Development and Implementation of Test Infrastructure and its Application during the Development and Prototyping of Service Hybrids for the Phase-2 Upgrade of the CMS Outer Tracker

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

M.Sc. Alexander Pauls

aus

Eschweiler

Berichter: Univ.-Prof. Dr. rer. nat. Lutz Feld
Univ.-Prof. Dr. rer. nat. Jörg Pretz

Tag der mündlichen Prüfung: 09.12.2024

Abstract

The Large Hadron Collider (LHC) at CERN will be upgraded to the HL-LHC to increase its instantaneous luminosity to more than $5 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The new Outer Tracker will replace the strip tracker of the CMS detector as part of the Phase-2 upgrades that prepare the experiment for future data-taking at the HL-LHC. The beginning of the Outer Tracker installation is planned for the year 2026 and the Outer Tracker will be able to provide tracking information to the first level of the CMS trigger.

This thesis is concerned with one of the two module types used in the Outer Tracker: the 2S module with two silicon strip sensors. The module electronics are implemented on dedicated PCBs, so-called hybrids. The Service Hybrids connect the 2S modules with the services outside the detector. They also feature two DC-DC converters and a serializer ASIC and opto-electrical converter for multi-gigabit electrical and optical communication. The silicon sensor bias voltage is also filtered and distributed by the Service Hybrid. The hybrid design is challenging in terms of material and component selection and mechanical constraints. The application in the harsh environment of the tracking detector requires thorough and extensive qualification and testing of all detector components.

The scope of this work covers the development and prototyping of the Service Hybrids and their test infrastructure as well as the transition to and preparation for the series production. The author qualified several Service Hybrid versions, made contributions to the development of the test system, and commissioned new testing hardware. In addition, studies of the Service Hybrids' properties as a function of temperature and under the influence of a magnetic field were performed. The specifics of the different hybrid versions and their production are summarized in this thesis. Failures are discussed and the collected measurement data are analyzed.

The obtained results benefited the research and development of the Phase-2 Outer Tracker. The contributions provided functional Service Hybrid prototypes for vital activities of the upgrade project, confidence in the Service Hybrid performance in the future detector, and an understanding of encountered failures.

Zusammenfassung

Der Large Hadron Collider (LHC) am CERN wird zum HL-LHC umgebaut um seine instantane Luminosität auf über $5 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ zu erhöhen. Als Teil der Phase-2 Upgrades, die den CMS-Detektor auf die Datennahme am HL-LHC vorbereiten, wird der Streifen-Spurdetektor des Experimentes durch den neuen Outer Tracker ersetzt. Der Einbau des Outer Tracker, der Spurinformationen für die erste Triggerstufe liefern kann, ist für das Jahr 2026 geplant.

Betrachtet wird in dieser Arbeit eine der zwei im Outer Tracker verwendeten Modularten: das 2S-Modul mit zwei Siliziumstreifensensoren. Die Modulelektronik befindet sich auf speziellen Platinen, die Hybride genannt werden. Die Service Hybride verbinden die Module mit der Auslese- und Versorgungstechnik außerhalb des Detektors. Sie haben zwei DC-DC-Konverter und für elektrische und optische Multi-Gigabit-Auslese einen ASIC zur Serialisierung und einen Opto-Elektrischen-Konverter. Der Service Hybrid verteilt und filtert zudem die Verarmungsspannung der Siliziumsensoren. Das Hybriddesign ist herausfordernd aufgrund der Anforderungen an die Komponenten und Materialien und der mechanischen Einschränkungen. Alle Detektorkomponenten bedürfen wegen ihrer Anwendung in der anspruchsvollen Umgebung eines Spurdetektors einer gründlichen und ausführlichen Qualifizierung und Prüfung.

Diese Arbeit behandelt die Entwicklung und Prototypisierung der Service Hybride und ihrer Testumgebung sowie den Übergang und die Vorbereitung für die Serienfertigung. Mehrere Versionen der Service Hybride wurden durch den Verfasser qualifiziert, er lieferte Beiträge zur Entwicklung des Testsystems und nahm neue Testhardware in Betrieb. Außerdem wurden Studien des Verhaltens von Service Hybriden in Abhängigkeit des Magnetfeldes und der Temperatur durchgeführt. Die Besonderheiten der verschiedenen Hybridversionen und ihrer Produktion werden in dieser Arbeit zusammengefasst. Ausfälle werden diskutiert und die genommenen Daten analysiert.

Die gewonnenen Erkenntnisse haben die Entwicklung des Phase-2 Outer Tracker vorangebracht. Die Beiträge lieferten funktionale Service Hybrid-Prototypen für wichtige Anwendungen im Upgradeprojekt, bestätigen die Eignung der Service Hybride für den zukünftigen Detektor und identifizierten Fehlerursachen.

Contents

Introduction

The discovery of the Higgs boson by the ATLAS and CMS Collaborations in 2012 completed the elementary particle content of the Standard Model of particle physics. An abundance of studies at the LHC confirmed the Standard Model's precision and predictive power. However, many open questions remain unanswered, such as the description of gravity, the nature of Dark Matter, or the origin of the matter-antimatter asymmetry. An increasing amount of collision data is required to gain access to rare and statistically limited processes and channels. Since the center-of-mass energy cannot be increased much further, a higher instantaneous luminosity is driving the discovery potential at higher masses.

The CERN accelerator complex and the detectors at the LHC will therefore undergo a major upgrade program to maintain their leading role in the field and their physics potential. The LHC upgrade to the HL-LHC with a higher instantaneous luminosity increases the particle density and level of radiation the detectors need to withstand. A higher detector granularity, improved trigger capabilities, and more readout bandwidth of the detector will allow the CMS Collaboration to profit from the accelerator upgrade. The Phase-2 Outer Tracker upgrade will fully replace the current CMS Phase-0 strip tracker. A major driver for the new Outer Tracker design is the goal to provide tracking information for each collision to the first trigger level. It requires the rejection of low transverse momentum charged particle tracks by the detector itself. The Outer Tracker is therefore structured into modules with two closely spaced sensors. The correlation of hits on both sensors allows the rejection of low-momentum particles on-module for the data supplied to the first trigger level. Still, the necessary bandwidth imposes the usage of a multi-gigabit optical link on every Outer Tracker module. The power consumption of the tracker increases by about a factor of two and requires a power distribution with a two-step DC-DC conversion scheme.

One of the two module design variants that will be used in the Outer Tracker is the 2S module. It has two silicon strip sensors. The printed circuit board that hosts all connections of the 2S module to the detector services is the Service Hybrid. It also hosts the lpGBT, a multi-gigabit transceiver ASIC, the VTRx$^+$ module, which is an opto-electrical converter, and two DC-DC converter ASICs, bPOL12V and bPOL2V5.

Within the scope of this work, important milestones of the prototyping phase of the Phase-2 Outer Tracker upgrade were achieved. The I. Physikalisches Institut

B of RWTH Aachen University designed and produced the first prototypes of the Service Hybrid. Tight geometric and material restrictions, multi-gigabit data lines, and the new implementation of a two-step DC-DC conversion scheme required a comprehensive development and prototyping phase. The Service Hybrid design does not allow for redundancy and is essential to the module functionality. It necessitated extensive evaluations of the design and the prototypes. A quality control and quality assurance program during series production is mandatory for the upgrade's success. The Service Hybrid test system evolved in parallel to the hybrid development and hybrids were qualified and distributed for module assembly. The first Service Hybrid design with the final ASICs was manufactured by CERN, which also set up a quality control and assurance program for the production of the required hybrids. The RWTH Aachen University contributed to the project by commissioning and testing the prototype hybrids and via the development of dedicated testing hardware and software. Additional studies of the Service Hybrid response to the presence of a magnetic field and to temperature changes, as well as the hybrids' influence on the 2S module strip noise, were performed. The efforts covered in this thesis span more than five years and also include hybrids from a kick-off and pre-series production run that followed the prototyping phase. Test hardware in sufficient numbers for the full Service Hybrid production was tested in Aachen, delivered to CERN, and installed at CERN.

The Phase-2 Outer Tracker upgrade is a large project with many contributing institutes, scientists, engineers, and students. Naturally, certain aspects discussed in this thesis that are relevant to its context are the outcome of the Collaboration's effort. Where relevant, such content is indicated as well as the author's contributions.

The author's contributions cover a wide range starting from performing measurements on Service Hybrids with new and preexisting setups, over investigations of hardware failures, design and commissioning of new test hardware and systems, to writing software for data-taking and analysis. These efforts required a deep understanding of the requirements and specifications of the 2S module's hybrids and the respective ASICs, the complex software framework for the Outer Tracker, the testing hardware, and the developments in the upgrade project.

This thesis has the following structure: Chapter 1 is an introduction to the LHC, the CMS detector, and the Phase-2 Outer Tracker upgrade. After the context is provided, more detailed information is given on the core subjects of this work, namely the 2S module, the Service Hybrid, and the test system for the hybrid production. The first chapter focuses on the final design, while Chapter 2 presents the variants of prototypes that were manufactured during the research and development phase of the Outer Tracker. Technical details of the different designs and their ASICs are provided, too. Chapter 3 focuses on the testing hardware for the first Service Hybrid prototypes that supported an optical readout and the hybrids' performance. The design of the 2S-SEH test card that is used for qualification during the Service Hybrid production is presented in Chapter 4. The test routines, the card commissioning,

and the production of a larger quantity are discussed. The test cards were used to assess batches of the final Service Hybrid prototypes. The results are summarized in Chapter 5. Here, the influence of different temperatures is studied, too. The performance of a final Service Hybrid prototype was also tested in the magnetic field of the M1 magnet at CERN. Chapter 6 covers the test results. Chapter 7 provides information on the Service Hybrids that were produced in the scope of the hybrid production kick-off. Two different designs and properties that have an impact on the 2S module strip noise were studied. As the last part of this thesis hybrids from the pre-series were evaluated with the test card, a climatic chamber, and the production test setup at CERN. The analysis of the data is given in Chapter 8.

1 Introduction to the Phase-2 Upgrade of the CMS Outer Tracker

1.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [1] is the most powerful particle accelerator in the world. It is operated by the European Organization for Nuclear Research (CERN[1]). The accelerator is installed in a 27 km long tunnel that is 45 to 170 m deep underneath the Swiss-French border region near the city of Geneva. The tunnel previously housed the Large Electron-Positron Collider which was in operation from 1989 to 2000. The LHC's two counter-rotating proton beams currently reach an energy of 6.8 TeV. Their collisions have a center-of-mass energy of up-to 13.6 TeV and a typical instantaneous luminosity of about $2 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [2]. The bunch crossing frequency of 40 MHz corresponds to a minimal time difference of 25 ns between collisions. The LHC can also collide heavy ions such as lead during special physics runs.

The acceleration chain of the protons starts with hydrogen atoms. They are first negatively ionized and then injected into the Linear Accelerator 4, where the hydrogen anions reach an energy of 160 MeV. After the electrons are stripped from the protons, the subsequent Proton Synchrotron Booster, the Proton Synchrotron, and the Super Proton Synchrotron accelerate them to 450 GeV. The LHC's radio frequency (RF) cavities accelerate them to the final beam energy. A total of 1232 15 m long dipole magnets with a magnetic field of up to 7.74 T and more than 400 quadrupole magnets guide 2808 bunches through the ring and focus them at the interaction points [3].

The LHC beams are collided at four interaction points. Each of them features one of the four large experiments. The ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) experiments are multi-purpose detectors. With their almost hermetic design, they can perform precision measurements of standard model (SM) parameters, studies of the Higgs boson, and search for new physics beyond the Standard Model. The LHCb (Large Hadron Collider beauty) experiment is a one-

[1] Conseil européen pour la recherche nucléaire

arm forward spectrometer that performs b-hadron measurements and investigates charge-parity-violation and rare decays. The ALICE (A Large Ion Collider Experiment) experiment focuses on heavy ion physics. The LHC was approved in 1994 and commissioned with the first beam in 2008. The LHC Run 1 started in 2009 with the first collisions. By its conclusion in 2012, about $30\,\mathrm{fb}^{-1}$ total integrated luminosity of proton-proton collision data were delivered to the CMS experiment. The center-of-mass energy was 7 and later 8 TeV. Run 1 was followed by a first Long Shutdown (LS1). The Run 2 (2015-2018) delivered $165\,\mathrm{fb}^{-1}$ for data-taking at 13 TeV. LS1 and Run 2 were used to install the first generation of detector upgrades, referred to as the Phase-1 upgrade program. After the second long shutdown (LS2, 2019-2021), the third LHC run started in 2022. Until the end of 2025, $250\,\mathrm{fb}^{-1}$ of integrated luminosity is planed to be be delivered at the center-of-mass energy of 13.6 TeV [4].

1.2 The CMS detector

One experiment to study the LHC collisions is the CMS detector. It is located on the French side of the LHC ring in the village Cessy, at the so-called Point 5. By design, it can study various physics phenomena. This includes the discovery of the Higgs boson in 2012 [5], precision measurements of the SM, and an abundance of searches for new physics beyond the SM [6].

A schematic drawing of the CMS detector (after the Phase-1 upgrade) is shown in Figure 1.1. In the following, the state of the detector during the LHC Run 3 is described. A full description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, is reported in Reference [7].

The CMS detector is 22 m long, has a diameter of 15 m, and weighs 14,000 tons. The nearly hermetic detector can trigger on and identify electrons, muons, photons, and (charged and neutral) hadrons. A superconducting solenoid with 6 m internal diameter and a length of 12.5 m provides a magnetic field of 3.8 T. A silicon pixel and strip tracker, a lead tungstate crystal electromagnetic calorimeter (ECAL), and a brass and scintillator hadron calorimeter (HCAL) are placed inside the magnet. Each of these sub-detectors has a barrel and two endcap sections. The pseudorapidity ($|\eta|$) coverage provided by the barrel and endcap detectors is extended by forward calorimeters. Gas-ionization detectors embedded in the steel flux-return yoke outside the solenoid measure muons.

