

Politecnico di Torino

Politecnico di Torino

Master's Degree in Electronic Engineering

A.Y.2025/2026

Application of V-Cycle to the development of a next-gen BJB board for sustainable mobility.

Supervisor: Guido Masera

Candidate: Battista Bettoni

Acknowledgements

To see the world, things dangerous to come to, to see behind walls, draw closer, to find each other, and to feel. That is the purpose of life.

Thanks to people whose actions inspire me every day, reminding me how real those words are.

Table of Contents

List of Figures VII

List of Tables IX

1 Introduction 1
- 1.1 Goal . . . 2
- 1.2 Thesis structure . . . 3

2 BMS overview 5
- 2.1 What is a BMS? . . . 5
- 2.2 BMS functions . . . 7
  - 2.2.1 Direct measurement . . . 8
  - 2.2.2 State estimation . . . 9
  - 2.2.3 Power and current limits . . . 9
  - 2.2.4 Safety monitoring . . . 10
- 2.3 State of the art on BMS technology . . . 10
  - 2.3.1 De-localizing reaction to the "edge" of the system . . . 10
  - 2.3.2 Electrochemical Impedance Spectroscopy (EIS) . . . 11
  - 2.3.3 Pyro-fuses . . . 12
- 2.4 Battery Junction Box: structure and state of the art . . . 14
  - 2.4.1 Functional definition and partitioning . . . 14
  - 2.4.2 Reference architectures . . . 14
  - 2.4.3 Comparison . . . 18

3 V-Cycle: a process, two norms 21
- 3.1 Purpose of V-Cycle and FuSa concept . . . 21
  - 3.1.1 The V-shape . . . 22
- 3.2 The Cycle . . . 23
  - 3.2.1 Comparing ISO 26262 and EN 50126 . . . 26

4 Defining HW architecture of the BJB (ML3) 29
4.1 End-of-line calibration . . . 30
4.2 HV positive voltage measurement . . . 30
4.3 HV bipolar voltage measurement . . . 34
4.4 Current sensing . . . 38
4.5 Port mapping on the MC33772C . . . 41
4.5.1 High-voltage chains . . . 41
4.5.2 Current chain . . . 42

5 Isolation measurement chain 45
5.1 State of the art on isolation measurement . . . 45
5.1.1 Bender iso175 . . . 45
5.1.2 NXP reference design . . . 46
5.2 Customer specifications and design premises . . . 47
5.3 Method and general design idea . . . 48
5.4 Model description and requirements . . . 49
5.4.1 Operating envelope and the role of the contactors . . . 50
5.4.2 Derived requirements . . . 51
5.4.3 How much uncertainty is “too much uncertainty”? . . . 51
5.5 Minimum required isolation . . . 52
5.6 Formal description . . . 53
5.6.1 Physical mapping . . . 53
5.6.2 System mapping: solving for the insulation . . . 55
5.6.3 Virtual mapping: removing the battery-voltage dependency . . . 56
5.7 Uncertainty of measurement . . . 58
5.7.1 Contributions to the budget . . . 58
5.7.2 Propagating the uncertainty . . . 60
5.7.3 Deterministic versus probabilistic interpretation . . . 62
5.8 Analysis and safety mechanisms . . . 62
5.8.1 Lower bounds on the readable voltage . . . 62
5.8.2 Uncertainty over the resistance space . . . 63
5.8.3 Isolation over the voltage space . . . 65
5.8.4 Isolation over the virtual space . . . 67
5.9 Summary . . . 69

6 Validation and results 70
6.1 Test environment . . . 70
6.2 Test campaign . . . 71
6.3 Results . . . 73
6.3.1 Test points and device baseline . . . 73
6.3.2 The safe operating area as specified . . . 74

6.3.3 The lower edge 74
6.3.4 The unbonded edge: SM4 74
6.3.5 Comparison 77
6.3.6 False alarms 77

7 Conclusions and future work 80

Bibliography 82

List of Figures

1.1 V-Cycle overall structure. 3

2.1 Centralized BMS topology. 6
2.2 Modular BMS topology. 6
2.3 Infineon's BMS general block diagram. 7
2.4 Project-specific Podium BMS architecture. 7
2.5 BMS main functions. 8
2.6 EIS equivalent circuit 12
2.7 Impedance diagram at different SoH of a li-ion battery. 13
2.8 EV current zones over time 13
2.9 RD772 reference design board 15
2.10 RDA777T2 reference design board 16
2.11 ADBMS2950 reference design board 17
2.12 IVT and BJB size comparison 18

3.1 Automotive V-Cycle overall structure. 21
3.2 Podium's V-Cycle 23
3.3 V-Cycle in ISO 26262 and EN 50126 26

4.1 Block diagram of the BJB. 30
4.2 Schematic of the HV positive measurement chain. 32
4.3 Schematic of the HV bipolar measurement chain. 35
4.4 Redundant current measurement chain 39
4.5 Current measurement channel of the MC33772C 44

5.1 Bender iso175 – Board for isolation measurement. 46
5.2 NXP – Isolation measurement reference design. 47
5.3 Isolation measurement model. 49
5.4 Image of the application over the assumed range 57
5.5 Relative uncertainty of the HV sensing chain 59
5.6 Uncertainty budget with 0.1% and 1% resistors 60
5.7 Maximum relative uncertainty over the resistance plane 64

5.8 Normalised sensitivity of V_1 65
5.9 Normalised sensitivity of V_2 65
5.10 Uncertainty of R_{isoP} over the resistance plane 65
5.11 Uncertainty of R_{isoN} over the resistance plane 65
5.12 Positive insulation estimate over the voltage plane 66
5.13 Negative insulation estimate over the voltage plane 66
5.14 Uncertainty of R_{isoP} over the voltage plane 66
5.15 Uncertainty of R_{isoN} over the voltage plane 66
5.16 Positive estimate minus its uncertainty 66
5.17 Negative estimate minus its uncertainty 66
5.18 Estimated $vvbat_1$ 67
5.19 Estimated $vvbat_2$ 67
5.20 Reachable insulation pairs 68
5.21 Safe operating area over the virtual plane 69

6.1 dSPACE SCALEXIO HIL system. 71
6.2 Test sequence and verdicts 72
6.3 Test points in the virtual plane 73
6.4 Lower edge against the exact threshold locus 75
6.5 Bounds on the unbonded edge 76
6.6 Miss rate and false-alarm rate 78

List of Tables

2.1 Processing architecture and integration class. 18
2.2 Measurement channels and conversion 19
2.3 Current-measurement path 19
2.4 Communication and environmental. 20

3.1 Podium's process against the two standards 28

4.1 Characteristics of the `GPIO/ANx` port 42
4.2 Characteristics of the `ISENSE` port and of ADC2 43

5.1 Bender iso175 specifications 46
5.2 RDA777T2 insulation measurement circuit 47
5.3 Boundary parameters of the ADC 59

6.1 Test plan 71
6.2 Scores of the criteria 77

Chapter 1

Introduction

The emphasis on sustainability in the society of the 21st century is growing to an adult and mature state, where the market is slowly adapting to offer real answers to the anxiety of a society more and more sensitive to preserving the beautiful diversity of our planet.

Even though recent events have momentarily reduced media coverage to Climate Change and mitigation techniques, this is a trend with an important critical mass, very difficult to stop, and for good reasons.

Few people in the transportation industry would deny that electrification is a great opportunity to bring new life to a mature and slow-moving market, an opportunity that European manufacturers have largely missed.

The pillars of electrification when it comes to the transportation are two: electric motor and battery, both elements are not exactly new, as the first electric motor is more than 150 years old (older than internal combustion engines) and some even consider Baghdad Battery (2000+ y.o) to be the oldest example of an electric battery.

The modern era of batteries began with the commercialization of lithium-ion cells in early 90s, their significantly higher energy density, lower weight, and cycle life enabled broad adoption across multiple applications, from cell phones to vehicles and more. Over the last decade, advances in cell chemistry, manufacturing processes, and battery management systems (BMS) have dramatically increased versatility while reducing costs. Developments such as nickel-rich cathodes, lithium iron phosphate (LFP) cells, and cell-to-pack integration have become major milestones in battery technology.

The improvements in BMS (Battery Management System) technology inside battery packs is drastically reducing safety concerns connected to improper use, from cell to pack perspective, while improvements in modelling of cells such as the transition to Electrochemical models instead of ECM (Equivalent Circuit Model), are allowing to extract more and more performance in terms of peak and continuous

discharge but also charge while maintaining high safety standards.

In order to ensure highest safety and functional standards, the development of products intended for automotive markets, must comply with extremely demanding standards, that precisely rule the methods and procedures of design of equipments, when it comes to both automotive and railway applications, the same procedures apply, in the so called **V-Cycle**.

If China is leading electrification in large scale production in the automotive industry, players of the Old Continent are bringing this "green sensibility" to niche markets like railway, aerospace and motorsport, unlocking new performances, challenges and possibilities.

It is precisely within this demanding and fast-evolving, highly specialized, landscape that specialized engineering companies play a crucial role, this is the environment in which Podium Advanced Engineering operates.

Podium Advanced Engineering is an Italian high-performance engineering company specializing in the design and development of cutting-edge solutions for motorsport and automotive applications. With a strong focus on innovation and technical excellence, Podium supports teams and manufacturers in the development of advanced electronic and mechanical systems, bridging the gap between prototype research and competitive racing environments, the company fuses at its core the passion for race, speed and performance to attention to nature, respect and innovation.

1.1 Goal

The present thesis guides the reader through the complete development process of the BJB (Battery Junction Box), a critical sub-system within the Battery Management System (BMS) developed internally by the Podium Advanced Engineering EE team.

The thesis follows a specific railway project, the most ambitious in the BMS team, aiming to comply with the stringent railway norm **EN 50126**, in the context of a wider project to realise a baseline for future evolutions of the "Podium BMS".

The work is structured according to the V-Cycle methodology, a systems engineering framework that organises the development process into two symmetrical phases. The descending left branch represents the decomposition phase, in which requirements and specifications are progressively broken down from the system level through subsystem and component levels, translating high-level functional needs into concrete design choices. The ascending right branch, conversely, represents the verification and validation phase, in which each level of implementation is tested and verified against the specifications defined in the corresponding descending stage, progressively reintegrating the validated components back into the broader

system.

Following this structure, the thesis opens with a state-of-the-art review of product development methodologies and BJB design practices, before moving through the design and physical realisation of a first board prototype and its subsequent validation campaign.

Once the bottom of the descending branch is reached, the work climbs back up through testing and validation activities, ultimately returning to the system level.

Here, the conclusions drawn from the complete loop inside the lower level of the cycle, are used to propose a revised set of system-level requirements for a next-generation BJB board, introducing new functionalities, effectively closing the V-Cycle back upon itself and laying the groundwork for future iterations.

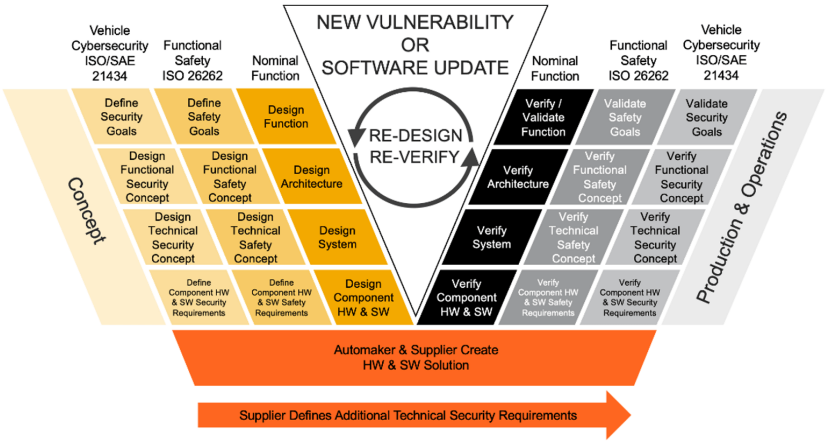


Figure 1.1: V-Cycle overall structure.

1.2 Thesis structure

Thesis is structured as follows:

- **Chapter 2** introduces State of the art on BMS technology and briefly presents Podium's BMS as well as some competitor solutions.
- **Chapter 3** introduces the V-Cycle and provides a brief literature review on the topic, analysing the differences among ISO 26262 [1] and EN 50126 [2].
- **Chapter 4** introduces choices of design of the BJB and derivation of HW requirements starting from System-level requirements, alongside with an introduction to the most relevant project phases and documents when it comes to FuSa.

- **Chapter 5** goes through design phase of isolation measurement chain of the BJB, analysed from a functional perspective.
- **Chapter 6** reports the test and validation setup and the results of the isolation measurement system.
- **Chapter 7** draws the conclusions of the work and proposes a revised set of system-level requirements for the next-generation BJB.

Chapter 2

BMS overview

This section aims to present the BMS as a system, explain its functions, working principles and to present a state of the art research in this field.

2.1 What is a BMS?

Battery management systems (BMSs) monitor, control, and protect battery packs from cell to pack level, ensuring efficient operation, enhancing performance and addressing the inconsistencies that lead to unstable performance and, in the worst cases, to early ageing and cell degradation.

Functional safety is critical as lithium-Ion batteries pose a significant safety hazard when operated outside their safe operating area, that is why, to meet demand of increased performances, monitoring and real time modelling of battery behaviour becomes more and more important.

BMS topologies are classified into two main categories:

- centralized
- modular

Their main difference lies in the physical location of the cell-sensing electronics:

In a **centralized BMS**, a single board measures every cell through sense wires, directly connected to it. It is cheap and compact but scales poorly and the wiring becomes fragile, noisy and messy as cell count grows, making it extremely cost-effective for smaller packs, but introducing substantial limitations as the amount of cells enlarges.

The **Modular BMSs** have small monitoring boards mounted to every cell or every few, usually clustering them in small groups, communicating back to the controller over a link. In a distributed topology, those boards called **CMUs** Cell

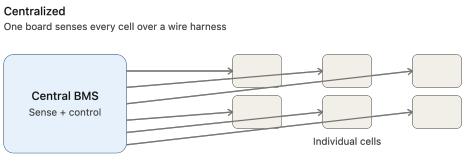


Figure 2.1: Centralized BMS topology.

Management Units sit close to their cells and connect to the controller called **BMU** Battery Management Unit, that acts as master, whereas the CMUs act as slaves of the system. This topology is now the backbone of most HV battery packs for traction, as it represents the state of the art for BMS tech and it is the one used by Podium as well. Modular topology is in fact proven to be the most efficient and

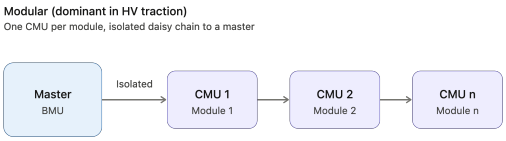


Figure 2.2: Modular BMS topology.

most scalable, being the reference for every good BMS design when it comes to HV batteries to be deployed in vehicles (RESS) Rechargeable Energy Storage Systems.

As said in the introduction, the BMS is responsible for all measurements from cell to pack level, where the latter function is assigned to a specific board, called the **BJB** Battery Junction Box, responsible for all HV (high-voltage) voltage measurements, as well as for measurement of current consumption and Coulomb Count (integrator of the current flow out of the battery). This board is also responsible for the isolation measurement from battery to chassis.

For certain applications it is necessary to add a **PDU** Power Delivery Unit or a **PIM** Power Input Module to power all LV boards.

A BMS holds countless inputs from sensors, but its only means of acting on the physical world is by driving the contactors: electromechanical switch used to establish and interrupt high-power electrical circuit, similar to a relay, they use a low-voltage, low-current control signal (usually through a PWM) to safely switch huge electrical loads (like ESS from the inverter/bus/charger).

The BMS is also responsible for reducing the inrush current when connecting and disconnecting the battery from its loads by making use of precharge circuitry, as well as communication to the charger when recharging the battery and to the **VCU**

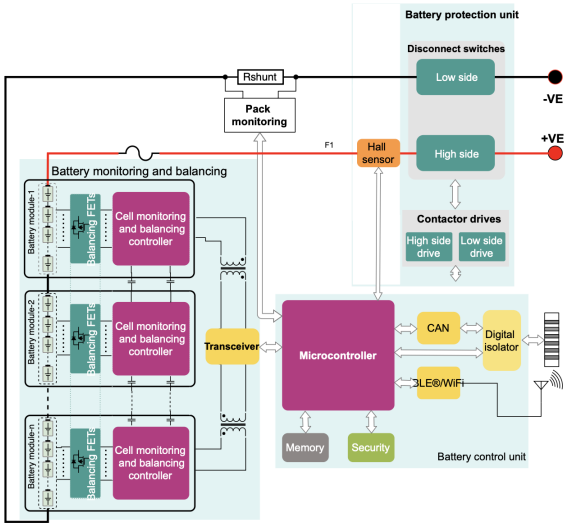


Figure 2.3: Infineon's BMS general block diagram.

Vehicle Control Unit, to communicate states of the battery pack. Usually in order to reduce the damage caused by over-current events BMS also implements a **fuse**, while last trends in automotive prefer to have pyro-fuses triggerable by external signals or even from BJB board using specific ASICs capable of an event-to-trigger time of few milliseconds.

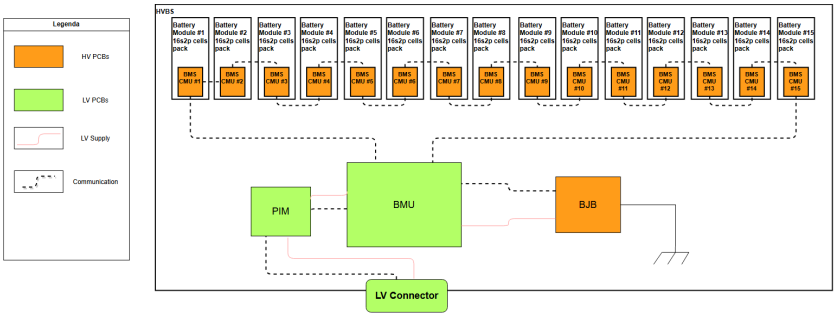


Figure 2.4: Project-specific Podium BMS architecture.

2.2 BMS functions

The most obvious functions of a BMS are the safety of the RESS, direct measurements of main physical quantities such as of voltage, current, temperature, and management of charge and discharge phases of the battery. Beyond that a series of estimators are deployed to derive more complex estimates, such as SoC, SoH,

SoP, SoE etc. Each of those parameters alone would be worth a full thesis, as the real innovation in BMS technology lays in the accuracy of estimation of those parameters, that allows the battery to run closer and closer to its chemical limit while retaining highest possible safety.

- Direct measurements
- State Estimation
- Power and current limits
- Safety monitoring

These macro areas can be further divided into specific functions, as the next figure shows:

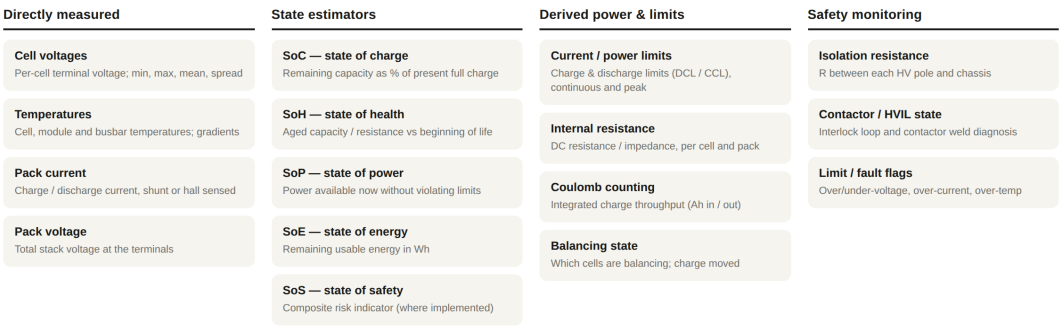


Figure 2.5: BMS main functions.

2.2.1 Direct measurement

Direct measurements are the foundation everything else is built on: no estimator is better than the raw quantities feeding it. They are conducted at two levels, cell and pack, and the modular architecture splits them across two different boards for reasons that are partly physical and partly safety related.

At cell level the job belongs to the CMUs. Each CMU sits on its module and measures the voltage of every series cell it supervises together with a set of temperatures distributed across the module. The reason this is done locally, and not by running a harness of sense wires back to a central board, is the same reason the modular topology exists at all: keeping the analog front-end next to the cells turns a mass of long, noise-prone analog leads into a single digital link, and that link is the only thing that has to cross the high-voltage isolation barrier. In

Podium's BMS, as in most HV designs, this front-end is not discrete circuitry but a dedicated battery cell controller ASIC. The reference is the NXP family (MC33774) [3], which performs the differential ADC conversions of cell voltages, current and temperature on-chip and ships the result upward over an isolated daisy chain. The communication uses what NXP calls the TPL (Transformer Physical Layer): a transformer-isolated two-wire bus that chains the cell controllers together and lets a single isolated link serve an entire stack of modules.

At pack level the job belongs to the BJB. Here the quantities of interest are the total HV bus voltage, the pack current, and the isolation resistance between each HV pole and chassis (the subject of Chapter 5). Current is read across a shunt – typically around $100\,\mu\Omega$ – and integrated over time to give the Coulomb count, that is the backbone of every SoC estimate.

2.2.2 State estimation

State estimation is where the intelligence of the BMS lives, all inside the central node of the system: the BMU. The CMUs and the BJB produce measurements; the BMU collects them over the isolated bus and turns them into the derived quantities, to update to the rest of the vehicle main states such as SoC, SoH, SoP, SoE. Estimation needs a global view of the whole pack, non-trivial compute, a model of the cell, specific both to the cell and to the configuration and pack and often a stored history, concentrating it in the master controller.

The CMUs stay close to the silicon and the physics; the BMU stays close to the model. CMUs and BJB can be considered the boards closest to the physics, as their main purpose is to provide measurements, so they are very much HW related, on the other side the BMU operates at an estimation level, leaning more on SW and modelling results.

This division also sets up a latency consideration that turns out to matter a great deal for safety, and it is worth being concrete about it.

2.2.3 Power and current limits

The power and current limits are situated between pure estimation and outright protection. From SoC, temperature and the estimated internal resistance, the BMS computes, cycle by cycle, how much current the pack may source or sink without pushing any cell past a voltage, current or thermal boundary – the charge and discharge current limits, in both a continuous and a short-duration events. These are produced by the BMU, because they depend on the same global, model-based view as the estimators, but they are enforced collaboratively: the BMU publishes the limits while the local protections on the CMU and BJB stand. This collaboration is one of the most active areas of research in the field, as anticipated

in the introduction: it allows the cells to be run closer and closer to their real limits, recomputed in real time.

2.2.4 Safety monitoring

The last function is the one that frames all the others, and it is genuinely distributed: there is no single board that "does safety". The CMU watches its cells for over- and under-voltage and over-temperature and can act on them locally; the BJB watches the pack for over-current and loss of isolation and can disconnect through the contactors or fire the pyro-fuse; the BMU supervises the whole, cross-checks the slaves, manages the contactor sequencing and the high-voltage interlock loop, and talks to the VCU. The contactors are the only means the entire system has of acting on the physical world (apart from pyros) – everything else is sensing and deciding – which is why their command path, and the diagnostics to confirm they are opened or closed and did not weld, is treated as safety-critical.

What the BMS does, badly understates the real engineering problem, because the hard part of this field is not making the functions work. It is proving that they cannot fail in a way that hurts someone – and that proof is what the rest of the engineering effort goes to. That is what *chapter 3* is about: FuSa, V-Cycle and good design practice.

2.3 State of the art on BMS technology

2.3.1 De-localizing reaction to the "edge" of the system

Routing a measurement to the BMU costs a considerable delay: the full round trip adds the conversion time, the response travelling back up the chain, and then the time for the BMU to actually react.

