were configured to record at 1 Hz. The forward-motion trace along the y-axis was analysed for around one minute to illustrate performance clearly.

Results

Figure 4.7 illustrates the longitudinal (y-axis) acceleration profiles of the prototype and the reference device. Although the overall trends align, their magnitudes and peak values diverge noticeably. Since both sensors operate at 1 Hz and employ downsampling, discrepancies of up to one second can occur between their respective averaging windows. Differences in downsampling and filtering methods between the two devices further explain the observed misalignment.

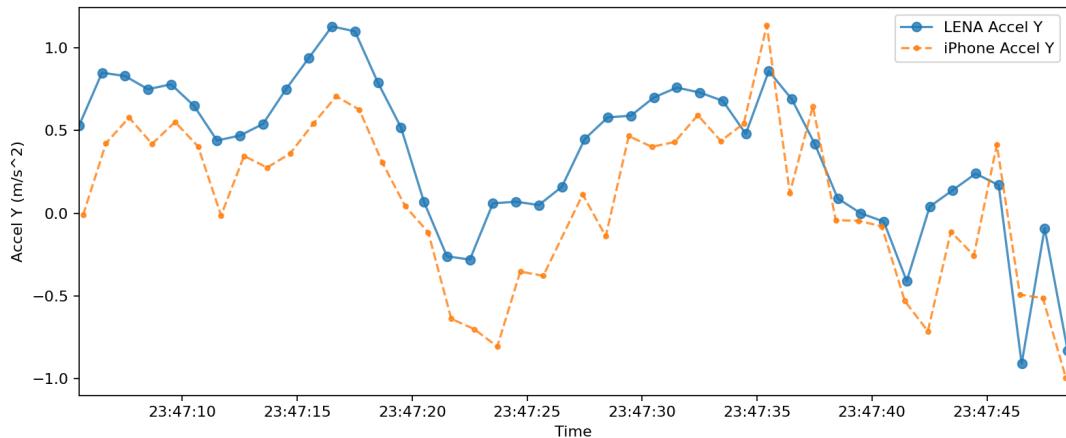


Figure 4.7: Comparison of longitudinal acceleration (y -axis) between prototype and reference device during road testing.

Typical accelerations during normal driving were small ($|a| \leq 0.1$ G). At this level, the accelerometer's inherent offset (35 mG) and cross-axis sensitivity introduced significant relative errors. Given the low utilisation of the 2 G full-scale range (about 5%), offset errors became the dominant contributor, reducing accuracy and correlation. Additional uncertainty arose from the non-rigid enclosure mounting, imperfect axis alignment, sampling offsets between both devices, and the unknown accuracy of the reference device.

The accelerometer interface functioned correctly; therefore, the design—referencing the manufacturer specifications—is assumed to be reliable within its specified ± 2 G range. The system FSR accuracy is 2.88% FSR at room temperature and 4.38% FSR across the full operating range [116]. This satisfies FR1.4 ($\leq 10\%$ error over the ± 2 G FSR) and NFR1.2 (accuracy maintained across the full temperature range). However, since typical vehicular accelerations were below 0.1 G, offset and alignment errors dominated the response, limiting effective accuracy at this smaller scale despite formal compliance. The requirement is therefore met, though the design could be improved to better mitigate offset-induced error.

Outcome

The accelerometer functioned reliably under vehicular conditions and satisfied FR1.4 (acceleration error $\leq 10\%$ across ± 2 G FSR). However, at low urban driving accelerations, offset effects dominated the response, limiting precision. Future revisions should incorporate calibration or lower-offset sensors.

4.2.4. GNSS

This subsection validates the accuracy and robustness of the GNSS receiver under automotive conditions. The test aims to ensure that the modem is robust to automotive speeds, maintains coverage throughout the drive, and maintains reasonable accuracy. Additionally, it tests acquisition and reacquisition of fix times to ensure the modem does not take unacceptably long and hinder the comprehensiveness of the data.

Method

The GNSS tests aim to ensure correct operation and reliability of the GNSS system within an automotive context.

1. *Driving Test:* This evaluates system reliability under high-speed conditions up to 120 km h^{-1} , over an extended 20-minute driving period with a roughly half-sky antenna view. This informs the system's robustness to typical automotive installation conditions, extended runtimes, and maximum speeds. This test uses an iPhone (2024 model) as a reference; its typical accuracy is not specified, and the 1 Hz update rate of both devices implies a 1 s maximum sampling mean offset. Because of this,

and because GNSS is not a fixed maximum-accuracy system, the test only seeks to achieve a reasonable accuracy conclusion to inform correct operation and reliability within an automotive context.

2. *Time-to-first-fix (TTFF) and reacquisition time:* This test measures the time to gain a valid fix after startup given acceptable line-of-sight after initial module power on, as well as the time to regain that fix after intermittent signal loss.

Results

Driving Test

During the highway driving test, the GNSS receiver maintained a continuous duty cycle across the entire trip, confirming reliable operation under sustained vehicle speeds and extended durations. The two-dimensional trajectory overlay in Figure 4.8 shows strong spatial alignment between the LENA module and the reference smartphone. Figure 4.9 illustrates the vehicle speed profile and the boolean GNSS fix validity flag over time, demonstrating consistent GNSS availability even at peak speeds of 120 km h^{-1} .

The horizontal (2D) error had a mean of 2.33 m, a standard deviation of 0.90 m, and a maximum deviation of 4.19 m, as shown in Figure 4.10. The three-dimensional (3D) error had a mean of 4.97 m, a standard deviation of 1.39 m, and a maximum deviation of 8.10 m, as illustrated in Figure 4.11. The higher 3D error primarily reflects GNSS altitude accuracy limitations. Figure 4.12 presents the corresponding velocity error profile, with a mean deviation of only 0.02 km h^{-1} , a standard deviation of 0.77 km h^{-1} , and a maximum of 4.23 km h^{-1} . Residual discrepancies arise from factors such as sample offsets, reception variability (antenna and signal processing variations in quality and placement), and small differences between geoid models. Overall, the GNSS receiver satisfied Functional Requirement FR1.3 by maintaining positional and velocity errors within 10 m and 5 km h^{-1} , respectively, under real-world highway conditions.

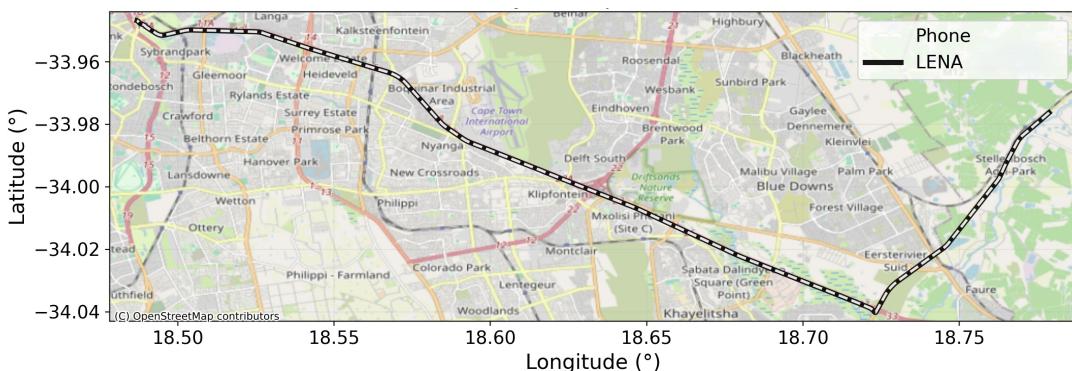


Figure 4.8: 2D map overlay of module and reference positions throughout the trip.

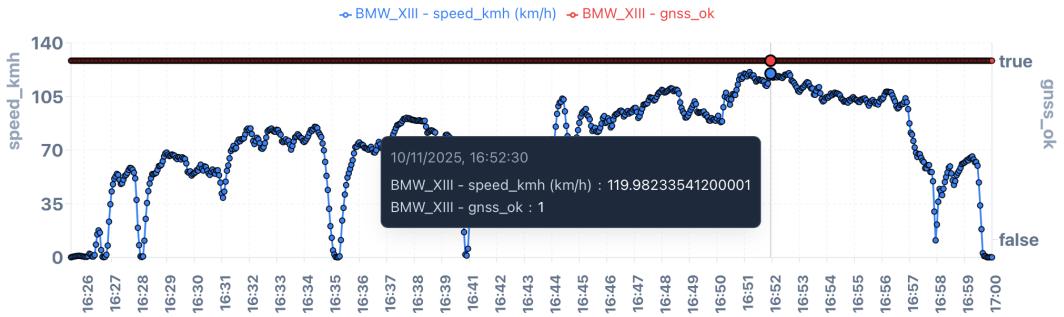


Figure 4.9: Vehicle speed and GNSS fix validity over time during the highway mobility test, demonstrating full coverage at high-speed.

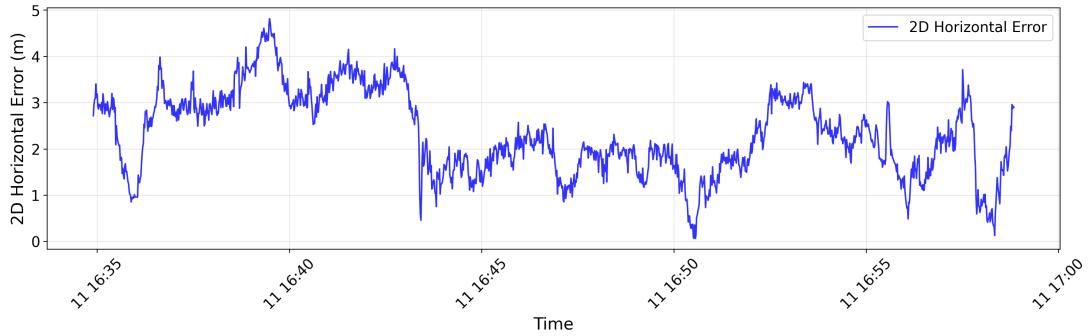


Figure 4.10: Two-dimensional position error of the prototype relative to the reference receiver during the highway mobility test.

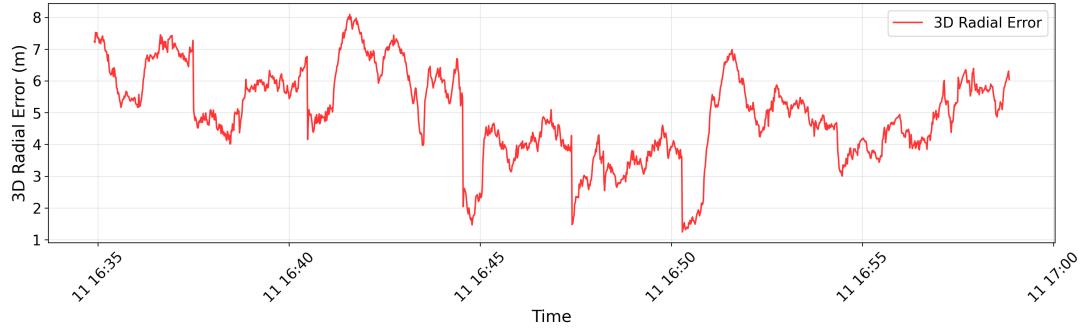


Figure 4.11: Three-dimensional position error of the prototype relative to the reference receiver during the highway mobility test.