A two-tier trigger system selects events of interest. Information from the muon detectors and the calorimeters is used by the first-level (Level-1 or L1) trigger. It selects events within a latency of 4 µs. The second level is the high-level trigger (HLT). It consists of a cluster of commercial processors that run a version of the full event reconstruction software. The Run 3 L1 trigger operates at a typical output rate of 110 kHz and the HLT runs at 5 kHz [9]. A particle flow approach is used for

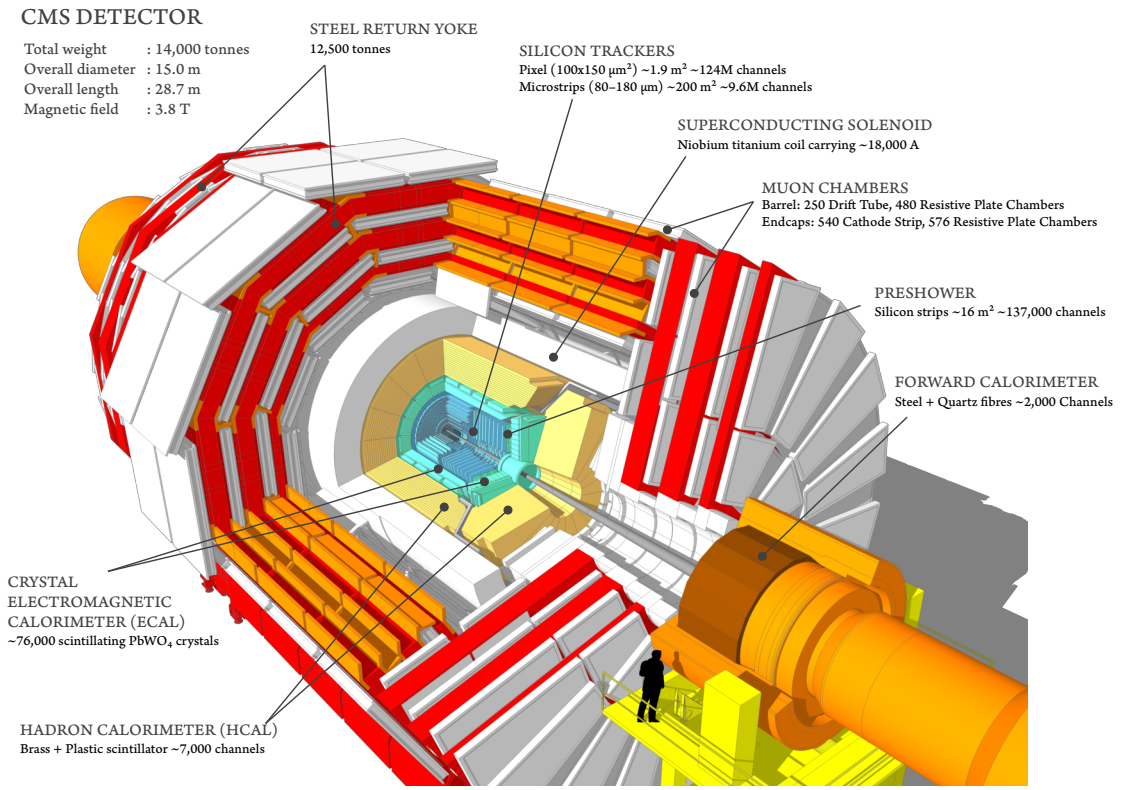

Figure 1.1: Illustration of the CMS detector after the installation of the Phase-1 upgrade [8].

the reconstruction of particles. They are identified by software that combines the information from the various detector components.

Charged particle trajectories are measured by the two innermost parts of the CMS detector. The pixel detector is closest to the interaction point. It has four barrel layers, three endcaps disks per side, and a total of 124 million readout channels. The strip tracker has ten barrel layers and nine disks per endcap. It has 9.3 million silicon micro-strips with a hit resolution of 20 µm. The coverage in $|\eta|$ extends to 2.5.

The ECAL identifies electrons and photons and measures their positions and energies. It consists of about 76,000 lead tungstate crystals ($PbWO_4$) and covers the range of $|\eta| < 3$. The crystals have lengths of about 23 cm, equivalent to 25 radiation lengths X_0. Their scintillation light is collected by avalanche photodiodes. The resolution for electrons from Z boson decays with a transverse momentum p_T of about 45 GeV ranges from 1.6% to 5%. The ECAL is also important for jet reconstruction, the measurement of the missing transverse momentum, and the determination of the electron and photon arrival time. Two-layered preshower detectors with 3 X_0

lead absorbers sit in front of the endcap ECAL. They reduce the rate of π^0 mesons falsely identified as photons.

The HCAL has a barrel, an endcap, an outer, and a forward section. The barrel and endcap sections surround the ECAL inside the magnet. They consist of brass plate absorbers that are interleaved with scintillating tiles. The readout is performed with silicon photomultipliers. The forward section outside the magnet increases the calorimeter coverage to $|\eta| < 5$. Its photomultiplier tubes read out quartz fibers in between steel absorbers. The outer section is made of plastic scintillators and the first layer of the iron flux-return yoke. The HCAL measures the energy of charged and neutral hadrons. The particle flow algorithm combines the information from the entire CMS detector to form jets. The jet energy resolution is about 10% at 100 GeV.

The muon detectors measure the momentum of muon tracks and are embedded in the flux-return yoke. The four sub-systems with a total of almost one million channels are all gas ionization detectors. The four different types are the drift tubes (DTs), the cathode strip chambers (CSCs), the resistive-plate chambers (RPCs), and the gas electron multiplier (GEM) detectors. The DTs are placed in the barrel region and cover up to $|\eta| < 1.2$ They are composed of drift chambers with rectangular cells. The CSCs cover the region $0.9 < |\eta| < 2.4$ in both endcaps. They consist of multi-wire proportional chambers with cathode strips. Both the barrel and endcap regions feature RPCs, which complement DTs and CSCs with a fast response time to identify unambiguously the bunch crossing corresponding to a muon trigger candidate. The GEM detectors are located in front of the inner ring of CSCs. They have a fast response and good spatial resolution [10].

The BRIL (Beam Radiation, Instrumentation, and Luminosity) project controls the background conditions of the CMS detector. It provides information on the beam and the radiation environment. In case of dangerous beam loss events, it has a beam dump functionality. Dedicated equipment delivers precise luminosity measurements [10].

1.3 The High Luminosity LHC and the CMS Phase-2 upgrades

The accelerator complex and the experiments undergo a constant evolution. Figure 1.2 gives an overview of the LHC schedule. The schedule also includes the time after the conclusion of Run 3. At the end of 2025, the LHC will have delivered about $400\,\text{fb}^{-1}$ at up to 13.6 TeV center-of-mass energy. In the following three years, the accelerator will be upgraded to the High Luminosity LHC (HL-LHC). The HL-LHC project started in 2011 and has a budget of about 1.1 billion Euros.

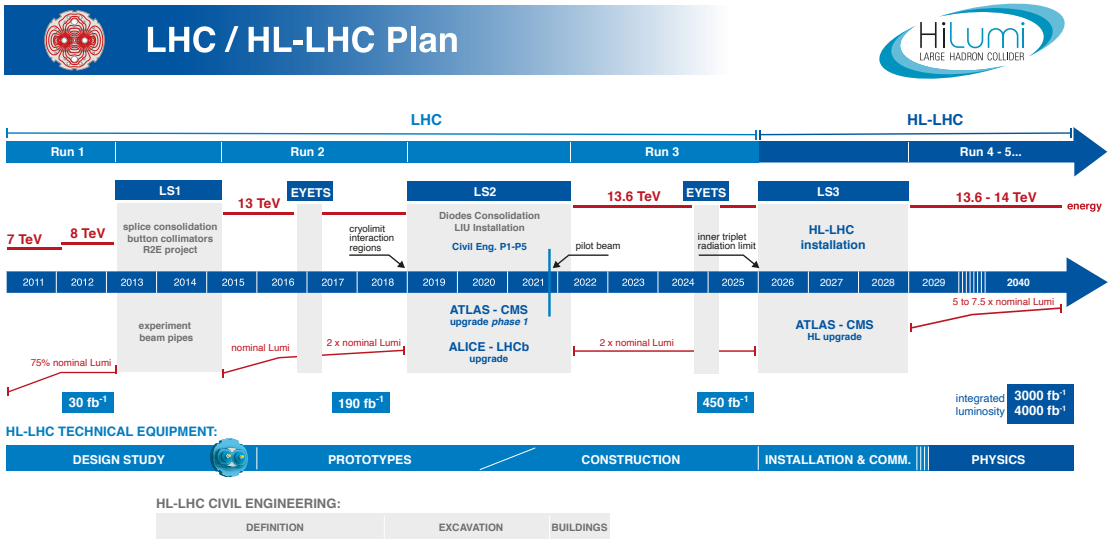

Figure 1.2: LHC and HL-LHC schedule. It provides information on the years and center-of-mass energies of past and future runs. The long shutdowns for upgrades of the instruments and the detectors are also shown [11].

A major part of the upgrade is the installation of Nb_3Sn focusing quadrupole magnets. They are responsible for beam focusing and are referred to as the inner triplet magnets. The magnets increase the beam aperture from 70 to 150 mm and allow a reduction of β^* by a factor of four. To limit the length increase to 9 m for a total magnet length of 32 m their magnetic field is increased to 10 T. The acceptance criteria for produced magnets is a nominal current of 16530 A at 1.9 K.

Additional underground galleries are installed 10 m above the existing LHC tunnel. They are accessible during beam operation and house the new infrastructure, powering, and protection equipment. The 19 magnets require up to 120 kA of DC (direct current). Several 80-100 m long superconducting cable links connect the accelerator magnets to their power converters. They are based on MgB_2 and ensure an efficient power connection.

The second important part of the upgrade is the installation of superconducting crab cavities. They will compensate for the adverse effect of the crossing angle on the luminosity. RF dipoles and double-quarter wave cavities deflect bunches in the horizontal and vertical crossing planes, respectively.

When the HL-LHC begins operation in 2029, the crab cavities will only be started to be exploited for the first year, β^* will be as low as 30 cm, and the bunch intensity of 1.8×10^{11} particles per bunch will be the same as at the end of Run 3. At the end of Run 4 in 2032, the intensity will have doubled and the β^* reduced to 20 cm.

Year	ppb [10^{11}]	Virtual lumi. [$10^{34}\text{cm}^{-2}\text{s}^{-1}$]	Days in physics	θ [μrad]	β^*_{start} [cm]	β^*_{end} [cm]	CC	Max. PU
2029	1.8	4.4	90	380	70	30	exp	116
2030	2.2	9.7	120	500	100	30	on	132
2031	2.2	11.3	160	500	100	25	on	132
2032	2.2	13.5	160	500	100	20	on	132
2033-34			Long shutdown 4					
2035	2.2	13.5	140	500	100	20	on	132
2036	2.2	16.9	170	500	100	15	on	132
2036	2.2	16.9	200	500	100	15	on	200

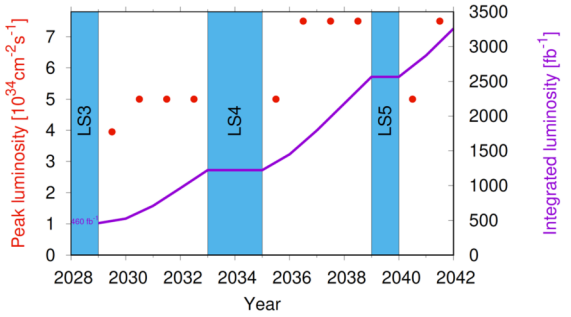

Figure 1.3: Left: intensity and performance ramp-up in Run 4 and Run 5 of the HL-LHC era. Right: prospected evolution of integrated luminosity in the HL-LHC era [11].

Figure 1.3 summarizes the operating conditions until 2036 (left) and the prospected evolution to 3000 fb^{-1} at the end of Run 6 in 2042 [4, 11]. The virtual luminosity is the theoretical luminosity without the leveling reduction factors that are achieved by the use of the crab cavities. The CMS experiment will experience an instantaneous luminosity of about $5 \times 10^{34}\,\text{cm}^{-2}\text{s}^{-1}$ and as many as 200 simultaneous collisions within the same bunch crossing (pileup).

The HL-LHC operating conditions expressed in Figure 1.3 by far exceed the design parameters of the present CMS detector. Most notably the instantaneous luminosity, the pileup conditions, and the expected radiation damage pose challenges. The third long shutdown (LS3) will be used to install major upgrades of almost all sub-detectors. The hardware and software for readout, triggering, and detector control will also be renewed. The detector upgrade is commonly referred to as the Phase-2 upgrade. The following paragraphs summarize the Phase-2 upgrade efforts for the most important sub-systems. The Phase-2 Outer Tracker upgrade is discussed separately and in more detail in the next section.

1.3.1 Inner Tracker

The innermost sub-detector of the CMS experiment is composed entirely of layers of pixelated silicon sensors and is often just referred to as the pixel detector.

The present pixel detector was installed in 2017 as part of the Phase-1 upgrade. It has about 124 million readout channels and an active area of 1.9 m^2. The size of each pixel is $100 \times 150\,\mu\text{m}^2$. The upcoming HL-LHC enforces high requirements on the pixel detector and makes a detector replacement necessary [12]. The new Inner Tracker (IT) must withstand a radiation dose of about 1 Grad and a fluence[2] of $2 \times 10^{16}\,\text{n}_{\text{eq}}\text{cm}^{-2}$. The proposed design features an increased granularity for better

[2]The radiation damage caused in silicon depends on the type of particle and its energy. The equivalent fluence of neutrons with a kinetic energy of 1 MeV is used as a reference.

track separation in the high-pileup conditions of the HL-LHC. It also has a reduced material budget and extends the tracking acceptance to $|\eta| \leq 4$. A high bandwidth is required to cope with a hit rate of up to $3.5\,\mathrm{GHz/cm^2}$.

The detector is separated into three sub-detectors: The barrel (TBPX), the forward disks (TFPX), and the endcaps (TEPX). It has a total active silicon area of $4.9\,\mathrm{m^2}$. The mechanical structure is designed for easy installation and removal during a technical stop at the end of a year. The main support is made of carbon fiber with integrated cooling pipes for CO_2 cooling.

All pixels are rectangular with a size of $25 \times 100\,\mu\mathrm{m^2}$. The planar n-in-p type sensors have a bulk thickness of $150\,\mu\mathrm{m}$. The innermost barrel layer has 3D pixels that can be depleted with a lower bias voltage and give more margin in terms of thermal stability. They are slightly less efficient (98%) than the planar sensors (> 99%) used in the rest of the detector. The sensors are read out by the CROCv2 chip which was designed by the RD53 Collaboration [13]. It has a pixel size of $50 \times 50\,\mu\mathrm{m^2}$ which requires a special connection scheme for the rectangular sensor pixels. It is produced in a 65 nm radiation-hard complementary metal–oxide–semiconductor (CMOS) technology. Its linear analog front-end has a low threshold of fewer than 1000 electrons and the time-over-threshold is proportional to the collected charge. It can cope with a trigger rate of 750 kHz, a readout latency of $12.5\,\mu\mathrm{s}$, and an output bandwidth of 1 to $4 \times 1.28\,\mathrm{Gb/s}$.

Pixel modules are an arrangement of either 1×2 or 2×2 readout chips (ROCs). The ROCs are bump-bonded to the sensors. A passive printed circuit board (PCB), the High Density Innerconnect, is glued to the other side of the sensor. Wire bonds connect it to the ROC and a Parylene coating provides the mechanical and spark protection. Aluminum-nitride structures under the ROCs support the modules, whose designs vary between the sub-detectors. The modules are connected with $1.28\,\mathrm{Gb/s}$ electrical links to portcards that provide optical links to the back-end. Each portcard hosts three low-power Gigabit Transceiver (lpGBT) [14] application-specific integrated circuits (ASICs) and Versatile Transceiver plus ($\mathrm{VTRx^+}$) modules [15] with $10\,\mathrm{Gb/s}$ optical links. Radiation-hard direct current to direct current (DC-DC) buck point-of-load converters (bPOL), namely the bPOL12V [16] and the bPOL2V5 [17], are required to provide power for the portcards. The number of electrical links (e-links) and the bandwidth per ROC are optimized for each layer of the different sub-detectors. The optical links are connected to the Serenity boards [18] at the back-end. Each board has two field-programmable gate arrays (FPGAs) and one central processing unit (CPU) and is the interface to the central CMS data acquisition (DAQ).

A serial powering scheme is implemented for the detector. Up to twelve modules are powered in series. It limits the current in the transmission lines and reduces the cable mass. An on-chip shunt low-dropout regulator (LDO) on each ROC regulates the voltage and dissipates excess power.

A part of the readout bandwidth is used for real-time cluster counting on FPGAs to measure the luminosity. The innermost ring of the last endcap disk is placed at $|\eta| > 4$. A separate chain reads out the modules with 825 kHz exclusively for the luminosity measurement.

The upgrade of the Outer Tracker is described in detail in Section 1.4.