And that is not even considering the polling rate of the BMU towards its slaves, worth at least 10 ms. It is evident that for reflex-fast tasks, where the safety of the battery, vehicle, and potentially occupants of the vehicle is critical, the ignition of a pyro-fuse cannot accept this kind of delay, requiring different strategies to be applied.

The architecture answers this in two ways, and both come down to not making the signal travel further than it has to. The first is to push the time-critical safety reactions down into the cell-controller silicon itself.

The second is to handle the most violent pack-level events on the BJB with dedicated hardware rather than software, so that the decision never has to make the round trip. The clearest example is over-current protection: where a classical melting fuse reacts on its own thermal time constant, recent automotive practice prefers a pyro-fuse, a pyrotechnic disconnect that can be fired by an external trigger

or directly by a dedicated ASIC on the BJB, reaching event-to-trigger times of a few milliseconds. A short circuit can drive the pack current to destructive levels far faster than a daisy-chained reading could be collected, evaluated by the master and acted upon, so the detection and the trigger are kept together, on the board where the current actually flows.

The general principle is the one that organizes the whole architecture: slower model-based reasoning goes to the BMU, fast reflex-like protection stays on the board that owns the quantity. The CMU protects single cells, the BJB protects the pack, and the BMU orchestrates all the system.

2.3.2 Electrochemical Impedance Spectroscopy (EIS)

Electrochemical Impedance Spectroscopy (EIS) is a non-destructive technique that involves applying a sinusoidal excitation signal to measure the response of an electrochemical system.

The technique allows to evaluate diffusion processes, mass transport, and charge transfer, providing essential insights to estimate the SoH (state of health).

Instead of watching how the terminal voltage responds to the load over time, EIS injects a small sinusoidal excitation – a current or a voltage – at a known frequency, measures the amplitude and phase of the response, and from the two computes the complex impedance $Z(\omega)$ at that frequency.

Sweeping frequency across a wide range from the order of mHz up to several kHz, results in the impedance spectrum of the cell. Because different physical processes inside the cell live on different timescales, they separate out in frequency: the fast, almost purely resistive conduction of the electrolyte and current collectors dominates at high frequency; the charge transfer across the electrode interfaces, with its associated double-layer capacitance, shows up as a semicircle in the mid band; and the slow diffusion of ions into the electrode material appears at low frequency as the characteristic 45° Warburg tail, typical of the lithium-ion diffusion process of the electrode.

Principle

The idea is to evolve classical "resistor representation" opening up to a frequency-dependent impedance. Instead of watching how the terminal voltage responds to the load over time, EIS injects a small mono-sinusoidal, multi-sinusoidal or chirp excitation at a known frequency, measures the amplitude and phase of the response, and from the two computes the complex impedance $Z(\omega)$ at that frequency.

The standard way to read all this is the **Nyquist plot**, which represents the negative imaginary part of the impedance against the real part, one point per frequency. On a lithium-ion cell the trace runs, left to right and high frequency

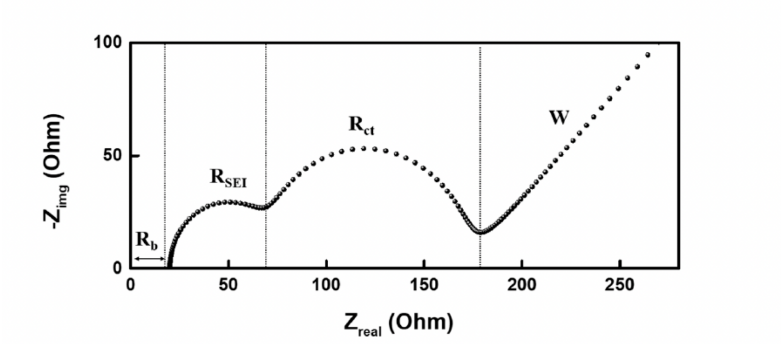


Figure 2.6: EIS – simplified equivalent circuit and impedance diagram.

to low, from a small inductive region above roughly 1 kHz, through one or more depressed semicircles between about 1 kHz and 0.5 Hz, and out along the Warburg line below that, the typical behaviour can be seen in Figure 2.6. The **Bode plot**, magnitude and phase against frequency, carries the same information in a form that makes the resistive, capacitive and diffusive regimes easy to separate by eye.

Implications of SoH

EIS parameters track ageing and incipient failure far more directly than a terminal-voltage measurement can. A concrete result reported in the literature on LFP cells is that the charge-transfer resistance and the Warburg coefficient stay essentially flat across the top of the cell's life, in the 95–100% SoH band, and then start to climb sharply once SoH falls below 90%, where capacity fade accelerates. An impedance feature, provides an early and quantitative estimate that a capacity-counting scheme would not provide until the damage is already done. Frequency resolution is what allows EIS detect temperature gradients, distinguish resistance growth from capacity fade, and flag micro-short circuits that a time-domain monitor would not distinguish. There is even a neat practical by-product: because the high-frequency impedance is strongly temperature-dependent, EIS can be used to estimate cell temperature internally, reducing the number of physical temperature sensors a pack needs [4], but this is a far from market-ready tech.

2.3.3 Pyro-fuses

The pyro-fuse, also called pyrotechnic safety switch, was introduced briefly above as the BJB's answer to fast over-current events. This has become ubiquitous in HV traction packs as comparison with the conventional fuse it supplements is a clean illustration of the reflex-vs-reasoning principle that rules the whole BMS architecture. In figure 2.8 it is possible to see the zones of current over time that

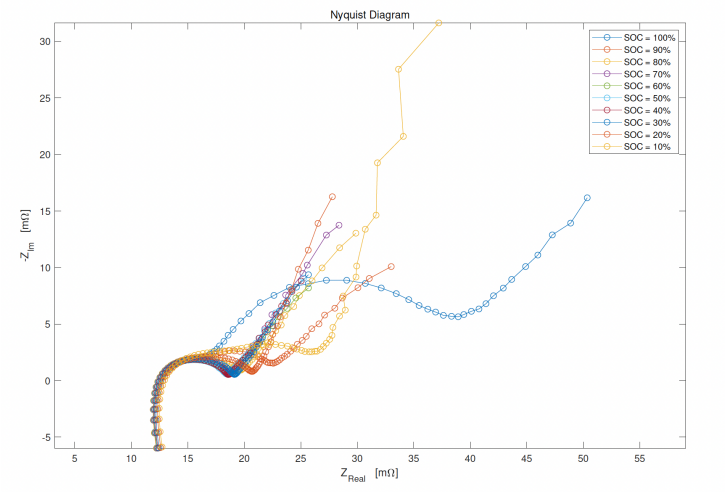


Figure 2.7: Impedance diagram at different SoH of a li-ion battery.

characterise an EV (electric vehicle).

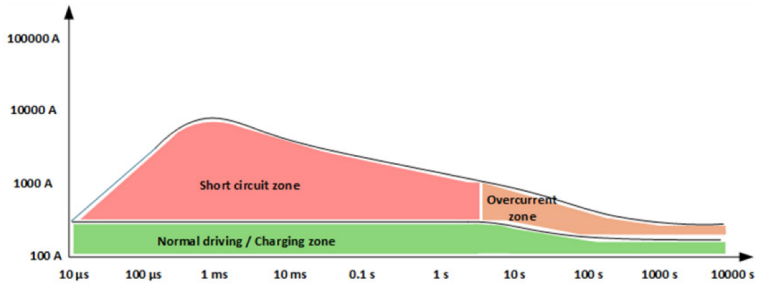


Figure 2.8: Basic example of characterising EV current zones over time [5].

Working principle

Where a melting fuse waits for the fault current to heat and sever a conductor, the pyro-fuse severs the conductor on command. Inside it sits a small pyrotechnic charge – the same family of technology as an airbag inflator with an igniter, typically a bridgewire, driven by a control signal.

In this way the triggering of the fuse is no longer bound to the actual current flowing the fuse itself, adding dof to the curve $I\,x\,t$ as it is possible to retrieve a higher peak current without triggering its ignition, while still limiting the curve where it matters most, abandoning the more or less linear relation among current and time to trigger the fuse and fitting its behaviour to actual needs of the pack.

2.4 Battery Junction Box: structure and state of the art

2.4.1 Functional definition and partitioning

The BJB is the subsystem that concentrates the pack-level power path and the pack-level safety functions. It is electrically distinct from the cell-monitoring chain: where the CMUs operate at module potential and resolve per-cell quantities, the BJB operates at full pack potential and resolves quantities defined across the pack terminals. Its function set is bounded and recurrent across designs:

1. **Current measurement** across a shunt in the main path, integrated to a Coulomb count.
2. **High-voltage measurement** of the bus and of the nodes across the contactors and fuses, for contactor-state and fuse-state diagnosis.
3. **Isolation-resistance measurement** between each HV pole and chassis.
4. **Temperature measurement** of the shunt and of the pre-charge resistor.
5. **Actuation**: pyrotechnic-switch (pyro-fuse) drive, contactor command, and the crash-signal / event inputs that gate them.
6. **Isolated communication** of the above to the BMU.

The BJB crosses the HV/LV galvanic barrier, this means its controller sits in the HV domain and its supply and communication crosses an isolation barrier realised as an isolated DC-DC converter and a transformer/capacitor-coupled isolated serial link [6].

Furthermore ASIL-D requirement for certain measurements forces *redundant* design: the quantities are acquired on two independent channels and, in the highest-integrity designs, processed by two independent signal paths. If current gen silicon uses double chips to implement redundancy, next gen ones are integrating at chip level redundant architectures in the same package, reducing space occupation and enhancing consistency as the chips now sit at same temperature.

2.4.2 Reference architectures

NXP UM11728 – dual MC33772C

The RD772 (UM11728) is the previous-generation NXP BJB reference design. It is not a specific BJB controller but a repurposed cell controller: it carries two MC33772C [7] six-channel cell-controller AFEs, each of which exposes a single

coulomb-counting current channel (CT1) and five general-purpose analog inputs multiplexed across the HV and diagnostic signals.

The two AFEs share ground and supply and are linked by a non-isolated TPL segment; only AFE1 crosses the isolation barrier to the BMU, through a 1:1 transformer isolating the TPL that runs up to $2\,\mathrm{Mbit\,s^{-1}}$.

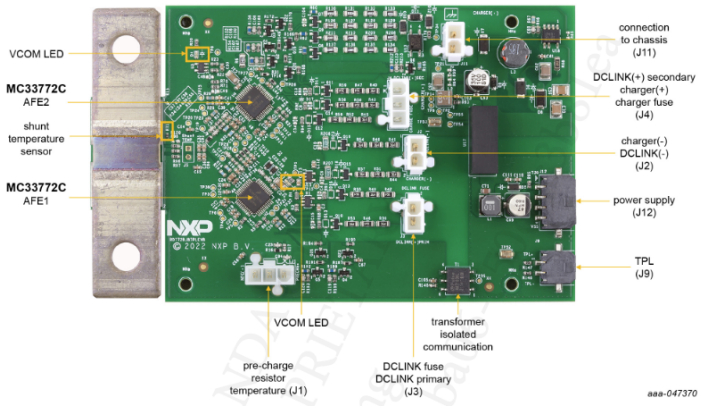


Figure 2.9: RD772 (UM11728) – NXP reference design based on two MC33772C.

The chip implements HV sensing through a voltage divider scaled to 4.85 V input, five positive inputs to 500 V and two bipolar inputs to ±500 V; a single 100 µΩ shunt to ±1,500 A, redundantly sensed by both AFEs; two NTC channels (shunt and pre-charge resistor); and the unbalanced two-/three-measurement isolation circuit, opto-isolated from the board. The acquisition is slow by construction: a per-input low-pass network gives 530 Hz to 600 Hz cut-offs, and a GPIO-enabled divider requires a blanking time of $\geq 2\,\mathrm{ms}$ before each conversion is valid. The key limitation, stated in the manual itself, is that the on-board shunt connection is not redundant as routed on the reference board; the redundant 4-wire connection must be added in the final application.

NXP UM12056 – single MC33777A

The RDA777T2 (UM12056) showcases the MC33777A [8], NXP's first single-chip BJB monitor and the direct successor to the dual-MC33772C.

The chip offers **dual-ADC-per-channel** front end: each current channel and carries a high-precision path (24-bit, signals to 1 kHz) and an independent high-speed path (16-bit, bandwidth to 125 kHz). The slow path serves coulomb counting and SoC; the fast path serves over-current detection and fuse emulation. Implementing redundancy principle that the RD772 implemented with two whole AFEs.

Board-level measurement capability from the chip is fully used: four current channels (doubling previous UM11728, and having ±300 mV) of range, eight positive HV inputs 0 V to 1,000 V and two bipolar inputs ±1,000 V, with positive/secondary pairs able to monitor the same node redundantly, as well as two NTC temperature channels to monitor shunt temperature.

The board implements new features already disclosed in previous chapter, such as isolated crash input; and *two independent pyrotechnic-switch outputs* each with its own energy-reservoir capacitor. The qualitative leap over RD772 is in the on-chip

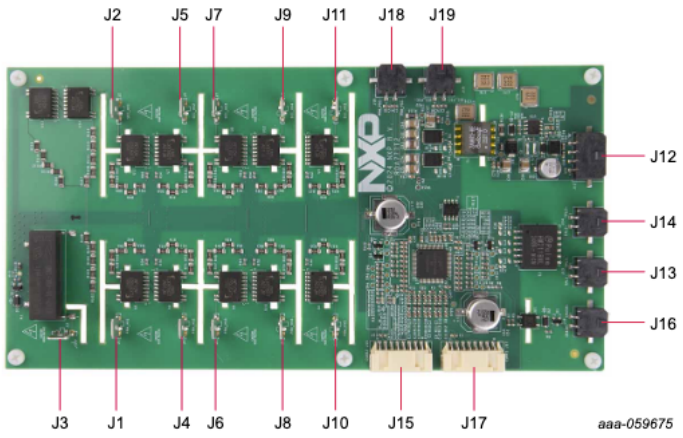


Figure 2.10: RDA777T2 (UM12056) – NXP reference design based on the MC33777A.

event engine: configurable logical combination of debounce or internal and external events (over-current, fuse emulation, over-voltage, crash) triggering reactions (pyro drive, MCU wake, GPIO) *without* a microcontroller in the loop, with ASIL D reached through redundant measurement and redundant signal processing internal to the device.

NXP states that integration allows an up-to-80% component-count reduction versus the discrete BJB-monitoring approach, while really optimistic that is for sure a simplification over their previous iteration. Communication is an isolated daisy chain (TPL3) or, alternatively, SPI.

Analog Devices ADBMS2950

AD's entry in the BJB/pack-monitor class is the ADBMS2950 (and the SPI-only ADBMS2952; the ADBMS2951 is a link monitor). It is a dedicated pack monitor: it measures pack current across a low-offset shunt with coulomb counting, measures pack voltage synchronously, and exposes 12 buffered high-impedance inputs for pack voltage, temperatures, HV-link voltages, chassis isolation, and contactor/fuse

state supervision [9, 10]. The approach of Analog Devices is to have less specialized and more versatile GPIOs compared to NXP.

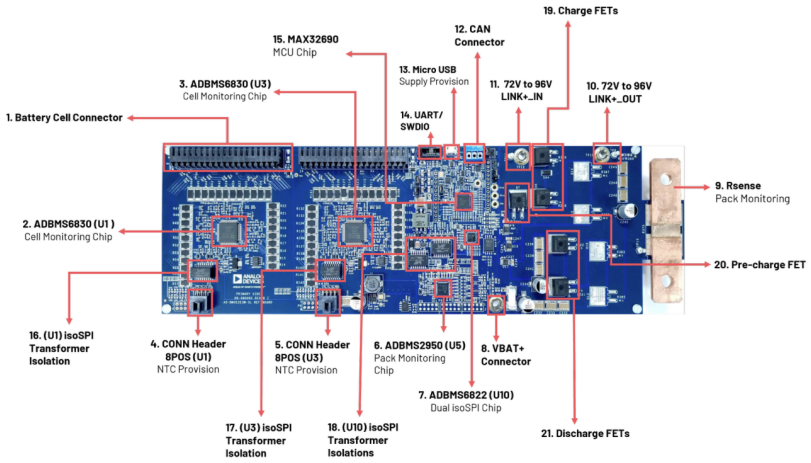


Figure 2.11: ADBMS2950 – Analog Devices reference design board.

Its over-current architecture is redundant to comply with ASIL-D requirements: fast over-current ADCs with triple-redundant digital threshold comparators reporting on dedicated alert line, plus PWM output options [9].

Actuation is handled through six GPO (open-drain or push-pull, dual-threshold read-back) able to drive HV transistors that disconnect the external dividers, and four GPIOs configurable as an I^2C or SPI controller.

Communication is built over isoSPI at $2\,\mathrm{Mbit\,s^{-1}}$, capacitor or transformer-coupled, with a daisy-chain option and a 4-wire SPI alternative; a second isoSPI port and a reversible-isoSPI mode provides a redundant communication path.

The ADBMS295x is developed for ISO 26262 and ASIL-D compliance [9], but architecturally the part differs from the NXP BJB monitor in two ways: it has no integrated pyro driver, (disconnect is delegated to external transistors via the GPOs), and it daisy-chains natively with ADI's ADBMS683x cell monitors over the same isoSPI fabric, making the pack monitor and the cell monitors one homogeneous communication domain [10].

Isabellenhütte IVT-S and IVT3-Pro shunt sensor modules

The IVT family is categorically different from the controller ICs above: it is a calibrated shunt-sensor module and a "finished" product, it enters this section as the company is entering the BJB market, as emerged in a meeting few months back with some commercials visiting Podium.

This is a shelf product very much used in motor-sport applications, very much appreciated for its extremely compact size and all-in-one solution, extremely easy

to integrate, thanks to its mounting point directly over the shunt current bar.

The IVT packs a shunt, a 16-bit ADC multiplexed over three or six HV channels (up to 1,000 V) and five further GPIOs, a microcontroller, temperature compensation and a galvanically isolated bus into a single plug-and-play unit, placed in the negative path of the pack. Compared to previous products from the same brand, it complies with ASIL for current measurement [11]. Communication is isolated

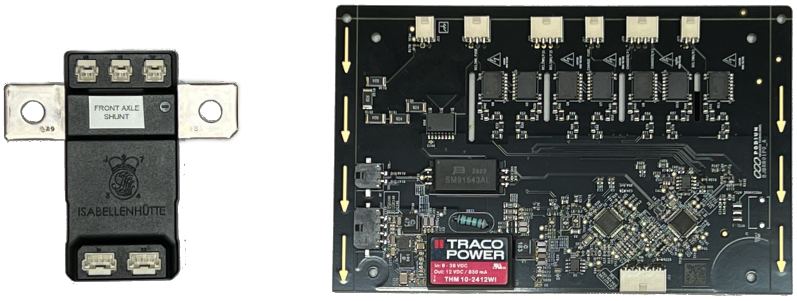


Figure 2.12: IVT size difference compared to Podium's current BJB.

CAN 2.0. The IVT3-Pro raises range to ±4,000 A from ±2,500 A and ±1,250 V from ±1,000 V, supports six, from 3, HV channels plus insulation-resistance monitoring, and is developed to ISO 26262 with an ASIL rating (up to ASIL-C), so it is not certified as ASIL-D compliant like NXP and AD counterparts.

2.4.3 Comparison

The following table summarises the main differences among the BJBs presented in the chapter.

It's worth noting that chips implementing EIS are now in the early adoption phase on the market, but for the scope of this thesis it was not possible to get one to work on, so it is not shown in the following table, as none of the presented solutions implements it:

Table 2.1: Processing architecture and integration class.

	MC33772C	MC33777A	ADBMS2950	IVT-S / IVT3-Pro
On-chip MCU	No (host computes)	event engine, No MCU	no (host computes)	integrated MCU
Redundancy	2 AFEs (board)	dual signal path (on chip)	triple OC comparators	incomplete (2 ADCs)
ASIL	ASIL-D	ASIL-D	ASIL-D	up to ASIL-C
Pyro driver	No	2× integrated	no	no
Package / form	2×LQFP (board)	LQFP64-EP	48-QFN	module

Three technical trends separate the generations rather than the vendors.

Table 2.2: Measurement channels and conversion. Figures from device datasheets / reference-design manuals.

	RD772	MC33777A	ADBMS2950	IVT-S
HV inputs	5 pos. (500 V) + 2 bip. (±500 V)	8 pos. (1,000 V) + 2 bip. (±1,000 V)	12 buffered HV/aux	up to 3 (1,000 V)
Isolation meas.	yes (unbalanced 2–3 meas.)	yes	yes (chassis input)	IVT3-Pro only
Temp. channels	2 (NTC)	2 (NTC)	via buffered inputs	internal + ext.
ADC (voltage)	16-bit, 154 µV/LSB	16-bit, 154 µV/LSB	buffered, synchronous	16-bit

Table 2.3: Current-measurement path. The dual-ADC entries distinguish the precision (coulomb-counting) path from the fast (over-current) path.

	MC33772C	MC33777A	ADBMS2950	IVT-S / IVT3-Pro
Method	shunt, 2 AFEs	shunt, 4 ch.	shunt, low-offset	shunt (integrated)
Precision ADC	16-bit	24-bit slow, 16-bit slow	lossless coulomb	16-bit fast
Fast ADC	—	16-bit, 125 kHz	fast OC ADC (3×)	Multiplexed ×3/×8
Range (board)	±1,500 A (100 µΩ)	ext. shunt (±300 mV)	ext. shunt	±100 A to ±2,500 A (IVT3-Pro: ±4,000 A)
Accuracy	ASIL-D class	ASIL-D class	low-offset	0.4 % rdg (IVT3: 0.1 %)
Coulomb count	yes (MC33772C)	yes	yes (lossless)	yes (Ah counter)

First trend is *the migration of redundancy from board to die*, next-gen boards, namely MC33777A and ADBMS2950, are implementing redundancy, discarding the double AFE approach in name of an internalized approach [9].

Second the *split of the current measurement into two ADCs of different character* a slow, high-resolution path optimised for coulomb-counting accuracy and a fast, lower-resolution path optimised for over-current latency, both for MC33777A and ADBMS2950. This directly addresses the latency argument of Section 2.3.1: the disconnect decision is taken on the fast path, locally, while the BMU still receives the accurate slow-path value for estimation.

Third, the *pyro drivers on the board*: NXP integrates the pyro drivers and an autonomous event engine on-chip, eliminating the microcontroller from the disconnect loop; AD keeps actuation external (GPO-driven transistors) but unifies the pack monitor and the cell monitors on one isoSPI fabric.

Finally last trend would be the introduction of EIS (electrochemical impedance spectroscopy), not analysed in this work, but certainly relevant in SoA discussion on pack-level monitoring for BMSs.

The thesis focuses on the adaptation of the current-generation Podium BJB board to the railway project. As it will be discussed by the end of the paper, an evolution to the MC33777-class chip is under discussion, while a new chip, the BMA8420 [12], is defining the next-generation reference for BJB design.

Table 2.4: Communication and environmental.