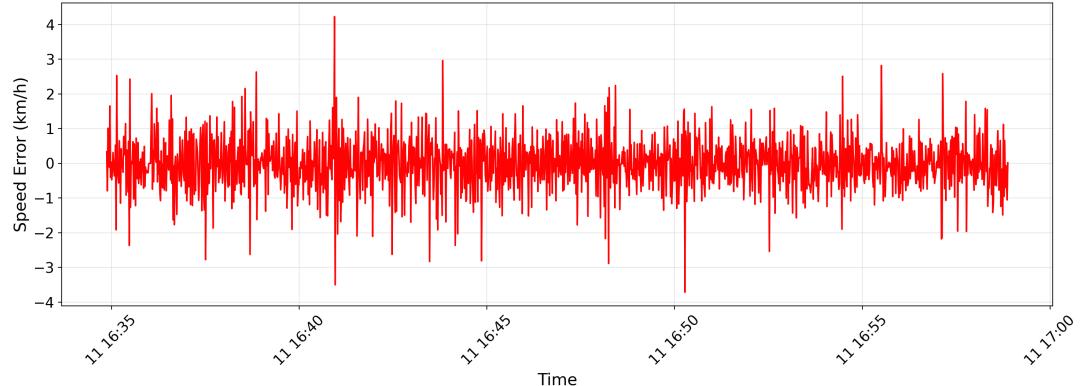


Figure 4.12: Velocity error during highway mobility test. Maximum error: 4.23 km h^{-1} .

Time-to-first-fix (TTFF) and reacquisition time

The TTFF varied between 21 and 39 s, remaining within the required 45 s, with variation attributable to environmental conditions. Reacquisition missed at most two 1 Hz cycles, so the maximum reacquisition time was ≤ 3 s, meeting the required maximum of 3 s. Both metrics were therefore compliant with FR2.7.

Outcome

GNSS performance met accuracy requirements for speed and 3D position within FR1.3 limits. TTFF and reacquisition times remained within FR2.7. The receiver and antenna maintained accuracy and coverage under automotive conditions, enabling reliable mobility profiling.

4.2.5. Temperature

This subsection evaluates the temperature sensing subsystem, which comprises a digital temperature sensor and its I²C interface. The objective is to confirm correct operation and reliable data acquisition under in-vehicle conditions, thereby validating the sensor's datasheet performance and the integrity of the digital interface.

Figure 4.13 shows the temperature profile recorded during the vehicle mobility test described in Section 4.2.4. The measured temperature varied consistently with minor ambient fluctuations, indicating correct sensor response and stable I²C interface operation. No communication errors or missing samples were observed, and timestamps remained aligned with the GNSS/RTC timebase. The observed deviations remained within the manufacturer-specified accuracy of ± 0.5 °C across the full temperature range. This verifies compliance with FR1.5 (ambient temperature sensing accuracy ≤ 1 °C) and NFR1.2 (accuracy maintained from -5 to $+55$ °C).

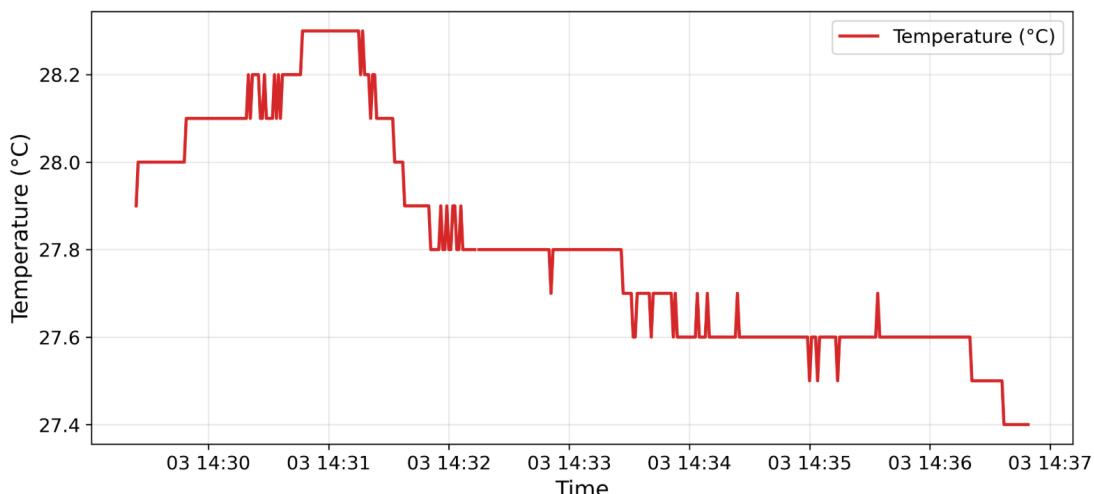


Figure 4.13: Ambient temperature recorded by the enclosure-mounted sensor throughout the vehicle drive test.

4.3. Data Handling and Communication

This section evaluates how the system processes, stores, and communicates the acquired data to local and remote interfaces. Reliable data handling is essential to ensure that measurements are not only preserved but also transmitted securely and made available with low latency to end users. Local and remote storage, LTE telemetry, SD logging, system timing, latency, and cloud security are assessed to verify end-to-end reliability and accessibility.

4.3.1. LTE Transmission and SD Logging

This subsection validates LTE connectivity under real-time operation, backlog flushing, and high-speed mobility. End-to-end latency and time-drift-induced data loss are excluded here and addressed separately in Section 4.3.3 and Section 4.3.2, respectively. The objective is to confirm that LTE transmission meets the requirements for 1 Hz cloud streaming (FR2.2), reliable retransmission of historical data (FR2.3), and robustness at vehicular speeds (NFR2.2).

Method

Four tests were conducted to evaluate LTE operation and reliability:

1. *Reconnection test*: Reconnection behaviour was evaluated by intentionally removing and reattaching the antenna for both short and long disconnect periods. These correspond to cases without and with an MQTT timeout (60 s) and were used to validate reconnection robustness and confirm that 1 Hz transmission was maintained under intermittent connectivity.
2. *Backlog flush test*: The maximum sustained backlog transmission rate was assessed to determine the throughput at which stored samples could be uploaded without module instability, validating the system's ability to clear large datasets accumulated during extended disconnections.
3. *Mobility stress test*: LTE connectivity and MQTT reliability were evaluated during a 20-minute drive up to 120 km h^{-1} to validate handover and coverage capability in the South African automotive context. The objective was to confirm stable connectivity independent of speed and conditions. SD and web application logs were compared to verify identical datasets, confirming 1 Hz throughput for both FR2.1 (local logging) and FR2.2 (cloud transmission).
4. *SD card storage test*: This test validates whether the SD card capacity is sufficient based on the empirically measured average payload size.

Results

Reconnection Test

When antenna disconnections were introduced, the module re-established the MQTT session and resumed transmission within approximately 2 s (a single missed sample at most) for both short and long outage periods. This demonstrates robust reconnection performance.

Backlog Flush Test

Backlog flushing was stable and sustained up to approximately 10 samples/s, with higher rates limited by the 20 ms AT command processing interval and UART congestion. Although firmware optimisation could increase throughput, the measured rate was deemed sufficient for clearing backlog datasets accumulated during typical connectivity outages. This is sufficient to ensure reliable retransmission of historical data (FR2.3).

Mobility Stress Test

During the 20-minute highway test up to 120 km h⁻¹, no MQTT downtime was observed once the connection was established, confirming compliance with NFR2.2. Figure 4.14 shows the MQTT connection status remaining true throughout the mobility test, overlaid with the vehicle speed profile. Both SD and frontend logs were identical and complete, confirming 1 Hz throughput for FR2.1 (local logging) and FR2.2 (cloud transmission).



Figure 4.14: Vehicle speed and LTE network validity during the mobility test, demonstrating full coverage at high-speed.

SD Card Storage Test

The mean payload size across the test dataset was 317 bytes per sample across the test dataset (27,000 samples). At 1 Hz with a 100% duty cycle, this corresponds to roughly 4.998 GB of storage per six months. Even when accounting for backlog data and additional CAN logs, the selected 16 GB card provides ample capacity, satisfying NFR2.1 (capacity beyond six months at 1 Hz). No corruption occurred during development, and the card remained fully readable on a PC, confirming FR2.1 (local logging).

Outcome

These results demonstrate that the LTE subsystem provides robust, high-coverage connectivity with rapid reconnection suitable for automotive applications. Both LTE transmission and SD logging operated reliably, maintaining complete 1 Hz data capture to both local and cloud storage in compliance with FR2.1 and FR2.2. Backlog flushing successfully cleared historical datasets (FR2.3), while operation at 120 km h^{-1} confirmed high-speed LTE reliability and compliance with NFR2.2. In addition, the SD card capacity was verified to be sufficient for over six months of continuous 1 Hz logging, satisfying NFR2.1.

4.3.2. System Timing

This subsection validates the timing subsystem, which combines GNSS time with an RTC fallback to ensure accurate timestamping and temporal integrity. It demonstrates compliance with FR2.4 (timekeeping with GNSS and RTC fallback), FR2.6 (startup time $\leq 10 \text{ s}$), NFR1.3 (Nyquist sampling), NFR1.4 (inter-channel and timestamp misalignment $\leq 1 \text{ s}$), and NFR2.3 (RTC drift $\leq 1 \text{ s/hour}$).

Method

Three aspects of system timing were validated against the specified requirements:

1. *Startup Time*: The interval from power-on to the first successful publish was measured, excluding GNSS TTFF. This was compared against FR2.6, which specifies a maximum startup time of 10 s.
2. *Timekeeping, Temporal Alignment, and Nyquist*: This test verified correct operation of both GNSS and RTC time sources (FR2.4), alignment of timestamps and all sensor channels (NFR1.4), and sufficient sampling rate to satisfy Nyquist criteria (NFR1.3). These requirements are satisfied in firmware design and all sensors intrinsically meet the Nyquist rate; testing aimed to confirm that no blocking, duplication, or omissions occur in the 1 Hz aggregated output that would compromise this design. Operation was evaluated over a 30 minute period using both GNSS and RTC time sources.
3. *RTC Drift*: The RTC offset drift from the GNSS time was measured until exactly 1 s deviation occurred after the initial GNSS synchronisation to validate NFR2.3 (RTC drift $\leq 1 \text{ s/hour}$). The objective was to verify that the RTC maintained correct operation and expected accuracy, thereby validating the manufacturer-specified drift budget and stray-capacitance assumptions.

Results

Startup Time

System startup (excluding GNSS TTFF) was dominated by LTE modem attachment and MQTT connection establishment. The modem required approximately 2 s to power on and an additional 2 s to complete the MQTT handshake. Including MCU boot and peripheral initialisation, the total startup time was approximately 6 s, remaining within the FR2.6 limit of 10 s. This ensures timely data acquisition upon vehicle startup.