1.3.2 Timing layer

The upgrades of the tracking detectors provide a spatial vertex separation down to 300 µm [12]. For smaller distances, vertices will merge in the high-pileup environment of the HL-LHC. However, even within one bunch crossing the collisions are separated in time. The spatial overlapping of vertices can be resolved with the precision timing of charged tracks. It will suppress pileup and maintain physics performance. The CMS Collaboration will add a timing layer that provides a time stamp with 30-50 ps resolution for every track [19].

The layer provides hermetic coverage up to $|\eta| \leq 3$. It is comprised of two sub-detectors: the barrel timing layer (BTL) in the barrel tracker support tube and the endcap timing layer (ETL) in front of the calorimeter endcap. The limited available space is a major design constraint.

The BTL covers an area of 38 m^2 with LYSO[3] scintillator crystals. The crystals are radiation tolerant, have a good light yield, and a fast rise time. Silicon photomultiplier (SiPM) photo-detectors collect the scintillator light. They are compact, fast, and insensitive to magnetic fields. A dedicated readout ASIC with dark-count-rate-suppression was derived from a PET (positron emission tomography) application. The data are merged on a concentrator card.

The ETL faces two challenges: it has to withstand the radiation and particle density of the forward CMS detector region and can only be 45 mm thick. The technology of choice is low-gain avalanche diodes (LGAD). They fulfill the requirements of radiation hardness, fast-rising pulses, and relatively thin sensors. An additional gain layer is processed directly on the sensors. A readout ASIC that meets the timing requirements is currently under development and qualification.

Both sub-detectors use the CERN-developed radiation-hard ASICs lpGBT, GBT-SCA (Gigabit Transceiver Slow Control Adapter) [20], and the bPOL12V as well as the $VTRx^+$ module in the readout. The cooling is performed with two-phased CO_2 cooling.

[3] Lutetium–yttrium oxyorthosilicate

1.3.3 Calorimetry

The readout speed and the radiation damage are the drivers of the Phase-2 upgrades of the calorimeters. The latter necessitates a complete replacement in the forward region, while an upgrade of the electronics is sufficient in the barrel region.

In the future, the barrel electromagnetic calorimeter still consists of $PbWO_4$ crystals. While the readout by avalanche photodiodes remains, the front-end electronics will be replaced. It will achieve 30 ps time resolution for electrons and photons of 30 GeV energy at 40 MHz [21, 22]. An additional front-end board and cooling to 9°C instead of 18°C will improve the performance. The back-end boards will also be replaced. The hadronic barrel calorimeter receives the same back-end board upgrade. It is the only hardware upgrade of this sub-system.

The High Granularity Calorimeter (HGCal) will replace the two endcap calorimeters [23]. It provides 3D spatial positions in addition to the energy measurement and 50 ps timing resolution. It is structured in an electromagnetic section (CE-E) with lead sheets, copper cooling plates, and copper-tungsten baseplates as absorbers and a hadronic section (CE-H) with steel absorbers. The active technologies are also mixed. About 26,000 silicon sensor modules with a total of 6 million channels cover the region closest to the interaction point. Their cells have a size of $120\,\mu m \times 0.5\,cm^2$ and $200\ (300)\,\mu m \times 1.1\,cm^2$. To cover an area of roughly $630\,m^2$, 12 different module types with (partly-)hexagonal sensors produced from 8-inch wafers are required. In the outer part of the hadronic section, 250,000 scintillator tiles read out by SiPMs are used. They cover an area of $370\,m^2$.

The CE-E has 26 layers with a total depth of 27.7 radiation length and 1.5 hadronic interaction lengths. The CE-H features 7 all-silicon layers and 14 layers with both technologies. The total depth is equivalent to more than 10 hadronic interaction lengths. CO_2 cooling facilitates detector operation at −30°C.

The same ROC can read out the silicon and the scintillator modules. It measures time-of-arrival, time-over threshold, and absolute charge. Concentrator chips for the data are connected to lpGBT ASICs and $VTRx^+$ modules for the optical connection to the back-end. CERN's bPOL12V and linPOL12V (linear point-of-load converter 12V) [24] ASICs are also used.

1.3.4 Muon system

The upgrade of the muon systems has four goals. The detectors must sustain higher trigger rates while maintaining performance and a higher amount of background. It will also add redundancy in the forward region and extend the acceptance in the very forward region [25]. The existing drift tubes, resistive plate chambers, and cathode strip chambers remain in the CMS detector after LS3.

The DTs receive new on-detector and back-end electronics to allow for the L1 trigger rate of up to 750 kHz. The trigger primitive generation[4] is moved to the back-end, which allows higher complexity and exploitation of the full DT time information. The system also uses the lpGBT and VTRx$^+$ module for the optical readout.

Following the modifications during LS2, the CSCs will only receive upgraded off-chamber optical data boards to cope with the higher trigger rate.

Similarly, the back-end infrastructure of the RPCs is improved. It will comply with the lpGBT standard and improve background rejection with 1.56 ns timing. Additional chambers with an improved RPC design are installed in the forward region. They are complementing the existing CSCs. For the same purpose, GEM stations were already installed during LS2 and prior to Run 3 in the layers closer to the interaction point.

The installation of the ME0 system during LS3 extends the muon coverage to the region $2.0 < |\eta| < 2.8$. The ME0 consists of six triple-GEM detectors and will be the closest muon detector to the interaction point.

1.4 Phase-2 Outer Tracker upgrade

This thesis concerns the Phase-2 upgrade of the Outer Tracker (OT). The currently operational strip tracker will be entirely replaced. The new tracker will be able to withstand the challenging radiation and pileup conditions of the HL-LHC. The L1 trigger rate of CMS will increase from currently 110 kHz to 750 kHz. Tracking information will be incorporated into the L1 decision algorithm. In addition, the trigger latency will be increased to 12.5 µs. The main driver of the tracker design is the requirement to provide track information to the L1 trigger. Figure 1.4 shows a schematic view of the layout of the new CMS tracking system. The OT consists of three sub-detectors and two different module types. The Tracker Barrel with PS (macro-pixel and strip) modules (TBPS) has a radial distance r from the collision point of $200\,\text{mm} < r < 600\,\text{mm}$. It is the innermost region of the OT. To ensure good tracking performance and a low material budget in the high η region, the PS modules are progressively tilted. The Tracker Barrel with 2S (strip and strip) modules (TB2S) surrounds the TBPS. A Tracker Endcap Double Disk section (TEDD) is placed at either end of the barrel structure. As in the barrel, the inner TEDD radial section is covered with PS modules, and the outer with 2S modules.

The new Outer Tracker has a front-end power consumption of about 100 kW. A two-phase CO_2 cooling system with a coolant temperature of -33°C will dissipate the heat.

[4] Creation of a trigger object with reduced data size and less information.

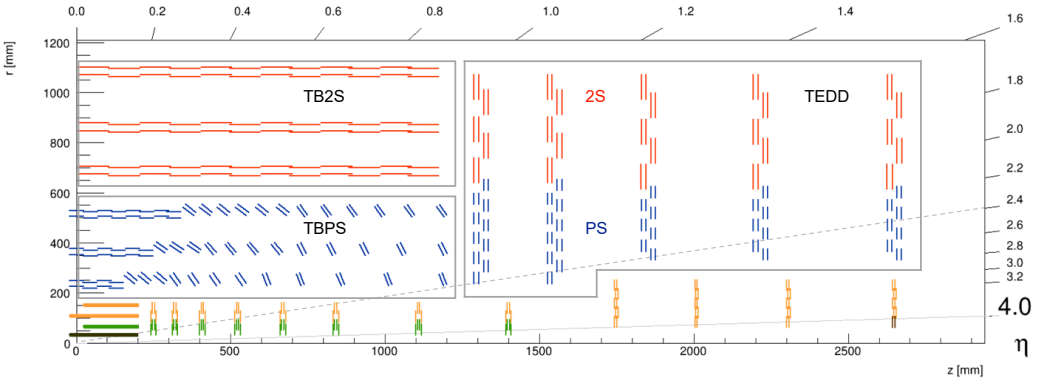

Figure 1.4: Schematic view of one-quarter of the future CMS tracker in the r-z plane. The Outer Tracker and its sub-detectors are highlighted. 2S module positions are shown in red, PS module positions in blue. The Inner Tracker modules are shown in yellow and green. The 3D pixel modules are drawn in black and the ones for the luminosity measurement in brown [12].

Both module types have two vertically stacked silicon sensors. This gives the modules the capability to identify particle candidates with a high transverse momentum p_{T} and provide their hit information to the Level-1 trigger. The principle is sketched in Figure 1.5. Charged particles have a bent trajectory as they pass through the magnetic field in the tracking volume. The curvature depends on the particle momentum. High-p_{T} particles have only a small offset between hit clusters in the two closely spaced sensors. The window size can be tuned to select particles above a 2 GeV threshold. The tuned size is a function of the position in the detector and the sensor spacing. Three different sensor spacings are used for the PS modules and two for the 2S modules. A correlation of two hit clusters in both sensors is called "stub". Only the information on the identified stubs is forwarded to the L1 trigger for each of the bunch crossings that occur with a rate of 40 MHz. Tracks are formed from the stubs at the back-end and utilized by the trigger. The full hit data is buffered on the module for up to 12.5 µs and transferred after an accept signal from the trigger is received.

The next paragraph provides more information on the PS module. As this thesis is concerned with the 2S module, it is discussed separately in Section 1.5. The bottom sensor in the PS module is a macro-pixel sensor with an active area of roughly $5 \times 10\,\mathrm{cm}^2$. Its 30,000 macro-pixels have a size of $100\,\mu\mathrm{m} \times 1.5\,\mathrm{mm}$. The upper sensor is a strip sensor of the same size with two rows of 2.5 cm long strips. The 960 strips per row have a pitch of 100 µm. Both sensors have n-in-p type polarity with an active thickness of 290 µm. The depletion voltage is about 350 V before irradiation and the voltage can be increased to about 800 V to maintain performance during HL-LHC operation. Aluminum-nitrite bridges separate the sensors by 1.6, 2.6, or 4.0 mm depending on a module's position in the tracker. The coefficient of thermal expansion of the composite material matches the one of silicon. Four flexible polyamide PCBs

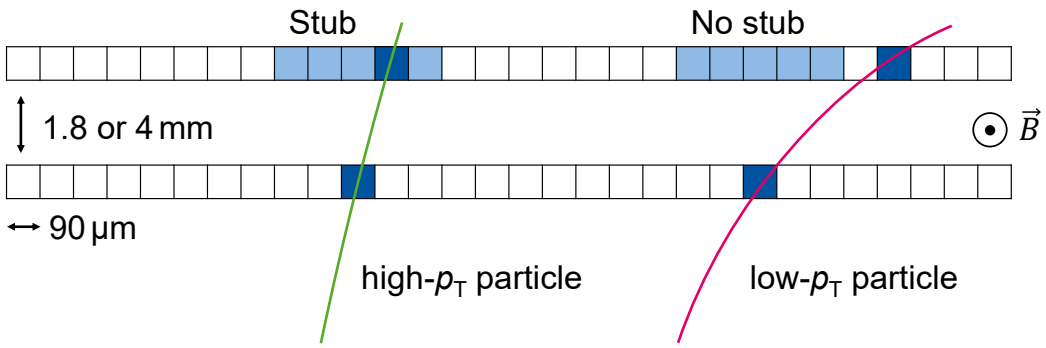

Figure 1.5: Illustration of the stub mechanism in the 2S module. The sketch is not to scale and exaggerates the bending. Here, no charge sharing between strips is assumed.

are placed around the sensors. They are referred to as "hybrids". The two Front-End Hybrids (FEHs) sit on the long side of the module and perpendicular to the strips [26]. Each hosts eight Strip Sensor ASICs (SSAs) [27], which read out the hits in the strip sensors. A total of 16 Macro-Pixel ASICs (MPAs) [28] are attached through bump bonds to the inner side of the macro-pixel sensor to read out its hits. They also perform the hit correlation and generation of the stubs. Wire bonds connect the MPAs to the FEHs. Each FEH also hosts one Concentrator Integrated Circuit (CIC). It receives data from the MPAs, aggregates them, and forwards them to the Readout Hybrid (ROH) on one of the short module sides. The ROH features an lpGBT for the data serialization and a VTRx$^+$ module for the optical communication with the back-end of the tracker. The remaining module side hosts the Power Hybrid (POH). It features one bPOL12V and two bPOL2V5 ASICs. Both chips are radiation-hard DC-DC converters and provide 2.5, 1.25, and 1.0 V from a module input voltage of about 10 V [12, 29].

1.5 Outer Tracker 2S module

Figure 1.6 shows an illustration of the 2S module design. Both strip sensors have a size of $10 \times 10\,\mathrm{cm}^2$. Each sensor is divided into two rows of 1016 strips with a pitch of 90 µm and a length of 5 cm. The sensor spacing is either 1.8 or 4 mm and is defined by aluminum carbon fiber (AlCF) bridges. Two 2S-FEHs are placed perpendicular to the strip orientation on two module sides. The sensor and the FEHs are connected through wire bonds. Eight CMS Binary Chips (CBCs) on each FEH read out the strips and generate stubs. The CIC is used on the 2S module to bundle the CBC data. The Service Hybrid (SEH) combines the functionality of the ROH and POH.

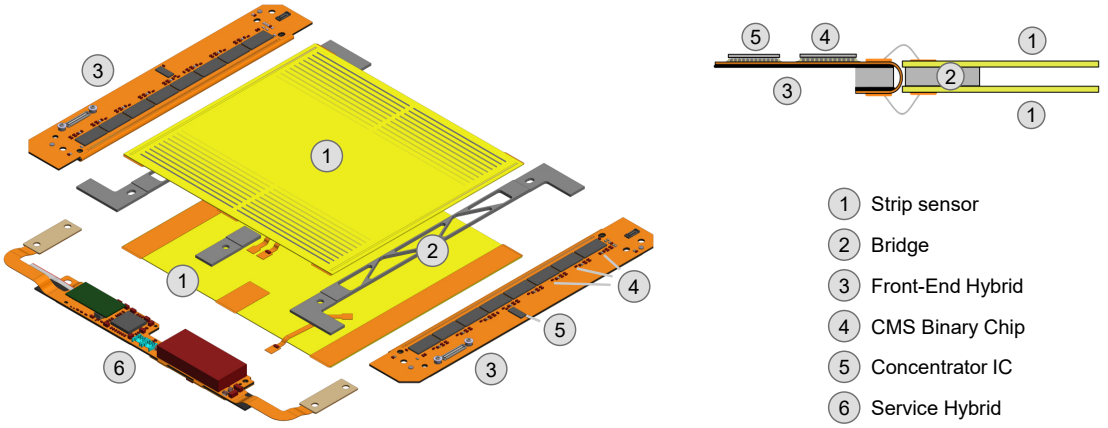

Figure 1.6: 2S module design drawing. On the left, the module is shown in an exploded view. The relevant components are labeled. The upper right corner shows a cross-section of the FEH and the adjacent sensor edge [29].

It also features the lpGBT and hosts a VTRx$^+$ module. Only one bPOL12V and one bPOL2V5 converter are required to provide 2.5 and 1.25 V.

1.5.1 Sensor and Front-End Hybrid

The active (physical) silicon sensor thickness is 290 µm (320 µm). The sensors are produced as n-in-p type with p-stop strip isolation and have an aluminum backside. A negative bias voltage is applied to the backside and depletes the p-bulk. Charged particles, which traverse the p-bulk, generate electron-hole pairs. N-strip implants collect the electrons and guide them to so-called DC-pads. The charge collection process induces the actual signal that is processed by the readout chip. The DC-pads are connected to the electrical ground with 1.5 MΩ bias resistors. The implants are AC-coupled to the actual aluminum strips on the sensor topside[5]. The electrical connection to the FEH is realized via wire bonds on the so-called AC-pads.