	MC33772C	MC33777A	ADBMS2950	IVT-S
Host link	isolated TPL	isolated TPL / SPI	isoSPI / 4-wire SPI	isolated CAN 2.0
Bit rate	$2\,\mathrm{Mbit\,s^{-1}}$	$2\,\mathrm{Mbit\,s^{-1}}$	$2\,\mathrm{Mbit\,s^{-1}}$	$500\,\mathrm{kbit\,s^{-1}}$ (typ.)
Daisy chain	yes (TPL)	yes (TPL)	yes (+ ADBMS683x)	yes (2nd CAN port)
Temp. range	automotive	$-40\,^{\circ}\mathrm{C}$ to $125\,^{\circ}\mathrm{C}$	AEC-Q100	$-40\,^{\circ}\mathrm{C}$ to $105\,^{\circ}\mathrm{C}$
Isolation	1.5 kV (DC-DC)	barrier (TPL3)	cap/transf. isoSPI	1 kV (pk 6 kV)

Chapter 3

V-Cycle: a process, two norms

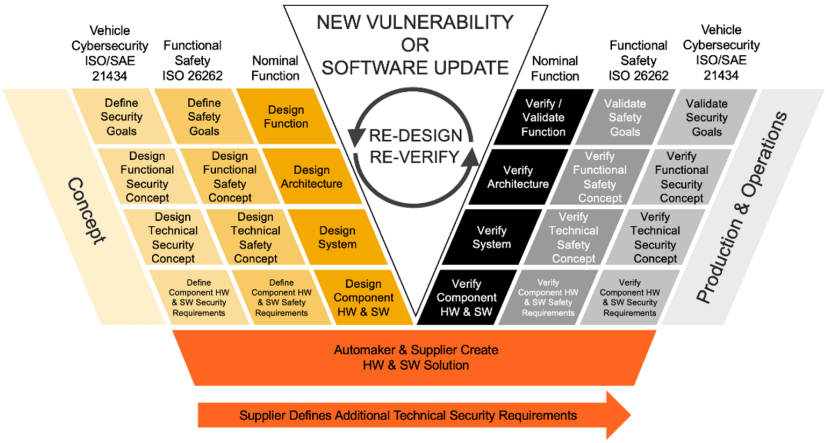


Figure 3.1: Automotive V-Cycle overall structure.

3.1 Purpose of V-Cycle and FuSa concept

While the primary objective of any engineering system is to perform its intended function, correct operation alone is not sufficient in safety-critical applications.

From an abstract perspective, a system that "works" defines a mapping from nominal operating conditions to the expected outputs, this part of the development flow is identified in the image 3.1 with the inner columns of *Nominal Function*.

Functional Safety (FuSa) does not change the mapping, instead, it augments it with a complete set of traceable relationships that explain why the mapping

exists, under which assumptions it remains valid, and how it behaves when those assumptions no longer hold.

The difference between a concept and a product lies in the two columns beside the Nominal Function one: there is an equally important body of work gathered under the term "Functional Safety". It is important to note that the part of cyber is not covered by this thesis but its presence inside the cycle underlines how important nowadays this is becoming.

Consequently, activities such as requirement traceability, hazard analysis, documentation, verification, and validation are essential elements of the engineering process. Standards such as ISO 26262 [1] for the automotive sector and EN 50126 [2] for railway applications provide frameworks that allow safety-related functions to be systematically specified, implemented, and demonstrated.

So V-Cycle can be seen as a development cycle defined per-industry that fits the FuSa concepts to the development process of a product, to give development guidelines. What is interesting is that the standard never touches technical specifications or choices; it only describes the full process to be followed to achieve high functional safety products.

This chapter presents how the V-Cycle is structured in Podium's method.

3.1.1 The V-shape

The V-Cycle is a visual/metaphoric representation of the activities of a development so that nothing is built without a reason recorded in advance and where previous steps are constantly called into question and eventually reworked if they do not stand up to scrutiny.

Reference norms are ISO 26262 for automotive (V structure is in figure 3.1) and EN 50126 for railway applications. The two norms share intents and main structure, broken down to their reference industry.

Just as a reminder, automotive and railway industries are widely different: first of all in the kind of customer, but also the type of product, their scale of production, cost and metrics of interest.

Two norms, one concept

Consider what a railway is, as an object of safety engineering. It is a small number of very large, very long-lived systems, built once and then operated for decades under continuous, professional, centrally organised maintenance. EN 50126 is built around this product category.

Cars on the other side are the opposite object: a mass-produced product, made in big numbers, owned and operated privately by an untrained individual, with no central maintenance authority and an operational life that, from the manufacturer's

standpoint, effectively ends when the warranty does. This is a faster, more lively market, where agility is necessary and is reflected in a more streamlined standard than the railway one: ISO 26262.

3.2 The Cycle

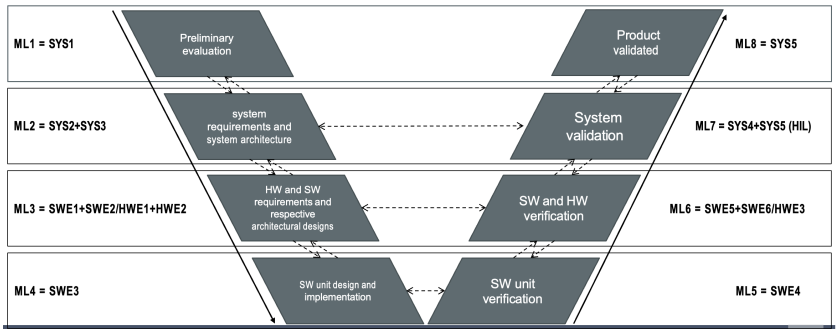


Figure 3.2: V-Cycle as derived by Podium, with the maturity levels.

The cycle is broken down in the descending and ascending phases, where each level of the descending has a correspondent ascending phase.

Podium internally divides the Cycle down in phases identified by a *Maturity Levels* ML, each maturity level represents a milestone associated with specific engineering activities, design deliverables and verification objectives. The process maintains full traceability between requirements, implementation and testing, ensuring compliance with Functional Safety principles regardless of the target application.

This subdivision is the keystone to the structure of this thesis.

ML1 – Preliminary Evaluation

The first maturity level corresponds to the initial feasibility assessment of the project. *Customer Requirements* are collected, the intended application is analysed and the overall project scope is defined.

Preliminary assumptions regarding the operating environment, system boundaries and technical feasibility are established. Although detailed requirements are not yet available, this phase provides the foundation for the subsequent development activities.

Within ISO 26262, this phase mainly corresponds to the *Concept Phase*, including the *Item Definition*, *Hazard Analysis* and *Risk Assessment (HARA)*, and the Functional Safety Concept. In EN 50126, the equivalent activities are the Concept phase together with the definition of the operational context and system boundaries.

ML2 – System Requirements and System Architecture

Once the project feasibility has been confirmed, the overall *System Requirements* are specified and organised into a coherent system architecture.

Functional, performance and safety requirements are collected and allocated to the various subsystems, while interfaces between hardware and software are identified. Verification planning also begins during this phase to ensure complete traceability throughout development.

This maturity level corresponds to the Product Development at the System Level described in ISO 26262 (Part 4), including the Technical Safety Concept and the definition of the System Architecture. In EN 50126, it is closely related to the specification of system requirements and the architectural allocation of system functions.

ML3 – Hardware and Software Requirements

The system architecture is progressively refined by decomposing the system-level requirements into dedicated hardware and software requirements. Separate hardware and software architectures are developed while maintaining complete traceability to the originating system requirements. At this stage, the functional allocation between electronic hardware and embedded software becomes explicit.

According to ISO 26262, this maturity level corresponds to the Hardware Development (Part 5) and Software Development (Part 6) activities concerning safety requirements specification and architectural design. Within EN 50126, this phase belongs to the system design and requirements allocation process, while implementation details are further addressed by companion standards such as EN 50128 [13] and EN 50129 [14].

ML4 – Software Unit Design and Implementation

The lowest point of the V-cycle represents the implementation phase. Individual software modules are designed, implemented and reviewed according to the requirements established during the previous maturity level. Although the Podium workflow emphasises software implementation at this stage, hardware development typically progresses in parallel through schematic design, component selection and PCB development.

This phase corresponds to the Software Unit Design and Software Implementation activities defined in ISO 26262. Within EN 50126, it belongs to the general design and implementation phase.

ML5 – Software Unit Verification

Following implementation, each software unit is verified independently to demonstrate compliance with its detailed specification. Typical verification activities include unit testing, static analysis, code inspection and software reviews. The objective is to detect implementation defects before integrating individual software modules into larger subsystems.

This maturity level directly corresponds to Software Unit Verification in ISO 26262, while EN 50126 considers it part of the general verification activities associated with software implementation.

ML6 – Hardware and Software Verification

After unit verification, hardware and software components are progressively integrated and verified. Particular attention is given to subsystem interfaces, communication protocols, timing behaviour and diagnostic functions. Hardware verification includes laboratory characterisation and functional testing, whereas software integration confirms the correct interaction among software modules.

Within ISO 26262, this phase corresponds to Hardware Integration and Verification together with Software Integration Testing and Embedded Software Verification. EN 50126 addresses these activities during the integration and verification stages of the system life-cycle.

ML7 – System Validation

At this stage the complete system is evaluated against the original system requirements. Within Podium, this activity is primarily performed through Hardware-in-the-Loop (HIL) testing, allowing the embedded controller to operate under realistic simulated operating conditions. Functional behaviour, diagnostic mechanisms and fault responses are assessed to verify that the integrated system satisfies both functional and safety requirements.

The equivalent activities in ISO 26262 include Item Integration and Safety Validation, while EN 50126 refers to this phase as System Validation.

ML8 – Product Validation

The final maturity level represents the validation of the complete product within its intended operational environment. Customer requirements, regulatory constraints and Functional Safety objectives are evaluated before product release. Unlike verification, which demonstrates compliance with technical specifications, validation confirms that the developed product fulfills its intended purpose under realistic operating conditions.

Within ISO 26262, this phase corresponds to the final Safety Validation and preparation for production. In EN 50126, it is associated with System Acceptance and the transition towards operation and maintenance.

3.2.1 Comparing ISO 26262 and EN 50126

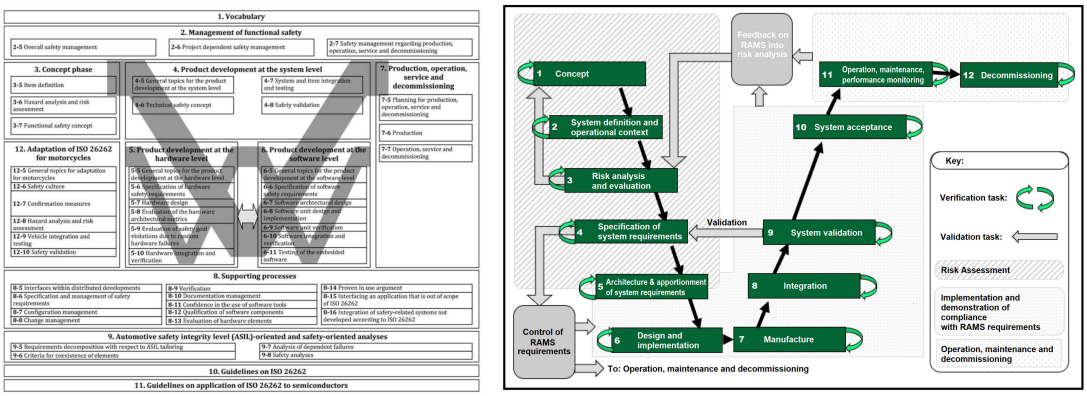


Figure 3.3: V-Cycle as described by ISO 26262 (left) and EN 50126 (right).

Although ISO 26262 and EN 50126 target different application domains, both adopt the V-cycle to ensure traceability between requirements, implementation, verification and validation. The main differences lie in the application rather than the development philosophy. ISO 26262 focuses on Functional Safety for automotive electronic systems through the ASIL framework, whereas EN 50126 adopts a broader RAMS-oriented approach covering the entire railway system lifecycle.

Automotive Safety Integrity Level (ASIL)

ISO 26262 classifies safety requirements according to the Automotive Safety Integrity Level (ASIL), ranging from ASIL A to ASIL D, where the latter represents the highest safety integrity. The ASIL is determined through the Hazard Analysis and Risk Assessment (HARA), considering the Severity (S), Exposure (E) and Controllability (C) of a hazardous event,

$$ASIL = f(S, E, C) \tag{3.1}$$

where $f(\cdot)$ denotes the decision matrix defined by ISO 26262 rather than an analytical expression.

Safety Integrity Level (SIL)

Within the railway domain, safety integrity is expressed through the Safety Integrity Level (SIL), ranging from SIL 1 to SIL 4. Unlike ASIL, SIL is associated with *quantitative safety targets* based on the acceptable probability of dangerous failures. For continuously operating systems, the relevant metric is the Probability of dangerous Failure per Hour (PFH),

$$PFH = \frac{N_{\text{dangerous}}}{T_{\text{operating}}} \tag{3.2}$$

where $N_{\text{dangerous}}$ is the number of dangerous failures observed during the operating time $T_{\text{operating}}$. The resulting PFH value is compared against the limits specified for each SIL level.

Podium's development process is therefore more closely aligned with the automotive approach of ISO 26262. However, maturity levels should not be interpreted as a direct implementation of the standard, but as an industrial adaptation that preserves the key Functional Safety principles of requirements traceability, systematic verification and progressive validation. Within this framework, the present thesis follows an automotive-oriented development process while extending its application and design considerations to a railway project.

Table 3.1: Conceptual mapping between Podium's development process, ISO 26262 and EN 50126.

ML	**Podium Development Process**	**ISO 26262**	**EN 50126**
ML1	Preliminary evaluation, project feasibility, system definition	Concept Phase: Item Definition, Hazard Analysis and Risk Assessment (HARA), FuSa	Concept, System Definition, Operational Context
ML2	System requirements specification and system architecture definition	Product dev. at System Level: Technical Safety, System Requirements, System Architecture	System Requirements Specification, Architecture and Allocation of Functions
ML3	HW and SW requirements together with their respective architectures	HW Development (Part 5): HW Safety Requirements and Architecture	System Design and Requirements Allocation (implementation further supported by EN 50128/50129)
ML4	SW unit design and implementation	SW Unit Design and Implementation	Design and Implementation
ML5	Software unit verification	Software Unit Verification	Verification Activities
ML6	HW and SW integration and verification	HW+SW Integration and Verification	Integration and Verification
ML7	System validation (typically through Hardware-in-the-Loop testing)	Item Integration and Safety Validation	System Validation
ML8	Final product validation and release	Final Safety Validation, Production Readiness	System Acceptance, Preparation for Operation

Chapter 4

Defining HW architecture of the BJB (ML3)

This chapter addresses the HW development and the requirements derived from the system level requirements evaluated in ML2. The result is the block diagram in Figure 4.1.

As seen in Chapter 3 the steps required to develop a certified product are many, and it has to be considered that in the development of this project, for the BMS alone, more than 20 engineers have been actively involved since 2024.

In this context, my work on the BJB board is part of a project already well underway, with ML1, ML2 and part of ML3 already accomplished.

This means that the choice of the silicon provider had already been made, coherently with the rest of the BMS: the project relies on NXP's previous-generation MC33772C [7], a multipurpose ASIC that can serve both as CMU and BJB chip, featuring TPL connectivity, advanced support as part of an agreement between Podium and NXP and numerous drivers.

The chip does not support the latest-generation features, as those were not required in the customer requirements, and any addition would have required additional validation, extremely expensive and time consuming.

The sections follow the system composer's block diagram. This block diagram represents all the components of the system, hence all functional blocks of the BJB. The tool used to describe the static architecture of the BMS is *System Composer* [15], a MathWorks tool that works within the Simulink framework and describes the static structure of the system, in the sense of the System Modeling Language (SysML): it models all the existing blocks and the general structure in the Block Definition Diagram (BDD), and the interfaces, data flow, parts and ports of the single blocks in the Internal Block Diagram (IBD).

System Architecture elements can be linked to Simulink or Stateflow models for

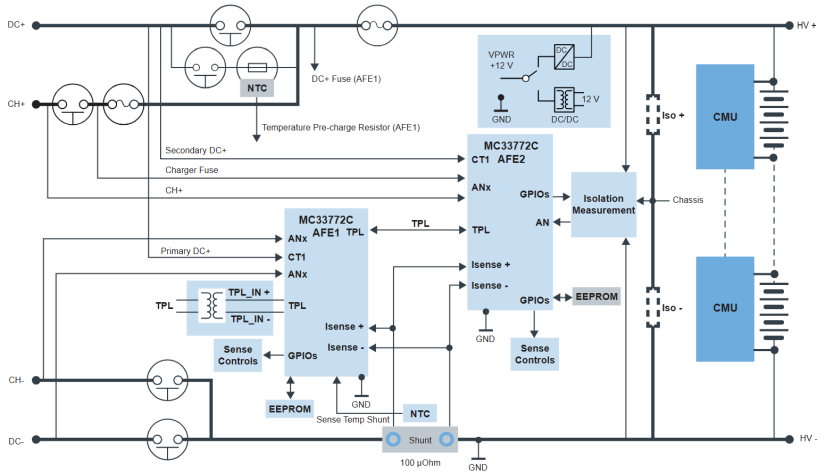


Figure 4.1: Block diagram of the BJB.

behaviour, and to requirements in Requirements Toolbox for traceability.

The approach to this chapter as well as others in the thesis is mainly descriptive of processes, the choice not to include numbers is intentional, as it keeps the sensitive data of the project confidential.

4.1 End-of-line calibration

Although this is more of a system-level requirement, during the development of the board an End-Of-Line calibration of the measurement chains was considered, in order to reduce the measurement uncertainty further than the figures stated by the component manufacturers.

This kind of calibration is unfortunately expensive and can also lead to misleading accuracy figures, since a large part of the value drift of electronic components comes from ageing and from thermal behaviour; hence, although an EEPROM was foreseen and finally placed on the PCB to store the chain configurations and calibrations, it was eventually chosen not to make use of it at all.

4.2 HV positive voltage measurement

The first functional block of the BJB is the one in charge of measuring the potentials of the high voltage nodes.

The block interfaces with the HV bus on one side, with the HV ground on the other, and lastly with the GPIO of the MC33772C chip.

The BJB's ground is still considered to be an HV point as it is connected (referred) to the battery's negative terminal (HV–), constituting the local ground of the whole HV domain of the board.

All the BJB's blocks are considered to operate in HV and components must be chosen accordingly, to comply with the stringent standards for surge and noise voltages, much higher than what can be experienced in the rest of the BMS, which operates in the LV domain, galvanically isolated from the HV domain.

The measurement is therefore single ended and, by construction, always positive: this is what distinguishes this block from the bipolar measurement described in the next section.

The nodes of interest are the ones that allow the BMS to supervise the power path, namely the battery side of the main positive contactor (*DC_pos voltage*), the load side of the same contactor (*DC_pos_fuse voltage*) and the pre-charge node (*DC_pos_prech*) voltage. This is what enables, at system level, the following functions:

- **Pre-charge supervision**: the DC-link voltage is monitored during the pre-charge phase and the main contactor is allowed to close only when the ratio between link and pack voltage exceeds the threshold defined in the system requirements (around 95%), within the allotted pre-charge time.
- **Contactor diagnostics**: the comparison between the two sides of the contactor, with the contactor commanded open, detects a welded contact; the same comparison with the contactor commanded closed detects a failure to close or an excessive contact resistance as required by the Failure Mode, Effects and Diagnostic Analysis (FMEDA) document.
- **Plausibility check**: the pack voltage measured by the BJB is compared with the sum of the cell voltages reported by the CMUs. The two quantities are acquired by independent channels, on independent boards (CMUs vs BJB), and their agreement is the diagnostic that covers a common cause failure of the cell measurement chain or of the BJB's measurement chain.
- **Interlock and safe state**: the absence of voltage on the load side is the condition that declares the HV bus discharged after a shutdown request.

Derived requirements

From the system level requirements (system + component) defined in ML2, the block inherits the so called *HW Requirements* that allow at component level to comply with system defined requirements.

The following is a short list of requirements, and not the full list, which counts more than 30 requirements just for this block:

- *Range* of measurement from 0 to V_{max}, where V_{max} is the maximum pack voltage (it shall be considered that the overshoot is minimal, hence V_{max} is only slightly higher than the nominal voltage of the pack);
- *Accuracy* of the measurement better than ±1% of the reading over the full operating temperature range and over the lifetime of the product above V_{batt_min}, and better than ±2.5 V below that value, in order to discriminate the pre-charge threshold and the contactor states without ambiguity;
- *Bandwidth* sufficient to follow the pre-charge exponential, whose time constant is set by the pre-charge resistor and the DC-link capacitance;
- *Leakage current* toward HV– low enough not to perturb the isolation resistance measurement performed by the insulation monitoring device and not to contribute appreciably to the self discharge of the pack at rest;

Implementation

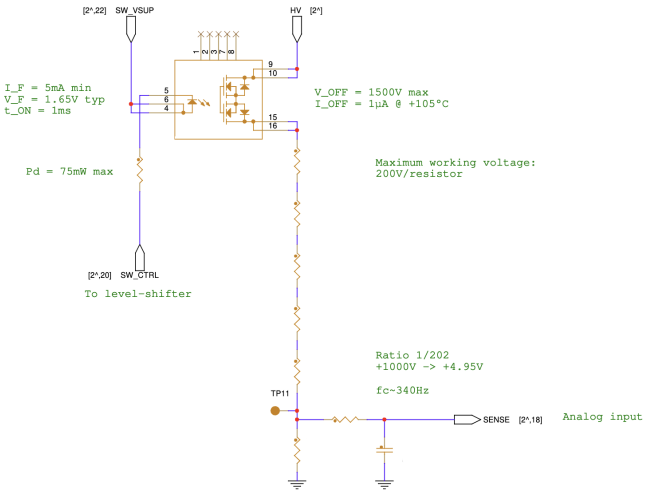


Figure 4.2: Schematic of the HV positive measurement chain.

Divider ratio and full scale Each node is scaled down by a resistive divider and acquired by one of the measurement channels of the MC33772C, which accepts an input span of 0 to 5 V and converts it with a ΣΔ ADC. The divider ratio is chosen as

$$k = \frac{R_{low}}{R_{high} + R_{low}} \leq \frac{V_{max_{ADC}}}{V_{max}} \tag{4.1}$$

so that the full scale of the ADC is approached but never exceeded in the worst case. The residual span is the trade-off between resolution and headroom for transients: in this project, while the battery sits at lower voltage, the maximum measurable voltage of the chain is 1000 V.

Series chain on the high side The upper branch R_{high} is a series chain. Three reasons justify this choice: the working voltage rating of a single thick film resistor is limited to a few hundred volts (200 V at best, usually), well below the pack voltage; the series chain distributes the voltage stress and the power dissipation; and it guarantees that a single resistor failure does not short the HV node to the AFE input.

As an additional benefit, the statistical uncertainty drops with the square root of the number of equal elements in the chain, so the overall uncertainty is marginally reduced as well.

Disconnection of the chain In order to isolate the measurement chain when not used, reducing current leakage, cross-talk to other channels, and exposure of the components to HV noise or damaging events, a photorelay is placed on top of the high-side branch of resistors, disconnecting the chain and withstanding up to 1500 V DC, more than double the rated voltage of the battery pack.

Sizing of the total resistance The total resistance is dimensioned as a compromise between the leakage requirement, which pushes it up, and the error caused by the input bias current of the AFE, which pushes it down, since this current flows in the Thévenin equivalent resistance seen at the divider tap and produces an offset that adds directly to the reading; this resistance is in the order of MΩ.

Filtering The tap is followed by a low-pass filter working as *Analog Front End (AFE)*, that serves both as anti-aliasing for the switched capacitor front end of the ADC and as attenuation of the common mode noise injected on the HV bus by the switching of the inverter. The cut-off is well above the pre-charge dynamics and well below the sampling rate of the GPIO.

Isolation and layout As anticipated, almost the whole BJB sits in the HV domain, so the isolation barrier towards the rest of the BMS is implemented on the communication interface: apart from the power supply of the board, the only element connecting it to the rest of the BMS is the TPL daisy chain, which is transformer coupled and crosses the barrier towards the BMU. The divider itself therefore lies entirely inside the HV domain, and the clearance and creepage distances of its layout are dimensioned according to the pollution degree (2) and

the working voltage prescribed by the standards; as this is a double certification board it must comply both with railway and automotive standards, widely different in terms of EMC.