Inter-Channel and Timestamp Alignment and Nyquist

Testing confirmed that the main loop executed without blocking, maintaining concurrent operation of all peripherals and buses during extended runtime. Both RTC and GNSS time sources were evaluated over several hours, with GNSS providing continuous correction. Under both sources, the system maintained a consistent 1 Hz output without missed or duplicated samples, confirming reliable timestamping and temporal alignment. The implementation therefore satisfies FR2.4 (timekeeping with GNSS and RTC fallback), NFR1.3 (Nyquist sampling), and NFR1.4 (inter-channel alignment ≤ 1 s).

RTC Drift

The RTC drifted by 1 s after approximately 12 hours without GNSS correction, corresponding to a drift rate of roughly 23.1 ppm. This is within the first-year estimated typical expectation of 30 ppm. This confirms the validity of the manufacturer-specified error budget and the stray-capacitance assumptions, supporting the worst-case budgeted lifetime drift of 70.5 ppm (0.306 s/hour) at room temperature and 115.5 ppm (0.4158 s/hour) across the full temperature range. As a result, NFR2.3 (RTC drift ≤ 1 s/hour) and NFR1.2 (accuracy maintained from -5 to $+55$ °C) are both satisfied. This also meets the latter part of FR2.4 (GNSS time with RTC fallback).

Outcome

The timing subsystem satisfied FR2.4, FR2.6, NFR1.3, NFR1.4, and NFR2.3. Startup time remained within specification, while the GNSS and RTC time sources ensured accurate timestamping. All sensor channels maintained temporal integrity and met the Nyquist criteria, and RTC drift remained well below the allowable limit, ensuring reliable timekeeping during GNSS outages.

4.3.3. Real-Time Latency and Historical Data Access

This subsection validates compliance with NFR2.4 (real-time latency ≤ 2 s) and FR2.5 (historical data access for raw datasets and downsampled views).

Method

Two aspects were validated:

Real-Time Latency Test: Real-time latency was defined as the delay between MCU transmission of a sample and its presentation on the frontend. This included transmission to the cloud, inter-service storage and query delays, delivery of the response to the frontend, and final rendering. Because AWS servers in the US-East region were used due to unavailability in Africa, an additional transmission delay was expected. Sensor averaging added a deliberate half-loop (0.5 s) delay to centre data around UTC time.

Latency was measured by configuring the frontend to query at 1.1111 Hz and subsequently calculating and printing the difference between the frontend system time and the latest received timestamp while the system was active. This slightly higher query rate caused natural oscillation between the minimum delay, representing the true end-to-end transmission latency, and the maximum delay, which included one additional sample interval (1 s) due to misalignment between the request time and the moment a new sample was written to storage.

Historical Data Access Test: Historical access was validated by executing large queries to confirm complete retrieval of raw datasets and correct downsampling for time-series plots and GNSS maps, consistent with FR2.5. Tests included queries of up to 27,000 samples to assess scalability and system behaviour under heavy loads.

Results

For real-time queries, a fetch of 60 samples was used, as this produced no appreciable difference from single-sample retrieval. The observed latency of 0.9–1.9 s reflects temporal offsets between sample storage and frontend requests, where the lower bound corresponds to no offset between the two. Excluding the 0.5 s sensor averaging delay, the true transmission latency is 0.4 s, satisfying NFR2.4 (≤ 2 s typical). As dataset size scales into the hundreds of thousands, DynamoDB efficiency decreases, making Timestream—avoided here due to cost—the preferred alternative.

For historical queries (27,000 samples), DynamoDB retrieval took ~ 12 s to obtain the entire raw dataset. Despite this, all requested samples were retrieved and displayed, fulfilling FR2.5 for complete dataset access. In future, using AWS S3 object storage is preferred to access extensive raw logs that must not be downsampled.

Outcome

The system therefore met NFR2.4 for real-time latency, ensuring timely delivery of live data, and satisfied FR2.5 by providing complete raw dataset access. However, it is noted that future developments with larger budgets should opt for better suited solutions such as Timestream for real-time access, MQTT over WebSockets for lower latency, and S3 for

raw data storage, as discussed in Section 3.4.

4.3.4. Remote Interface Security

This test validates the security of the frontend login interface, the cloud API for web access, and the MQTT interface for device communication. It addresses NFR2.5 (secure remote interfaces).

Method

Five checks were performed:

1. Cognito login: secure user creation, login via the frontend, and token expiry handling.
2. API requests without a valid Cognito JWT.
3. MQTT connection attempts without the provisioned X.509 client certificate.

Results

1. User creation via Cognito succeeded, and frontend login generated a valid session JWT enabling authorised data requests. Figure 4.15 shows the Cognito-hosted login interface. The frontend correctly blocked invalid credentials, rejected expired tokens, and enforced re-login after the 5 minute access-token lifespan once the refresh token was revoked.
2. API calls without a valid or unexpired JWT were rejected.
3. MQTT connection attempts without the device's provisioned X.509 certificate were refused by AWS IoT Core (TLS authentication failure). Valid X.509 certificates enabled secure MCU communication.

Outcome

These validation checks confirm that the frontend login interface, cloud APIs, and MQTT connections are secure and operate as intended. This demonstrates compliance with NFR2.5 (secure remote interfaces).

4.3.5. Frontend and Data Interface

This section validates the frontend dashboard against **FR2.5** exclusively, which requires complete raw dataset downloading and downsampled data viewing via the frontend (maps and plots). Latency and security are considered within prior sections. The focus is on data correctness, query behaviour, retention, and interface functionality. Real-time performance and security considerations are addressed separately.

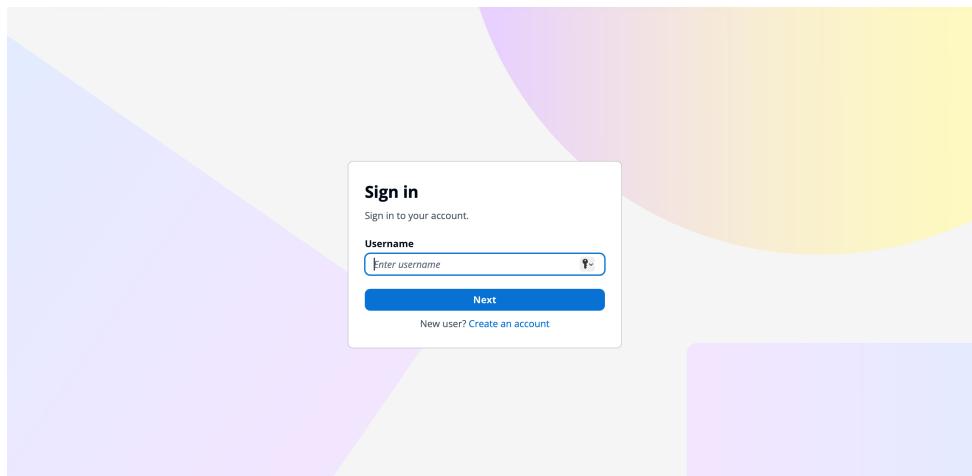


Figure 4.15: Cognito-hosted login interface presented to operators before accessing the telemetry dashboard.

Method

Testing was performed in four stages:

1. **Time-Series Chart:** Queries were tested across long (one month) and short (one minute) time spans to assess downsampled and per-second viewing correctness, and filtering by time range and data fields was exercised to confirm correct query behaviour and accurate data subsets. The auto-refresh functionality was also tested.
2. **Raw Logs (CSV Export):** The CSV export functionality was validated to ensure it encompassed all SD card logs correctly within the requested time frame.
3. **Live Map:** The live GNSS map was tested for correct display of the last known location, associated telemetry, and continuous auto-refresh functionality to allow real-time monitoring of the fleet.
4. **Command Interface:** The command panel was tested for issuing configuration instructions (e.g. WiFi credentials, LTE settings, sensor ranges) to devices, verifying that commands are successfully published and displayed in the console.

Results

Time-Series Chart

A 30 s dataset was retrieved successfully in under 10 s, limited largely by network latency, while a zoomed one-minute view demonstrated full-resolution query behaviour. Queries worked reliably across both long and short time spans, and the time-range selector and data field toggles produced accurate subsets of data, confirming correct query and filtering behaviour. The auto-refresh function operated reliably, as illustrated by the one-month overview in Figure 4.16, the zoomed detail in Figure 4.17, and the filter controls in Figure 4.18.

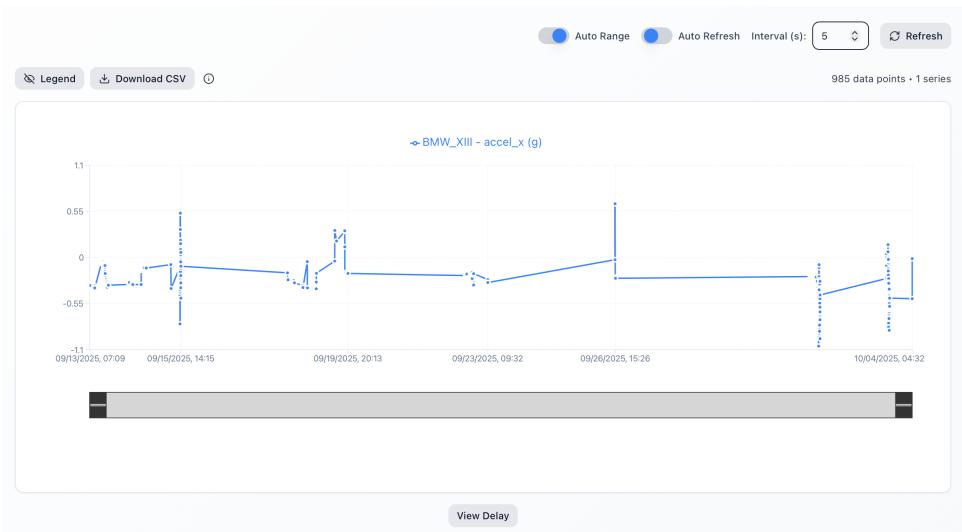


Figure 4.16: Dashboard view of a one-month telemetry query illustrating pagination controls and downsampled trend plotting.



Figure 4.17: Dashboard view of a one-minute telemetry window emphasising high-resolution chart detail.

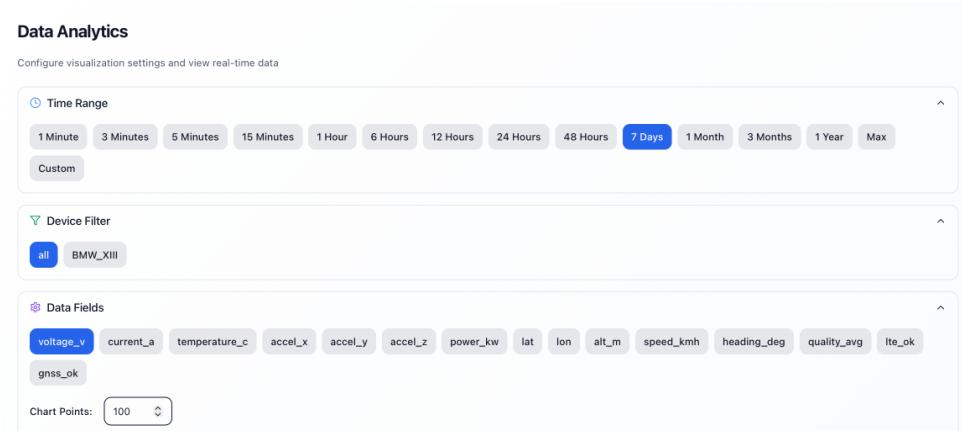


Figure 4.18: Frontend controls for selecting time ranges, devices, and telemetry channels before charts render.