The main supporting structures are the AlCF bridges, which are screwed to the cooling supports of the larger detector structures. The aluminum backsides of both sensors are glued to two long bridges and, depending on module position and cooling requirements, one or two short ones. Glued polyimide films (*Dupont Kapton® MT* [30]) isolate the bridges from the bias voltage on the backside. Each FEH is attached to a long bridge. Eight CBCs each read out 127 channels on both the top and the bottom sensor. As indicated in Figure 1.6, a section of the flexible FEH PCB is folded to the hybrid backside to allow the wire bond connection to the bot-

[5] AC (alternating current)

tom sensor. The signal of each strip is amplified and shaped before a comparator identifies hits.

A CBC starts the formation of a stub by rejecting over-sized clusters of hits in a sensor. The maximal allowed cluster width is four and adjustable via a register. The bottom sensor is the seed layer by default and the top sensor is the correlation layer. The CBC looks for clusters in the top sensor within a correlation window of the bottom sensor cluster. The width of the window of up to ± 7 channels about a central channel in the correlation layer and its offset (up to ± 3 channels for four independently programmable regions per chip) are adjustable to account for different module positions in the detector. The assignment of seed and correlation layer can also be swapped depending on the direction a module faces. The use of cluster positions results in the possibility of half-strip positions. Each CBC exchanges information about hits in its edge channels with its nearest neighbors. In total 30 active-high CMOS inputs and the same number of outputs per chip are used to exchange five seed layer hits and 25 correlation layer hits.

Each CBC is connected to the CIC with six 320 Mb/s e-links. Five are used to send up to three stubs per bunch crossing to the CIC. If more than three stubs are found, only the three with the lowest cluster position are forwarded. The CIC collects the stubs from all connected CBCs and sorts them by their p_{T}. A higher p_{T} corresponds to a lower bend and a lower digitized bend code. In the 2S module, stubs are collected for eight bunch crossings. The 16 stubs with the lowest bend are sent to the back-end. After the module receives an L1 trigger, the sixth e-link is used to send the full hit data ("L1 data"). The CBCs store their hit data for 12.8 µs in a 512 bunch crossing deep pipeline. The CIC bundles and compresses the data from the CBCs before they are forwarded to the back-end.

1.5.2 Service Hybrid

The focus of this thesis lies on the Service Hybrid. It is therefore discussed in detail in this section.

As the top-side view in Figure 1.7 indicates, the Service Hybrid resides on one module side between the two FEHs. It is connected to the FEHs via fine-pitch connectors (Panasonic A35S [31]) that are marked in the bottom-side view in Figure 1.8. A distinction of the two FEHs between a "right" and a "left" one is made based on the viewing perspective in Figure 1.7. The Service Hybrid serves as the entry point for all connections to the back-end. It has two receptacles (Hirose DF57-2P [32]) for two-pin cables, one low-voltage cable with a positive supply voltage and the ground potential, and the bias voltage cable. The third back-end connection is the connector for the VTRx$^+$ module. The VTRx$^+$ module is assembled with five optical fibers. One fiber transmits data from the back-end to the module ("downlink") and four from the module to the back-end ("uplink"). The SEH only uses one uplink fiber.

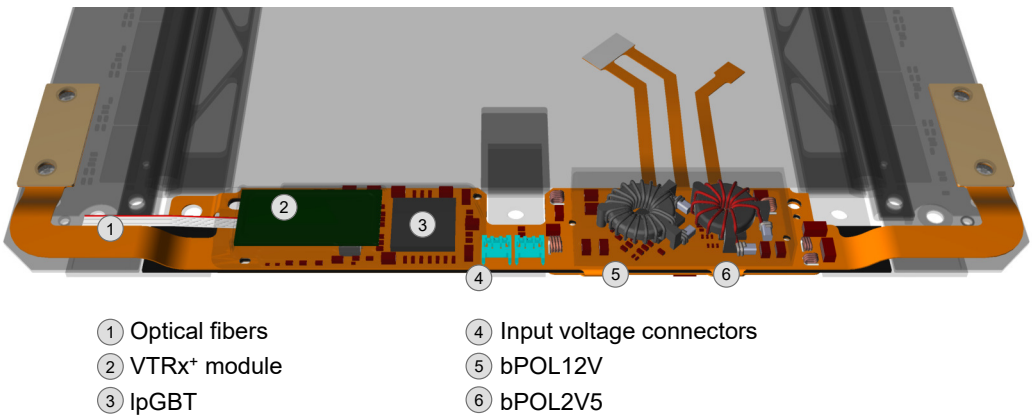

Figure 1.7: Illustration of the Service Hybrid assembled into a 2S module. The most important components on the top side are labeled. The rest of the module and the DC-DC shield are grayed out. On the right, the SEH is connected to the right FEH, and on the left to the left FEH.

The lpGBT ASIC and the VTRx$^+$ module occupy one-half of the hybrid. The latter performs the optical-electrical conversion of the high-speed links that connect it to the lpGBT. The other half features the two DC-DC converter ASICs bPOL12V and bPOL2V5 and their accompanying passive components. An aluminum shield protects the module from radiated noise. A fold-over section to the backside contains a filter for the sensor bias voltage. Two connectors on the hybrid fold-over connect two flexible flat cables to the sensor backsides to provide the filtered bias voltage. The end of the cable is glued to the sensor and wire bonds provide the electrical connection. A similar cable with a thermistor is glued to the upper sensor.

The flexible hybrid PCB has a carbon fiber stiffener. In the module, the hybrid is glued to three bridges. One of the screws that mount a module to its cooling contacts electrically connects the pads of a grounding flap to the larger sub-structure. The hybrid stiffener is grounded via inserts pressed into the carbon fiber material.

An overview of the functional components of the Service Hybrid and their connections is shown in Figure 1.9. The following sections give additional information on the functionalities of the Service Hybrid and its components.

1.5.2.1 Powering

The power supplies for the OT are placed in the underground CMS service cavern. Up to twelve modules share a power supply channel for the positive supply voltage and the ground return. However, each module has a dedicated supply voltage wire and can be switched on or off individually. The cables that connect the power supplies with the modules are more than 60 m long. Voltage drops and the result-

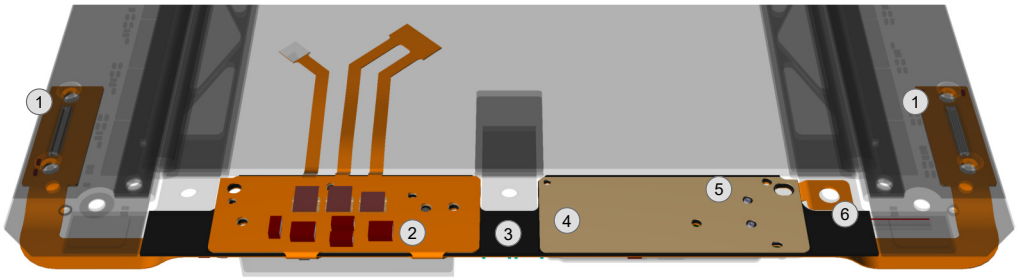

1 FEH connectors
2 Bias voltage RC filter
3 Carbon fiber stiffener
4 Compensator
5 Grounding inserts
6 Grounding flap

Figure 1.8: Illustration of the Service Hybrid assembled into a 2S module. The most important components on the bottom side are labeled. The rest of the module and the DC-DC shield are grayed out.

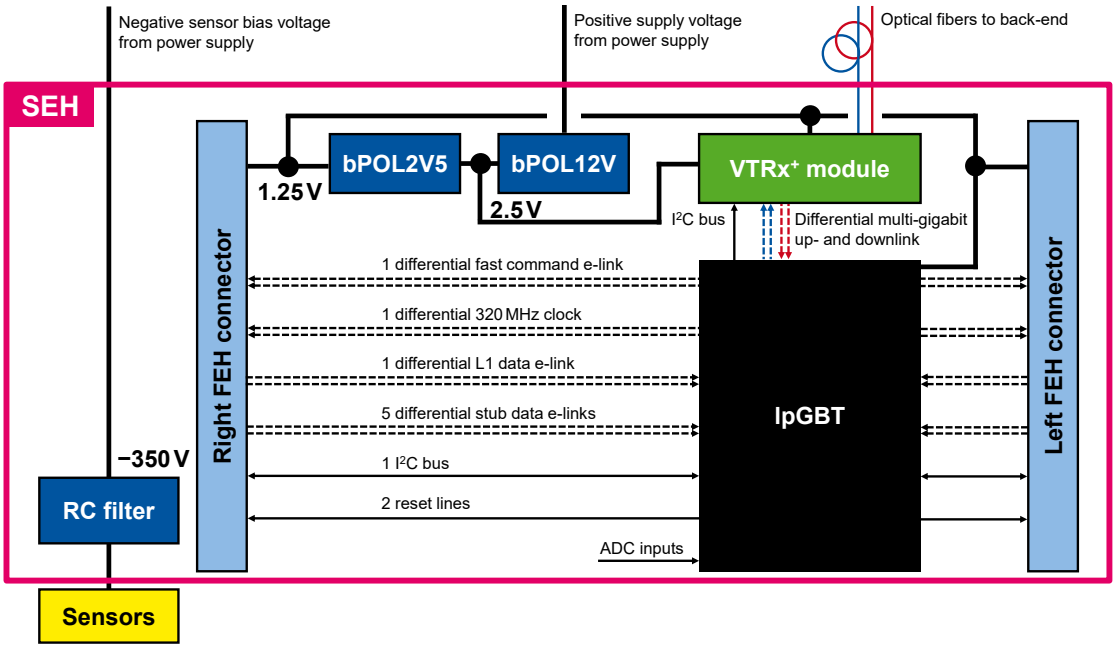

Figure 1.9: Overview of the functional Service Hybrid components and their interconnections. The ground connections, the passive components, and configuration inputs are not shown.

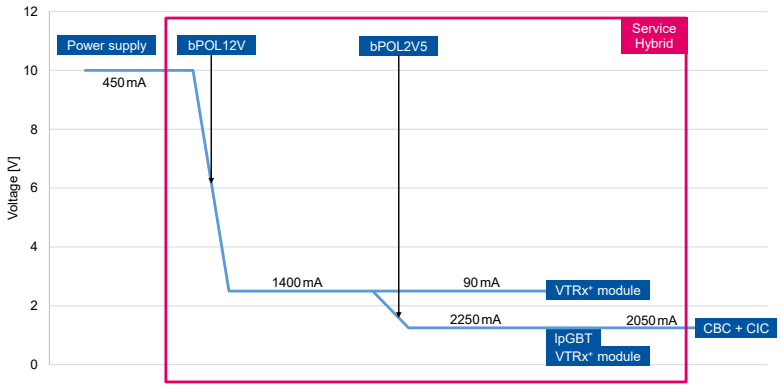

Figure 1.10: Illustration of the voltage distribution scheme on the Service Hybrid. A power supply provides roughly 10 V of input voltage to the hybrid. The bPOL12V converts the voltage to 2.5 V required by the VTRx$^+$ module. In a second step, the bPOL2V5 reduces this further to 1.25 V to power the lpGBT, the VTRx$^+$ module (combined roughly 200 mA), and the CBCs and CICs on the FEHs. The shown currents are typical consumptions.

ing power loss depend quadratically on the transported current. The FE electronic power consumption of the Outer Tracker will be around 100 kW compared to a consumption of 60 kW for the current strip tracker. Direct powering would result in almost four times more cable losses. Since the available cable space is limited, the OT is not powered directly from the power supply output. Instead, a two-stepped powering scheme with DC-DC converters is implemented.

The on-module powering scheme of a Service Hybrid is sketched in Figure 1.10. The current required by a module results in voltage drops along the supply cable. As a result of these Ohmic losses, a module receives an input voltage of about 10 V. The bPOL12V converter reduces this voltage to 2.5 V, which is required by the VTRx$^+$ module. These 2.5 V are also the input voltage for the bPOL2V5 converter in the second DC-DC stage. Its output voltage of about 1.25 V is the supply voltage of the other ASICs (lpGBT, the CICs, and the CBCs) and is also required by the VTRx$^+$ module.

The same power supply unit that provides the DC-DC converter input voltage for twelve modules also provides the channels with their sensor bias voltage. Up to four modules share a power supply channel. Voltages between 0 and -800 V are possible, while -350 V is a typical value. Individual cables forward the voltage to each module. The supply voltage ground connection is used as the return path. The bias voltage distribution features removable jumpers, which allow the disconnection of any module [33]. After irradiation, the bias voltage leakage current through the sensor is of the order of a few milliamperes, and voltage drops on the cables are not a concern. The Service Hybrid hosts a low-pass RC filter for the bias voltage on

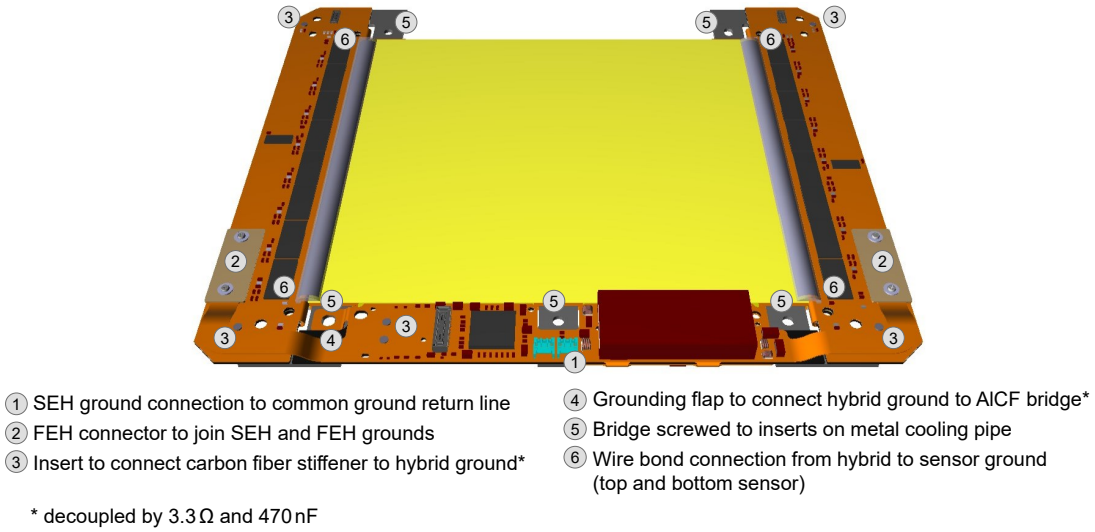

Figure 1.11: Illustration of the interconnections between different ground potentials on a 2S module.

its fold-over section. Additional resistors in the bias voltage flex cables and behind the filter limit currents between the filter and the sensors. The resistors also limit current flow between both sensors.

The relevant ground connections on the 2S module are indicated in Figure 1.11. The DC-DC input and output voltages and the sensor bias voltage are referenced to the same ground potential on the SEH that is provided by the common ground return line. The ground potentials of the SEH and FEHs are joined at the level of the FEH connectors. Inserts electrically connect each carbon fiber stiffener to its hybrid ground. On each hybrid, a resistor (3.3 Ω) and a capacitance (470 nF) decouple stiffener and hybrid ground. A glue layer isolates the stiffeners from the bridges and the grounding flap is used to provide a discrete connection between the SEH stiffener and the left bridge. The module bridges are screwed to metal cooling inserts. Although the inserts are glued to the pipe, an electrical connection between the pipe and the insert is present. The connection is verified during production. Wire bonds connect ground pads on the FEHs with counterparts on both sensors.