Error budget and diagnostics

Contributions to the budget The goal when tuning the accuracy of the block is to make it as independent as possible from the divider uncertainty, but also to set the ADC to a working point where its relative uncertainty is not high. The contributions that have been budgeted are the initial tolerance of the resistors, their temperature coefficient, their long term drift and the self heating of the upper branch, to which the gain and offset error of the AFE channel and the effect of the input bias current are added.

Further detail on the uncertainty budget can be found in Section 5.7.1.

Choice of the components Resistors with 0.1% tolerance and low temperature coefficient have been chosen for the divider, as ratio matching between the two branches is extremely important since no end-of-line measurement is foreseen to trim the ratio itself.

Failure modes and diagnostic coverage The failure mode of the block requires an additional remark: an open circuit in the upper branch drives the GPIO input to zero, which is indistinguishable from a genuinely discharged node. The fault is therefore *misleading*, and the diagnostic coverage of this failure is delegated to the plausibility checks listed above, in particular to the comparison against the sum of the cell voltages and to the coherence between the nodes during the pre-charge sequence.

4.3 HV bipolar voltage measurement

The second measurement block shares topology, components and layout rules with the one described in Section 4.2, and everything stated there about the series chain on the high side, the sizing of the total resistance against leakage, the filtering stage and the clearance and creepage rules applies unchanged. What distinguishes it is the sign of the measured quantity.

The nodes handled by this block belong to the portion of the power path whose potential can occasionally drop slightly below the HV– ground. This can happen because of transients following an opening under load, or noise, or other nonidealities of the ground plane.

The functional need is therefore to have a bipolar measurement chain, and to remap it onto the input span of the ADC, which is unipolar and bounded between 0 and 5 V. The whole section is about the single element that performs this mapping: a 2.5 V reference injected into the low branch of the divider.

Derived requirements

Only the requirements that differ from the previous block are reported:

- *Range* of measurement symmetric around zero, from $-V_{max}$ to $+V_{max}$, with the same worst case magnitude of the unipolar chain;
- *Accuracy* expressed in absolute terms over the whole range rather than as a $\pm 2\%$ of the reading or ± 5 V, since around zero a relative specification degenerates;
- *Sign discrimination* around zero guaranteed within the resolution of the chain;
- *Stability of the offset* over temperature and lifetime, since no end-of-line trimming is performed and the offset is not observable from the measurement itself.

Implementation

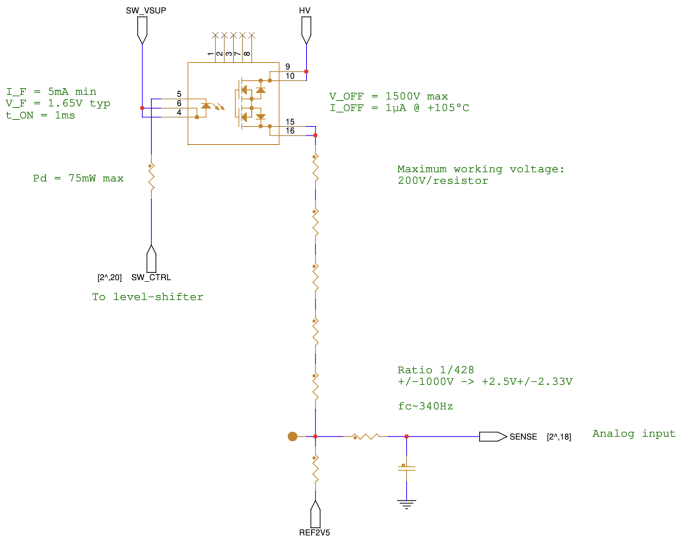


Figure 4.3: Schematic of the HV bipolar measurement chain.

Level shifting of the divider The low branch of the divider is not returned to HV– but to a reference potential V_{ref}. The voltage seen by the GPIO becomes

$$V_{ADC} = k\,V_{in} + (1-k)\,V_{ref}, \qquad k = \frac{R_{low}}{R_{high} + R_{low}} \tag{4.2}$$

and the quantity of interest is recovered by the software as

$$V_{in} = \frac{V_{ADC} - (1-k)\,V_{ref}}{k}. \tag{4.3}$$

Since $k \ll 1$, the second term of (4.2) is essentially V_{ref} itself: the reference sets the code corresponding to a null input, and the attenuated signal rides on top of it in both directions.

Choice of the offset value Placing V_{ref} at half of the ADC span is the choice that makes the two polarities symmetric, and it is the one adopted here. The price is paid on the attenuation ratio: the condition on k becomes

$$k \leq \frac{V_{max_{ADC}}/2}{V_{max}}, \tag{4.4}$$

that is, half of the ratio allowed in the unipolar case for the same V_{max}, hence half of the volts per LSB resolution. An asymmetric offset would recover part of that resolution whenever the negative excursion is known to be much smaller than the positive one, but it requires a dedicated reference value, it makes the chain no longer reusable between the two polarities and it complicates the cancellation described below.

The symmetric choice was preferred as it keeps the block identical to the positive one except for a single net, but also allows for more versatile operation in future projects where the GND reference and HV– of the battery might sit at different potentials.

Generation of the reference The reference is derived from the same regulated output that the MC33772C provides for ratiometric sensing, and which also sets the full scale of the GPIO conversion; if the offset is a fixed fraction of the ADC reference, a drift of that reference moves the measured code and the offset code in the same direction, and the two effects partially cancel in (4.3). A reference generated independently, however accurate, would not enjoy this cancellation and its drift would appear in full in the reading.

Loading of the reference node Unlike the unipolar chain, where the low branch always returns current to ground, here the reference node must both sink and source, since the current in the divider reverses with the sign of the input. Its output impedance R_s appears in series with R_{low}: it perturbs the ratio k and, more importantly, it makes the offset dependent on the measured voltage, turning what should be a pure offset into a gain error. The requirement is therefore $R_s \ll R_{low}$ over the whole current excursion imposed by $\pm V_{max}$, which in practice calls for a buffered reference rather than a plain resistive partition of the supply rail. Local decoupling of the reference node is equally critical, because noise present there is transferred to the ADC input with unity gain, while noise present at the HV node is attenuated by k.

Bidirectional protection and disconnection The photorelay used to disconnect the chain must block both polarities, which is guaranteed by the back to back MOSFET structure of the selected device; a switch built around a single MOSFET would conduct through its body diode for one sign of the input and would defeat the purpose of the disconnection. For the same reason the clamping network at the AFE input is bidirectional, as the node can now be pushed both above the supply rail and below the local ground, and the filter capacitors must be non polarised.

Error budget and diagnostics

Propagation of the reference error The dominant new term of the budget is the uncertainty of V_{ref}. Differentiating (4.3),

$$\frac{\partial V_{in}}{\partial V_{ref}} = -\frac{1-k}{k} \simeq -\frac{1}{k}, \tag{4.5}$$

which means that the reference error is referred back to the input amplified by the inverse of the attenuation ratio, the same factor that protects the chain from the noise of the HV bus. An error of a few millivolts on the reference is therefore worth volts at the input. This is the single figure that dictates the accuracy of the block and the reason why the tolerance, the temperature coefficient and the long term drift required for the reference are far more stringent than those required for any resistor of the divider.

Nature of the resulting error The consequence of (4.5) is that the error of this block is an offset, constant over the range, rather than a fraction of the reading. Near full scale the block is comparable to the unipolar one; near zero, where the unipolar chain would be intrinsically accurate, this one is at its worst in relative terms. Expressing the accuracy requirement in absolute terms, as done above, is

the direct translation of this behavior, and the plausibility thresholds that consume this measurement have been set accordingly.

Live zero and diagnostic coverage The offset brings one benefit that partly compensates its cost: here the null input corresponds to mid scale, and both extremes of the range are outside the plausible window: a tap collapsed to ground, an open low branch pulled to the rail, or a loss of the reference all produce a code that no legitimate operating condition can produce. The failure that was *misleading* in Section 4.2 becomes *detectable* by the channel itself, without relying on external plausibility checks.

4.4 Current sensing

Current is the only quantity of the battery that cannot be reconstructed from the others, and it feeds two functions with opposite needs: the estimation of the state of charge, which integrates it over hours and therefore lives on the accuracy around zero, and the protection of the pack, which reacts in milliseconds and therefore impacts the range and the latency. Most of the design choices of this block are the arbitration between the two.

The measurement is performed with a shunt. Hall effect transducers have an advantage for what concerns isolation, as it is provided "for free", but in this case as established in Section 4.2 the whole BJB already sits in the HV domain and the barrier is crossed by the TPL daisy chain.

Furthermore, Hall effect current sensors are not well regarded by some vehicle manufacturers, to the point that this became a stakeholder requirement during the development of the board.

The shunt is sized so that the peak detectable current is far greater than the maximum the system can withstand: this comes from the System Requirement that over-current events be correctly detected, in order to act on fuses, pyro-fuses and contactors in time to save the battery itself.

The value of the shunt is the usual compromise between the amplitude of the differential signal and the power dissipated on the element, and it is connected in Kelvin configuration, since its resistance is comparable to that of the joints that carry the current to it. A thermistor coupled to the shunt is acquired ratiometrically on a dedicated analog input: it corrects the residual temperature coefficient of the element, and it is at the same time the only quantity that observes the thermal state of a component carrying the full pack current.

Derived requirements

Only a subset of the requirements allocated to the block is reported:

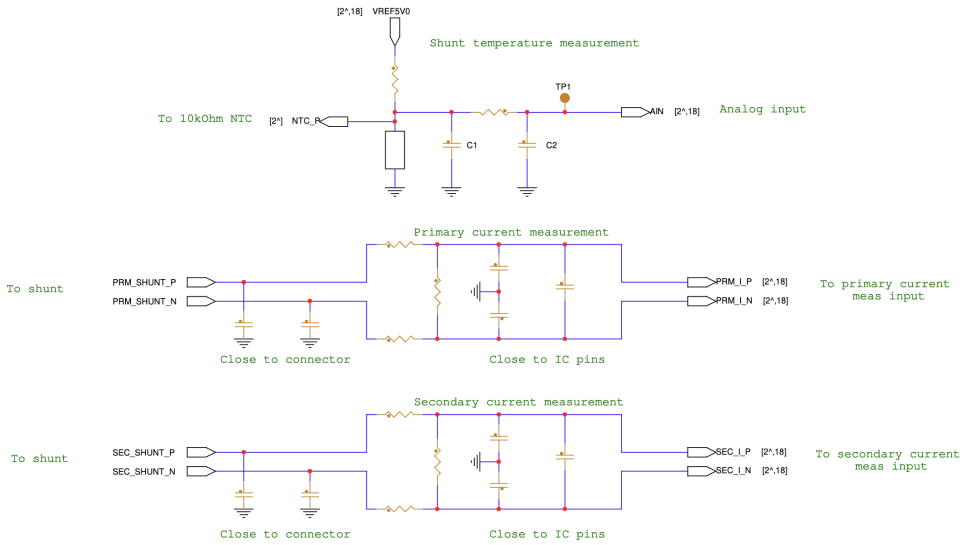


Figure 4.4: Schematic of the redundant current measurement chain, with the shunt temperature measurement.

- *Range* covering the peak current in both directions, with headroom for the fault currents that must be measured;
- *Offset stability* over temperature and lifetime, being the term that the integration accumulates;
- *Bandwidth and latency* compatible with the reaction time of the overcurrent protection (tens of milliseconds), sampled coherently with the voltage channels;
- *Independence* of the two chains, as required by the decomposition of the safety goal this measurement contributes to.

Implementation

Two stage filtering The filter is split between the two ends of the chain, as annotated on the schematic. The first stage, at the entrance of the board, diverts to ground the common mode noise picked up by the harness from the switching of the inverter, before it propagates along the board. The second stage, at the pins of the acquisition device, is the anti-aliasing filter proper and provides at the same time the charge reservoir that the switched capacitor input of the $\Sigma\Delta$ converter demands: a converter of this kind draws a current spike at every sampling instant, and without a local capacitor able to supply it the settling would depend on the impedance of the whole chain.

Symmetry of the two legs The series resistors of the two legs are equal and the differential capacitor is a single component between the lines rather than two capacitors to ground. Both choices preserve the common mode rejection, which holds only when a disturbance reaches the two inputs identically: any mismatch converts a fraction of the common mode into a differential signal that cannot be distinguished from a real current. Two capacitors to ground would make the rejection depend on the tolerance matching of two separate parts, an order of magnitude worse than what a single differential capacitor guarantees by construction. The capacitors to ground of the first stage sit before the series resistors, at a low impedance level, where their mismatch is far less effective. The series resistors also limit the current forced into the internal protection structures when the input is driven outside the supply rails, and a resistor across the pair defines the input potential so that an open sense line cannot leave it floating.

Redundancy of the chains Two chains, primary and secondary, acquire the same shunt and are independent from the sense taps onward, each with its own filtering network and its own acquisition input. This allows the safety requirement allocated to this measurement to be met without demanding from a single chain a diagnostic coverage it cannot provide: the comparison of two readings sharing no active component is a diagnostic on both. It also opens the possibility, configuring the two channels with different gains, of covering the low current region with the resolution required by the coulomb counting and the high current region with the range required by the protection, without asking one channel to do both, but this additional feature was not developed, in order to preserve the original intent of the redundant current measurement: compliance with the ASIL D safety level.

Error budget and diagnostics

Contributions to the budget The budgeted terms are, summarized:

- tolerance of the shunt and its residual temperature coefficient after correction
- error of the temperature channel performing that correction
- offset and gain error of the acquisition channel with their drift
- finite common mode rejection together with the mismatch (already discussed)
- noise integrated over the residual bandwidth

Their weight is not uniform: at high current the multiplicative terms dominate, around zero the additive ones do. To these adds a contribution specific to this block, the thermoelectric voltage of the junctions between the alloy of the shunt

and the copper of the sense connection, which cancels only as long as the two junctions are at the same temperature.

Diagnostic coverage and its limits In this case diagnostics are layered:

- the comparison between the two chains covers the failures of a single chain and is the diagnostic on which the decomposition of the safety requirement
- the resistor across the pair makes an open sense line observable
- the plausibility against the voltages of the power path and the currents reported by the loads covers what is common to both chains
- the shunt temperature is checked against the dissipation implied by the measured current rests

The two chains share the shunt, the connector and the harness: a drift of the shunt value, a mechanical failure of its joints or a disconnection of the harness affects both identically and produces two readings that agree with each other and are both wrong. However, this kind of event is estimated to have a Mean Time To Failure (MTTF) high enough to justify the absence of redundancy on the shunt itself.

4.5 Port mapping on the MC33772C

The selection of the right input/output port for each measurement chain is decisive in reaching the accuracy and quality goals set by the requirements.

The device offers three families of analog inputs. The `CT` terminals are a stack-referenced differential ladder, tied to the cell chain and unavailable for anything else.

The `GPIO` ports, configured as analog inputs (`ANx`), are single-ended channels referred to `AGND` and multiplexed into ADC1-A.

The `ISENSE` pair is a differential low-level input feeding a programmable gain amplifier (PGA) and its own converter, ADC2. Each choice fixes the reference, the input window and the available diagnostics before any external component is sized.

4.5.1 High-voltage chains

Pack-level voltages are referred to ground and must be divided down, so the `ANx` ports are the only option. Of the two acquisition modes, the ratiometric one cancels the source error only when the sensing network is biased from `VCOM`, which is the case of the NTC dividers and not of a high-voltage divider; the absolute mode is used instead, and is also the more accurate of the two (Table 4.1).

The bipolar chain needs a mid-scale offset, so only half of the window is available per polarity: k doubles and the stability of the offset enters the budget. The bipolar chains therefore drive the divider sizing.

Two constraints follow from the port itself. The input leakage flows through the Thévenin impedance of the divider, which keeps the low-side branch in the few-kΩ range. Each `ANx` also carries over- and under-threshold registers routed to the `FAULT` pin, which makes the port a monitoring resource and not only an acquisition channel.

4.5.2 Current chain

The shunt drop is differential and low-level, and has to be resolved over multiple orders of magnitude. This rules out a single-ended port and a fixed-gain chain. `ISENSE`$\pm$ provides a Kelvin differential input, automatic gain selection over $\{4, 16, 64, 256\}$ and a converter running continuously.

The gain is transparent to the user scale: the register resolution stays at V_{2RES} for any gain, so the PGA adds a factor 64 of instantaneous range without a per-range scale factor to handle in software, and the $\pm 152\,\mathrm{mV}$ span is read on a single linear axis.

Across the ranges, each gain has its own calibration factor with cold and hot terms interpolated on die temperature, so the transitions do not produce steps in the transfer characteristic and the $\pm 0.5\,\%$ gain error holds over the whole span, provided ADC2 runs at full resolution.

Table 4.1: Relevant characteristics of the `GPIO/ANx` analog input port (MC33772C, Rev. 4.0) [16].

Symbol	Parameter	Min	Max	Unit
V_{ADC}	Input range (ratiometric)	`AGND`	`VCOM`	V
V_{ANx_ERR}	Channel error, ratiometric	−16.0	16.0	mV
V_{ANx_ERR}	Channel error, absolute (1.0 V to 4.5 V)	−10.0	10.0	mV
V_{ANx_ERR}	Channel error, absolute, after ageing	−11.0	11.0	mV
$V_{CT_ANx_RES}$	Resolution, 15-bit register	–	0.153	mV/LSB
I_{IL}	Input leakage current	−100	100	nA
$V_{OL(TH)}$	Open-pin detection threshold	0.10	0.23	V
R_{OPENPD}	Open-detection pull-down	3.8	6.2	kΩ
V_{GPIOx}	Absolute maximum pin voltage	−0.3	`VCOM`+0.5	V

Table 4.2: Relevant characteristics of the `ISENSE` differential port and of the ADC2 current chain (MC33772C, Rev. 4.0) [16].

Symbol	Parameter	Min/Typ	Max	Unit
V_{IND}	Differential input range	−150	150	mV
V_{INC}	Common-mode input range	−300	300	mV
V_{2RES}	Register resolution (any gain)	0.6	–	µV/LSB
$V_{ISENSEX(OFFSET)}$	Input-referred offset error	–	0.5	µV
$I_{GAINERR}$	Gain error including nonlinearity	−0.5	0.5	%
$I_{SENSEX(BIAS)}$	Common-mode input bias current	−100	100	nA
$I_{SENSE(DIF)}$	Differential input bias current	−5.0	5.0	nA
V_{I_NOISE}	Noise, 16-bit conversion	3.01	–	µV rms
t_{ICONV}	Conversion time incl. PGA settling	37.67	–	µs
t_{AZC_SETTLE}	Auto-zero settling time	200	–	µs

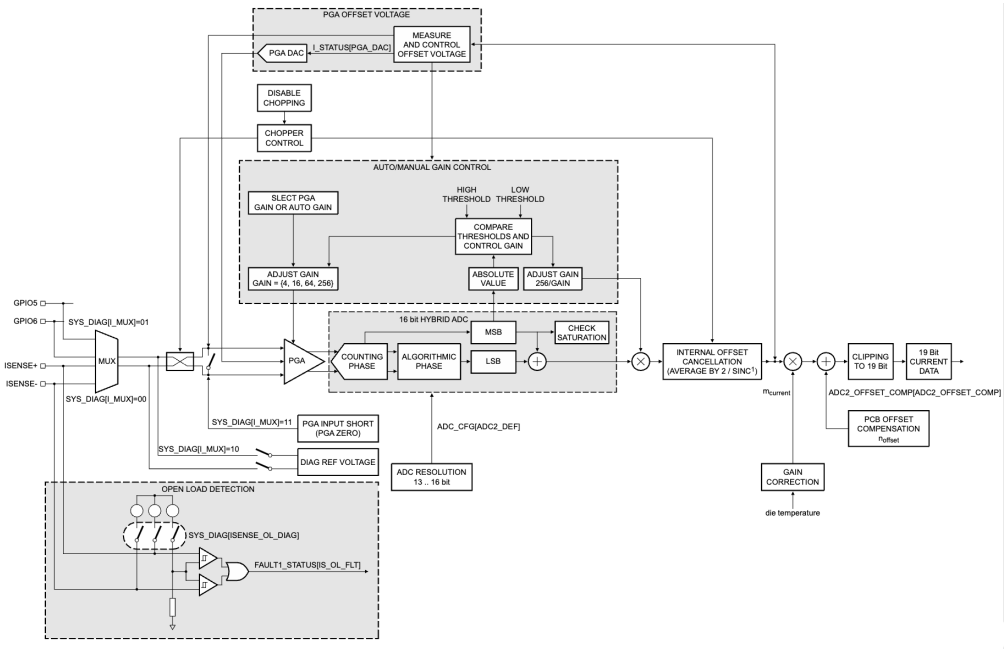


Figure 4.5: Current measurement channel: input multiplexer, chopper, automatic-gain PGA, ADC2 and the digital offset and gain correction stages.

Chapter 5

Isolation measurement chain

A specific chapter is required for the isolation measurement, as this is the core of the thesis since it is the part of the BMS that according to literature and internal knowledge has turned out to be the most challenging HW wise.

Automotive and railway systems, while different in many ways, share similar needs when it comes to isolation measurement, so the general functioning of the system can be treated consistently among the two.

5.1 State of the art on isolation measurement

When it comes to insulation monitoring device for ungrounded drive systems in electric vehicles, there are different strategies to choose from.

Generally the idea is to measure resistance either injecting a known current in a partially unknown system (*active measurement*) or in the other case by measuring a voltage fall across a voltage partition, from a known source voltage: *passive measurement*.

The active method is more complex, but offers faster readouts and potentially greater range, the passive one is simpler and while more limited is extremely reliable, furthermore in the case of the battery pack the voltage source is "for free" as measurement of its voltage is already available and computed as seen in previous chapter. This clearly adds an important bias to the technology of choice.

5.1.1 Bender iso175

Bender is the world leader when it comes to isolation measurement systems, from automotive to industrial fields. Bender is a producer well known inside Podium, as numerous projects implement their measurement boards. This is an active measurement board, that provides both a fast and a precise isolation estimation.

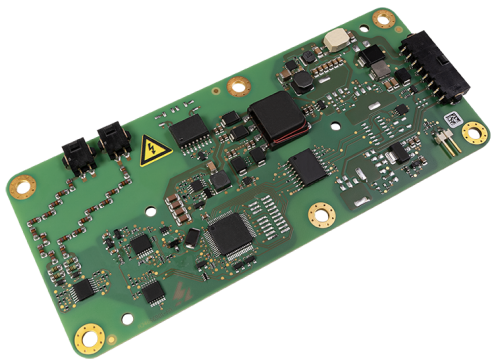


Figure 5.1: Bender iso175 – Board for isolation measurement.

Table 5.1: Main specifications of the Bender iso175 insulation monitoring device [17].

Parameter	**Standard**	**Fast Startup**
Supply	12/24 V DC	Same
DC voltage	0–1000 V DC (6000 V impulses)	Same
AC voltage	0–480 V_{RMS} AC	Same
Measurable R	0–50 MΩ	Same
Output range	0–35 MΩ	Same
High-cap mode	Up to 10 μF	Up to 10 μF
Response time	$\leq 30s$	Reduced (not specified)

The company ensures less than 60 s for a precise estimation, while also offering a fast estimation based on exponential prediction, with growing and growing confidence while time of measurement increases.

That board ensures a ±12% tolerance.

5.1.2 NXP reference design

NXP's RDA777T2 board implements an unbalance measurement system based on the voltage unbalancing principle. A high-impedance resistor divider scales the HV bus voltage. Optically isolated switches connect calibrated resistor branches to chassis. By measuring the resulting voltage variations, the positive and negative insulation resistances can be estimated independently in a two unknowns in two equations fashion.