Raw Logs (CSV Export)

The CSV exported file encompassed all SD card logs within the one month time frame, matching the SD logs exactly and ensuring reproducibility for further analysis. This demonstrates the ability to quickly import system logs for detailed offline examination.

Live Map

The live GNSS map correctly displayed the last known location and related timestamp, as well as relevant telemetry. The auto-refresh function worked continuously, enabling real-time monitoring of the fleet, and the map correctly tied telemetry data to the GNSS points, as shown in Figure 4.19. It also showed critical telemetry metrics useful for fleet management.

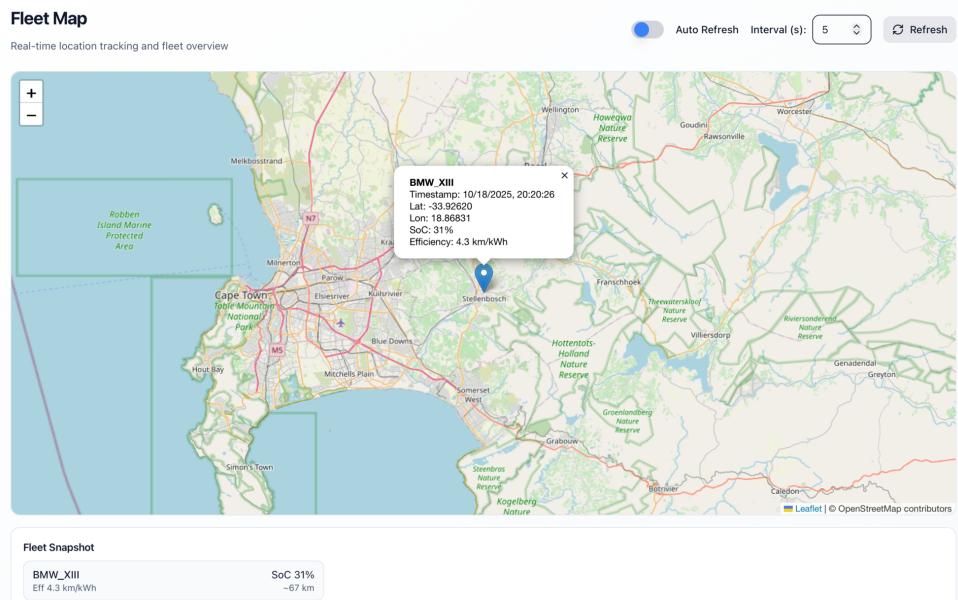


Figure 4.19: Frontend GNSS map displaying latest device positions with interactive metadata pop-ups.

Command Interface

The frontend successfully issued configuration commands such as `set_wifi`, `set_lte_config`, and `set_voltage_range` to the MQTT topic. These commands were successfully published and displayed in the console, but are not implemented on the MCU per requirement scope; Figure 4.20 shows the console output while Figure 4.21 summarises the available command templates.

Outcome

The frontend successfully met FR2.5 by enabling both complete raw data retrieval and downsampled visualisation. Time-series queries functioned correctly across short and long durations, with accurate filtering and dynamic refresh confirming correct data aggregation and query handling. The CSV export produced identical records to the SD-card logs,

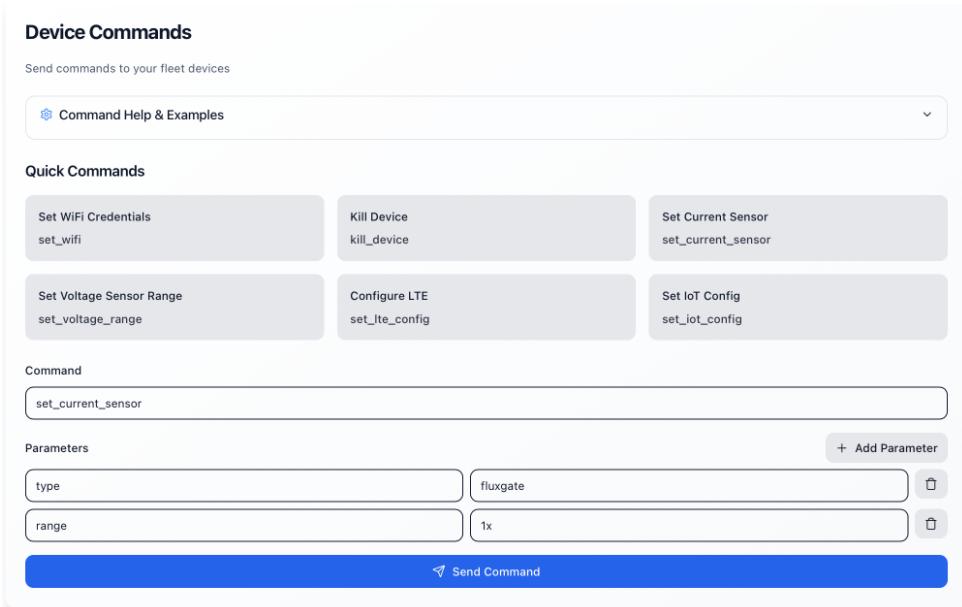


Figure 4.20: Device command console allowing remote configuration of telematics units.

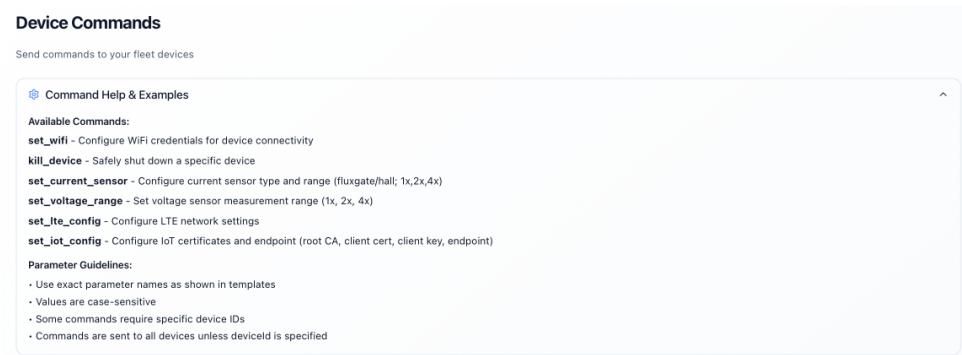


Figure 4.21: Command help and parameter templates confirming correct interface operation.

ensuring full dataset accessibility. The live GNSS map displayed correct positional and telemetry data with continuous auto-refresh, while the command console verified bidirectional communication through successful MQTT message publication. These results confirm full compliance with the required frontend functionality for data access, visualisation, and remote control.

4.4. System Retrofitability, Generalisability and Robustness

This section considers whether the system can be integrated into existing vehicles and applied across diverse fleet contexts. Retrofitability and generalisability are crucial for adoption, while robustness ensures reliable operation under the demanding conditions of real-world automotive use. Together these factors determine the practicality and resilience of the platform in its intended context.

4.4.1. Calibration and Protocol Independence

All tests were performed without calibration and without per-vehicle CAN protocol knowledge (except where explicitly required for proving CAN acquisition or for using CAN data as a reference). The system satisfies NFR3.7 (accuracy without calibration and protocol independence), demonstrating that performance is maintained across vehicles without per-device calibration or dependence on specific CAN mappings.

Non-Invasive Measurement

This subsection validates the non-invasive nature of the voltage, current, and CAN acquisition methods. The EV bench setup described in Section 4.2.2 was used to confirm that no permanent modifications were required to install the sensors. Voltage was measured by tapping into existing terminal connections, current by both the fluxgate sensor—requiring temporary line removal for insertion and refastening—and the GMW clip-on probe, as well as CAN by means of the read-only CSS clip-on inductive coupler. No cuts or splices were required, and both the current and CAN sensors were confirmed to fit the cable diameters used in the EV bench setup. Additionally, per design, the CAN coupler may operate reliably on conductors between 18–30 AWG (0.81–0.25 mm diameter), while the current probes may accommodate up to 27 mm diameter cables with the GMW clip-on probe, adequately covering typical automotive and HV conductors. The CAN clip-on sensor is shown in Figure 3.12, and the current sensors in Figure 3.4. This demonstrates compliance with NFR1.1 (non-intrusive CAN acquisition), NFR3.2 (no permanent vehicle alterations), and NFR3.6 (inductive CAN and current probe fit).

4.4.2. PCB Sizing

The PCB measures 127.4 mm \times 97.4 mm \times 30.2 mm with mounting holes at 90.0 mm \times 120.0 mm, confirmed with digital calipers, and remaining within the 15 cm \times 15 cm \times 5 cm required dimensions per NFR3.1.

4.4.3. Power Usage

This subsection addresses the requirement for power efficiency to minimise parasitic loading on the host EV.

Method

Maximum typical and peak power consumption of the telematics unit was characterised under normal operating conditions, with all subsystems enabled and both current sensors active. The test used LTE Cat 1 bis operation, as 2G RAT functionality is being phased out by 2027. The fluxgate's current draw scales with the measured primary current; however, the available test source was limited to approximately 60 A, well below the transducer's full-scale rating. This test verifies compliance with NFR3.9 by confirming that total power draw remains within design limits and that the regulators maintain correct operation under representative load conditions.

Results

At 12.00 V input, the steady-state supply current with all subsystems active and the higher-load fluxgate connected was typically 0.21 A, with transient peaks up to 0.39 A. This corresponds to approximately 2.52 W typical and 4.68 W peak, both within the design allowances of 25 W (typical) and 50 W (peak), respectively. The design power budget—accounting for worst-case estimates, higher currents, 2G operation, and conservative margins—is therefore considered valid. Accordingly, the system draw is assumed to be 3.41 W typical and 32.18 W peak when operating at full rated load.

Outcome

The system satisfies NFR3.9, with total power draw within limits across operational modes.

4.4.4. Mechanical Robustness

This test validates the enclosure's capability to protect the internal electronics against dust, water, vibration, and mechanical shock, while confirming its compact suitability for in-vehicle installation. The selected enclosure, rated to IP67, was verified through a basic water-immersion test and seen to maintain integrity. As noted in the Section 3.2.4, the cable glands were not installed on this prototype. The in-vehicle test, conducted with

the PCB mounted and the lid removed, showed that vibration did not cause instability or disconnection of wiring or components under high-speed automotive conditions. This subsection therefore demonstrates compliance with NFR3.5 (enclosure protection) and NFR3.8 (vibration and shock resistance). Figure 4.22 illustrates the assembled enclosure with the mounted PCB. Modifications are described in Section 3.2.5.



Figure 4.22: Fully assembled enclosure with the mounted PCB.

4.4.5. Power Rails

This subsection validates correct operation under non-ideal automotive supply conditions. The tests imposed adverse conditions through high acceleration events and monitored all system operations.