A schematic illustration of the grounding scheme of the 2S module that also includes the common ground return line and the mechanical structure is shown in Figure 1.12. The return line is connected to the main detector ground at its patch panel 0, which is placed at the outer edge of an OT substructure. The pipes of the CO_2 cooling loops are also referenced there.

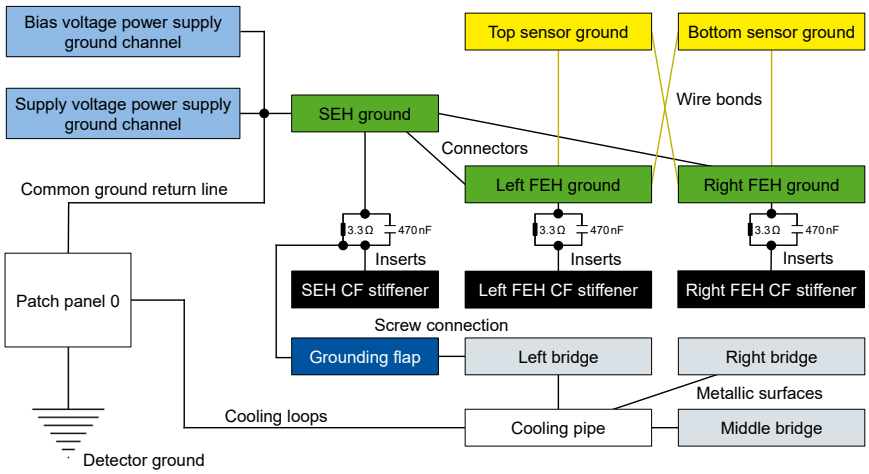

Figure 1.12: Schematic illustration of the grounding scheme of a 2S module.

Buck DC-DC converter

As their name implies, the bPOL12V [16] and the bPOL2V5 [17] are both buck converters. This type of DC-DC converter reduces the output voltage by periodically connecting and disconnecting the load and the power supply. They require an inductor to store energy in its magnetic field and to preserve a constant current flow. Additional capacitors or filters reduce voltage ripple at the input and output.

Figure 1.13 illustrates the basic operation principle. It describes the conversion of an input voltage V_{in} to the output voltage V_{out}. T_1 and T_2 are transistors and switched with a fixed frequency, f_s. As passive components, the converter uses an inductor L and an output filter capacitance C_{out}. In the on-state, T_1 is conducting and T_2 is open. The top sketch shows the current flow from the power supply to the load. At the same time, a magnetic field builds up in the inductor resulting in a reduced voltage at the load. Next, both transistors are switched. In the off-state (lower sketch), T_1 is open and T_2 conducting. The magnetic field is dissipated to maintain the current flow through the load. The switching period is much faster than the decay time of the magnetic field.

A control circuit drives the transistors and regulates the output voltage. For a lossless DC-DC converter the conversion ratio r corresponds to the inverse of the duty cycle D, which is the fraction of the time the device T_1 is in the on-state (t_{on}):

$$D = \frac{t_{\text{on}}}{T} \quad \text{with} \quad t_{\text{on}} + t_{\text{off}} = T = \frac{1}{f_s} \quad , \tag{1.1}$$

$$r = \frac{V_{\text{in}}}{V_{\text{out}}} = \frac{1}{D} = \frac{1}{t_{\text{on}} \cdot f_s} \quad . \tag{1.2}$$

A feedback loop modifies the duty cycle to maintain the set value of the output voltage irrespective of the load. This is possible independently of the input voltage.

Buck converters can have a switching frequency in the range of a few kHz to several MHz. The choice is a trade-off between higher output ripple and thus the need for larger inductors for smaller frequencies, compared to more switching and driving losses of the transistors and lower conversion efficiency for a higher frequency.

We define the conversion efficiency η as the ratio of the converter output to input power:

$$\eta = \frac{P_{\text{out}}}{P_{\text{in}}} = \frac{V_{\text{out}} \cdot I_{\text{out}}}{V_{\text{in}} \cdot I_{\text{in}}} \quad . \tag{1.3}$$

Here, V_{out} is the output voltage of the DC-DC converter, I_{out} is the output current, V_{in} is the supply voltage, and I_{in} is the input current.

Compared to commercial DC-DC converters, the bPOL devices need to fulfill several requirements for their usage in the OT. They need to withstand the radiation environment both in terms of the total ionizing dose (TID) of up to 100 Mrad and single-event upsets. The magnetic field of 3.8 T inside the tracking volume saturates ferrite components, so air-core inductors must be used. Electromagnetic emissions potentially degrade the module performance. The devices need to be well shielded while maintaining low mass and small volume, which is important for tracking performance [34, 35].

1.5.2.2 Digital readout

The lpGBT in conjunction with the VTRx$^+$ module perform the optical readout of the module. Two multi-gigabit differential e-links connect the two components. The VTRx$^+$ module receives optical data from the back-end and transfers it electrically to the lpGBT and vice versa. This section discusses 2S module-specific settings. More details on the lpGBT and VTRx$^+$ module are given in Section 2.2.2.

On each module, the lpGBT receives one 64-bit downlink frame per bunch crossing at 2.56 Gb/s. The lpGBT recovers its clock and optimal sampling phase from the incoming data. Each frame is composed of a four-bit header, two Internal Control (IC) bits, two External Control (EC) bits, a 32-bit data field, and 24 forward error correction (FEC) bits. The lpGBT searches for the header in consecutive frames and uses this to align each frame correctly. The two IC bits in each frame are used for control and monitoring of the lpGBT and subsequently the whole module. Bits from subsequent frames are concatenated and form the words of a write-read protocol to access the lpGBT registers. The EC bits are made available on an 80 Mb/s differential e-link. This feature is not used in the Outer Tracker. The 32-bit data field is

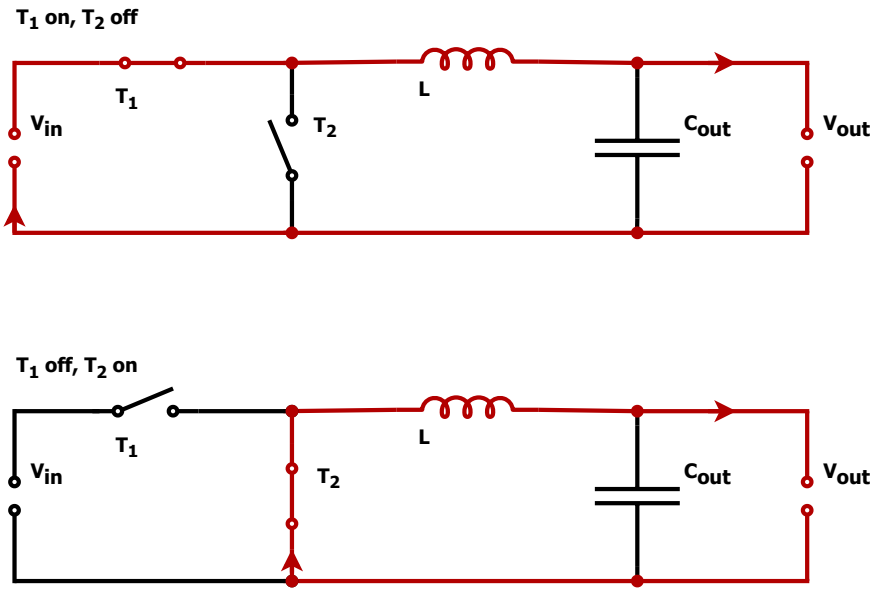

Figure 1.13: Working principle of a buck converter [34]. Red parts of the circuit indicate the current path in a given state. Top: on-state; the load is connected to the input. Bottom: off-state; the magnetic field of the inductor maintains the current through the load.

transmitted to the four output groups (eight bits per group and frame). The FEC code detects and corrects transmission errors.

To ensure sufficient transitions for the clock and phase recovery the IC, EC, and data bits are encoded by a recursive scrambling equation to achieve a DC-balanced frame. The scrambled bits of one frame are used to code/decode the data bits of the next frame. After the scrambler, the value of the FEC is calculated with a Reed-Solomon code. The resulting frame is interleaved and the FEC is subsequently able to correct up to twelve consecutive corrupted bits. The CBC and CIC ASICs feature a fast control interface that receives one eight-bit "fast command" on the outgoing lpGBT e-links per bunch crossing. They are sampled from a 320 MHz differential clock that the lpGBT provides for each hybrid. Each hybrid receives its own fast commands since the latency might vary between the sides. In the OT front-end chips, phase tuning is only possible by inversion of a side's sampling clock. This corresponds to a $\pm 180^\circ$ phase shift

The possible fast commands are listed in Table 1.1. They all start with the unique synchronization sequence 110_{bin}. The ASICs use it to produce their 40 MHz on-chip reference clock. The following four bits initiate a pre-defined function. The Fast Reset sends a reset to most core ASIC functionalities. The Trigger initiates a trigger signal to the pipeline control logic. The Test Pulse Trigger produces a test pulse

Fast command	B7	B6	B5	B4	B3	B2	B1	B0
Fast Reset	1	1	0	1	0	0	0	1
Trigger	1	1	0	0	1	0	0	1
Test Pulse Trigger	1	1	0	0	0	1	0	1
Orbit Reset	1	1	0	0	0	0	1	1
Orbit Reset & Fast Reset	1	1	0	1	0	0	1	1
Orbit Reset & Trigger	1	1	0	0	1	0	1	1
Orbit Reset & Test Pulse Trigger	1	1	0	0	0	1	1	1
Idle	1	1	0	0	0	0	0	1

Table 1.1: Overview of the allowed fast commands on the downlink e-links.

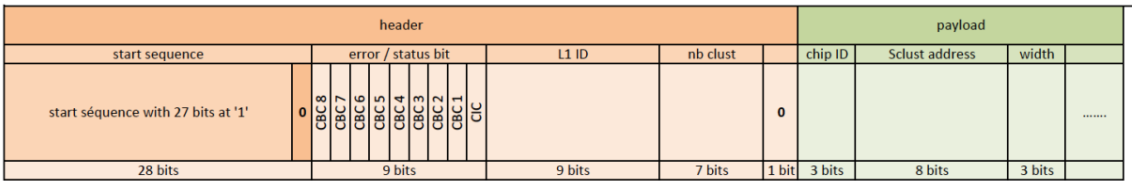

Figure 1.14: Overview of the CIC's L1 output frame. The sparsified frame suppresses unhit strips. The lpGBT receives it on an uplink e-link [36].

on the CBC. The Orbit Reset defaults various chip-internal counters. It is the only signal that can be sent at the same time as the other commands.

The lpGBT downlink has only one possible format, the uplink has four. The 2S module uses the one with a speed of 5.12 Gb/s and the capability to correct up to five consecutive wrong bits (FEC5). The 128-bit frame consists of two header bits, two IC and two EC bits, 112 data bits, and ten FEC bits. Scrambling, forward error correction, and interleaving are implemented similarly as for the downlink.

The module data that is injected can be divided into three parts. First, the two bits of the IC return frame; second, hit pipeline data sent by the front-end ASICs after an L1 trigger is received; and third, the stub data that are continuously sent. The lpGBT receives the latter two on its seven e-link receiver groups[6]. Each CIC sends L1 data on one 320 Mb/s e-link to the lpGBT. The frame contains sparsified (zero-suppressed) hits combined to clusters. After a header with a length of 52 bits (start sequence, status field, L1 identifier (ID), cluster multiplicity, and a padding bit), several 14-bit clusters are broadcasted (3-bit chip ID, 8-bit strip address, 3-bit cluster width). Figure 1.14 shows the composition of the frame. When no pipeline data is shipped, the CIC sends a continuous 10_{bin} pattern.

[6]Depending on the configured speed, multiple channels per group are available. Only twelve e-links are required to read out a 2S module.

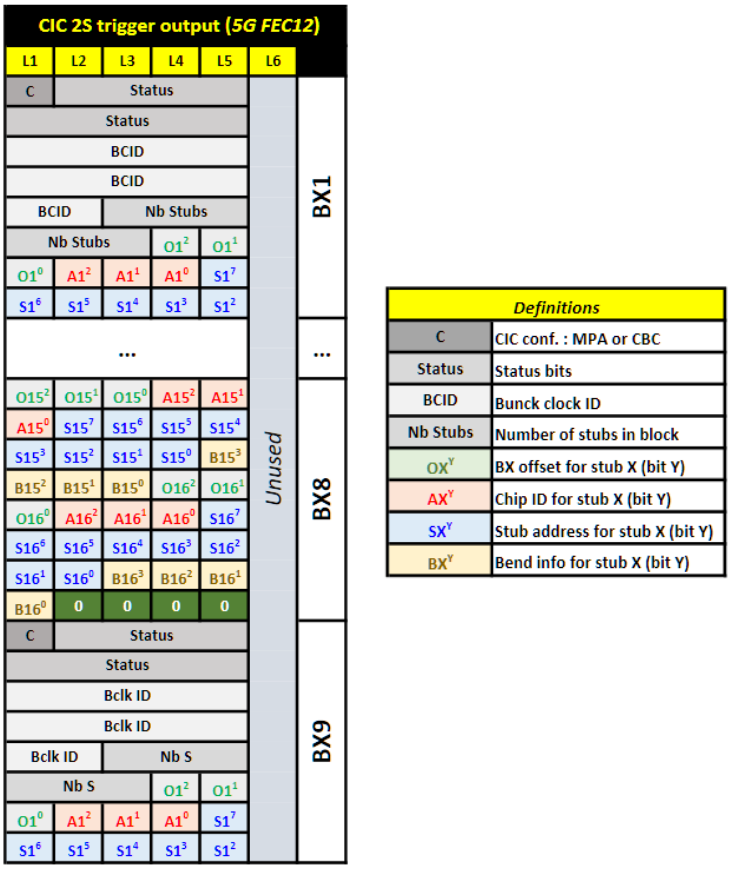

Figure 1.15: Summary of the CIC's stub output frame. One frame is transmitted every eight bunch crossings. The lpGBT receives it in parallel on five uplink e-links [36].

The CIC's stub data are split on five 320 Mb/s e-links. One block stretches over eight bunch crossings. A 28-bit header (front-end type, nine status bits, twelve bits for the bunch crossing ID, and six bits for the stub multiplicity) precedes 16 stub packets. Each packet encodes a stub into 18 bits (three bits for the timing offset within the eight bunch crossings, three bits for the FE chip ID, eight bits for the stub address on the bottom sensor, and four bits for the stub bend code). Figure 1.15 shows the block structure. The sixth stub data output is only used in the PS module. Therefore only twelve out of the 14 available upstream e-links are used in the 2S module. A new block is sent every eight bunch crossings, so no specific idle pattern is required.

1.5.2.3 Slow control features

The two IC bits in each frame per bunch crossing are used for control and monitoring of the lpGBT and subsequently the whole module. Bits from subsequent frames are combined and form the words of a write-read protocol to access the lpGBT registers.

The protocol bandwidth is limited to 80 Mb/s, and the communication stretches over many bunch crossings. It is referred to as "slow control" to distinguish it from the fast command interface. All slow control features are controllable by the manipulation of specific registers [14]. They can be grouped into three categories, and are discussed in more detail below:

- the I^2C interfaces,
- the analog-to-digital converter (ADC) block,
- the general purpose input/output (GPIO) ports.

I^2C interfaces

I^2C stands for inter-integrated circuit and is a single-ended serial communication bus. A controller is linked to one or several targets with a data and a clock line.