The high input impedance minimises the measurement current and power

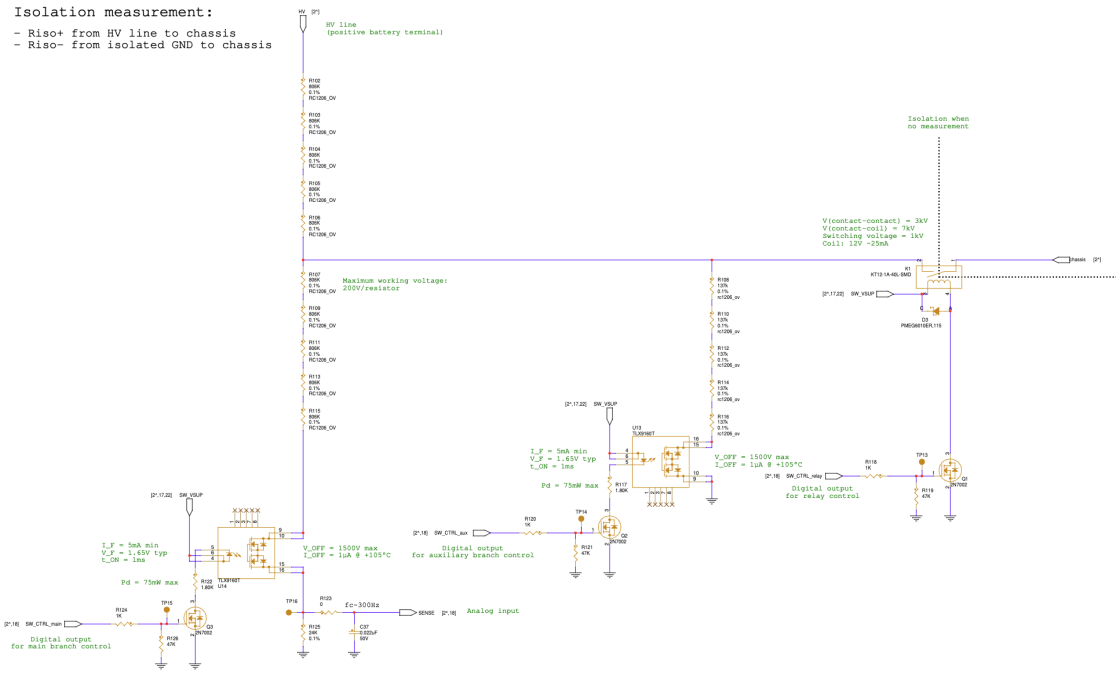


Figure 5.2: NXP – Isolation measurement reference design.

Table 5.2: Main specifications of RDA777T2 insulation measurement circuit.

Parameter	**Value**
Rated HV bus voltage	Up to 1000 V_{DC}
Main divider resistance	$\approx$ 8.06 MΩ
Auxiliary branch resistance	$\approx$ 548 kΩ
Divider resistor tolerance	$\pm$0.1%
Estimated divider current @ 1000 V	$\approx$ 124 μA
Estimated divider power @ 1000 V	$\approx$ 124 mW
Relay isolation voltage	3 kV (contacts), 7 kV (coil-contact)
OptoMOS output blocking voltage	1500 V

dissipation, making the solution suitable for railway applications

5.2 Customer specifications and design premises

Customer requirements do not specify the range of operation (input resistance) that the system shall be capable of estimating correctly, so to bridge this “short circuit” in the requirements structure, a “reasonable” range is estimated from systemists in

a specific paper called *Assumptions of Use*.

For what concerns the parameter of *bjbData_isomeas_rIso* the a-priori choice was to develop the system to be effective over a range from $100\,\Omega$ up to $100\,\mathrm{M}\Omega$.

Customer requires, following a de-facto rule born in automotive industry, that sets the under-isolation threshold at $500\,\Omega\,\mathrm{V}^{-1}$.

As seen in the previous chapters, the device selected for the acquisition is the NXP MC33772C [7], whose analog inputs and switched GPIOs are the resources available for the isolation measurement.

5.3 Method and general design idea

As seen in picture 4.1, the isolation measurement is the block that sits through positive and negative poles of the battery. From that image the idea is to make a symmetrical measurement system, exploiting the two unknowns in two equations system seen from NXP's design.

The model chosen for this project, as already subtly mentioned in the introduction to the chapter is the *unbalance model*, using a double voltage partition system.

The idea behind this design is to reduce measurement points from 2 to just one: as to compute two unknowns, two independent equations are needed, two measurements are performed, among the two a known perturbation to the system is introduced, then the measurement is repeated.

This allows to effectively measure two different quantities using only one ADC and one measurement chain, reducing space on PCB and costs.

Compared to NXP's reference design the idea is to have a "symmetrical" system, simpler, as it removes the need to disconnect the entire system by the use of a relay, in fact both the high and low branches are disconnected by means of an opto-mos.

The system is designed to have a partition of $175:1$ to ensure the input of the ASIC's ADC are not saturated. This partition constant k is lower then the one chosen for voltage measurement branches, reflecting the urge of a more tailored architecture to the battery needs as this measure is highly impacted by range and sensitivity of the ADC, as this chapter will hopefully highlight.

This chapter approaches the problem from a *functional* standpoint. Rather than opening with a netlist, it follows the chain of reasoning that leads from the physical arrangement of resistors to a number the Battery Management System can act on.

A methodological note applies throughout. The analysis was carried out almost entirely in MATLAB, and every uncertainty figure quoted in this chapter is a type B estimate in the sense of the GUM [18], computed with a coverage factor $k_p = 2$ unless stated otherwise.

Where a result could be reached by more than one route, it was checked against

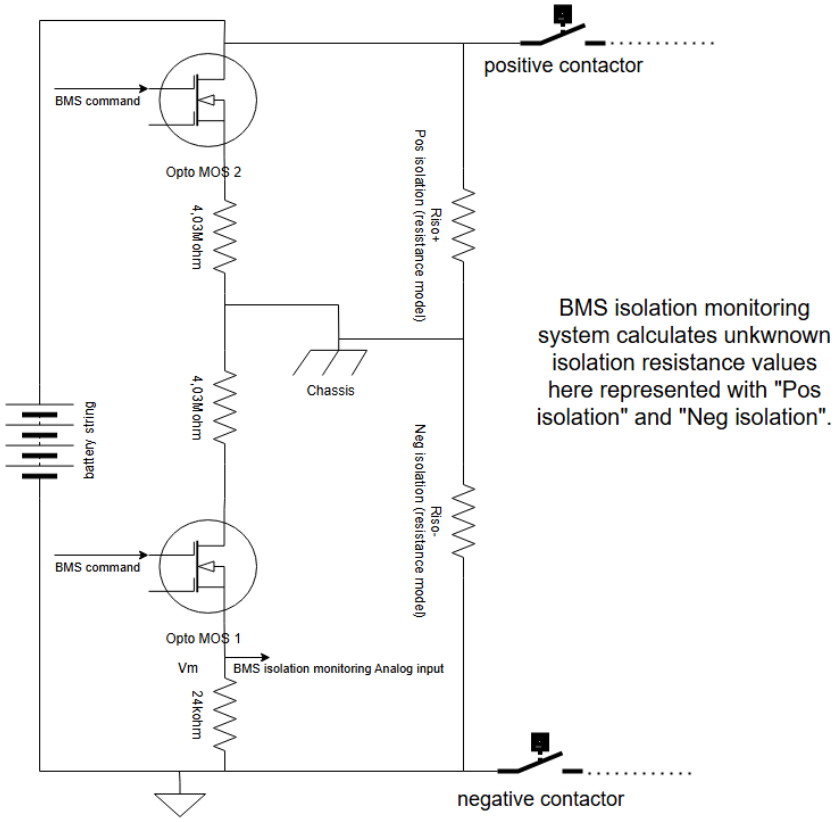


Figure 5.3: Isolation measurement model.

an independent derivation before being trusted; the dead ends and abandoned analytic attempts are not reproduced here.

5.4 Model description and requirements

Two known resistors set the divider on each side of the bus – R_1 in parallel with the positive insulation path and R_2 in series with the sense resistor R_m on the negative side and the insulation resistances R_{isoP} and R_{isoN} appear as the unknown parallel elements.

The voltage of interest, V, is taken as the drop across the ends of R_m and is scaled into the 0–5 V window accepted by the ADC1-A/B inputs of the NXP `MC33772C`. The design takes the economical route discussed in the previous section: rather than two sensing chains, it derives the two independent equations it needs from a single chain by taking two readings of V under two switch states:

The first, with the R_1 branch open, and the second, with the corresponding switch closed, present the front end with two electrically distinct systems while

using only one ADC and one measurement path, saving both board area and cost.

Throughout the remainder of the chapter the two readings are written as

- V_1 – the value of V measured with switch SW_1 *open*;
- V_2 – the value of V measured with switch SW_1 *closed*.

To comply with railway applications, the connection and disconnection of the chassis-to-V_{batt+} and chassis-to-V_{batt-} paths must satisfy the functional safety provisions of EN 50126 [2].

To meet this, two opto-MOS devices (`TPSI2140-Q1`) [19], each rated to isolate up to a surge voltage of 5000 V, are placed in the switching paths.

A consequence of the very high resistances of isolation of the pack is that the measurement is slow. The parasitic capacitances of the harness and the pack, multiplied by insulation resistances in the GΩ range, produce extremely long time constants, so the bus needs seconds (even hundreds of seconds) to reach steady state before each reading is valid.

5.4.1 Operating envelope and the role of the contactors

The pack is built so that the battery negative pole V_{BATT-} is referred to the train rails, and therefore to chassis, whenever the contactors are closed. An insulation measurement against chassis only carries meaning while the pack is floating, so the chain operates exclusively with the contactors *open*, that is, with the battery disconnected from the traction bus. This places the block at a precise point of the start-up sequence:

- **Precondition to power-up**: a confirmed-good isolation reading is one of the conditions that authorise the closing of the contactors to power the DC_LINK.
- **Exclusive operating window**: once the contactors close, V_{BATT-} and chassis become the same node, and the measurement is no longer defined. The function is therefore available only in the disconnected state.
- **Settling before validity**: the parasitic capacitances of the harness and of the pack, multiplied by insulation resistances in the GΩ range, give time constants in the order of seconds, so a reading is declared valid only after the bus has reached steady state.
- **Max measurement time**: the system must be operative in max $10s$ from operator's start-up procedure.

5.4.2 Derived requirements

The accuracy demanded of the result originates in the client specification and is propagated first into the System Requirements and then into the Component Requirements. From these, the block inherits the following constraints:

- relative uncertainty of the estimated insulation resistance below 35% of the computed value;
- measurement range from hundreds of kΩ to 50/100MΩ, covered by a single acquisition chain;
- execution restricted to the contactor-open state, with the result available before the closing sequence is authorised;
- settling time compatible with the start-up timing budget of the vehicle.

5.4.3 How much uncertainty is “too much uncertainty”?

The 35% figure looks generous read on its own, but stops looking generous once it is put against the range it has to hold across: the same divider and the same converter are asked to resolve both ends of the interval, this specification is derived from the “assumptions of use” paper, as the customer is not specifying a figure for this.

In order to explain the choice of 35% here’s an example:

Close to the admissible threshold, the uncertainty governs whether a reading can be trusted as a basis for a decision. A measurement can look compliant on its nominal value while its uncertainty still leaves a non-negligible probability that the true resistance sits below the limit. The danger sharpens when the relative uncertainty is large compared with the gap between the estimate and the threshold.

A numerical example makes the point. Suppose the system estimates $\hat{R} = 100\,\mathrm{M\Omega}$ against a threshold TH, with coverage factor $k_p = 2$, and compare a 95% relative uncertainty against a 35% one:

$$\text{Case 1:}\quad u_c = \frac{U(k_p)}{k_p} = \frac{95\%}{2} = 47.5\%,$$
$$k_{p,\max} = \frac{\hat{R} - TH}{u_c} = 2.098,$$
$$P_{\text{monolateral}} = 1.8\%,$$

$$\text{Case 2:}\quad u_c = \frac{U(k_p)}{k_p} = \frac{35\%}{2} = 17.5\%,$$
$$k_{p,\max} = \frac{\hat{R} - TH}{u_c} = 5.7,$$
$$P_{\text{monolateral}} = 1.2 \cdot 10^{-8}\%.$$

The margins differ by orders of magnitude. With 95% uncertainty the lower bound of the interval reaches the threshold at a coverage factor of barely 2.1, leaving a one-sided probability near 1.8% that the real insulation is sub-threshold; with 35% uncertainty that probability collapses to the order of 10^{-8}%.

The lesson is that a measurement can be nominally far above the limit and still be unsafe to act on, which is precisely why a maximum admissible uncertainty has to be enforced and any reading exceeding it discarded – at the price of some false rejections.

5.5 Minimum required isolation

The threshold the estimate is judged against is not fixed by an open standard. As a design rule it is taken at roughly $500\,\Omega/\text{V}$, a value that has the weight of a *de facto* standard in the automotive industry through the harmonised document LV123, agreed among the major German OEMs. Two thresholds are defined, a warning level and a minimum working level:

$$TH_{ISOwarning} = 1090\,\Omega/\text{V}, \qquad TH_{ISOmin} = 500\,\Omega/\text{V}.$$

For the project at hand, with a nominal working voltage $V_{BATTnom} = 660\,\text{V}$, these scale to

$$R_{ISOwarning} = 718\,\text{k}\Omega, \tag{5.1}$$

$$R_{ISOmin} = 500\,\Omega/\text{V} \cdot 660\,\text{V} = 324\,\text{k}\Omega. \tag{5.2}$$

5.6 Formal description

The chain can be described, electrically, as a double voltage divider. Reasoning about it informally is error-prone, so to avoid this mistake the section gives it a formal representation. Three views of the same system are built in: a *physical* view that maps insulation resistances to the voltages actually read, a *system* view that inverts that mapping to recover the resistances from the readings, and a *virtual* view that re-parameterises the problem to remove its dependency on the battery voltage.

5.6.1 Physical mapping

In its most general form the measurement is a function of seven quantities. Writing the subscript "e" for the *extended* model – more variables than the phenomenon strictly requires – and "p" for *physical*, the forward map is

$$f \ : \ D_{Pe} \longrightarrow C_{Pe}$$

with

$$D_{Pe} = \{(R_1, R_2, R_m, R_{isoP}, R_{isoN}, V_{BATT1}, V_{BATT2}) \in (\mathbb{R}^+)^7\}$$
$$C_{Pe} = \{(V_1, V_2) \in (\mathbb{R}^+)^2\}$$

Each reading is itself a function of six of those quantities – the three fixed resistors, the two insulation resistances, and the relevant battery voltage:

$$V_1 : \ D_{Pe1} \longrightarrow C_{Pe1}$$
$$R_1, R_2, R_m, R_{isoP}, R_{isoN}, V_{BATT1} \longmapsto f(R_1, R_2, R_m, R_{isoP}, R_{isoN}, V_{BATT1})$$

$$V_2 : \ D_{Pe2} \longrightarrow C_{Pe2}$$
$$R_1, R_2, R_m, R_{isoP}, R_{isoN}, V_{BATT2} \longmapsto f(R_1, R_2, R_m, R_{isoP}, R_{isoN}, V_{BATT2})$$

This extended form carries more structure than is useful, because it does not separate the genuine variables from the design constants. The three resistors R_1, R_2 and R_m are not subject to evaluation: they are considered fixed by the design (and time/temperature invariant), chosen following the silicon vendor's guidelines, and do not change between measurements. Removing them from the list of free variables collapses the description onto a cleaner four-dimensional domain,

$$F_p \ : \ D_P \longrightarrow C_P$$

with

$$D_P = \{(R_{isoP}, R_{isoN}, V_{BATT1}, V_{BATT2}) \in (\mathbb{R}^+)^4\}$$
$$C_P = \{(V_1, V_2) \in (\mathbb{R}^+)^2\}$$

which factors per reading as

$$F_p(R_{isoP}, R_{isoN}, V_{BATT1}, V_{BATT2}) = \Big(f_{p1}(R_{isoP}, R_{isoN}, V_{BATT1}),\ f_{p2}(R_{isoP}, R_{isoN}, V_{BATT2})\Big).$$

A note on treating V_{BATT1} and V_{BATT2}

The two battery-voltage entries deserve comment, because the temptation is to set $V_{BATT1} = V_{BATT2}$ and be done with it. That is the easiest option, and it is **not viable**. The two readings are separated in time by the settling interval, which for this system is long; over that interval the pack voltage can drift slightly, or even appreciably.

Carrying a single battery voltage for both readings folds that drift directly into the result, and – as the inverse formulas below make explicit, the term $(V_1 V_{BATT2} - V_2 V_{BATT1})$ sits in the numerator of both insulation estimates, so when V_1 and V_2 are close the estimate becomes acutely sensitive to any mismatch between the two voltages.

Keeping V_{BATT1} and V_{BATT2} as separate, independently measured quantities is what allows the system to be treated as time invariant. They are best understood not as extra unknowns but as *correction coefficients* that absorb the drift between the two readings.

Electro-technical description

Starting from the electrotechnical description of the divider, the two readings can be written explicitly. The positive insulation appears in parallel with R_1 on the high side, and the negative insulation appears in parallel with the $(R_2 + R_m)$ branch on the low side:

$$R_{hs} = \frac{R_{isoP} \cdot R_1}{R_1 + R_{isoP}} \tag{5.3a}$$

$$R_{ls} = \frac{R_{isoN} \cdot (R_2 + R_m)}{R_{isoN} + R_2 + R_m} \tag{5.3b}$$

$$V_1 = V_{BATT1} \cdot \frac{R_{ls}}{R_{isoP} + R_{ls}} \cdot \frac{R_m}{R_2 + R_m} \tag{5.4a}$$

$$V_2 = V_{BATT2} \cdot \frac{R_{ls}}{R_{hs} + R_{ls}} \cdot \frac{R_m}{R_2 + R_m} \tag{5.4b}$$

Here R_2 is of the order of the MΩ, which is what gives the chain adequate sensitivity and linearity across the working range of interest discussed in Section 5.5. Substituting (5.3a)–(5.3b) into (5.4a)–(5.4b) condenses the two readings into:

$$V_1 = V_{BATT1} \cdot \frac{\frac{R_{isoN} \cdot (R_2 + R_m)}{R_{isoN} + R_2 + R_m}}{R_{isoP} + \frac{R_{isoN} \cdot (R_2 + R_m)}{R_{isoN} + R_2 + R_m}} \cdot \frac{R_m}{R_2 + R_m} \tag{5.5a}$$

$$V_2 = V_{BATT2} \cdot \frac{\frac{R_{isoN} \cdot (R_2 + R_m)}{R_{isoN} + R_2 + R_m}}{\frac{R_{isoP} \cdot R_1}{R_1 + R_{isoP}} + \frac{R_{isoN} \cdot (R_2 + R_m)}{R_{isoN} + R_2 + R_m}} \cdot \frac{R_m}{R_2 + R_m} \tag{5.5b}$$

Equations (5.5a)–(5.5b) are the physical behaviour of the system: given the insulation that physically exists, they predict the voltages the ADC will read. They do not, by themselves, tell the BMS what it actually wants to know, which is the reverse – the insulation implied by a pair of readings. That inversion is the subject of the next subsection.

5.6.2 System mapping: solving for the insulation

From the BMS standpoint the inputs are the measured quantities V_1, V_2, V_{BATT1}, V_{BATT2} and the desired outputs are R_{isoP} and R_{isoN}. Treating the two battery voltages as fixed parameters reduces the forward map to a function of the two insulation resistances alone,

$$F_{p(V_{BATT1}, V_{BATT2})} : (\mathbb{R}^+)^2 \to (\mathbb{R}^+)^2,$$

$$F_{p(V_{BATT1}, V_{BATT2})}(R_{isoN}, R_{isoP}) = F_p(R_{isoP}, R_{isoN}, V_{BATT1}, V_{BATT2}),$$

whose inverse is exactly the system-side map of interest,

$$F_{BMS(V_{BATT1}, V_{BATT2})} = F^{-1}_{p(V_{BATT1}, V_{BATT2})}(V_1, V_2) = (R_{isoN}, R_{isoP}).$$

Before inverting, the inverse has to be shown to exist. Local invertibility is guaranteed wherever the Jacobian of the forward map is non-singular,

$$J = \det \begin{pmatrix} \frac{\partial V_1}{\partial R_{isoN}} & \frac{\partial V_1}{\partial R_{isoP}} \\ \frac{\partial V_2}{\partial R_{isoN}} & \frac{\partial V_2}{\partial R_{isoP}} \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \neq 0.$$

For this system the condition holds everywhere on the domain of interest, since all the partial-derivative terms are strictly positive and therefore

$$J = ad - -bc > 0 \quad \forall \Big((V_1, V_2) \subset C_p,\ (R_{isoN}, R_{isoP}) \subset C_p \Big).$$

With invertibility established, the system map

$$F_{BMS(V_{BATT1}, V_{BATT2})} \ : \ D_{BMS} \longrightarrow C_{BMS}$$

$$D_{BMS} = \{(V_1, V_2) \in (\mathbb{R}^+)^2\}$$
$$C_{BMS} = \{(R_{isoN}, R_{isoP}) \in (\mathbb{R}^+)^2\}$$

is obtained by solving (5.5a)–(5.5b) for the two insulation resistances. The algebra is heavy and was carried out with the MATLAB Symbolic Math Toolbox; the closed forms are

$$R_{isoP} = f_{BMS1} = \frac{R_1 R_m \cdot (V_1 V_{BATT2} - V_2 V_{BATT1})}{V_1 \cdot (R_2 V_2 + R_m V_2 - R_m V_{BATT2})} \tag{5.6a}$$

$$R_{isoN} = \frac{R_1 R_m \cdot (R_2 + R_m) \cdot (V_2 V_{BATT1} - V_1 V_{BATT2})}{V_1 V_2 (R_2 + R_m)^2 - R_m^2 (V_1 V_{BATT2} + V_2 V_{BATT1} - V_{BATT1} V_{BATT2}) + R_1 R_m (V_1 V_{BATT2} - V} \tag{5.6b}$$

The numerator $(V_1 V_{BATT2} - V_2 V_{BATT1})$ in (5.6a) is the formal reason the single-voltage shortcut of Section 5.6.1 was rejected: it nearly vanishes when the two readings coincide, amplifying any inconsistency between the two battery voltages.