Method

1. The system was tested from 8–20 V with \approx 50 mV ripple in a lab environment.
2. It was then evaluated under in-vehicle driving at highway speeds to rapidly charge and discharge the vehicle battery, thereby inducing transient voltage spikes and dips.
3. The system was tested using the EV bench setup to test system reliability on a real electric vehicle platform.

Results

All subsystems operated correctly across the swept range and during highway speeds, maintaining low-noise and correct DC-point operation. The EV bench test, performed in Section 4.2.2, also confirmed correct operation under real EV conditions. Sensitive subsystems maintained operation under all testing conditions, demonstrating robust performance.

Outcome

The system satisfies NFR3.4 (robust operation on non-ideal 12 V supply), maintaining functionality under both transient and steady-state deviations.

Chapter Summary

Table 4.2 provides a results table that consolidates each requirement, its specification, and the measured outcome at room and full temperature where applicable. Other than the first row, room temperature is used to refer to 25 °C and full to refer to the full –5 °C to 55 °C operating range.

Table 4.2: Summary of verification outcomes and traceability back to the functional and non-functional requirements.

ID	Req. and Spec	Measurements or Notes	Evidence
<i>Sensing</i>			
FR1.1	HVDC voltage accuracy $\leq 1\%$ FS to 880 V; safety per NFR3.3.	0.062% FS (25 °C); 0.442% FS (–5 °C to 55 °C); basic insulation 1 kV.	4.2.1
FR1.2 (Fluxgate)	HVDC current $\leq 1.5\%$ FS; range $\pm 1\text{ kA}$.	Error: $\leq 0.3\%$ FS (room); $\leq 0.5\%$ FS (full); non-contact with range $\pm 1500\text{ A}$.	4.2.2
FR1.2 (GMW)	HVDC current $\leq 1.5\%$ FS; range $\pm 1\text{ kA}$.	Error: 1.05% FS (room); 1.36% FS (full); non-contact with range $\pm 1000\text{ A}$.	4.2.2
FR1.3	GNSS position $\leq 10\text{ m}$, speed $\leq 5\text{ km h}^{-1}$.	Typical: $\approx 1.5\text{ m}$ position; $\approx 0.05\text{ m/s}$ speed.	4.2.4
FR1.4	Accelerometer error $\leq 10\%$ over $\pm 2g$.	2.88% FSR (room); 4.38% FSR (full).	4.2.3
FR1.5	Ambient temperature accuracy $\leq 1\text{ }^\circ\text{C}$ across –5 °C to 55 °C.	$\pm 0.5\text{ }^\circ\text{C}$ (room and full).	4.2.5
FR1.6	Acquire vehicle CAN frames.	EV bench test.	4.2.2
<i>Data Handling and Communication</i>			
FR2.1	Local logging to removable storage at 1 Hz.	Validated.	4.3.1, 4.2.2
FR2.2	Cloud transmission at 1 Hz over LTE.	Validated.	4.3.1, 4.3.2

Continued on next page

ID	Req. and Spec	Notes / Measured (room / max)	Evidence
FR2.3	Retransmit missed historical data after outages.	Validated up to 10 samples/s.	4.3.1
FR2.4	Timekeeping with GNSS and RTC fallback.	Validated.	4.3.2
FR2.5	Frontend and data interface (raw and downsampled views).	Validated.	4.3.5
FR2.6	Startup to first sample ≤ 10 s.	≈ 6 s (incl. MQTT, excl. GNSS TTFF).	4.3.2
FR2.7	GNSS TTFF ≤ 45 s (cold), reacquisition ≤ 3 s.	Cold ≤ 39 s; reacquisition ≤ 3 s.	4.2.4
<i>Non-Functional: Sensing</i>			
NFR1.1	Non-intrusive CAN.	Inductive tapping; validated during EV testing.	4.4.1, 4.2.2
NFR1.2	Accuracy maintained across -5°C to 55°C .	Verified for all sensors and part choices using Assumption 1(4.1.3).	4.1.3
NFR1.3	Nyquist sampling for each channel.	Validated.	4.3.2
NFR1.4	Inter-channel alignment ≤ 1 s.	Validated.	4.3.2
<i>Non-Functional: Data Handling and Communication</i>			
NFR2.1	SD capacity ≥ 6 months.	≥ 13 months including headroom.	4.3.1
NFR2.2	LTE reliability at automotive speeds.	Validated up to 120 km h^{-1} .	4.3.1
NFR2.3	RTC drift ≤ 1 s/hour.	0.306 s/hour (room) and 0.4158 s/hour (full).	4.3.2
NFR2.4	Real-time latency ≤ 2 s.	0.4 s MCU to frontend.	4.3.3
NFR2.5	Secure remote interfaces.	X.509; TLS; JWT via Cognito.	4.3.4
<i>Non-Functional: Retrofittability, Generalisability and Robustness</i>			
NFR3.1	PCB dimensions within $\leq 150 \times 150 \times 50$ mm.	127.4 mm \times 97.4 mm \times 30.2 mm.	4.4.2
NFR3.2	No permanent vehicle modifications.	terminal taps (voltage) and inductive clip-ons (current and CAN); no cuts/splices.	4.4.1
NFR3.3	HV insulation and safety margin.	Basic insulation to 1 kV; isolated ADC 1.2 kV; ≈ 1.85 kV transient tolerance.	4.2.1, 4.2.2
NFR3.4	Operate from non-ideal 12 V supply.	Stable on high-speed road test.	4.4.5, 4.2.2
NFR3.5	Enclosure protection.	IP67 enclosure.	4.4.4
NFR3.6	Probe fit for inductive sensing.	Current clamps ≤ 27 mm; CAN 18-30 AWG.	4.4.1, 4.2.2
NFR3.7	No calibration required.	No calibration performed; requirements met.	4.4.1
NFR3.8	Vibration/shock resistance.	Stable on high-speed road test.	4.4.4

Continued on next page

ID	Req. and Spec	Notes / Measured (room / max)	Evidence
NFR3.9	Power requirements within limits.	Typical \approx 3.4 W; peaks < 33 W.	4.4.3

The chapter's testing encompassed laboratory characterisation, in-vehicle driving, and a stationary EV bench. Across these contexts, the system acquired accurate, time-aligned electromobility data, transmitted reliably over LTE, and logged locally with sufficient capacity. The cloud system and web application provided secure storage, data access, and analytical functionality. Installation was non-invasive and compact, with appropriate electrical isolation and protection from high-voltage hazards.

Limitations of the present validation include reliance on datasheet specifications where high-precision references, high-power systems, or environmental chambers were unavailable (Assumption 1; see Section 4.1.3). The cloud system's US-based servers introduced latency for the South African test site, and its storage architecture lacked scalability for large-scale, fleet-level operation. Accelerometer accuracy during low-acceleration events was constrained by offset errors.

Overall, the system satisfies all stated functional and non-functional requirements under realistic conditions and demonstrates feasibility for retrofit, vehicle-agnostic, fleet-scale deployment. This enables the acquisition of high-quality electromobility data without permanent vehicle modifications or CAN access.

Chapter 5

Conclusion

The transition to EVs offers the potential for substantial reductions in transport costs and urban air pollution. Realising these benefits, however, depends on sound planning and viability assessments supported by accurate, locally relevant data. In South Africa, large-scale datasets linking how EVs are driven to how they consume energy remain scarce. Access to vehicle sensor data is restricted by proprietary CAN protocols, while no universal, direct-sensing telematics solution exists.

This work therefore set out to design and validate a retrofit, vehicle-agnostic telematics platform capable of acquiring accurate, time-aligned power and mobility data for fleet-scale profiling and decision-making. The developed prototype was evaluated through laboratory characterisation, a stationary EV bench, and on-road testing, demonstrating compliance with all sensing, data-handling, and retrofit requirements under realistic operating conditions.

Power sensing achieved sub-percent full-scale accuracy across the full temperature range for both voltage and current measurements, supporting DC voltages up to 880 V and bidirectional currents of ± 1.5 kA, with full electrical isolation and protection. Mobility sensing through GNSS systems achieved typical positional accuracy of 1.5 m and velocity accuracy of 0.05 m s^{-1} , enabling reliable motion profiling. Accelerometer and temperature sensing met their respective accuracy requirements for meaningful inference, and vehicle CAN data acquisition capability provided additional application generalisability.

Time was synchronised via GNSS and maintained by a battery-backed RTC, ensuring reliable 1 Hz timestamping. Measurements were logged locally to SD storage and simultaneously streamed to a secure cloud platform over LTE, with automatic retransmission of any missed records for complete archival. The web application enabled visualisation and export with end-to-end latency below 1 s.

The system remained compact, non-invasive, and robust under automotive conditions. Installation required no cuts or modifications to the vehicle, with inductive current and CAN sensing via clip-on connectors. Power consumption remained within limits while maintaining stable operation under vibration and non-ideal vehicle supply conditions.

The outcomes of this work demonstrate the feasibility of a fully functional, retrofit, and vehicle-agnostic telematics platform capable of producing high-fidelity electromobility datasets under local operating conditions. The device provides a practical and reliable

means of generating locally relevant data that address the current shortage of empirical evidence needed to support collective research and informed state intervention in the development of electric transport in South Africa.

Future Work

Future development should advance the prototype toward a production-ready system with increased safety, reliability, and scalability across the hardware, firmware, cloud, and web-application layers.

At the hardware level, potential improvements include reinforced insulation and other HVDC safety enhancements, AC sensing capability, higher-bandwidth electromobility data acquisition, CAN-FD-compliant components, multi-antenna RATs for increased coverage, accelerometers with lower inherent offset, a high-stability RTC to maintain timing during extended outages, widened apertures on current sensors to accommodate varied bus geometries, and standardised interface cabling for faster and more reliable installation. Additionally, firmware development improvements may involve secure over-the-air updates, reliable handling of inbound MQTT messages, improved robustness under edge-case conditions, optional sensor calibration, and at-rest encryption for both SD and MCU flash storage. Cloud infrastructure development may leverage more scalable storage architectures and WebSocket-based MQTT for higher throughput, supported by enhanced observability and regional deployment to reduce latency. Finally, the web application layer may be expanded with a more capable analytical and visualisation toolset to improve accessibility and insight generation.

Together, these enhancements would enable a robust, production-grade deployment supporting data-driven transport electrification across South Africa.