Three I^2C controllers are implemented in the lpGBT. Controller 0 (2) serves the left (right) side FEH. The CIC and CBCs share the same bus and are identified via their I^2C address. The I^2C bus is used to configure and monitor the configuration registers of the front-end ASICs. The pull-up resistors are placed on the FEH. An I^2C frequency of 1 MHz is used.

The bus of controller 1 is connected to the $VTRx^+$ connector and therefore the $VTRx^+$ module. Only the uplink is configurable via I^2C. The pull-up resistors for Controller 1 are placed on the SEH. It was found that a reliable operation is achieved with 400 kHz operation frequency.

ADC block

The lpGBT features a fully differential 10-bit successive-approximation ADC with a 16-channel input multiplexer. The multiplexer has eight external inputs and eight internal signal connections. In some cases, voltage dividers are used to translate a signal to the ADC measurement range. The basic schematic is shown in Figure 1.16 (left). A given voltage V_{in} is reduced to V_{out} by the ratio of the resistance R_2 between V_{out} and ground to the total resistance between V_{in} and ground:

$$V_{\text{out}} = V_{\text{in}} \cdot \frac{R_2}{R_1 + R_2} \quad . \tag{1.4}$$

The input pins are used to monitor various aspects of the module. The list below states the name of the lpGBT input pin and the signal name on the hybrid:

ADC0 → AMUX_L:
All eight left side CBC ASICs share a connection of their 17:1 analog multiplexer (AMUX) to this input. Via I^2C the desired CBC and analog bias signals

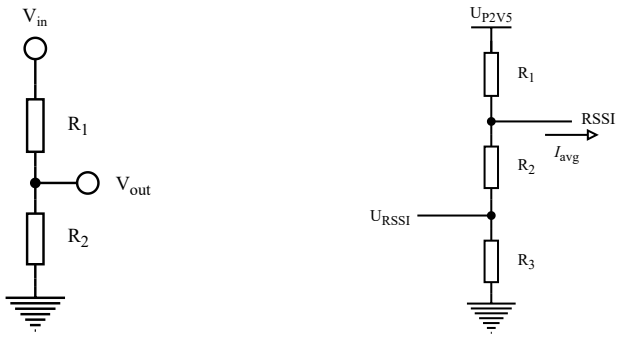

Figure 1.16: Left: general layout of a voltage divider. Right: schematic drawing of the circuit to measure the RSSI signal.

(e.g. the comparator threshold voltage) can be selected and measured. Only one CBC and only one signal can be monitored at a time.

ADC1 → VMON_P1V25_L:
This input is used to monitor the output voltage of the second stage DC-DC converter. The lpGBT can also measure its supply voltage from the internally connected signals. ADC1 is not connected directly to the power plane and senses the voltage at the left FEH connector and therefore provides information on the voltage drop to the point furthest away from the DC-DC output filter. A voltage divider (factor 0.645) adjusts the voltage level to the ADC's allowed input range.

ADC2 → VMIN:
The module's input voltage is monitored over this pin. It provides insights into the voltage drop between the service cavern and the tracker. A voltage divider (factor 0.049) adjusts the voltage level to the ADC's allowed input range.

ADC3 → AMUX_R:
See ADC0. The same functionality is available for the right-side CBCs.

ADC4 → TEMPP:
This input is connected with the thermistor that is glued to the sensor (1 kΩ negative temperature coefficient (NTC) thermistor). The lpGBT features a programmable current source on each external input that is used to measure the temperature-dependent resistance. For a given programmed value the current digital-to-analog converter (DAC) will increase the output voltage until the set current is reached. The conversion factor is roughly 3.6 µA per DAC unit. The resulting voltage is measured and converted to a resistance from which a temperature can be calculated.

ADC5 → VTRX+_RSSI_ADC:
The VTRx$^+$ module provides information about the average optical power of the light it receives from the back-end. This is done in the form of the received

signal strength current (RSSI). The used circuit is sketched in Figure 1.16 (right). A current sink draws a current equivalent to the average photocurrent (I_{avg}) on the RSSI output that is connected to the 2.5 V supply voltage rail (U_{P2V5}). A voltage divider adjusts the voltage level to the ADC's allowed input range. The voltage measured at the ADC5 input ($U_{\text{RSSI}}^{\text{VTRx+}}$) is given by

$$U_{\text{RSSI}}^{\text{VTRx+}} = \frac{R_3}{R_2 + R_3} \cdot (U_{\text{P2V5}} - R_1 I_{\text{avg}}) \ . \tag{1.5}$$

Typical values for the photodiode reponsivity are 0.45 to 0.50 A/W [37].

ADC6 → PTAT_BPOL2V5:
PTAT stands for Proportional To Absolute Temperature. The bPOL2V5 ASIC outputs a voltage proportional to its internal chip temperature on this pin. The slope is roughly 4 mV/K. The absolute value of the PTAT voltage has a wide sample-to-sample variability (up to 200 mV). It is best suited to monitor sudden changes during operation. An initial calibration and consideration of radiation-induced changes is needed for absolute measurements.

ADC7 → PTAT_BPOL12V:
The PTAT of the bPOL12V functions similarly to the PTAT_BPOL2V5. The slope is roughly 4.85 mV/K. The bPOL12V has the wide sample-to-sample variability of up to 200 mV, too. Due to the different voltage classes of the internal components, the absolute value is higher than for the bPOL2V5. A voltage divider (factor 0.623) adjusts the voltage level to the ADC's allowed input range. The observed slope will therefore be about 3 mV/K.

ADC measurements and calibration

The ADC is fully differential. It has essentially two inputs and measures the voltage difference between those, commonly referred to as the positive (V_{pos}) and negative (V_{neg}) input. The voltage difference is multiplied by an adjustable gain[7], G. The analog to digital conversion is performed for voltages between $-V_{\text{ref}}$ and $+V_{\text{ref}}$. V_{ref} is an internal reference voltage with a nominal value of 1 V that can be tuned via an 8-bit register[8]. The basic conversion formula to the ADC units U_{ADC} is

$$U_{\text{ADC}} = \frac{V_{\text{pos}} - V_{\text{neg}}}{2V_{\text{ref}}} \cdot G \cdot 2^{10} + \Delta \ \ , \tag{1.6}$$

where the offset Δ is defined as $U_{\text{ADC}}(V_{\text{pos}} = V_{\text{neg}})$. For an ideal ADC Δ is 512 (middle of the measurement range).

[7] The nominal values are 2, 8, 16, and 32. However, the actual factor is always smaller.
[8] 6-bit for the lpGBT version 0

In total 16 different signals can be chosen for both ADC inputs. The choices include the aforementioned eight external inputs and eight internal signals. In addition, a connection to ground potential is also possible. We discuss in more detail the choice of an external input for V_{pos} and the internal signal $V_{\text{ref}}/2$ for V_{neg} and the lowest gain setting. This choice represents the general ADC use case of measuring a positive voltage between 0 and 1 V. In this case, Formula 1.6 can be used to obtain

$$V_{\text{pos}} = 1\,\text{V} \cdot \left(\frac{U_{\text{ADC}} - \Delta}{512 \cdot G} + \frac{1}{2} \right) \quad . \tag{1.7}$$

G can be determined by measuring U_{ADC} when V_{pos} is connected to $V_{\text{ref}}/2$, and V_{neg} is connected to ground:

$$G = \frac{U_{\text{ADC}}\left(V_{\text{ref}}/2, \text{GND}\right) - \Delta}{512} \cdot 2 \quad . \tag{1.8}$$

Both G and Δ are independent of the V_{ref} tuning. The V_{ref} tuning can be performed by measuring or applying a known external voltage. The value of the tuning register is varied until the value for V_{pos} obtained via Equation 1.7 matches the expectation. This approach requires no further knowledge of the chip or the environment. It needs to be repeated whenever the temperature or other conditions change.

The lpGBT designers also provide calibration data for all their production ASICs. The necessary formulae follow the same logic as the previously presented approach. However, they also require a temperature measurement of the lpGBT junction temperature. The temperature is necessary to tune the voltage reference and to correct higher-order effects in the gain and offset. The calibration data include calibration constants for the reference voltage tuning, gain, offset, and their temperature dependence [14].

GPIO ports

The GPIO ports of the lpGBT are used for additional slow control applications. Four separate pins are connected to reset lines for the right and left side CBCs and CICs. The CIC is reset by an active low signal, and the CBC by a high signal. An additional pin is connected to the active low chip reset of the VTRx$^+$ uplink. In the lpGBT, the pins need to be configured as outputs.

Both DC-DC converters provide a Power Good (PG) signal to GPIO pins configured as inputs. The PG signals can indicate problems with the DC-DC converters as long as the provided power is sufficient to read out the module. Possible problems that are indicated by a low PG include an output voltage outside a $\pm\,6.5\%$ window around the set nominal, and in case the under-voltage, over-temperature, or over-current protections are triggered [16, 17]

One GPIO line ends blind in the left side FEH connector. It was intended for the control of an e-fusing circuit in the hybrid test system.

1.6 Hybrid quality assurance and quality control

The following section covers the challenges related to the quality assurance and quality control (QA/QC) of the Outer Tracker hybrids. It will present the implemented approach to tackle these challenges. The main focus is the 2S Service Hybrid, but key aspects can be transferred to the other types of hybrids.

The module design outlined in the previous section does not include redundancy in any of its on-detector components. To fully read out a single 2S strip the entire electronic chain from the sensor over the wire bond and the analog CBC input to the digital e-links of the CIC and lpGBT and the optical link needs to be functional. A malfunction of any one component directly translates to inefficiencies further down the chain. ASICs undergo production testing before they are mounted on hybrids. After their assembly, they cannot be removed[9]. The same is true for a hybrid after it has been glued on its module. Therefore, it is vital to thoroughly test each component. While single inefficient channels are tolerable, a bad digital line cannot be recovered due to the parallel nature of the data streams. A complete qualification must include production, performance, and reliability issues, which are increasingly difficult to spot.

Without considering any spares, the full Phase-2 Outer Tracker requires for its assembly 45,192 hybrids. Because of different sensor spacings and data rates, there are a total of 18 variants of the five hybrid designs. To produce the required amount plus an additional 5% in reserve, the peak production rate might need to be as high as 4,400 hybrids per month at a single contractor. A QA/QC system needs high throughput and needs to provide quick feedback in order to stop production in case of serious problems. In addition, it must be flexible to adapt to varying production rates of different hybrid variants throughout the entire production.

An important requirement is a hybrid operation at the expected temperature environment of the tracker down to -35°C. If a climatic chamber is used to test at this temperature, the test equipment does not need to meet the hybrid requirements in terms of power and space restriction and radiation hardness, but it must be partially operated at cold, too.

[9] While it is technically possible to replace ASICs and frequently done during prototyping, the long-term reliability of reworked components always remains questionable and they will not be used in the detector.

1.6.1 Production and workflow

In light of the sheer amount of required hybrids and the demanding specifications, hybrid production is carried out in industry. To produce a Service Hybrid, a flex manufacturer provides the flexible printed circuits. The company performs a surface treatment to achieve sufficient adhesion of the bottom surface for the gluing of the carbon fiber stiffener. The flexible circuit and the stiffener are baked out before gluing. After the lamination of the carbon fiber stiffener, two coefficient of thermal expansion (CTE) compensators are laminated under the carbon fiber. Lastly, the stiffeners above the FEH connectors made of FR4 (a class of glass-reinforced epoxy laminate material) are laminated on the top layer. Subsequently, all passive and active components are soldered by an assembly company. This includes the connectors, ASICs, and the converter coils. Underfill is applied for the bPOL2V5 and lpGBT. The sensor bias voltage part of the flex is folded and glued to the compensator. The grounding inserts are mounted and the circuit and assembly are cleaned. The conformal coating in the DC-DC shield area and the bias voltage filter is applied. As a last step, the shield is soldered by hand and a sticker with the ID matrix code is attached.

All of these steps come with a set of sub-steps and extensive requirements that must be obeyed by the contractor.

After the hybrid is produced, the contractor performs 20 passive thermal cycles between +50°C and −35°C. Afterwards, all hybrids pass through a warm and cold functional test at +40°C and −35°C. Only hybrids that pass this test and are found fully functional are delivered to CERN. CERN carries out a visual inspection of the hybrids. The Hungarian Wigner Research Centre for Physics assists CERN with 1/3 of the workload.

During the visual inspection, the readability of the ID code is checked. The hybrid flex must be free of wrinkles and the hybrid must not have a bow. The hybrid is inspected in terms of damages to the flex, stiffeners, compensators, or components and for cleanliness of the circuit and the connectors. The alignment of the various layers is checked with a pin test of the guiding holes. The fold-over needs to be accurate and the endpoint of the tails' adhesive in the correct location. The adhesive aspects between the flex and carbon fiber stiffener, the fold-over and compensator, and the compensator and carbon fiber stiffener are controlled. One verifies the quality of the soldering and if all visible components are correctly mounted and aligned. Visually inspected are also the quality of the underfill, the conformal coating, and the mounting of the grounding inserts. Observed problems are differentiated between repairable and not repairable. Hybrids with repairable features return to the contractor, where they are reworked, and re-join the production at the stage of the cold functional test.

After the inspection, all hybrids are subjected to another cold functional test at CERN (all 2S hybrids and 1/3 of the PS hybrids) or Italian National Institutes for

Nuclear Physics (INFN) in Catania (2/3 of PS-FEH) and Genova (2/3 of the PS-POH and PS-ROH). To assess the long-term reliability, a sample of the hybrids is selected for long-term thermal cyclings. A system for prototype testing (Chapter 3) was adapted for this task for the SEH [38]. It uses the so-called "test board". Only hybrids that pass the visual inspection and the functional tests are accepted and stay with the CMS Collaboration. All others return to the contractor.

1.6.2 Requirements of the functional testing

The same contractor is expected to produce multiple hybrid variants in parallel. The variants will not be produced, and thus are not ready for testing, at constant rates. The deliveries to the CMS Collaboration and the subsequent testing will be similarly affected. This calls for a test system with adjustable throughput and as many common parts as possible. It should also be compact and scalable.

Different solutions for each hybrid type are required because of the different connector interfaces and geometries. The tight requirements for the hybrids in a tracking detector do not allow for integrated test features on the level of the hybrid itself (neither mechanically nor electrically). The hybrid will need to be mounted on, fixed to, and connected to additional hardware. These manual actions must be easy, reliable, and fast to perform.

The functional testing will be performed by non-experts. The test system therefore needs to be useable without special training and without frequent maintenance. Gathered results need to be presented clearly and unambiguously, preferably as either "pass" or "fail". Nevertheless, a test should be thorough and results need to be stored for expert review or later analysis.

It is foreseen to perform the functional test at the lowest and highest expected detector temperatures inside a climatic chamber. For the "cold" case a temperature of -35°C is reached when the tracker is not powered and the CO_2 cooling is running. The "warm" extreme would be a tracker operation in a powered state with no cooling. An ambient temperature of $+40$°C is targeted to emulate this rare condition. The components need to start up and operate under these conditions and so must all components of test system inside the chamber. To reduce the required time, the transition between different temperatures should be as fast as possible without presenting a thermal shock to the components. The set target rate is 3 K/min. The humidity in the ambient needs to be controlled to avoid condensation. A climatic chamber can provide the required characteristics and has enough space to fit a system to cycle many hybrids at the same time.

The time required to mount and cycle hybrids can be recovered partially by using the period, while the climatic chamber is operated, to (un-)mount the (previous) next set of hybrids. Hardware with a mounted hybrid should therefore be removable from the system and available in large enough quantities.

If the time to perform a test on a hybrid is of the order of a few minutes, the testing time is dominated by the temperature cycling. In that case, full parallel testing of all hybrids in the climatic chamber is unnecessary. A serial approach that shares resources between hybrids makes the system cheaper and easier to implement.