5.6.3 Virtual mapping: removing the battery-voltage dependency

The system map of the previous subsection is valid only for a fixed pair (V_{BATT1}, V_{BATT2}). A description that holds across all battery voltages is more useful, particularly when defining safety rules that must apply regardless of the pack's state of charge. The dependency can be removed by dividing each reading by its own battery voltage, which introduces two normalised quantities

$$vvbat_1 = \frac{V_1}{V_{BATT1}}, \qquad vvbat_2 = \frac{V_2}{V_{BATT2}}.$$

Because the measurement model depends on the voltages only through these ratios, replacing the raw voltages by their normalised counterparts does not reduce the effective number of degrees of freedom; it is a reparameterisation of the input space in independent variables, in normalised form (5.5a)–(5.5b) become

$$vvbat_1 = \frac{V_1}{V_{BATT1}} = \frac{\frac{R_{isoN}\cdot(R_2+R_m)}{R_{isoN}+R_2+R_m}}{R_{isoP} + \frac{R_{isoN}\cdot(R_2+R_m)}{R_{isoN}+R_2+R_m}} \cdot \frac{R_m}{R_2+R_m} \tag{5.7a}$$

$$vvbat_2 = \frac{V_2}{V_{BATT2}} = \frac{\frac{R_{isoN}\cdot(R_2+R_m)}{R_{isoN}+R_2+R_m}}{\frac{R_{isoP}\cdot R_1}{R_1+R_{isoP}} + \frac{R_{isoN}\cdot(R_2+R_m)}{R_{isoN}+R_2+R_m}} \cdot \frac{R_m}{R_2+R_m} \tag{5.7b}$$

defining the virtual map

$$F_v \ : \ D_v \subset (\mathbb{R}^+)^2 \longrightarrow C_v \subset (\mathbb{R}^+)^2$$

$$D_v = \{(vvbat_1, vvbat_2) \in (\mathbb{R}^+)^2\}$$
$$C_v = C_s = \{(R_{isoP}, R_{isoN}) \in (\mathbb{R}^+)^2\}$$

The virtual domain has no direct physical meaning, but its extent is fully determined by the physical limits of the system, the ADC window and the battery-voltage range:

$$[vvbat_{min}, vvbat_{max}] = \left[\frac{V_{ADCmin}}{V_{BATTmax}}, \ \frac{V_{ADCmax}}{V_{BATTmin}}\right].$$

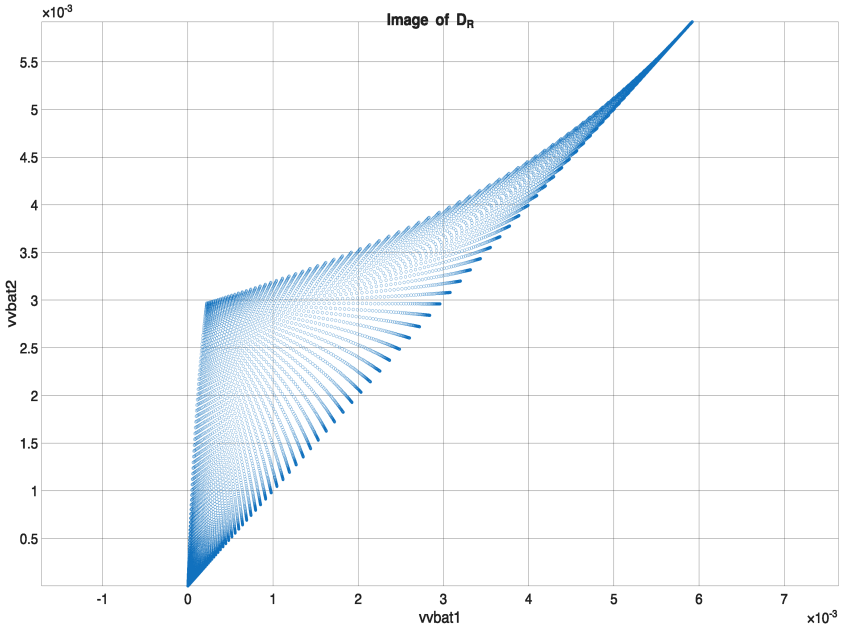


Figure 5.4: Matlab simulation – Image I of the application, in the range assumed in customer requirements.

Solving (5.7a)–(5.7b) for the insulation resistances, again with MATLAB, gives the battery-voltage-independent estimates

$$R_{isoP} = \frac{R_1 R_m (vvbat_1 - vvbat_2)}{vvbat_1 (R_2\, vvbat_2 - R_m + R_m\, vvbat_2)} \tag{5.8a}$$

$$R_{isoN} = \frac{(R_2 + R_m) \cdot R_1 R_m \cdot (vvbat_1 - vvbat_2)}{R_m^2 (vvbat_1 + vvbat_2 - 1) - R_1 R_m (vvbat_1 - vvbat_2) + R_2 R_m (vvbat_1 + vvbat_2) - R_2^2\, vvb} \tag{5.8b}$$

These are the working formulas used for the analysis and the safety mechanisms in Section 5.8, since they let every conclusion be stated once, for all battery voltages, rather than re-derived at each operating point.

5.7 Uncertainty of measurement

For a safety feature such as isolation monitoring, confidence in the number matters as much as the number itself. Precision is important, but it is secondary to a complete accounting of where the measurement can go wrong, because an unmodelled error source degrades the quality of the decision far more than a slightly noisier reading would.

Good practice, once the system is modelled, is to add *safety mechanisms* that flag a result as suspect when the conditions for a trustworthy measurement are not met; the virtual domain of Section 5.6.3 was introduced partly because it offers a clean space in which to express those flags.

The measurement model, the uncertainty budget and the safety mechanisms are three pieces of one coherent whole, and they have to agree with one another to avoid producing contradictory outputs.

5.7.1 Contributions to the budget

The contributors are the parameters listed in Section 5.6.1, each of which carries its own uncertainty.

ADC. The non-idealities of the ADC are the first to consider. They split conventionally into static and dynamic terms: on the static side, gain and offset error, quantization noise, integral and differential non-linearity (INL, DNL), reference-voltage stability and thermal noise; on the dynamic side, jitter, conversion time and sample-and-hold effects. The *dynamic terms can be discarded* for this application as sampling speed is slow (milliseconds). The `MC33772C` datasheet does not break out individual non-linearity figures; instead it guarantees that the total error stays below a stated bound, summarised in Table 5.3.

	V range	Resolution	LSB value	E_{typ}	E_{max}	Noise
ADC1-A	0–4.85 V	15 bit	152.59 μV/LSB	$\pm$0.8 mV	$\pm$8 mV	600 μV

Table 5.3: Key boundary parameters of the ADC, extracted from the datasheet [16].

Resistors. The chain resistors R_m, R_1 and R_2 contribute linearly to the budget, which makes them far easier to model than the ADC. The application is niche and only three resistors are involved, so all three are specified at 0.1% tolerance.

Battery voltage. The uncertainty of V_{BATT1} and V_{BATT2} enters through the high-voltage sensing chains. Both the `HVp-meas` and `HVbp-meas` chains of the `BJB8R02P00` were modelled in MATLAB, sweeping several precision-resistor configurations, including parallel arrangements on the low side intended to drive the sensing uncertainty down. Figure 5.5 compares the relative and absolute uncertainty of the chain at 400 V, 800 V and 1000 V for the different configurations.

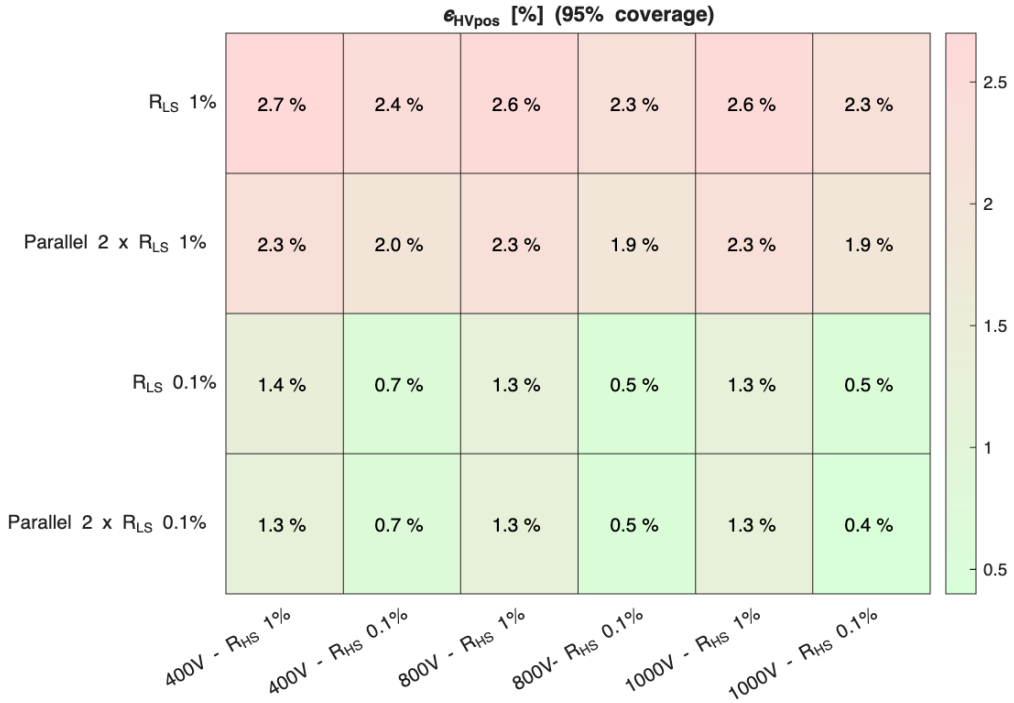


Figure 5.5: Relative uncertainty of the measurement chain for several parameter sets.

The expected result holds – the most precise configuration is the one combining

the tightest resistors with a parallel low-side arrangement – but the design decision is not about the absolute optimum. It is about the sweet spot between uncertainty, cost and complexity, and that lands on 0.1% resistors *without* resorting to a parallel R_m. Figure 5.6 shows why: with 0.1% resistors the budget is dominated by the ADC error, whereas at 1% the resistors take over and dominate, so refining the resistors past 0.1% buys little while the ADC term is still in control.

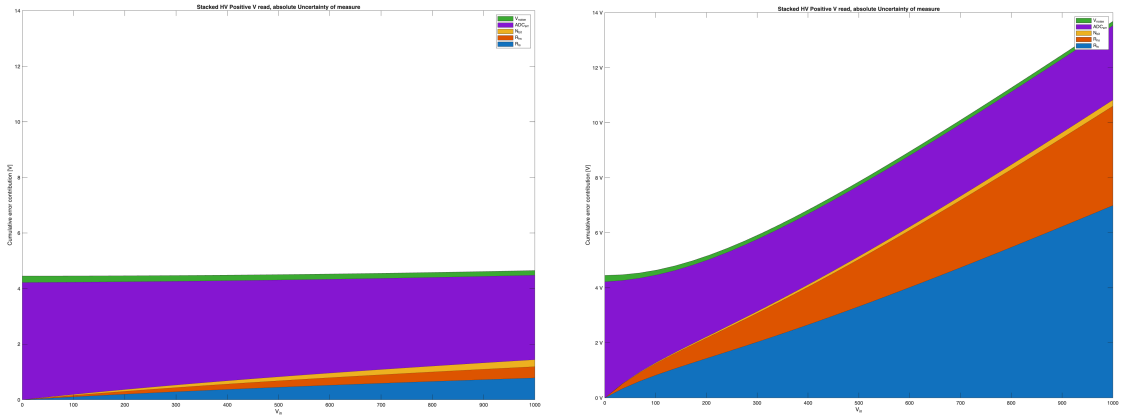


Figure 5.6: On the left, with 0.1% resistors, the budget is dominated by the ADC error; on the right, with 1% resistors, the resistor tolerance dominates.

In the chosen configuration the resistor contributions evaluate to $U(R_{hs}) = 0.052\%$ (for the five high-side resistors at 0.1%) and $U(R_{ls}) = 0.12\%$ (no parallel low side). To keep the number of inputs to the uncertainty function manageable, $U(V_{BATT})$ is not carried as a function of the actual battery voltage but fixed at its worst case over the nominal range, roughly $[360\,\mathrm{V}, 700\,\mathrm{V}]$. With coverage factor $k_p = 2$ this gives the constant

$$U_{V_{BATT}} = 4.5\,\mathrm{V}. \tag{5.9}$$

5.7.2 Propagating the uncertainty

A fully analytic propagation was attempted and abandoned: the closed-form insulation expressions are simply too unwieldy to differentiate by hand with any confidence. The working approach propagates the budget numerically through the model. Formally the uncertainty is a map

$$U_s : g_s(R_1, R_2, R_m, V_1, V_2, V_{BATT1}, V_{BATT2}) \longrightarrow (U_{R_{isoN}}, U_{R_{isoP}}) \tag{5.10}$$

which, transported to the virtual domain and treating R_1, R_2, R_m as parameters, becomes the lower-dimensional

$$U_v \;:\; g_v(R_1, R_2, R_m, vvbat_1, vvbat_2) \in (\mathbb{R}^+)^5 \longrightarrow (U_{R_{isoN}}, U_{R_{isoP}}) \in (\mathbb{R}^+)^2 \tag{5.11}$$

This change of space is lossless: it preserves the degrees of freedom of the full problem while removing the battery-voltage dependency, so the budget can be stated once for the whole domain

$$D_v = \{(vvbat_1, vvbat_2) \in (\mathbb{R}^+)^2\}$$
$$C_v = \{(R_{isoN}, R_{isoP}) \in (\mathbb{R}^+)^2\}$$

Only type B contributions are considered, since no repeated measurement is taken that would justify a type A term. The combined standard uncertainties of the two estimates follow from the usual quadratic propagation,

$$U_{R_{isoN}} = \sqrt{\nabla f_{v1}^T \, \Sigma_v \, \nabla f_{v1}} \tag{5.12a}$$

$$U_{R_{isoP}} = \sqrt{\nabla f_{v2}^T \, \Sigma_v \, \nabla f_{v2}} \tag{5.12b}$$

where Σ_v is the covariance matrix of the inputs. Its diagonal holds the variance of each contribution; the off-diagonal terms hold the covariances. In this application the cross-correlations are all taken as zero,

$$\Sigma_v = \begin{pmatrix} u^2(R_1) & 0 & 0 & 0 & 0 \\ 0 & u^2(R_2) & 0 & 0 & 0 \\ 0 & 0 & u^2(R_m) & 0 & 0 \\ 0 & 0 & 0 & u^2(vvbat_1) & \mathrm{cov}(vvbat_1, vvbat_2) \\ 0 & 0 & 0 & \mathrm{cov}(vvbat_2, vvbat_1) & u^2(vvbat_2) \end{pmatrix} \tag{5.13}$$

Intuition would suggest correlation between $vvbat_1$ and $vvbat_2$to be high: both readings come from the same ADC, separated by an interval short enough that the converter's error model can be treated as time invariant, so a shared gain or offset error would correlate them. The problem is that the datasheet does not decompose the total error into its gain, offset, INL and DNL parts, so there is no defensible basis on which to claim a common source. Lacking that information, the covariance is conservatively set to zero,

$$\mathrm{cov}(vvbat_1, vvbat_2) = \mathrm{cov}(vvbat_2, vvbat_1) = 0.$$

The gradients in (5.12a)–(5.12b) are taken with respect to the same five inputs,

$$\nabla f_{v1} = \frac{\partial f_{v1}}{\partial R_1}\hat{R}_1 + \frac{\partial f_{v1}}{\partial R_2}\hat{R}_2 + \frac{\partial f_{v1}}{\partial R_m}\hat{R}_m + \frac{\partial f_{v1}}{\partial vvbat_1}\widehat{vvbat}_1 + \frac{\partial f_{v1}}{\partial vvbat_2}\widehat{vvbat}_2 \tag{5.14a}$$

$$\nabla f_{v2} = \frac{\partial f_{v2}}{\partial R_1}\hat{R}_1 + \frac{\partial f_{v2}}{\partial R_2}\hat{R}_2 + \frac{\partial f_{v2}}{\partial R_m}\hat{R}_m + \frac{\partial f_{v2}}{\partial vvbat_1}\widehat{vvbat}_1 + \frac{\partial f_{v2}}{\partial vvbat_2}\widehat{vvbat}_2 \quad (5.14b)$$

5.7.3 Deterministic versus probabilistic interpretation

The uncertainty attached to an insulation estimate cannot be read as a plain tolerance band; it has to be interpreted against the chosen confidence level and the decision rule applied to the measurement. This matters most when the lower edge of the confidence interval drifts down towards the isolation threshold. In that regime even a modest excursion beyond the assumed coverage can correspond to a genuine under-isolation, and the robustness of the decision erodes.

Two evaluation models bracket the situation. The *deterministic* model combines the maximum deviation of every contributor and so returns a worst-case upper bound on the error. The *probabilistic* model propagates the statistical contributions through the model, allowing partial cancellation when uncertainties pull in opposite directions, and so generally returns a smaller figure. The catch is that the cancellation only helps when many comparable contributors are present. When the budget is instead dominated by a handful of terms – which is the case for this measurement – the probabilistic estimate creeps up towards the deterministic one, and in the limit can even exceed it, which is physically absurd. To keep the two consistent, the extended probabilistic uncertainty is clamped at the deterministic bound,

$$U(x) = \min\Big(U(x), \Delta x\Big), \quad U(x) = u(x) \cdot k_p, \quad k_p = 2,$$

with the deterministic uncertainty acting as a hard logical ceiling on the probabilistic one.

5.8 Analysis and safety mechanisms

With the model, its inverse and its uncertainty budget in place, the chain can be characterised across its whole operating space, and the points where it cannot be trusted can be turned into explicit safety mechanisms. The analysis proceeds in three coordinate systems – the resistance space (R_{isoN}, R_{isoP}), the voltage space (V_1, V_2) and the virtual space $(vvbat_1, vvbat_2)$ – because each makes a different property easiest to see.

5.8.1 Lower bounds on the readable voltage

Two safety mechanisms set the floor below which a reading is meaningless. The first, **SM1**, fixes the voltage at which the signal-to-noise ratio falls below an acceptable

target. Defining the total SNR from the quantisation and noise terms,

$$SNR_t = 10\log_{10}\left(\frac{SNR_{q(lin)} \cdot SNR_{n(lin)}}{SNR_{q(lin)} + SNR_{n(lin)}}\right) = 20\log_{10}\left(\frac{V_{in}}{V_q + V_n}\right),$$

and requiring $SNR_t > 10\,\mathrm{dB}$ gives

$$SNR_t = 20\log_{10}\left(\frac{V_{in}}{V_q + V_n}\right) > 10\,\mathrm{dB} \longrightarrow V_{in} > 2.4\,\mathrm{mV}. \tag{5.15}$$

The second floor guards against readings comparable to, or below, the ADC's own declared error. As a rule of thumb the input should never fall below twice the typical error, so that the implied uncertainty stays under 50%; this is admittedly a crude criterion, and the more rigorous probabilistic treatment supersedes it, but as a hard floor it gives

$$V_{in_{MIN}} > 2 \cdot E_{typ}/\sqrt{3} \sim 1\,\mathrm{mV}. \tag{5.16}$$

5.8.2 Uncertainty over the resistance space

Sweeping the resistance domain and mapping it forward through the physical model F_p, then attaching the uncertainty of (5.10), produces the relative-uncertainty surface of Figure 5.7. The construction is: generate the grid over (R_{isoN}, R_{isoP}) at nominal battery voltage, compute (V_1, V_2) from F_p, then evaluate U_s at each point. Formally the two stages are

$$\begin{aligned} F_s :\ & (\mathbb{R}^+)^2 \times (\mathbb{R}^+)^5 \longrightarrow \mathbb{R}^2 \\ \Big((R_{isoP}, R_{isoN}), (R_1, R_2, R_m, V_{BATT1}, V_{BATT2})\Big) & \longmapsto (V_1, V_2) \\ \\ U_s :\ & (\mathbb{R}^+)^5 \times \mathbb{R}^2 \longrightarrow (\mathbb{R}^+)^2 \\ \Big((R_1, R_2, R_m, V_{BATT1}, V_{BATT2}), (V_1, V_2)\Big) & \longmapsto (U_{R_{isoN}}, U_{R_{isoP}}) \end{aligned} \tag{5.17}$$

Fixing V_{BATT} rather than carrying it as a variable could in principle narrow the validity of the result, but taking the worst case $V_{BATT1,2} = 360\,\mathrm{V}$ makes the estimate conservative: every other operating point is at least as good.

The result is characteristic of unbalanced topology: the chain is at its best when the two insulations are of comparable magnitude, and degrades as they diverge, because the divider output clamps near a rail when one side dominates. Figures 5.8–5.9 show the underlying readings, and the correspondence with the uncertainty maps is direct: low uncertainty coincides with the regions where the readings change most steeply with resistance.

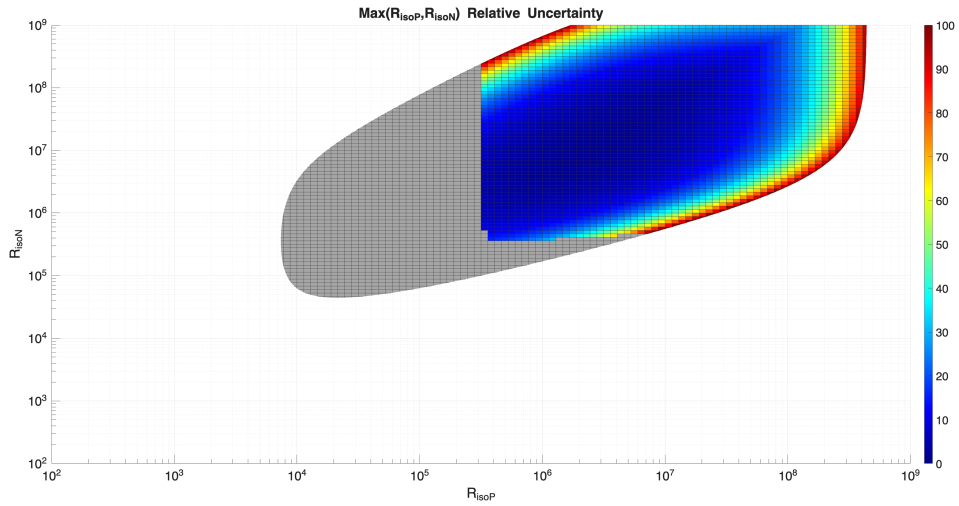


Figure 5.7: Maximum relative uncertainty in the resistance domain. The greyed region marks where, after subtracting uncertainty, the estimate falls below threshold, i.e. $\min(R_{isoP} - U(R_{isoP}),\ R_{isoN} - U(R_{isoN})) < TH$.

Outcome 1 – balance of positive and negative isolation When the positive-to-chassis and negative-to-chassis insulation differ widely, by more than about three orders of magnitude, the uncertainty of at least one of the two estimates degrades severely, potentially to the point of being unusable.

The link between sensitivity and uncertainty is made explicit by the normalised sensitivity maps of Figures 5.8 and 5.9. These plot the derivatives of the readings normalised on a logarithmic resistance scale, which gives a kind of universal sensitivity expressing how readily the chain detects a step from one resistance value to an adjacent one. Normalising this way is necessary: a plain derivative would be swamped by the large absolute resistances and could not be compared uniformly across the decades. There is no absolute scale on these plots, since the quantity is in effect $V/\log(\Delta R)$.

Outcome 2 – sensitivity drives uncertainty Regions of constant colour in Figures 5.8 and 5.9 are low-sensitivity regions, and they coincide with the regions of greatest measurement uncertainty. Sensitivity peaks on the bisector of the plane, where the chain is symmetric ($R_1 = R_2$).

Outcome 3 – why the two poles differ The sensitivity on V_1 (Fig. 5.8) correlates strongly with the uncertainty of R_{isoP} (Fig. 5.10): when V_1 is read, R_{isoP} is the only high-side resistor, giving good linearity and sensitivity, especially where $R_{ls} = R_{isoN} \parallel (R_2 + R_m)$ is comparable to $R_{hs} \equiv R_{isoP}$. The picture is weaker

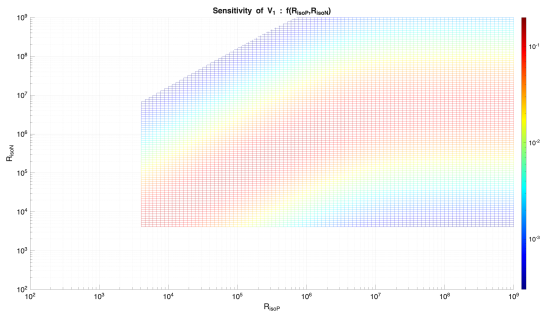


Figure 5.8: Resistance-normalised sensitivity of the $V_1 = f(R_{isoN}, R_{isoP})$ measurement.

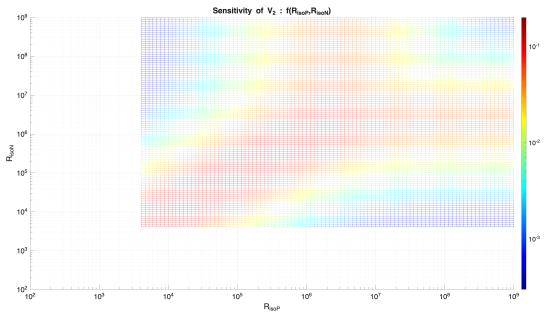


Figure 5.9: Resistance-normalised sensitivity of the $V_2 = f(R_{isoN}, R_{isoP})$ measurement.