Bibliography

- [1] E. T. Inc. (2025) Ev power electronics: Purpose of key components. Accessed: 23 Oct 2025. [Online]. Available: <https://www.exro.com/industry-insights/ev-power-electronics-explained>
- [2] Electronics Tutorials. (2025) Hall effect sensor and how magnets make it work. Accessed: 26 October 2025. [Online]. Available: <https://www.electronics-tutorials.ws/electromagnetism/hall-effect.html>
- [3] D. A/S. (2020) How a zero flux current transducer works: new video from danisense. Accessed: 2025-09-13. [Online]. Available: <https://danisense.com/news/fluxgate-video/>
- [4] “Contactless can bus reader,” CSS Electronics, Aarhus, Denmark, 2025, accessed: 2025-08-13. [Online]. Available: <https://www.csselectronics.com/products/contactless-can-bus-reader-adapter#void>
- [5] J. Finnerty. (2025, February 25) Why norway leads the world in ev adoption. GRIDSERVER Newsroom. [Online]. Available: <https://www.gridserve.com/why-norway-leads-the-world-in-ev-adoption/>
- [6] I. E. Agency, *Global EV Outlook 2024: Outlook for Electric Mobility*, 2024. [Online]. Available: <https://www.iea.org/reports/global-ev-outlook-2024/outlook-for-electric-mobility>
- [7] P. Reddy and J. Radmore, “Electric vehicles: Market intelligence report 2024,” GreenCape, Cape Town, South Africa, Tech. Rep., 2024, accessed October 2025. [Online]. Available: <https://greencape.co.za/wp-content/uploads/2024/08/Electric-Vehicles-digital.pdf>
- [8] S. Lacock, A. A. du Plessis, and M. J. Booysen, “Using driving-cycle data to retrofit and electrify sub-saharan africa’s existing minibus taxis for a circular economy,” *World Electric Vehicle Journal*, vol. 14, no. 10, 2023. [Online]. Available: <https://www.mdpi.com/2032-6653/14/10/296>
- [9] K. Collett and S. Hirmer, “Data needed to decarbonize paratransit in sub-saharan africa,” *Nature Sustainability*, vol. 4, no. 8, pp. 642–655, 2021. [Online]. Available: https://www.researchgate.net/publication/351569881_Data_needed_to_decarbonize

https://www.semanticscience.org/resource/paratransit_in_Sub-Saharan_Africa/link/60be5f3ba6fdcc22eae893e9/download?tp=eyJjb250ZXh0Ijp7InBhZ2UiOiJwdWJsaWNhdGlvbiIsInByZXZpb3VzUGFnZS16bnVsbH19

[10] W. Lacock, C. Hull, K. Collett, and M. Booysen, “High fidelity estimates of paratransit energy consumption from per-second gps tracking data,” *Transportation Research Part D: Transport and Environment*, vol. 118, p. 103695, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1361920923000925>

[11] Z. Bi, G. Xu, G. Xu, C. Wang, and S. Zhang, “Bit-level automotive controller area network message reverse framework based on linear regression,” *Sensors*, vol. 22, no. 3, p. 981, 2022.

[12] S. Kothari, “Here’s how california could pave the way for an industry standard ev diagnostic system,” Nov. 2023. [Online]. Available: <https://insideevs.com/news/696065/ev-diagnostic-system-arriving-by-2026/>

[13] L. Kang. (2025, Apr.) Byd to officially launch 5-minute charging han 1 sedan, tang 1 suv on apr 9. Accessed: 2025-08-24. [Online]. Available: <https://cnevpost.com/2025/04/07/byd-to-launch-han-tang-l-apr-9/>

[14] I. N. Laboratory, “Plugged in: How americans charge their electric vehicles (ev project summary report),” Idaho National Lab, Tech. Rep., 2015. [Online]. Available: <https://avt.inl.gov/sites/default/files/pdf/arra/PluggedInSummaryReport.pdf>

[15] Z. Guo, L. You, R. Zhu, Y. Zhang, and C. Yuen, “A city-scale and harmonized dataset for global electric vehicle charging demand analysis,” *Scientific Data*, vol. 12, no. 1254, 2025. [Online]. Available: <https://www.nature.com/articles/s41597-025-05584-7>

[16] W. Achariyaviriya, W. Wongsapai, D. Rinchumphu, N. Tippayawong, K. Yaibuathet, and P. Suttakul, “A comparative study of vehicle powertrain efficiency: Data-driven analyzing energy consumption and environmental impact,” *Transportation Engineering*, vol. 18, 2024.

[17] G. Team. (2025) The ultimate 2025 guide to fleet telematics: Maximize safety & efficiency. Accessed: 28 October 2025. [Online]. Available: <https://www.geotab.com/blog/fleet-telematics-safety-efficiency-guide/>

[18] J. Pearce, “Electric vehicle telemetry,” 2010. [Online]. Available: <https://roblab.org/theses/2010-REV-Telemetry-Pearce.pdf>

[19] P. Kaur and A. Gupta, “Real-time battery monitoring and predictive maintenance for electric scooters using iot technology,” *International Journal of Advanced*

Research in Engineering and Technology (IJARET), vol. 10, no. 6, pp. 846–855, 2019. [Online]. Available: <https://iaeme.com/Home/issue/IJARET?Volume=10&Issue=6>

[20] A. F. Hasibuan and A. Ma’arif, “Battery usage monitoring system internet of things-based electric cars (iot) and radio telemetry,” *Buletin Ilmiah Sarjana Teknik Elektro*, vol. 6, no. 3, pp. 214–222, 2024. [Online]. Available: <https://doi.org/10.12928/biste.v6i3.11524>

[21] I. S. Priya, A. Abinesh, S. Aswin, B. M. Kumar, and S. Prabin, “Iot based ev battery parameters monitoring system,” *International Journal of Advanced Trends in Engineering and Management*, vol. 2, no. 4, pp. 28–35, 2023.

[22] A. Abinaya, N. Vivetha, G. Shibin, P. Y. Prasanna, and N. Sowmiya, “Monitoring the ev battery management system by using iot,” *International Journal of Innovative Research in Technology*, vol. 11, no. 12, pp. 6541–6546, 2025.

[23] A. Jadhav, H. Rathod, and V. Kanade, “Smart battery parameter monitoring system for electric vehicles,” *International Journal of Advances in Engineering and Management (IJAEM)*, vol. 5, no. 5, pp. 685–688, 2023.

[24] A. Pérez-Gonzalez, A. F. Villa-Salazar, I. N. Gómez-Miranda, J. D. Velásquez-Gómez, and A. F. Romero-Maya, “An iov-based real-time telemetry and monitoring system for electric racing vehicles: Design, implementation, and field validation,” *Preprints.org*, 2025, available under CC BY 4.0 License.

[25] J. Winkelholz, M. Hitzemann, A. Nitschke, A. Zygmanowski, and S. Zimmermann, “Resistive high-voltage probe with frequency compensation by planar compensation electrode integrated in printed circuit board design,” *Electronics*, vol. 11, no. 21, p. 3446, Oct 2022. [Online]. Available: <https://doi.org/10.3390/electronics11213446>

[26] K. Zhu, W. K. Lee, and P. W. T. Pong, “Non-contact voltage monitoring of HVDC transmission lines based on electromagnetic fields,” *IEEE Sensors Journal*, vol. 19, no. 8, pp. 3123–3132, Apr. 2019.

[27] *Voltage Transducer DVC 1000-P*, LEM International SA, February 2018, datasheet, Version 0. [Online]. Available: <https://www.lem.com>

[28] J. Fattakhov, “Maximize ev battery longevity with integrated resistor dividers,” *Electronic Design*, 2025, march 3, 2025. Accessed 18 Aug 2025. [Online]. Available: <https://www.electronicdesign.com/markets/automotive/article/55269327/texas-instruments-maximize-ev-battery-longevity-with-integrated-resistor-dividers>

- [29] *IEC 60664-1: Insulation coordination for equipment within low-voltage systems – Part 1: Principles, requirements and tests*, International Electrotechnical Commission (IEC) Std., 2007, older standard, superseded by IEC 62368-1 in this context.
- [30] *IEC 62368-1: Audio/video, information and communication technology equipment – Part 1: Safety requirements*, International Electrotechnical Commission (IEC) Std., 2018, primary safety standard governing creepage, clearance, and insulation requirements.
- [31] B. Schweber, “Faq on creepage and clearance: Part 2,” *Power Electronic Tips*, 2023, accessed 18 Aug 2025. [Online]. Available: <https://www.powerelectronicstips.com/faq-on-creepage-and-clearance-part-2/>
- [32] M. P. Corporation, “Clearance and creepage distance in pcb design: Essential guidelines for high-voltage insulation,” *Minmax Technical Blog*, 2023, accessed 18 Aug 2025. [Online]. Available: <https://www.minmaxpower.com/en/resource/clearance-creepage-pcb-design>
- [33] L. Gao, C. Jiang, Z. Guo, B. Han, S. Sun, T. Cao, and H. Huang, “High sensitivity dc current sensor based on microfiber mach-zehnder interferometer combined with copper wire,” *Optics and Lasers in Engineering*, vol. 193, p. 109111, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0143816625002969>
- [34] M. P. Systems. (2025) Magnetic position and angle sensors: Common types, key components, parameters, usage considerations and applications. Accessed: 2025-09-13. [Online]. Available: <https://www.monolithicpower.cn/cn/learning/resources/magnetic-sensors-common-types-key-components-parameters-usage-considerations-and-applications>
- [35] ——, “Current sensors: Types, key parameters, performance comparison, and common applications,” *Monolithic Power Systems Technical Article*, 2025, accessed 18 Aug 2025. [Online]. Available: <https://www.monolithicpower.com/en/learning/resources/current-sensors-types-key-parameters-performance-comparison-and-common-applications>
- [36] D. d.o.o., “How to measure current using current sensors,” *Dewesoft Blog*, 2024, accessed 18 Aug 2025. [Online]. Available: <https://dewesoft.com/blog/how-to-measure-current-using-current-sensors>
- [37] M. Soltero, “What are you sensing? pros and cons of four temperature sensor types,” Texas Instruments, Tech. Rep. SSZTCA2, 2015. [Online]. Available: <https://www.ti.com/lit/ta/ssztca2/ssztca2.pdf>

- [38] A. Pacanowsky. (2024) The most common types of temperature probe sensors and how to choose the right one. Accessed: 2025-08-18. [Online]. Available: <https://www.powerblanket.com/blog/the-most-common-types-of-temperature-probe-sensors-and-how-to-choose-the-right-one/>
- [39] S. L, *The Influence of Temperature on the Capacity of Lithium Ion Batteries with Different Anodes*, 2021. [Online]. Available: <https://www.mdpi.com/1996-1073/15/1/60>
- [40] Endevco, “Selecting the right accelerometer for your application,” Endevco (Meggitt Sensing Systems), Tech. Rep. TP327, 2019. [Online]. Available: https://www.endevco.com/contentStore/mktgContent/endevco/dlm_uploads/2019/02/TP327.pdf
- [41] O. E. . DwyerOmega. (2023) Accelerometers: How to choose. Accessed: 2025-08-18. [Online]. Available: <https://www.dwyeromega.com/en-us/resources/accelerometers-how-to-choose>
- [42] Q. C. Sun, R. Odolinski, J. C. Xia, J. K. Foster, T. Falkmer, and H. Lee, “Validating the efficacy of gps tracking vehicle movement for driving behaviour assessment,” *Travel Behaviour and Society*, vol. 6, pp. 32–43, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214367X15300120>
- [43] I. Analog Devices, “Which adc architecture is right for your application?” *Analog Dialogue*, 2024, accessed 18 Aug 2025. [Online]. Available: <https://www.analog.com/en/resources/analog-dialogue/articles/the-right-adc-architecture.html>
- [44] S. Electronics, “How do microcontrollers work and how do i choose one?” 2025, accessed 18 Aug 2025. [Online]. Available: <https://www.suntsu.com/knowledge-center/how-microcontrollers-work-and-choosing-the-right-mcu/>
- [45] KitBuilder, “Are you choosing the right microcontroller? comprehensive analysis of arduino vs esp32 vs stm32,” 2022, accessed 18 Aug 2025. [Online]. Available: <https://kitbuilder.io/blog/arduino-vs-esp32-vs-stm32-comparison>
- [46] GeeksforGeeks. (2025, July) Advantages and disadvantages of microcontroller. Accessed: 2025-08-18. [Online]. Available: <https://www.geeksforgeeks.org/computer-organization-architecture/advantages-and-disadvantages-of-microcontroller/>
- [47] T. A. Wild, G. Wilbs, D. K. Dechmann, J. E. Kohles, N. Linde, N. Richter, S. Sfantzourakis, H. Nicolaou, E. Erokstrova *et al.* (2024, October) Time synchronisation for millisecond-precision on bio-loggers. [Online]. Available: <https://movementecologyjournal.biomedcentral.com/articles/10.1186/s40462-024-00512-7>