The test system needs to verify the integrity of all analog and digital features of a hybrid. The highest amount of, especially differential, data lines that need to be externally accessed and tested are present on the PS-FEH. Its test system demands the most resources and thus sets the dimensions.

To reduce the time and effort required for the development, the system should be compatible with existing software and firmware that is present in the tracker upgrade project.

1.6.3 Common testing hardware and infrastructure

Figure 1.17 shows a schematic overview of the hybrid test system that accounts for the previously outlined requirements. The components are explained below. Hybrids are mechanically installed on and electrically connected to interface cards called "test cards". Each test card is specifically designed for a certain hybrid flavor but follows common mechanical and interface specifications that allow a connection to a multiplexer backplane inside a crate. The crate can hold three interconnected backplanes and a single backplane can support four test cards. Thus, twelve hybrids can be tested in series.

The test cards and the backplane are project-specific developments. The electrical, and in the case of the SEH and ROH the additional optical, readout is performed by Micro Telecommunications Computing Architecture (µTCA) components that were developed for the DAQ of the current CMS detector.

The shared constituents of the test system are presented in detail below. The system and all required boards were developed by CERN, except for the 2S-SEH test card. The 2S-SEH test card was designed by the I. Physikalisches Institut B of RWTH Aachen University and more details on it are given in Chapter 4.

1.6.3.1 Test cards

All test cards are housed by a standard 3U-size 19-inch crate. The imposed mechanical constraint is a PCB size of $100.00 \times 228.00\,\mathrm{mm}^2$. The only electrical connection is one FPGA Mezzanine Card (FMC) connector (Samtec Searay [39]), which is placed on the center-line of the short side and connects to the backplane. Two independent guide post holes are placed symmetrically to this center line at a distance of 84 mm from each other. The last common mechanical feature of all cards is two 3.5 mm holes for the attachment of a front panel on the opposite side of the PCB. Figure 1.18 shows the basic dimensions of all test cards. Components can be present on both

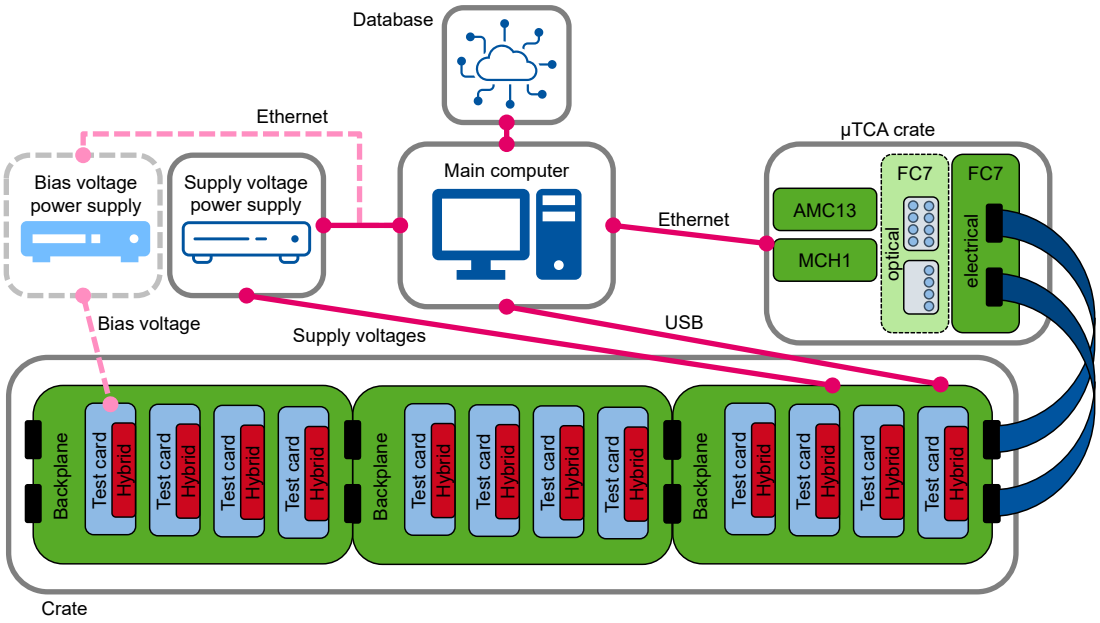

Figure 1.17: Schematic overview of the hybrid test system and its relevant components. Dashed and lighter color connections and components are only required for certain hybrids. The optical fiber and bias voltage connections for every test card are not shown for the sake of readability. Details are given in the text.

sides, but to avoid mechanical clashes, the maximum height in a given area must be less than 20.3 mm.

The hybrid under test is hosted on the test card with a dedicated fixture. Jumper cables allow access to the signals present on the hybrid connectors.

The test cards are powered with 3.3 V and −3.3 V from the backplane and receive an additional voltage to power the hybrid under test. All required digital signal transmissions and analog measurements are performed by the test cards.

The main interface to control the test cards is USB. The USB communication is hosted by a single-chip USB-to-SPI bridge (Silicon Labs CP2130 [41]). It is common to all test card flavors and features one-time programmable fuses. The main test computer uses the fusible product string (the USB standard also refers to it as Product ID) to identify the card type and serial number. Subsequently, the computer controls the test cards from the USB interface.

A test card is inserted upright into a crate's frame and guided by rails to connect with a backplane. Additional accessories like fixtures, cables, and jumpers vary between test card types and depend on the specific needs.

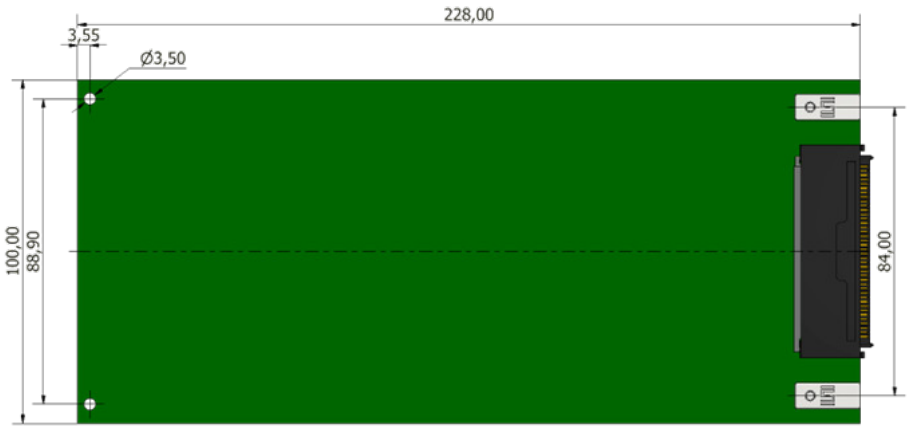

Figure 1.18: Basic test card dimensions [40]. All dimensions are given in millimeters.

1.6.3.2 Backplane

The backplanes are installed on the backside of the crate. One can accept four test cards and up to three backplanes can be connected in series by right-angle connectors on both sides. This enables the installation of small-scale test setups and reduces repair costs. Guideposts on both sides of the test card connector ensure the proper insertion of the cards.

Only the first backplane in the chain is connected to the µTCA board, the power supply, and the USB of the main computer. The backplanes multiplex all required high-speed input and output signals and the USB link and distribute the power. No hybrid-specific test functionalities are implemented.

Because it is the simplest layout and requires the least amount of components, a tree architecture with 2:1 multiplexers was selected [42]. Logic signals from the µTCA board are used to select the proper backplane and test card. After selection, the multiplexing is transparent to the main computer and test card and does not influence test results.

Figure 1.19 shows a perspective view of a crate with three backplanes and eight generic test cards for illustration.

1.6.3.3 µTCA components

Since its beginning, the tracker prototyping effort used readout hardware that is designed for and used by the current CMS detector. The comparable requirements in terms of data rate and interfaces made it a good choice for the test system. The relevant components for the hybrid test system are based on the µTCA and AMC standards [43, 44]. µTCA stands for Micro Telecommunications Computing Architecture and the standard “defines a backplane which can accommodate up to 12 Advanced Mezzanine Card (AMC) modules, along with up to two MicroTCA

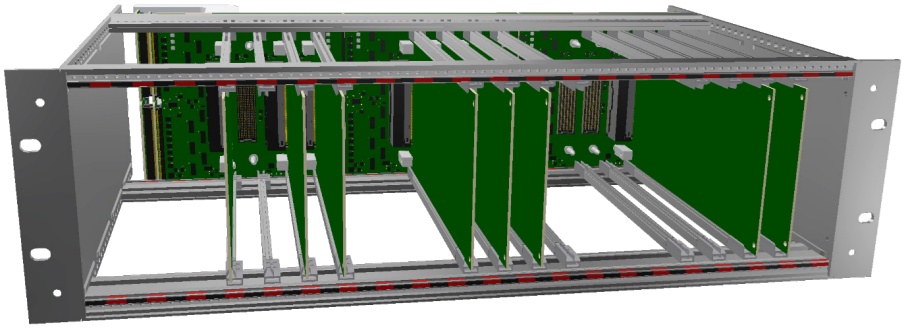

Figure 1.19: Perspective view of a crate for the hybrid production test system. The three interconnected backplanes are visible at the backside of the crate and several test cards are inserted [40].

Carrier Hub (MCH) modules" [45]. In a telecommunication application, the two MCH sites (MCH1 and MCH2) provide redundant backups of a commercial module that provides management and Ethernet service to all modules in a crate.

The FC7 is a µTCA compatible full-size, double-width AMC and is designed around the Xilinx Kintex-7 FPGA and the FMC standard [46]. It features two 400-pin Samtec SEARAY sockets that provide up to 20 high-speed serial links. The boards target generic data acquisition and system control uses. The FPGA design is split into two parts. A system core implements all basic infrastructure such as an Ethernet link based on IPBus (a simple IP-based control protocol designed for controlling xTCA-based hardware [47]), the configuration of the clocking circuitry, the control of the powering circuitry, SD card programming, board identification and various monitoring tasks. The second part is the user logic that is written individually for a specific application. The programming firmware is stored on an SD card and loaded into the FPGA. The firmware framework is called *µDTC* and used in several subprojects of the CMS Phase-2 tracker upgrade [48].

The CMS experiment uses the MCH2 site to house an AMC13 (the thirteenth AMC) module [45]. It takes advantage of this connection to provide a central clock (LHC machine clock), timing, and DAQ services within the µTCA crate environment. The hybrid test system mainly uses the AMC13 to synchronize the clocks of several FC7s via a common signal on the µTCA backplane. This requires the installation of a fiber in loop-back configuration on the Trigger Timing and Control Small Formfactor Pluggable (SFP) on the AMC13's front panel.

The electrical FC7 connects to the backplane via two high-density cables (Samtec HDR-169473-**), which are plugged to the FMC connectors. To allow optical readout, an additional optical FC7 in conjunction with commercial FMCs using the common SFP+ standard is required. The available solutions provide access to eight optical channels on each FMC socket so that one optical FC7 is sufficient for a crate with 12 hybrids and test cards. It is important to set the clock source of both

the optical and electrical FC7 to the AMC13 clock. Test cards that do not require optical readout can function with the electrical FC7's internal clock.

The FPGA enables the electrical FC7 to control the hardware selection signals and to provide the communication platform for the Low Voltage Differential Signaling (LVDS) signals. Special test routines are also implemented in its firmware. The optical FC7 runs a firmware image that is also used for the 2S module and PS module readout in the 2S-SEH test card and PS-ROH test card application, respectively. The MCH1 and the Ethernet interface provide the connection to the main test computer, which also executes part of the data processing.

1.6.3.4 Test enviroment

In addition to the custom components in the 3U crate and the µTCA crate described above, several commercial components are required. Partly, they have already been introduced above. The mechanical parts of the 3U crate as well as the MCH1 and µTCA crate itself are generic. The optical FC7 is only required when testing SEHs or PS-ROHs. Its up to twelve fiber pairs for the optical link can be connected directly to the front panels of the respective test cards. They use commercial LC and MPO fiber variants [49, 50].

A standard PC equipped with a Linux system runs the test software. A *C++* based driver for each test card flavor is used to communicate with a test card via USB. Those drivers are integrated into the likewise *C++* based test procedure, which is written in the data acquisition framework for the Phase-2 tracker upgrade, *Ph2_ACF*. The framework is capable of communicating with the FC7s, which are set up with their appropriate firmware. The communication is based on Ethernet and the IPBus protocol. A graphical user interface (GUI) written in *Python* takes care of registering the hybrids together with the test cards they are tested on, selecting the card on the backplane, running the test procedure, and uploading the test data to the CMS construction database. Figure 1.20 gives an overview of the test software and interfaces. It also shows that the FC7 is interfaced with a hybrid's ASICs through the FMC cables, backplane, and test card.

The required voltages are provided by two power supplies. A Hameg HMP4040 four-channel power supply [52] can deliver up to 32 V and is connected directly to the backplane. The system differentiates between the supply voltages and the test voltage. The supply voltages are used to power the backplanes and test cards themselves. The test voltage powers the hybrid under test. There are two general operation modes of the test voltage, which are set via the firmware of the FC7 connected to the backplane, a "front-end" mode and a "power hybrid" mode. In the front-end mode, the backplane forwards a test voltage between 1.3 and 3.0 V to the test cards. Subsequent linear regulators on the 2S-FEH, PS-FEH, and PS-ROH test cards provide the appropriate voltage to their respective hybrids. In the

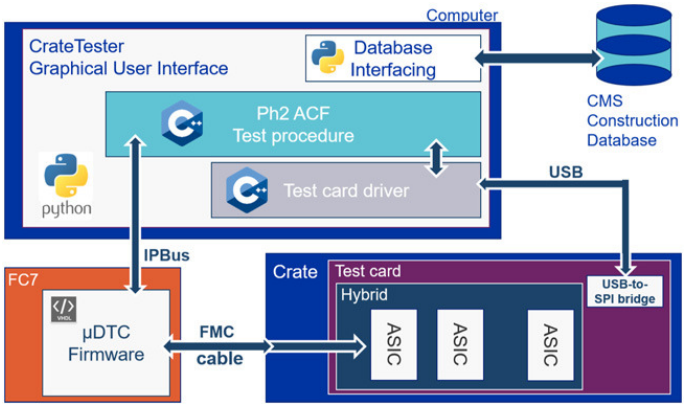

Figure 1.20: Simplified block diagram of the test system software and firmware [51]. It illustrates the Front-End Hybrid case without power supply control and additional optical readout.

power hybrid mode, the test voltage is directly applied to the hybrids on the 2S-SEH and PS-POH test card as input of the hybrids' first-stage DC-DC converters. Voltages between 4.5 and 12 V are possible. The test cards cannot be selected for values outside the limits. If already enabled, they will be disconnected. For reasons of safety, the default setting is the front-end mode. All test cards can be selected safely and identified from the product string in this mode.

The Service Hybrid and the right-side PS-FEH host their modules' bias voltage filter in addition to their other functionalities. To test the bias voltage circuits, an additional power supply is required. The iseg SHR [53] is used, which has four channels that each can provide up to 2 kV. The bias voltage cannot be provided from the backplane to the test cards. Instead, a coaxial high-voltage cable connects the SHR's output to a high-voltage compatible connector on the test card's front panel. Additional cables with a male and female connector can be used to daisy-chain the bias voltage of neighboring test cards. In this way, only one power supply channel and connecting cable are required for all test cards and hybrids in a crate that are tested one after another.

Both power supplies feature an Ethernet interface and remote-control capabilities.