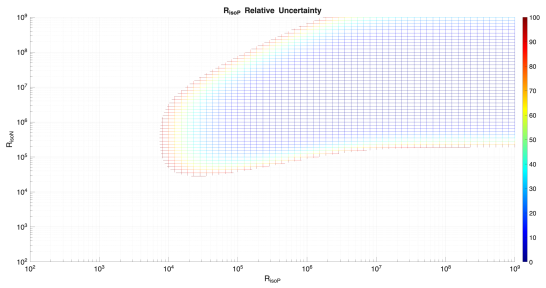


Figure 5.10: Relative uncertainty $U(R_{isoP})$ of the positive insulation estimate over the resistance plane (R_{isoN}, R_{isoP}).

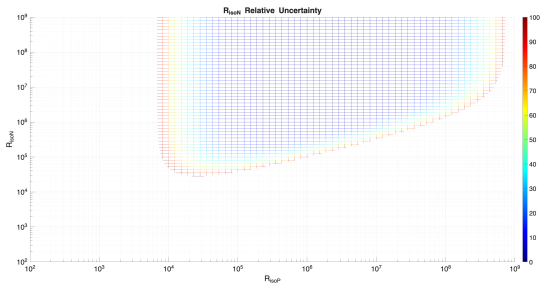


Figure 5.11: Relative uncertainty $U(R_{isoN})$ of the negative insulation estimate over the resistance plane (R_{isoN}, R_{isoP}).

for R_{isoN}, because the V_2 configuration has the extra variable R_1 inside R_{hs}; once R_{isoP} grows large, R_{hs} clamps near R_1 and the correlation between Figure 5.11 and Figure 5.9 weakens.

5.8.3 Isolation over the voltage space

Read directly in the voltage space (V_1, V_2), with the battery voltage fixed at $V_{BATT_x} = 660\,\mathrm{V}$, the estimates and their uncertainties make the input-to-insulation relationship intuitive. Figures 5.12–5.13 give the positive and negative estimates together with their 95% uncertainty and the uncertainty-corrected lower bound.

Some specifications require only an isolation *check* – a boolean from the BMS asserting that insulation is sufficient – while others want the measured value reported as well. For the project the requirement is only the boolean: a good-isolation signal that permits the contactors to close and the tram to power up. Even so, it is worth weighing how a value might be presented if it were required;

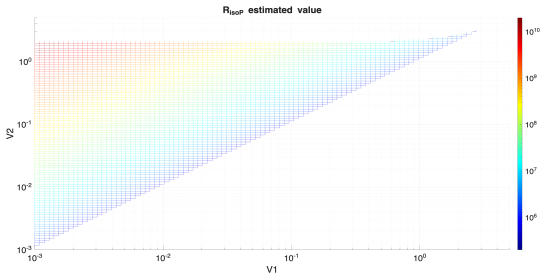


Figure 5.12: Positive insulation estimate $R_{isoP} = f(V_1, V_2)$.

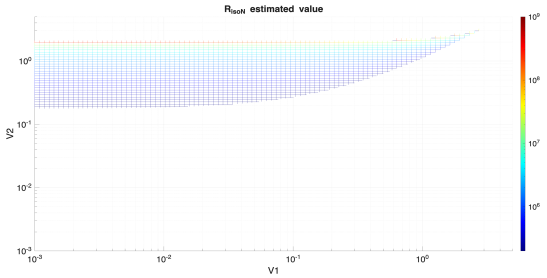


Figure 5.13: Negative insulation estimate $R_{isoN} = f(V_1, V_2)$.

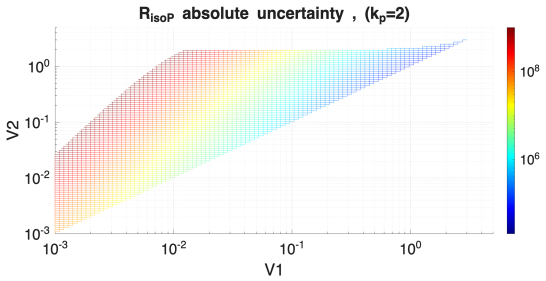


Figure 5.14: Positive uncertainty interval value $U_{Riso_P} = f(V_1, V_2)$.

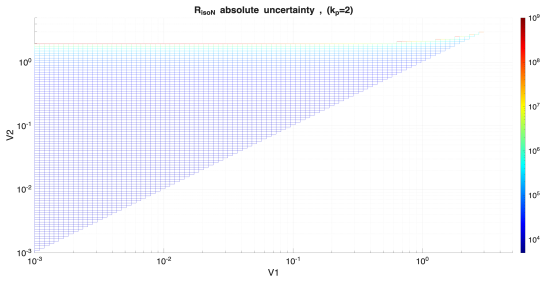


Figure 5.15: Negative uncertainty interval value $U_{Riso_N} = f(V_1, V_2)$.

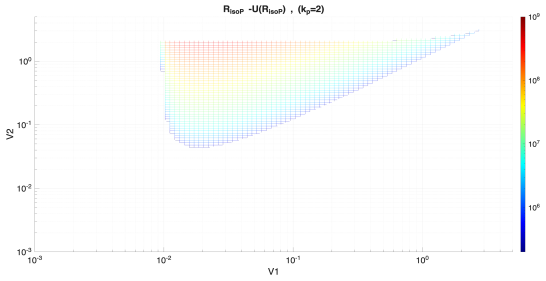


Figure 5.16: Positive resistance insulation ($k_p = 2$ uncertainty interval subtracted) $R_{isoP} - U(R_{isoP}) = f(V_1, V_2)$.

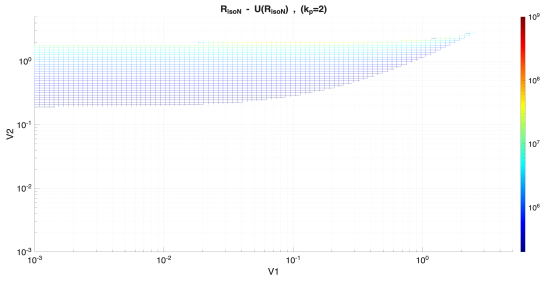


Figure 5.17: Negative resistance insulation ($k_p = 2$ uncertainty interval subtracted) $R_{isoN} - U(R_{isoN}) = f(V_1, V_2)$.

the usual approach is to report the estimate together with a confidence band, or to guarantee a fixed maximum percentage error.

5.8.4 Isolation over the virtual space

The virtual space $(vvbat_1, vvbat_2)$ is where the system is represented most completely, because it drops the battery-voltage dependency without sacrificing any validity – unlike the voltage-space view, which had to pin V_{BATT} to a constant.

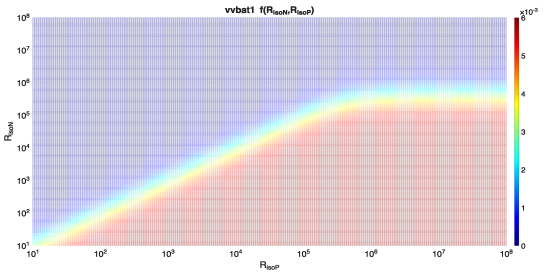


Figure 5.18: Estimated $vvbat_1$ versus the insulation resistances.

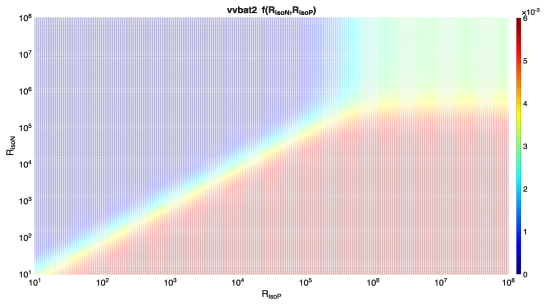


Figure 5.19: Estimated $vvbat_2$ versus the insulation resistances.

To validate the virtual representation as a basis for safety mechanisms, several quantities have to be characterised: the reachable subset of the codomain (which insulation pairs the map can actually produce), the inverse-image subset (which input pairs correspond to valid insulation), the uncertainty of the evaluation, the reduced subset where that uncertainty stays under a maximum, and finally the subset where the worse of the two insulations clears the threshold.

Reachable space. Working from the ADC input limits, only the half of the domain with $vvbat_1 < vvbat_2$ carries physical meaning, since the complementary region forces one or both insulation estimates negative. Restricting to that subdomain,

$$D_{vvbat_1} \subset D_{vvbat} \,\Big|\, \forall x \in D_{vvbat_1},\ vvbat_1 < vvbat_2,$$

the reachable insulation pairs are those of Figure 5.20.

Outcome 4 – discarding non-physical inputs All input pairs with $vvbat_1 > vvbat_2$ must be discarded: they have no physical meaning, since one or both insulation estimates come out negative.

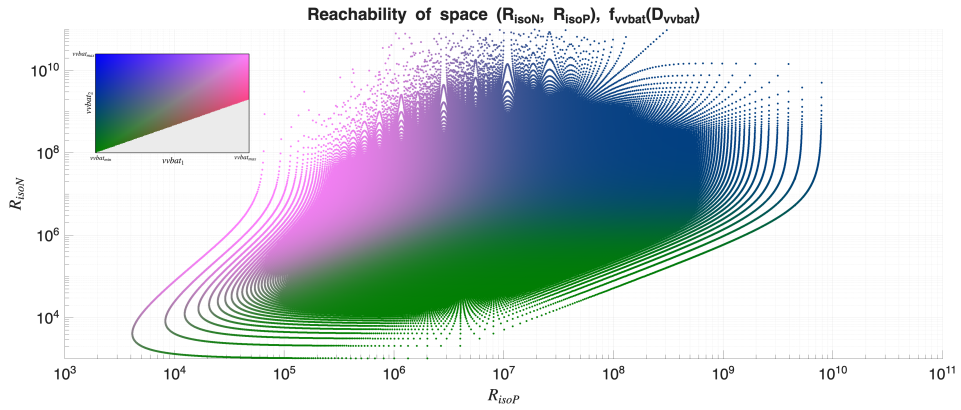


Figure 5.20: All reachable insulation pairs with positive value; colour encodes the source point in the $(vvbat_1, vvbat_2)$ domain.

Outcome 5 – where the chain can resolve The chain resolves small changes in R_{isoN}, R_{isoP} very well in the centre of its dynamic range – roughly $R_{isoN} \in [10^4, 10^9]\,\Omega$ and $R_{isoP} \in [10^5, 10^9]\,\Omega$ – but loses that ability toward the extremes, where many target states are simply unreachable and the readings poorly represent the true insulation.

Inverse image The region of the domain that yields valid insulation can also be found from the other direction, by starting from the valid subset of the co-domain and computing its preimage under the map,

$$\mathbb{I}\!\gg \left(F^{-1}_{vvbat}(R_{isoP}, R_{isoN})\right) \;\middle|\; (R_{isoP}, R_{isoN}) \in D_R.$$

The result coincides exactly with the physically-meaningful subset found above, which is a useful internal consistency check.

Safety Mechanism – safe operating area. The motivation for working in the virtual model is precisely to define a *safe operating area* that does not depend on battery voltage. The mechanism delimits a region of the $(vvbat_1, vvbat_2)$ plane within which sufficient isolation can be guaranteed for both poles, with the uncertainty already folded in.

Computing the extended uncertainty in real time on a resource-limited microcontroller is impractical, so the scheme separates the two concerns: the value reported is the uncertainty-free estimate, which is representative but not worst case, while compliance with the minimum isolation is enforced by membership in the safe operating area. Figure 5.21 shows the worse of R_{isoN} and R_{isoP} as a function of $f_{vvbat}(vvbat_1, vvbat_2)$, with the valid region marked.

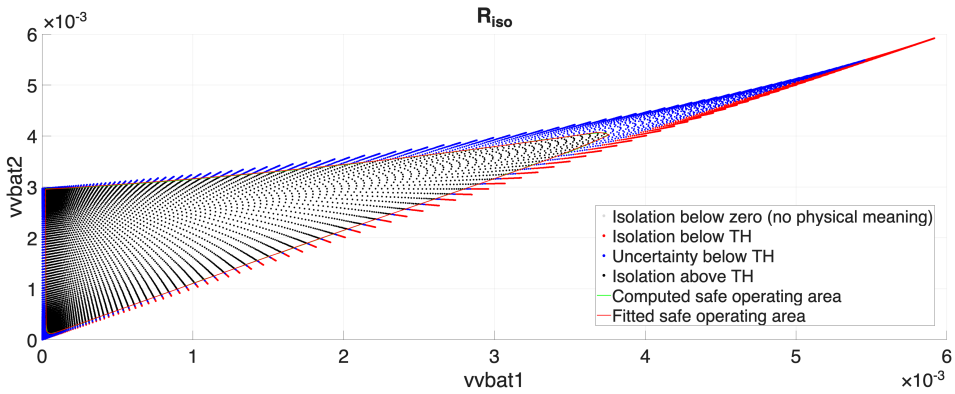


Figure 5.21: Worse of R_{isoN} and R_{isoP} over the virtual plane, with the valid points defining the safe operating area highlighted. The region is the intersection of the two preimages of the positive and negative isolation subsets, and it is what the safety mechanism uses to filter out the invalid operating areas of the $(vvbat_1, vvbat_2)$ plane.

5.9 Summary

This chapter has reconstructed the isolation measurement system from first principles. It placed the chosen unbalanced, single-channel, two-reading topology within the bigger family of isolation-measurement techniques; gave it a formal description in three equivalent coordinate systems; and built its uncertainty budget, using the resulting model to expose the topology's strengths and its limits, beyond components selection, deep in the topology itself.

Two conclusions carry into the test and validation work that follows:

1. The chain is excellent in the center of its dynamic range and where the two poles are balanced, but its uncertainty climbs steeply when the insulations diverge by more than about three orders of magnitude (Outcomes 1 and 5); and because the budget is dominated by a few contributors, the probabilistic and deterministic uncertainties nearly coincide, which removes much of the usual benefit of statistical propagation and forces an enforced ceiling on the admissible uncertainty.

2. The safe operating area is the practical response to those limits, and it is carried forward as a design input to the validation campaign.

This perimeter represents the end of the modelling process of the measurement chain, opening to the validation part into next chapter.

Chapter 6

Validation and results

The combined tests of development phase ML6 showed the computed isolation dropping to zero in some scenarios. The cause is the set of software watchdogs that prevent numerical cancellation and negative results: they react to the asymptotic behaviour of the inverse formulas, so the measurement becomes inconsistent exactly at the borders of the region found in Chapter 5. This chapter checks the chain and the safe operating area SM3 against a known reference, and uses the results to correct the area where the correction can be justified.

6.1 Test environment

The isolation measurement was validated in Hardware-in-the-Loop (HIL): the BJB runs as in the application, while a dSPACE SCALEXIO real-time system [20] (Figure 6.1) reproduces its electrical environment. This makes nominal and fault conditions repeatable without a complete pack and without provoking a real insulation failure.

Two limitations shape the results. Literature on insulation failure modes of traction batteries is scarce, so it is not known whether a real loss of insulation is typically symmetrical or confined to one pole. The HIL emulates the two insulation paths as independent channels adjustable from 5 kΩ to 5 MΩ, well below the 100 MΩ expected on a healthy pack (Section 5.4.1); a channel can also be disconnected, in which case the DUT sees the HIL's own insulation, of order 100–500 MΩ and not specified. An *open* channel is therefore a known state but not a known value, and none of the arguments below assumes a value for it.



Figure 6.1: dSPACE SCALEXIO HIL system.

6.2 Test campaign

Four failure modes combine the loss of insulation on the two poles: A, failure of R_{isoN} with R_{isoP} open; B, failure of R_{isoP} with R_{isoN} open; C, symmetrical failure; D, asymmetrical failure, run as a crossing sweep (D1) and as two one-sided sweeps with the healthy pole bonded at 5 MΩ instead of open (D2, D3). Table 6.1 gives the plan; Figure 6.2 the executed sequence.

Test	Applied R_{isoN}	Applied R_{isoP}	Points (sound / faulty)
A	5 MΩ → 5 kΩ	open	120 (74 / 46)
B	open	5 MΩ → 5 kΩ	110 (78 / 32)
C	5 MΩ → 5 kΩ	$= R_{isoN}$	247 (155 / 92)
D1	5 kΩ → 5 MΩ	5 MΩ → 5 kΩ	
D2	5 MΩ	3 MΩ → 5 kΩ	378 (178 / 200)
D3	3 MΩ → 5 kΩ	5 MΩ	
–	open	open	38 (38 / 0)

Table 6.1: Test plan. C was executed twice; the last row collects the intervals between tests. A point is faulty when either applied resistance is below R_{ISOmin}.

The BJB was logged with NXP FreeMASTER [21] and the applied bonding by the HIL, and the two logs were synchronised offline. A point is retained when the BJB publishes a new estimate while the applied pair is unchanged since the previous update, which discards the transients between settings. The multiplexed

analog input is low-pass filtered at 1 Hz, and $vvbat_1$, $vvbat_2$ are formed with the battery voltages read by the BJB.

The campaign was run with the logging set to the maximum possible speed. FreeMASTER does not guarantee constant intervals between measurements, so an ad-hoc script was responsible for synchronising the two logging devices (HIL setting parameters and FreeMASTER collecting readouts) plus a specific algorithm was developed to reduce measurements too unstable during transitions. Out of the hundreds of thousands of samples, fewer than 900 were retained as valid points: 523 sound and 370 faulty.

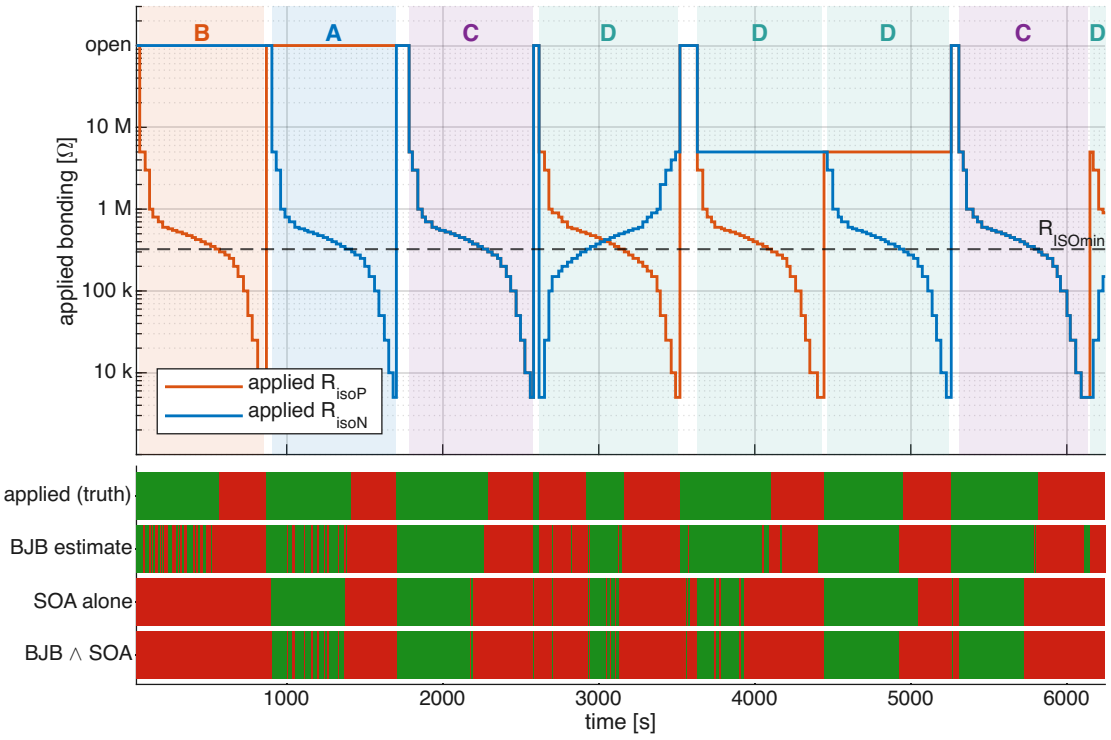


Figure 6.2: Test sequence. Top: applied bonding, tests shaded. Bottom: verdicts (green sound, red fault) from the applied values, the BJB estimate, the safe operating area alone and the two combined as in (6.1).

Every criterion is scored against the applied pair, with the insulation fault as the positive class: a missed fault (FN) is a sound verdict on a faulty point, a false alarm (FP) the opposite. They are reported over their own populations, miss rate $= FN/$faulty and false-alarm rate $= FP/$sound, and never combined, for the reasons of Section 6.3.6. From Chapter 5 (Outcomes 1 and 5) the chain is expected to struggle when the two insulations differ widely, which tests A and B produce by construction.

6.3 Results

6.3.1 Test points and device baseline

Figure 6.3 places the points in the virtual plane with the safe operating area and the boundaries of Section 5.8.4. Each test occupies a recognisable region: A lies on the left edge, where $vvbat_1$ drops to the noise floor; B and the open-open intervals lie on the upper edge, the locus $R_{isoN} \to \infty$; C and D follow the lower oblique edge, where the threshold is crossed. Some points fall beyond the image boundary, as expected from borders drawn inside the image by the uncertainty, from the unmodelled HIL insulation and from the non-idealities of the setup.

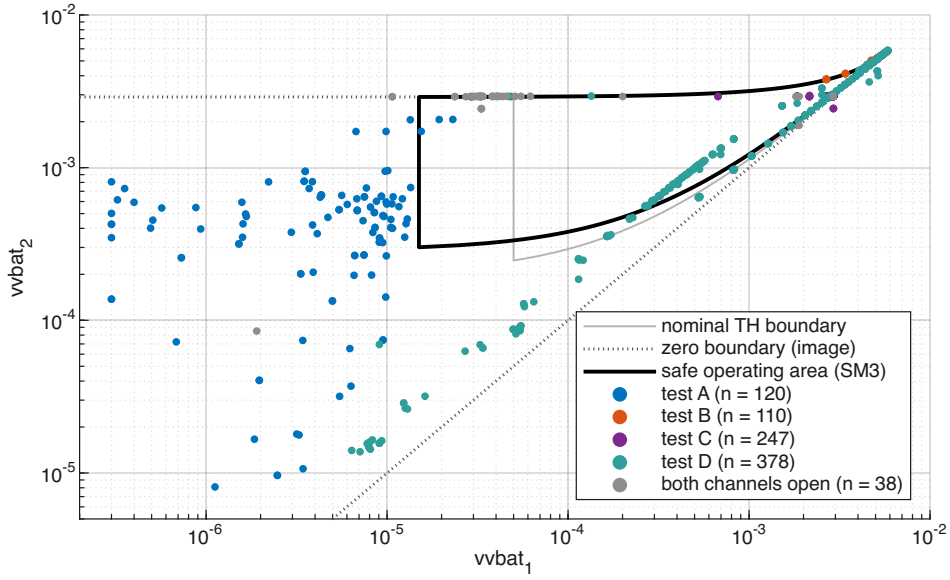


Figure 6.3: Test points against the safe operating area, the nominal TH boundary and the zero (image) boundary.

The baseline is the BJB estimate compared with R_{ISOmin}: 14 missed faults and 81 false alarms (miss rate 3.78 %, false-alarm rate 15.5 %). Every fault of A and B is detected; the false alarms fall on their healthy halves (15 in A, 41 in B) and are the symptom described above: the estimate of the open pole diverges and the watchdog clamps it to zero (10 of 15 in A, 36 of 41 in B). The missed faults all come from C and D. Ten are readings at 5 kΩ, where the device reports 180 MΩ to 1.4 GΩ; four are estimation errors, the largest 150 kΩ reported as 453 kΩ.