- [48] A. Devices. (2022) How to level shift 1-wire systems. Accessed: 2025-08-18. [Online]. Available: <https://www.analog.com/en/resources/design-notes/how-to-level-shift-1wire-systems.html>
- [49] B. Automotive. (2025, March) Beijer measures first vehicles with only can fd: Kia ev3 and audi q6 e-tron. Accessed: 2025-08-21. [Online]. Available: [https://www.beijer.com/en/news/beijer-meet-eerste-voertuigen-met-uitsluitend-can-fd-kia-ev3-en-audi-q6-e-tron/#:~:text=CAN%20FD%20\(Controller%20Area%20Network,for%20their%20complex%20electronic%20systems.](https://www.beijer.com/en/news/beijer-meet-eerste-voertuigen-met-uitsluitend-can-fd-kia-ev3-en-audi-q6-e-tron/#:~:text=CAN%20FD%20(Controller%20Area%20Network,for%20their%20complex%20electronic%20systems.)
- [50] A. Ltd. (2023) Gnss module performance: 4 key metrics. Accessed: 24 October 2025. [Online]. Available: <https://blog.antenova.com/gnss-module-performance-4-key-metrics>
- [51] R. Security. (2019) Accuracy, precision and reliability of gps tracking and gps trackers. [Online]. Available: <https://www.rewiresecurity.co.uk/blog/gps-tracking-tracker-satellite-accuracy-precision>
- [52] M. Illidge, *Changes for South African 2G and 3G switch-off*, 2025. [Online]. Available: <https://mybroadband.co.za/news/cellular/601109-changes-for-south-african-2g-and-3g-switch-off.html>
- [53] M. Bosson, “Lte cat 1 bis: An ultimate guide with practical examples,” May 2023. [Online]. Available: <https://onomondo.com/blog/what-is-lte-cat-1-bis-and-why-is-it-good-for-iot/>
- [54] J. Rosende, “5g iot: what it means for lte-m, nb-iot, and your deployment,” Nov 2024, accessed 26 Aug 2025. [Online]. Available: <https://onomondo.com/blog/what-is-5g-iot-lte-m-nb-iot-and-your-deployment/>
- [55] *What is NB-IoT? Current NB-IoT status and settings for South Africa*, SIM Control, 2025. [Online]. Available: <https://flickswitch.freshdesk.com/support/solutions/articles/23000013586-what-is-nb-iot-current-nb-iot-status-and-settings-for-south-africa>
- [56] N. S. DevAcademy, “Lte-m and nb-iot – cellular iot fundamentals, lesson 1,” 2023, accessed: 2025-08-26. [Online]. Available: <https://academy.nordicsemi.com/courses/cellular-iot-fundamentals>
- [57] J. Rosende, *An extended guide to NB-IoT: Practical insights for IoT projects*, 2024. [Online]. Available: <https://onomondo.com/blog/what-is-nb-iot-definitive-guide/#:~:text=Latency:%20NB%2DIoT%20has%20a,critical%20real%2Dtime%20data%20transmissions>

- [58] O. Akbarzadeh and H. Attar, “Investigating the effect of latency in a 5g infrastructure for applications requiring ultra-low latency,” *2022 International Engineering Conference on Electrical, Energy, and Artificial Intelligence (EICEEAII)*, pp. 1–6, 2022, latency benchmarks for 2G–5G networks.
- [59] I. Grigorik, *High Performance Browser Networking*. O'Reilly Media, 2013, discusses TCP/IP, TLS, HTTP/2, cellular latency and throughput characteristics. [Online]. Available: <https://hpbn.co/>
- [60] u-blox, *LENA-R8 Series: Multi-mode LTE Cat 1 bis Modules*, 2022. [Online]. Available: https://content.u-blox.com/sites/default/files/documents/LENA-R8_Datasheet_UBX-22003110.pdf
- [61] u-blox AG, *SARA-U2 series – HSPA modules with 2G fallback*, 2017. [Online]. Available: <https://www.alldatasheet.com/datasheet-pdf/view/1242691/U-BLOX/SARA-U260.html>
- [62] Q. W. Solutions, *RG500U Series Hardware Design*, 2024, released documentation, 5G module hardware design guide. [Online]. Available: https://www.quectel.com/download/quectel_rm500u_series_hardware_design_v1-1/
- [63] (2024) Lte network latency compared with 2g, 3g & wifi. CableFree. Accessed: 12 Oct 2025. [Online]. Available: <https://www.cablefree.net/wirelesstechnology/4glte/lte-network-latency/>
- [64] everythingRF, “Gnss antennas: A resource of whitepapers and articles on gnss antenna,” 2025. [Online]. Available: https://cdn.everythingrf.com/live/GNSS_Antenna_eBook_2024_638705533460814107_638761398462817853.pdf
- [65] A. Castaldo, “Switching regulator fundamentals,” Texas Instruments, Tech. Rep. SNVA559C, 2012, revised February 2019. [Online]. Available: <https://www.ti.com/lit/an/snva559c/snva559c.pdf>
- [66] S. Electric. (2024) All you need to know about voltage regulators & its application. Accessed: 2025-08-20. [Online]. Available: <https://eshop.se.com/in/blog/post/all-you-need-to-know-about-voltage-regulators-its-application.html>
- [67] Z. Peterson, “What is pcb star grounding and why would anyone use it?” *Altium Resources*, September 21 2021, accessed: 2025-08-20. [Online]. Available: <https://resources.altium.com/p/using-pcb-star-grounding-can-keep-your-design-shining>
- [68] ——. (2022, January) Using an ipc-2152 calculator: Designing to standards. [Online]. Available: <https://resources.altium.com/p/using-ipc-2152-calculator-designing-standards>

- [69] E. Next. (2025) Your guide to electric car battery capacity and lifespan. Accessed: 2025-08-20. [Online]. Available: <https://www.eonnext.com/electric-vehicles/guides/costs-and-benefits-of-electric-cars/battery-capacity-and-lifespan>
- [70] P. Mell and T. Grance, “The nist definition of cloud computing,” National Institute of Standards and Technology, Tech. Rep. SP 800-145, 2011. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>
- [71] U. Hunkeler, H. L. Truong, and A. Stanford-Clark, “Mqtt-s—a publish/subscribe protocol for wireless sensor networks,” in *Proc. 3rd Int. Conf. on Communication Systems Software and Middleware*, 2008.
- [72] Z. Shelby, K. Hartke, and C. Bormann, “The constrained application protocol (coap),” 2014. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7252>
- [73] I. Fette and A. Melnikov, “The websocket protocol,” 2011. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6455>
- [74] “Triggers and bindings in azure functions,” 2025. [Online]. Available: <https://learn.microsoft.com/azure/azure-functions/functions-triggers-bindings>
- [75] “What is aws lambda?” 2025. [Online]. Available: <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>
- [76] “Kubernetes concepts overview,” 2024. [Online]. Available: <https://kubernetes.io/docs/concepts/overview/>
- [77] “X.509 certificate authentication for iot hub,” 2025. [Online]. Available: <https://learn.microsoft.com/azure/iot-hub/tutorial-x509-introduction>
- [78] “What is mutual tls?” 2025. [Online]. Available: <https://www.cloudflare.com/learning/access-management/what-is-mutual-tls/>
- [79] “What is the principle of least privilege (polp)?” 2025. [Online]. Available: <https://www.cloudflare.com/learning/access-management/principle-of-least-privilege/>
- [80] 5 Top Cloud Service Providers in 2025 Compared, datacamp, 2025. [Online]. Available: <https://www.datacamp.com/blog/top-cloud-service-providers-compared>
- [81] B. Beyer, C. Jones, N. Petoff, and J. L. Murphy, “Site reliability engineering: Monitoring distributed systems (four golden signals),” 2016. [Online]. Available: <https://sre.google/sre-book/monitoring-distributed-systems/>
- [82] “Aws iot core pricing,” 2025. [Online]. Available: <https://aws.amazon.com/iot-core/pricing/>

- [83] Z. Shi, J. Dong, and Y. Gan, “Democratizing digital transformation: A multisector study of low-code adoption patterns, limitations, and emerging paradigms,” *Applied Sciences*, vol. 15, no. 12, p. 6481, 2025.
- [84] J. Swacha and A. Kulpa, “Evolution of popularity and multiaspectual comparison of widely used web development frameworks,” *Electronics*, vol. 12, no. 17, p. 3563, 2023.
- [85] A. O’Dwyer, “Benchmark comparison of javascript frameworks: React, angular, vue,” Trinity College Dublin, Tech. Rep., 2021. [Online]. Available: <https://www.scss.tcd.ie/publications/theses/diss/2021/TCD-SCSS-DISSERTATION-2021-020.pdf>
- [86] C. Nordstrom and A. Dixelius, “Comparisons of server-side rendering and client-side rendering for web pages,” Master’s thesis, Uppsala University, Sweden, 2023. [Online]. Available: <https://www.diva-portal.org/smash/get/diva2:1797261/FULLTEXT02.pdf>
- [87] Cloudflare, “What is a cdn? how do cdns work?” 2025. [Online]. Available: <https://www.cloudflare.com/learning/cdn/what-is-a-cdn/>
- [88] Netlify, “Https (ssl) – netlify documentation,” 2025. [Online]. Available: <https://docs.netlify.com/manage/domains/secure-domains-with-https/https-ssl/>
- [89] Vercel, “Working with ssl certificates,” 2025. [Online]. Available: <https://vercel.com/docs/domains/working-with-ssl>
- [90] D. Verma and P. Aland, “A comparative review of server rendering and client side rendering in web development,” *International Journal of Scientific Research & Engineering Trends*, vol. 10, no. 2, pp. 521–526, 2024. [Online]. Available: https://ijsret.com/wp-content/uploads/2024/03/IJSRET_V10_issue2_218.pdf
- [91] Vercel, “Incremental static regeneration (isr) – next.js documentation,” 2025. [Online]. Available: <https://nextjs.org/docs/pages/guides/incremental-static-regeneration>
- [92] R. Appelqvist and N. Örnmyr, “Performance comparison of xhr polling, long polling, server-sent events and websockets,” Blekinge Institute of Technology, Tech. Rep., 2017. [Online]. Available: <https://www.diva-portal.org/smash/get/diva2:1133465/FULLTEXT01.pdf>
- [93] W. Lasocha, “Comparison of websocket and http protocol performance,” *Journal of Computer Sciences Institute*, vol. 19, pp. 87–92, 2021. [Online]. Available: <https://ph.pollub.pl/index.php/jcsi/article/view/2452>

- [94] M. W. Docs. (2025) Responsive web design – learn web development. [Online]. Available: https://developer.mozilla.org/en-US/docs/Learn_web_development/Core/CSS_layout/Responsive_Design
- [95] W. WAI, “Understanding success criterion 1.4.3: Contrast (minimum),” 2025. [Online]. Available: <https://www.w3.org/WAI/WCAG22/Understanding/contrast-minimum.html>
- [96] M. W. Docs, “Wai-aria basics,” 2025. [Online]. Available: https://developer.mozilla.org/en-US/docs/Learn_web_development/Core/Accessibility/WAI-ARIA_basics
- [97] ——. (2025) Set-cookie header. [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/HTTP/Reference/Headers/Set-Cookie>
- [98] OWASP, “Session management cheat sheet,” 2024. [Online]. Available: https://cheatsheetseries.owasp.org/cheatsheets/Session_Management_Cheat_Sheet.html
- [99] A. Monga. (2023) Securing cookies: why you should always set `httpOnly`. [Online]. Available: <https://medium.com/@ajay.monga73/securing-cookies-why-you-should-always-set-httpOnly-92489cbf76c1>
- [100] Wikipedia. (2025) Secure cookie. [Online]. Available: https://en.wikipedia.org/wiki/Secure_cookie
- [101] A. Lock. (2025) Understanding `samesite` cookies. [Online]. Available: <https://andrewlock.net/understanding-samesite-cookies/>
- [102] M. W. Docs, “Cross-origin resource sharing (cors) – guide,” 2025. [Online]. Available: <https://developer.mozilla.org/en-US/docs/Web/HTTP/Guides/CORS>
- [103] OWASP, “Content security policy cheat sheet,” 2021. [Online]. Available: https://cheatsheetseries.owasp.org/cheatsheets/Content_Security_Policy_Cheat_Sheet.html
- [104] ——. (2025) Http security response headers cheat sheet. [Online]. Available: https://cheatsheetseries.owasp.org/cheatsheets/HTTP_Headers_Cheat_Sheet.html
- [105] ——. (2024) Clickjacking defense cheat sheet. [Online]. Available: https://cheatsheetseries.owasp.org/cheatsheets/Clickjacking_Defense_Cheat_Sheet.html
- [106] PortSwigger. (2025) Clickjacking. [Online]. Available: <https://portswigger.net/web-security/clickjacking>
- [107] Geotab. (2025) Electric vehicle reporting and monitoring – user guide. [Online]. Available: <https://support.geotab.com/mygeotab/doc/ev-reporting>

[108] ——. (2025) Electric vehicle fleet management solutions. [Online]. Available: <https://www.geotab.com/ie/fleet-management-solutions/electric-vehicles/>

[109] *u-blox 8 / u-blox M8 Receiver Description*, u-blox AG, 2023. [Online]. Available: https://content.u-blox.com/sites/default/files/products/documents/u-blox8-M8_ReceiverDescrProtSpec_UBX-13003221.pdf

[110] *AMC131M01 1-Channel, 64-kSPS, Simultaneous-Sampling, 24-Bit, Reinforced Isolated Delta-Sigma ADC With Integrated DC/DC Converter*, Texas Instruments Incorporated, Sep. 2023, rev. SBASA57. [Online]. Available: <https://www.ti.com/lit/ds/symlink/amc131m01.pdf>

[111] *AMC131M03 Evaluation Module User's Guide*, Texas Instruments, April 2023, <https://www.ti.com/lit/ug/sbau402/sbau402.pdf>.

[112] L. I. SA, *CAB-SF 1500 Series Automotive Current Transducer – Fluxgate Technology*, 2023, model CAB-SF 1500-000. Nominal current: ± 1500 A. Power supply: +12 V. Output: High-speed CAN (500 kbps). Mounting: Busbar or panel. [Online]. Available: https://www.lem.com/sites/default/files/products_datasheets/cab_150_0_generic_public_v4.pdf

[113] G. Associates, “Cpc-1000-27-bp2 dc-ac current probe clip, 27 mm aperture,” 2018, nominal range ± 1000 A, maximum ± 1100 A. Accuracy $\pm 1.0\%$ full scale, non-linearity $\pm 0.2\%$. Output ± 2.00 V full scale, bandwidth DC to 75 kHz, supply 3.5–5.5 V, ± 85 mA. Operating temperature -40°C to $+100^{\circ}\text{C}$. [Online]. Available: https://gmw.com/wp-content/uploads/2019/08/CPC-XXXX-27-BP2_04_18_18.pdf

[114] *ADS1220 4-Channel, 2-kSPS, Low-Power, 24-Bit ADC With Integrated PGA and Reference*, Texas Instruments Incorporated, Aug. 2016, sBAS601C. [Online]. Available: <https://www.ti.com/lit/ds/symlink/ads1220.pdf>

[115] *DS18B20: Programmable Resolution 1-Wire Digital Thermometer*, Maxim Integrated (Analog Devices), 2019. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/ds18b20.pdf>

[116] *ADXL343 3-Axis, ± 2 g/ ± 4 g/ ± 8 g/ ± 16 g Digital MEMS Accelerometer, ADXL343BCCZ-RL7*, Analog Devices Inc., 2024, rev. A. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/ADXL343.pdf>

[117] E. Systems, “Esp32-s2 series datasheet,” version 1.7, accessed: 2025-09-08. [Online]. Available: https://www.espressif.com/sites/default/files/documentation/esp32-s2_datasheet_en.pdf

- [118] *MCP7940N Battery-Backed I²C Real-Time Clock/Calendar with SRAM*, Microchip Technology Inc., October 2011, datasheet. [Online]. Available: <https://ww1.microchip.com/downloads/aemDocuments/documents/MPD/ProductDocuments/DataSheets/MCP7940N-Battery-Backed-I2C-RTCC-with-SRAM-20005010J.pdf>
- [119] E. Toyocom, “Sc-32s series: 32.768 khz smd quartz crystal unit,” accessed: 2025-09-08. [Online]. Available: https://www.mouser.co.za/datasheet/3/394/1/file_PRODUCT_MASTER_50812_GRAPHIC03.pdf
- [120] STMicroelectronics, “Guidelines for oscillator design on stm8af/al/s and stm32 mcus/mpus,” STMicroelectronics, Tech. Rep. AN2867 Rev. 23, January 2025. [Online]. Available: https://www.st.com/resource/en/application_note/cd00221665-oscillator-design-guide-for-stm8af-al-s-stm32-mcus-and-mpus-stmicroelectronics.pdf
- [121] S. Arar, “Assessing the effect of load capacitance on the frequency of a quartz crystal,” *All About Circuits*, June 2021, accessed: 17 Oct. 2025. [Online]. Available: <https://www.allaboutcircuits.com/technical-articles/assessing-effect-of-load-capacitance-on-frequency-of-quartz-crystal/>
- [122] *MCP2515 Stand-Alone CAN Controller with SPI Interface*, *MCP2515-I/ML*, Microchip Technology Inc., Jun. 2005, rev. J. [Online]. Available: <https://ww1.microchip.com/downloads/en/DeviceDoc/MCP2515-Stand-Alone-CAN-Controller-with-SPI-20001801J.pdf>
- [123] *TCAN3413 3.3-V CAN FD Transceiver with Flexible IO and Standby Mode*, *TCAN3413DRBR*, Texas Instruments Incorporated, Mar. 2023, sLLSFB5A. [Online]. Available: <https://www.ti.com/lit/ds/symlink/tcan3413.pdf>
- [124] MobileStore. (2024, May) Comparing mobile providers in south africa. Accessed: 2025-08-22. [Online]. Available: <https://mobilestore.gumtree.co.za/blog/view/comparing-mobile-providers-in-south-africa>
- [125] Kyocera AVX, *X9000984-4GDSMB External Hinged EU 4G/LTE Antenna*, Nov. 2023, datasheet. [Online]. Available: https://datasheets.kyocera-avx.com/ethertronics/AVX-E_X9000984.pdf
- [126] TE Connectivity / Linx Technologies, *ANT-GNRM-L125A Series L1/L2/L5 Magnetic Mount Active GNSS Antennas*, Oct. 2022, datasheet. [Online]. Available: https://www.te.com/commerce/DocumentDelivery/DDEController?Action=srchtrrv&DocNm=ANT-GNRM-L125A_Series&DocType=Data+Sheet&DocLang=English&DocFormat=pdf&PartCntxt=ANT-GNRM-L125A-3

- [127] L. Gough, “Experiment: microsd card power consumption & spi performance,” *Gough Lui Blog*, February 27 2021, accessed: 2025-08-21. [Online]. Available: <https://goughlui.com/2021/02/27/experiment-microsd-card-power-consumption-spi-performance/>
- [128] T. Instruments, “Lmzm23601 36-v, 1-a step-down dc/dc power module.” [Online]. Available: <https://www.ti.com/lit/ds/symlink/lmzm23601.pdf>
- [129] I. Analog Devices, “Lt3972imse#pbf 33v, 3.5a, 2.4mhz step-down switching regulator with $75\ \mu\text{A}$ quiescent current,” 2025, low Ripple Burst Mode® operation, 3.5A max output, $I_Q = 75\ \mu\text{A}$ at 12V_{in} to 3.3V_{out} , 15mV_{pp} noise, line regulation 0.01%/V. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/3972fa.pdf>
- [130] TDK-Lambda, “Ccg30-24-12s isolated 30 w dc-dc converter (9-36 v in, 12 v out).” [Online]. Available: https://product.tdk.com/system/files/dam/doc/product/power/switching-power/dc-dc-converter/specification/ccg15-24-xxs_spc.pdf
- [131] *TMR 1211 Isolated DC/DC Converter, 5 V Output, 2 W, SIP-8 Package*, Traco Power, 2021, tMR 2 Series Datasheet. [Online]. Available: <https://www.tracopower.com/products/tmr2.pdf>
- [132] *LT3060IDC-2.5 Low Noise, 100 mA Low Dropout Linear Regulator*, Analog Devices Inc., 2020, LT3060 Series Datasheet. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/lt3060.pdf>
- [133] *LTC1550L Low Noise, Switched Capacitor Regulated Voltage Inverter, Fixed -2.5 V Output*, Analog Devices Inc. (Linear Technology), 2013, LTC1550L/LTC1551L Datasheet. [Online]. Available: <https://www.analog.com/media/en/technical-documentation/data-sheets/15501lf.pdf>
- [134] *18024115401H SPI Digital Isolator 5000Vrms 4 Channel 100Mbps 150kV/μs CMTI*, Würth Elektronik eiSos GmbH & Co. KG, Jan. 2024, datasheet. [Online]. Available: <https://www.we-online.com/components/products/datasheet/18024115401H.pdf>