6.3.2 The safe operating area as specified

SM3 (Section 5.8.4) keeps the device estimate as the reported value and enforces compliance through membership of the reading in the area $\mathcal{S}$. The verdict is therefore the conjunction

$$\text{fault} \iff \hat{R}_{iso} < R_{ISOmin} \ \vee \ (vvbat_1, vvbat_2) \notin \mathcal{S}, \tag{6.1}$$

with $\hat{R}_{iso}$ the smaller device estimate. On its own $\mathcal{S}$ misses 13 faults and the device 14, but only one fault is missed by both, so (6.1) misses one. The criteria fail on different points: the device fails where its inversion leaves the model, and those readings fall outside $\mathcal{S}$; the area fails on its lower edge, where the device is right. The diversity concerns the processing only, since both use the same readings, so a fault of the front end is a common cause.

6.3.3 The lower edge

Twelve of the thirteen faults missed by $\mathcal{S}$ are in D3: $R_{isoP} = 5\,\mathrm{M\Omega}$ with R_{isoN} at 250–300 $\mathrm{k\Omega}$. The readings are not at fault, since they agree with (5.7a)–(5.7b) within 1.4–3.7 %; the edge is. The lower edge should be the locus $\min(R_{isoN}, R_{isoP}) = R_{ISOmin}$, while the area uses a polynomial approximation of it. Figure 6.4 compares the two: the edge is within 1 % of the locus near and beyond the corner ($vvbat_1 \approx 2.8 \times 10^{-3}$), and falls to -6 % at 7.7×10^{-4}, -13 % at 3.3×10^{-4} and -27 % at the left border. The discrepancy is clearest at the intercept: as $vvbat_1 \to 0$ the exact locus tends to $vvbat_2 = 3.973 \times 10^{-4}$ (Section 6.3.4), whereas the edge meets the left border at 3.01×10^{-4}. An edge below the locus accepts sub-threshold readings, so the approximation errs in the unsafe direction, and increasingly so as $vvbat_1$ decreases. The D3 readings sit exactly in that band. On its own, the area would need its lower edge refitted on the exact locus.

6.3.4 The unbonded edge: SM4

The largest group of false alarms of $\mathcal{S}$ is test A: 71 of 74 sound points are rejected, because 117 of its 120 points lie left of the area, where the inversion is useless. On that edge a further classification can still be derived, with each insulation bounded by a different reading. Take $a = R_m/(R_2 + R_m) = 5.8196 \times 10^{-3}$ and $x_L = 1.5 \times 10^{-5}$ the left border of the area; the extension covers $vvbat_1 < x_L$.

Positive insulation. Solving (5.7a),

$$R_{isoP} = R_{ls}\left(\frac{a}{vvbat_1} - 1\right), \qquad R_{ls} = \frac{R_{isoN}(R_2 + R_m)}{R_{isoN} + R_2 + R_m}. \tag{6.2}$$

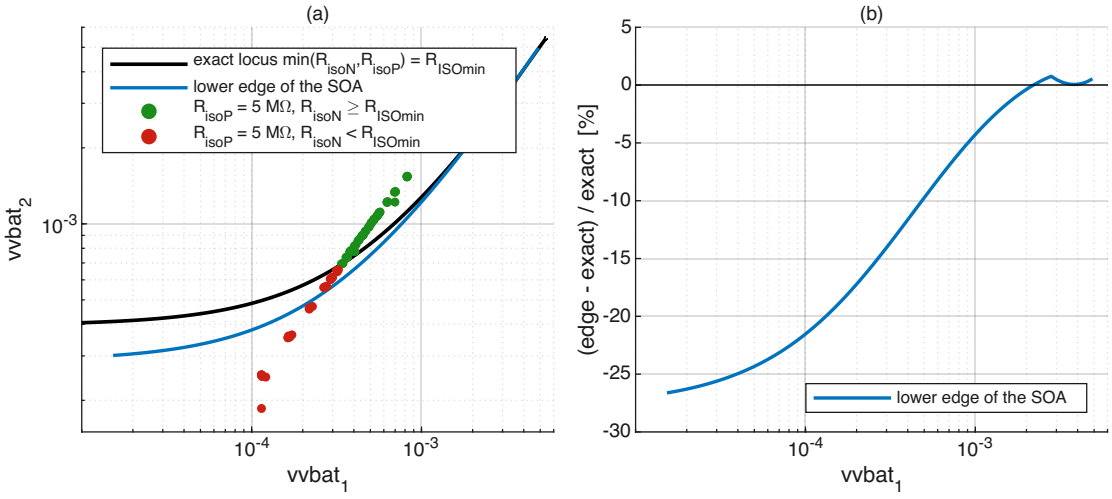


Figure 6.4: (a) Exact threshold locus against the lower edge of the safe operating area, with the D3 points. (b) Relative error of the edge.

The estimate diverges as $vvbat_1 \to 0$, but the condition $R_{isoP} > R_{ISOmin}$ does not. Solving it for R_{isoN},

$$R_{isoN} > (R_2 + R_m)\,\frac{b}{1-b}, \qquad b = \frac{R_{ISOmin}}{R_2 + R_m} \cdot \frac{vvbat_1}{a - vvbat_1}. \tag{6.3}$$

The bound increases with $vvbat_1$, so its worst case over the domain is at x_L: $R_{isoN} > 837\,\Omega$, almost 400 times below R_{ISOmin}. Equivalently, at $R_{isoN} = R_{ISOmin}$ a reading at the border guarantees $R_{isoP} \geq 1.16 \times 10^8\,\Omega$. With the channel open the first reading is noise, of the order of

$$\frac{u_{ADC}}{V_{BATTmin}} = \frac{8\,\mathrm{mV}/\sqrt{3}}{360\,\mathrm{V}} = 1.283 \times 10^{-5} < x_L, \tag{6.4}$$

and 114 of the 120 readings of A lie below this value.

Negative insulation. For $R_{isoP} \to \infty$, $R_{hs} \to R_1$ and (5.7b) depends on R_{isoN} alone:

$$\lim_{R_{isoP}\to\infty} vvbat_2 = \frac{R_{ls}}{R_1 + R_{ls}}\,a, \qquad vvbat_2\Big|_{R_{isoN}=R_{ISOmin}} = 3.973 \times 10^{-4}. \tag{6.5}$$

The readings of A follow (6.5) over three decades of bonding (Figure 6.5b). The limit is not the worst case, because the HIL insulation is finite (R_{hs} is already 4 % below R_1 at 100 MΩ). The joint inversion of (5.7a)–(5.7b), with $p = vvbat_1/a$ and $q = vvbat_2/a$, involves no R_{isoP}:

$$R_{ls} = R_1 \frac{q}{1-q}\left[1 - \frac{p\,(1-q)}{q\,(1-p)}\right], \qquad R_{isoN} = \frac{R_{ls}\,(R_2 + R_m)}{(R_2 + R_m) - R_{ls}}. \tag{6.6}$$

The bracket decreases with p, so for fixed $vvbat_2$ the implied R_{isoN} decreases as $vvbat_1$ grows: at x_L the value (6.5) corresponds to only 311.7 kΩ. Solving (6.6) for $R_{isoN} = R_{ISOmin}$ at $vvbat_1 = x_L$ gives the cut valid over the whole domain,

$$vvbat_{2,min} = 4.103 \times 10^{-4}. \tag{6.7}$$

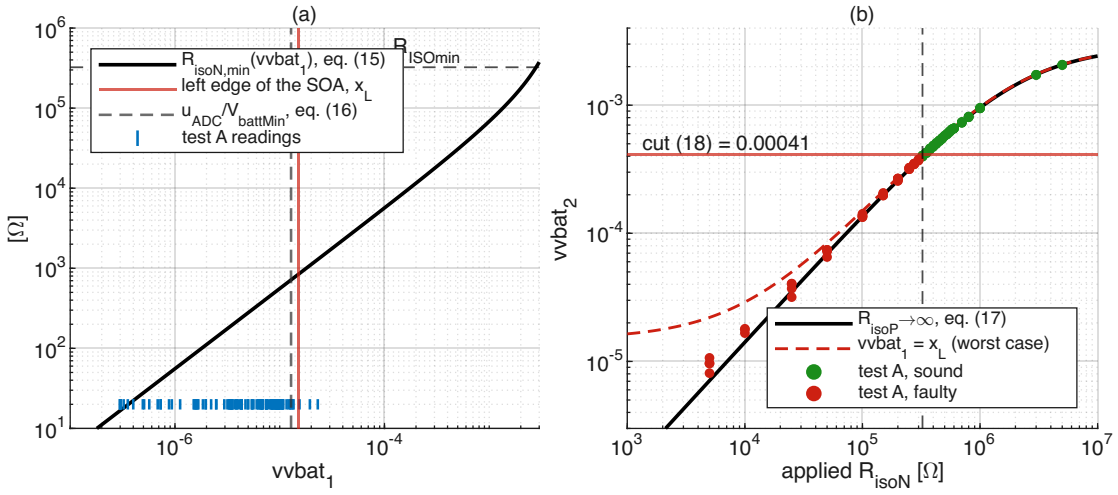


Figure 6.5: (a) Bound (6.3) versus the first reading, with x_L, the noise level (6.4) and the readings of test A. (b) Readings of A against the limit (6.5), the worst case at $vvbat_1 = x_L$ and the cut (6.7).

On the data the cut is not critical: the highest faulty reading left of x_L is 3.770×10^{-4} (300 kΩ) and the lowest valid one is at 3.965×10^{-4} (325 kΩ), so any value inside that gap separates the campaign. The cut (6.7) is adopted because it makes $R_{isoN} \geq R_{ISOmin}$ a consequence of the reading rather than an observation on this campaign. It costs seven false alarms at 325 kΩ for the area alone, and two for the combined criterion (6.1).

SM4 – validity on the unbonded edge. A reading with $vvbat_1 < x_L$ is accepted as sound when $vvbat_2 \geq vvbat_{2,min}$ (6.7). This guarantees $R_{isoN} \geq R_{ISOmin}$, and then (6.3), satisfied for any $R_{isoN} > 837\,\Omega$, guarantees $R_{isoP} \geq R_{ISOmin}$.

SM4 removes 66 false alarms of the area alone (64 of them in test A) and 55 of (6.1), without accepting any faulty point. One limitation remains: SM4 relies on the model, and the 5 kΩ readings that violate it land within a factor two to nine of x_L, at 15–71 mV, which SM1 and SM2 do not reject. None crossed the border here,

but the argument does not exclude it. In addition, the strip could be closed above at the $R_{isoN} \to \infty$ value of (6.5), 2.92×10^{-3}, which no reading exceeded.

6.3.5 Comparison

Table 6.2 adds a logistic classifier from an earlier proof of concept. Training and inference on the same samples bias the score optimistically, so the classifier is scored out of fold over five folds. Some bias remains: the features were selected on those same folds, and all folds come from one campaign whose sweep points are strongly correlated. Its figures are therefore an upper bound, whereas the other criteria are derived from the model, this model while being extremely hard to bring on the certification process, offers an interesting glimpse at what a future research might work on: ML based feature extraction with supervised or unsupervised learning.

Criterion	TP	TN	FP	FN	Accuracy	Miss rate	False-alarm rate
Logistic classifier (5-fold)	322	501	22	48	92.2 %	12.97 %	4.2 %
BJB estimate	356	442	81	14	89.4 %	3.78 %	15.5 %
$\mathcal{S}$ alone	357	266	257	13	69.8 %	3.51 %	49.1 %
$\mathcal{S}$ with SM4, alone	357	332	191	13	77.2 %	3.51 %	36.5 %
BJB $\wedge$ $\mathcal{S}$	369	261	262	1	70.5 %	0.27 %	50.1 %
BJB $\wedge$ $\mathcal{S}$ with SM4	369	316	207	1	76.7 %	0.27 %	39.6 %

Table 6.2: Scores over 893 points (523 sound, 370 faulty); $\wedge$ is the verdict (6.1).

Ranked by accuracy, the classifier comes first and the final mechanism last. Ranked by miss rate, the order reverses: one missed fault in 370 for the mechanism, 14 for the device, 48 for the classifier, which minimises a loss that weighs a missed fault and a false alarm equally.

It is important to underline that the single remaining miss is the first sample after the step to 900 kΩ / 150 kΩ in D1: its first reading is 34 % below the model, while the six following samples agree within 3 % and are rejected. It is a settling transient, and a stability filter spanning more than one update would have discarded it. The final software release foresees exactly this: a verdict is confirmed only when at least two consecutive samples agree, so that a single reading taken during a transition cannot drive the decision.

6.3.6 False alarms

The measurement runs with the contactors open, and a good result is a precondition for closing them (Section 5.4.1). A false alarm keeps the contactors open: the vehicle is not powered, at a cost in availability that is visible and recoverable.

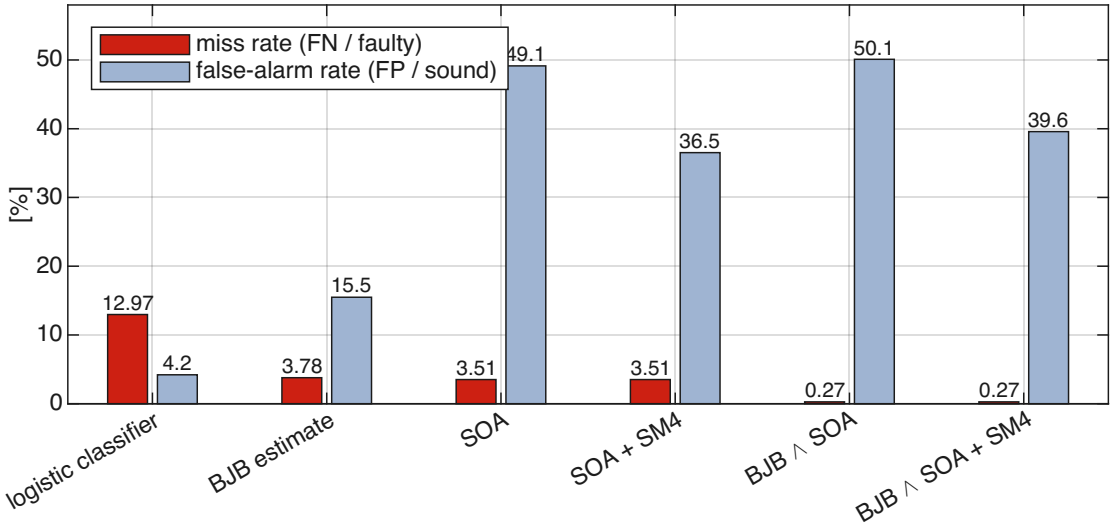


Figure 6.6: Miss rate and false-alarm rate of the criteria of Table 6.2.

A missed fault lets them close on a degraded barrier. Once V_{BATT-} is referred to the rails, a low positive insulation opens a leakage path through the chassis at full pack voltage, and nothing reports it. EN 50126 treats availability and safety as distinct RAMS attributes, and here the choice favours safety: the criterion is tuned for coverage, $1 - \text{miss rate} = 99.7\,\%$ against $96.2\,\%$ for the device.

The 207 residual false alarms have identifiable origins. 111 come from test B and the open-open intervals, on the upper edge where the area coincides with the image border $R_{isoN} \to \infty$; this is the mirror case of SM4 with the readings exchanged, and a mechanism of the same kind is the natural next step.

50 are readings of C and D at $vvbat_1 > 1.3 \times 10^{-3}$ that lie above the exact locus but are rejected by the lower edge, which there sits up to $0.8\,\%$ above it.

35 are the device's own false alarms (15 in A, 20 in C and D), 9 lie below the exact locus within the resolution of the chain, and 2 are the $325\,\text{k}\Omega$ readings rejected by the cut of SM4.

False alarm disclaimer Numbers shown in Table 6.2 might look extremely high, reflecting a measurement chain that does not measure the real isolation of the pack very accurately, but the tests were always conducted at the edge of the operating region, where the chain is at its weakest; the campaign therefore says little about its behaviour in nominal conditions, which the test plan did not cover. Test points were chosen at the edge because, even where the pack is not technically under-isolated, that is already a worrying condition as far as isolation is concerned.

So particularly with false alarm rates, the number is high particularly in the condition of Test B, where the area of validity of measurement is compressed (right

tip of the SOA triangle). The fact that a measurement falls close to that area is alone a concern even without a failure in isolation, as far as the applicable standards are concerned.

It shall also be considered that uncertainty of measurement imposes a shrinking of the area to be considered to be sure to operate into, while on the "reality side of things" uncertainty allows also for measurements expanded compared to the nominal condition, this makes for a higher rate of false alarms, natural to this kind of safety-first approach.

Outcome 6 – validation Used as specified, on top of the device estimate, the safe operating area with SM4 reduces the missed faults of the campaign from 14 to 1 (coverage 99.7 %), at a false-alarm rate of 39.6 %. The campaign exposed a lower edge that lies below the exact threshold locus, by up to 27 % at low $vvbat_1$, and a gap on the unbonded edge, closed by SM4 through the exact bounds (6.3) and (6.6)–(6.7).

Chapter 7

Conclusions and future work

This thesis presented the design, implementation and validation of the Battery Junction Box developed for a railway and automotive project inside Podium Advanced Engineering, following the V-Cycle from the system requirements down to the hardware architecture and back up through verification, in its first ISO 26262 compliant BMS.

Starting from customer and system requirements the thesis follows the development of the BJB board, responsible of pack-level voltage and current measurement, and gives particular attention to the isolation measurement chain, the pack-level safety function on which the authorization to close the contactors depends.

The experimental results obtained in Hardware-in-the-Loop, across single, symmetrical and asymmetrical insulation failures, demonstrate that the device estimate alone detects single-pole failure but tests on the battery return some concern about the assumptions made before the development phase of the product, as the system misses some faults.

The chain's model is used to re-engineer an additional criterion, namely the Safe Operating Area, an area in a virtual voltage-over-voltage domain operating as an additional condition rather than as an alternative criterion. This criterion is developed with the purpose of reducing to zero by design the false-negatives, hence the under-isolation events unnoticed by the system. This approach eliminates completely False-negatives from a precious 4%, but comes at the cost of a raise in the False-Positives, and consequentially a reduction in the overall accuracy from 89% down to 77% of the test campaign, conducted at the edge of the safety conditions (an uniformly distributed test over the full operating range of the system would return an accuracy estimated to be higher then 99.9%).

The campaign also corrected the design, exposing a boundary that lay on the unsafe side of the exact threshold locus, and a region where the inversion is undefined and every sound reading was refused. There the exact bounds of the forward model showed that a verdict remains possible without an estimate,

and led to a fourth safety mechanism. A statistical classifier trained on the same data reached the best accuracy of all criteria while missing one every eight faults, confirming that a safety function cannot be selected on a loss that weighs both errors equally.

Future work could focus on bonding the healthy pole, for which the campaign already provides evidence, on extending the same reasoning to the opposite edge of the operating area, and on the migration to a new generation of junction box silicon, more accurate, while the tests allowed to have better understanding of real world isolation values inside the battery pack, fundamental for a future version of the isolation measurement system.

While simple, the logistic classifier opened to great possibilities and demands a in-depth analysis of the tool.

Overall, this work contributes to the development of safe battery management systems by carrying a safety function through a complete V-Cycle and exploiting its defining strength: the evidence gathered in validation returns as revised system requirements, closing the cycle upon itself.

Bibliography

[1] *ISO 26262-1:2018 – Road vehicles – Functional safety – Part 1: Vocabulary*. International Organization for Standardization. 2018. URL: https://www.iso.org/standard/68383.html (cit. on pp. 3, 22).

[2] *EN 50126-1:2017 – Railway Applications – The Specification and Demonstration of Reliability, Availability, Maintainability and Safety (RAMS) – Part 1: Generic RAMS Process*. CENELEC. 2017 (cit. on pp. 3, 22, 50).

[3] NXP Semiconductors. *MC33774 – 18 Channel Li-Ion Battery Cell Controller IC ASIL D*. URL: https://www.nxp.com/products/MC33774 (cit. on p. 9).

[4] NXP Semiconductors. *NXP Improves Battery Health Monitoring with EIS-Capable Battery Management Chipset*. https://www.globenewswire.com/news-release/2025/10/29/3176157/0/en/NXP-Improves-Battery-Health-Monitoring-with-EIS-Capable-Battery-Management-Chipset.html. 2025 (cit. on p. 12).

[5] Sensata Technologies. *PyroFuse: Pyrotechnic Circuit Breakers for HV Systems (whitepaper)*. https://www.sensata.com/sites/default/files/a/sensata-el-ev-pyrofuse-whitepaper.pdf (cit. on p. 13).

[6] NXP Semiconductors. *UM12056: RDA777T2 battery junction box reference design, Rev. 1.0*. https://www.nxp.com/docs/en/user-manual/UM12056.pdf. 2025 (cit. on p. 14).

[7] NXP Semiconductors. *MC33772C – 6-Channel Li-Ion Battery Cell Controller IC*. URL: https://www.nxp.com/products/MC33772C (cit. on pp. 14, 29, 48).

[8] NXP Semiconductors. *MC33777: Battery Junction Box Monitor IC*. https://www.nxp.com/products/MC33777 (cit. on p. 15).

[9] Analog Devices. *ADBMS2950/2951/2952 Battery Pack Monitors – Data Sheet, Rev. SpC*. https://www.analog.com/media/en/technical-documentation/data-sheets/adbms2950-2951-2952.pdf (cit. on pp. 17, 19).

[10] Analog Devices. *ADBMS2950 Battery Pack Monitor – Product page*. https://www.analog.com/en/products/adbms2950.html (cit. on p. 17).

[11] Isabellenhütte. *IVT-S – Current, voltage and temperature measurement system*. https://www.isabellenhuette.com/solutions/products/ivt-s (cit. on p. 18).

[12] NXP Semiconductors. *BMA8420 – Battery Junction Box Monitor IC with EIS Capability*. URL: https://www.nxp.com/products/BMA8420 (cit. on p. 19).

[13] *EN 50128:2011 – Railway applications – Communication, signalling and processing systems – Software for railway control and protection systems*. CENELEC. 2011 (cit. on p. 24).

[14] *EN 50129:2018 – Railway applications – Communication, signalling and processing systems – Safety related electronic systems for signalling*. CENELEC. 2018 (cit. on p. 24).

[15] The MathWorks, Inc. *System Composer*. URL: https://www.mathworks.com/products/system-composer.html (cit. on p. 29).

[16] *MC33772C Battery cell controller IC – Product data sheet, Rev. 4.1*. NXP Semiconductors. 2026. URL: https://www.nxp.com/docs/en/data-sheet/MC33772C-DS.pdf (cit. on pp. 42, 43, 59).

[17] Bender GmbH & Co. KG. *ISOMETER iso175 – Insulation monitoring device*. URL: https://www.bender.de/en/products/insulation-monitoring/isometer-iso175/ (cit. on p. 46).

[18] *JCGM 100:2008 – Evaluation of measurement data – Guide to the expression of uncertainty in measurement (GUM)*. Joint Committee for Guides in Metrology. 2008. DOI: 10.59161/JCGM100-2008E. URL: https://www.bipm.org/documents/20126/2071204/JCGM_100_2008_E.pdf (cit. on p. 48).

[19] *TPSI2140-Q1 1200-V, 50-mA, Automotive Isolated Switch With 2-mA Avalanche Rating – Data sheet*. Texas Instruments. URL: https://www.ti.com/product/TPSI2140-Q1 (cit. on p. 50).

[20] dSPACE GmbH. *SCALEXIO – Modular real-time platform*. URL: https://www.dspace.com/en/pub/home/products/hw/simulator_hardware/scalexio.cfm (cit. on p. 70).

[21] NXP Semiconductors. *FreeMASTER Run-Time Debugging Tool*. URL: https://www.nxp.com/design/design-center/software/development-software/freemaster-run-time-debugging-tool:FREEMASTER (cit. on p. 71).