Both the GUI itself and the test procedures it executes must be able to control the power supply to set voltage levels, toggle channels, or read back measurements. A software-based *TCP* (Transmission Control Protocol [54]) server is used for the implementation. An *XML* settings file contains information on the type of device and interface and defines the configurable channels and their identifier names. Using the identifier, the GUI or the test software can send requests to the server, which handles the actual communication with the power supplies based on the information in the *XML* file.

To enable the thermal cycling, the crate needs to be installed in a climatic chamber. The selected model at CERN is an EXCAL2 4025-HE [55]. The chamber is temperature and humidity-controlled and large enough to house three crates. To run parallel tests on all three crates, three electrical FC7, two Hameg HMP 4040, and one iseg SHR are necessary. Each crate that tests hybrids with a lpGBT requires an additional optical FC7.

CERN has collected and installed all the required components to operate three crates in the climatic chamber. Most test cards and their accessories are available. The system along with its software was commissioned and is being finalized for hybrid production [56].

The 2S-SEH test card was designed, prototyped, and ultimately produced by the I. Physikalisches Institut B of RWTH Aachen University. The produced cards have been extensively qualified and then provided to CERN for hybrid production. The development in the framework of this thesis also contributed to the various software and firmware efforts and the hybrid prototyping campaigns in general.

2 Service Hybrid prototypes

This chapter provides information on the various generations of Service Hybrid prototypes that were assembled in the context of the Phase-2 upgrade of the Outer Tracker. The first section briefly summarizes the history of the hybrid design and important milestones on the way to the optical readout of prototype modules. The first SEH used for the assembly of optical 2S modules is introduced in the next section. The section also provides information on the used ASICs and the hybrid design. The last section covers the first prototype with the final chipset. It discusses the design, the ASICs, and their configurations.

2.1 Summary of the prototype generations

In 2015, the technical proposal of the CMS Phase-2 upgrade already foresaw the Service Hybrid as part of the 2S module, its tasks, and the main components [57]. In particular, it included the use of ASICs that are centrally developed by CERN to fulfill a certain role and can be used by different experiments and projects. One refers to such chips as "common ASICs". An example of a common ASIC is the FEAST2 ASIC, which was designed by the CERN Group for Electronics Systems for Experiments. It is a radiation and magnetic field tolerant point-of-load buck DC-DC converter [58] and for example used by the CMS Collaboration to power the Phase-1 pixel detector. The pixel DC-DC converter boards that feature the FEAST2 have been operating since 2017 and were designed, produced, and tested by the I. Physikalisches Institut B of RWTH Aachen University [35, 59].

At the time of the technical proposal, it was evident that DC-DC converters would be required to power a future tracker. The use of the lpGBT and VTRx$^+$ module for the optical readout of the modules was also planned. Although no functional ASIC prototypes were available, it was already clear that more than one supply voltage was required on the module. The CMS Collaboration decided to implement a "two-step DC-DC conversion scheme, with two DC-DC converters working in series (...), taking into account system efficiency, cabling aspects, and required development effort" [59].

Following this decision, the first Service Hybrid prototype (v1) was designed by the I. Physikalisches Institut B of RWTH Aachen University. Shape and dimensions followed the proposed design, yet no interface to a FEH prototype was present. It

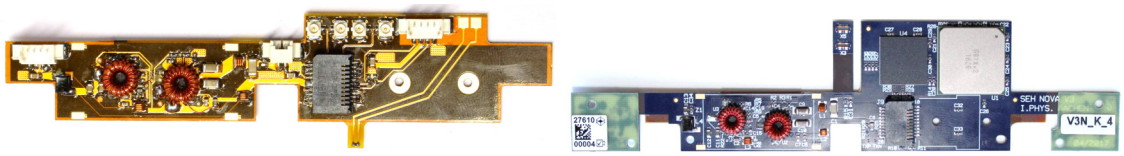

Figure 2.1: Early Service Hybrid prototypes. Left: version v1. Right: version v3. The DC-DC converter shield is not assembled on both hybrids. See text for details.

included a Versatile Transceiver (VTRx) module receptacle, connectors for the high-speed link qualification, and a sensor bias voltage filter circuit. The main purpose of these SEH prototypes was the first implementation of the two-step DC-DC conversion scheme. The design used the FEAST2 for the first stage and a commercial chip in the second stage. Figure 2.1, left, shows a picture of a v1 prototype.

It was demonstrated that the approach of a two-step scheme in general worked [59]. The second SEH generation featured three variants with different bias voltage RC filter circuits. It also had more practical connectors. It was already possible to power a full-sized 2S module with this generation. It still required soldered connections and an electrical readout with additional hardware [12, 60].

A major step forward was the development of the third prototype generation. Figure 2.1, right, shows a picture of the first hybrid produced within that generation (v3). It targeted the first optical readout of the 2S module. Therefore, common ASICs that are predecessors of the lpGBT needed to be integrated into the Service Hybrid. It also required a decision on a data connector between the Service Hybrid and the Front-End Hybrids. For this purpose, two different variants were tested. The powering and the optical readout were successfully demonstrated [61]. Unfortunately, a layout bug prevented its deployment on a large scale.

The bug was fixed in a second iteration of the third prototype generation. Those hybrids got version number v3.1 and their design and components are described in the next section. Section 3.2 provides more details on their performance and measurements in a prototype test system.

The first prototypes of the common ASICs for the SEH became available in early 2020. Another prototype design was required that included them. This design and its prototyping campaign are summarized in the last section of this chapter. Chapter 5 presents the characterization of those hybrids. They were designed and produced under the supervision of CERN, who took over this responsibility from the I. Physikalisches Institut B of RWTH Aachen University.

2.2 Development of the SEH v3.1

The main goal of the production of version v3.1 of the Service Hybrid was to use them in prototype 2S modules and exercise their optical readout. In 2020, FEH hybrids with eight CBC ASICs were available that could be read out electronically via interface boards. They did not yet host a CIC ASIC, but a CIC could be added via a mezzanine card. After the first experiences with third-generation Service Hybrids, they were equipped with a 50-pin socket (Panasonic A35P [31]) as the designated interface to the SEH. The interface imposed the same requirements in terms of provided and accepted signals as in the final module. Since the lpGBT, VTRx$^+$ module, and both bPOLs were not yet available, replacement ASICs were needed. Instead of the bPOLs the same DC-DC converter design that had been used in earlier versions could be used. The lpGBT was substituted by the Gigabit Transceiver ASIC (GBTx) and the GBT-SCA. The latter is necessary because the GBTx does not provide the I^2C controller interfaces that are required to control the front-end ASICs. The opto-electrical conversion was performed by the VTRx module. The bias voltage circuit was implemented according to the experience gained with the second generation. While the footprint of the DC-DC converters fitted on one half of the baseline SEH design, the combination of GBTx, GBT-SCA, and VTRx module is much larger than the allowed space. The PCB in Figure 2.1, right, had to be increased compared to the baseline model. The following sections give more details on the PCB design, the ASICs, and the components.

2.2.1 Design

Figure 2.2 shows a picture of an assembled v3.1 prototype hybrid. It is made of a flexible polyimide PCB with four layers and a FR4 stiffener. The reference copper thickness is 17 µm with staggered vias. The ASICs, the input voltage connectors, and the VTRx connector are placed on a rigid part with the 0.55 mm thick FR4 stiffener glued to the bottom. The stiffener has a 70 µm copper layer on both sides. The connectors to the FEH ("FEH connectors") are placed on the bottom side at both ends of the hybrid and have a 0.3 mm FR4 stiffener on the top side. The bias voltage filter is placed on the top of a fold-over section and glued to the bottom side of the stiffener below the DC-DC converters. The hybrid is 134.5 mm long and 15 mm wide. Except for the area where the GBTx, GBT-SCA, and the VTRx connector are placed, which is wider (34 mm) than the baseline model. The VTRx module can be connected with its fibers pointing away from the DC-DC converters. A small tail section between the FEH connector and the GBTx and second stage DC-DC converter output filter, respectively, is not glued to the bottom stiffener to remain flexible and compensate a height difference between the SEH and FEH. The DC-DC converters are covered by a 29 mm by 14 mm large and 5 mm high aluminum shield. Parts of their input and output filters are placed outside the shield to avoid radiated noise coupling. The shield is manufactured from about 150 µm thick foil

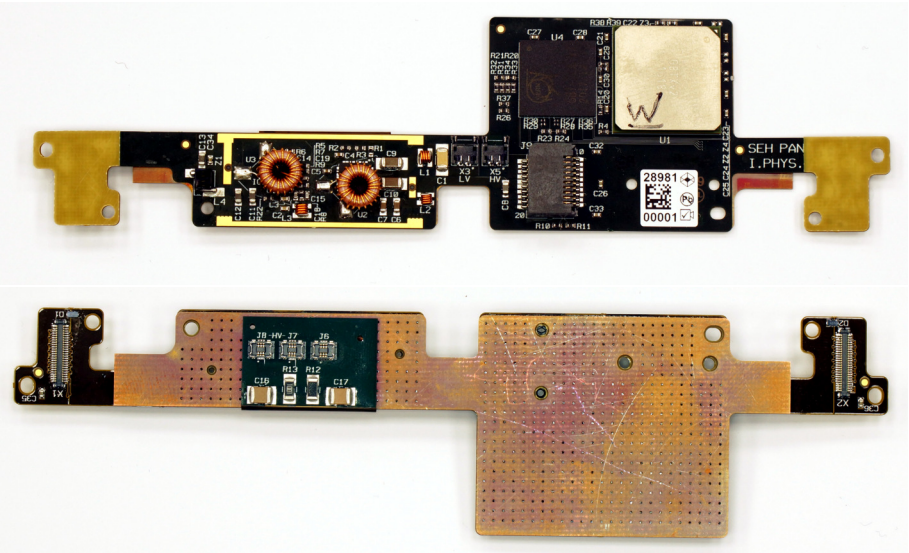

Figure 2.2: Pictures of the top and bottom side of a Service Hybrid v3.1 prototype. The DC-DC converter shield is not assembled.

that is folded and has soldered edges. It is plated with nickel and tin and soldered all around to the PCB with small openings where signals and traces are routed on the top layer.

For correct operation, the ASICs require additional passive components (in general resistors, capacitors, and inductances). They are chosen according to the recommendations in their respective datasheets. More details on the ASICs are given below. The respective sections also discuss the selections of passive components that need to be changed by the designers to configure the ASICs.

2.2.2 ASICs

GBTx

The GBTx ASIC is part of the larger Gigabit Transceiver chipset, which also includes the GBT-SCA, a Gigabit Trans-Impedance Amplifier for the optical receiver (GBTIA), and a Gigabit Laser Driver (GBLD) [20, 62, 63]. The latter parts are integrated into the VTRx module. In all of those chips, total dose and other radiation effects are mitigated following special layout techniques in a commercial 130 nm CMOS technology. The GBTx is designed to provide a high-speed (3.2-4.48 Gb/s) bidirectional optical link. Analog to the lpGBT link scheme it combines the three distinct data paths for timing, trigger and fast control, data acquisition, and slow control on a single optical link. The GBTx "implements all the needed functions of the data and timing transceiver" [62]. On the downlink, its clock and data recovery (CDR) circuit receives high-speed serial data from the GBTIA. For the correct sam-

pling of the incoming data stream, it recovers the high-speed clock from the data. Subsequently, the data stream is converted to parallel form and decoded (error correction and descrambling). On the uplink, the reverse operations are performed. The data to be sent are DC-balanced via scrambling, encoded with a forward error correction, and serialized. The GBTx configuration can be set through several hardwired pins and a large amount of registers. It is therefore highly flexible and in turn complex in its operation. The next paragraph briefly covers the most important features used in the SEH application.

In the SEH, the GBTx is hardwired to operate as a transceiver with bidirectional data communication. Differential e-links connect to the FEH connectors and their e-port receivers feed the data to the serializer. Transmitter e-ports receive data from the CDR circuit which are forwarded to the FEH. The e-links are set up for low-voltage differential signaling, whose settings can be optimized. The GBTx's clock manager also provides differential clocks needed for the front-end ASICs. A control and monitoring logic is implemented in the GBTx to provide the configured operation mode. E-fuses set the initial configuration. After the optical link is established, it is used to configure the chip. The reset during start-up is latched by the second DC-DC converter stage's Power Good signal. An I^2C interface to control the laser driver on the VTRx module is present. The recommended supply voltage is 1.5 V [62]. In addition to the ASIC itself, the GBTx package assembly includes six decoupling capacitors, a crystal oscillator, and a 0.5 mm thick copper heat sink. The total size of the substrate is $17 \times 17 \times 2.97\,\mathrm{mm}^3$. Its ball grid array (BGA) matrix has 20×20 balls with a pitch of 0.8 mm and a diameter of 0.5 mm.

GBT-SCA

The GBTx does not provide all interfaces that are required to operate a 2S module. The missing interfaces are all related to the slow control data path. The GBT-SCA complements those needs. Like the GBTx, the GBT-SCA is a highly flexible and powerful chip. The relevant features are briefly covered below.

The GBT-SCA provides up to 16 independent I^2C interfaces, two of which are used by the SEH. It also features a 12-bit ADC with 31 multiplexed analog inputs. Seven inputs are used in the SEH. Four of the 32 GPIO interfaces are used to provide reset signals for the front-end ASICs, and two are used to measure the DC-DC converter Power Good signals. The connection to and control by the GBTx is realized by an 80 Mb/s e-link. The EC field of the GBTx high-speed link is used to feed and receive the data. The supply voltage is 1.5 V. The size of the ASIC is $12.15 \times 12.15 \times 1.7\,\mathrm{mm}^3$. The GBT-SCA features a 14×14 BGA matrix with 0.8 mm pitch and a ball diameter of 0.5 mm [20].

GBTIA and GBLD (VTRx module)

The optical readout components that were used in the original LHC detectors featured a diverse set of solutions. For the Phase-1 and 2 upgrades, the dispersed development efforts were regrouped in the Versatile Link (VL) and Versatile Link plus (VL^+) projects. Figure 2.3 gives an overview of the Versatile Link architecture. The VTRx module is the result of the VL effort and is targeted for installation in the Phase-1 upgrades with the radiation resistance necessary for calorimeter usage. Its design was driven by a generalist vision to provide "a bidirectional front-end transceiver implemented as a customized module based on the well defined standard SFP+ format" [63].

The VTRx module is responsible for the opto-electrical conversion of the high-speed data link. It is an additional PCB with active, passive, and mechanical components. The ASICs are part of the GBT chipset, while the laser and photodiodes are selected from commercial components after their radiation hardness was qualified. The optical module connects to the hybrid via a variant of an SFP edge connector and to the back end via a multi-mode LC fiber pair. A radiation hard laser driver (GBLD) and a vertical cavity surface emitting laser diode are used to transmit the uplink data. The GBLD can be configured via an I^2C interface from the GBTx. The optical signal from the back-end is received by a GaAs photodiode and a limiting amplifier in combination with the GBTIA. Those components are placed on a flexible subassembly and forward the data electrically to the GBTx. The supply voltage of the VTRx module is 2.5 V. The received signal strength can be measured analog to the RSSI voltage of the $VTRx^+$ (Figure 1.16, right) [64].

FEAST2

The FEAST2 is a radiation-tolerant synchronous buck DC-DC converter [65]. It is produced in a 350 nm process and can sustain a total ionizing dose of more than 200 Mrad and an integrated particle fluence of $5 \times 10^{14}\,\mathrm{n_{eq}/cm^2}$. Due to its radiation hardness, it is used for instance to power the current CMS pixel detector. For the v3.1 prototypes, the radiation hardness is not important, as a radiation-tolerant variant for the second stage is missing. Still, the bPOL12V is largely based on the FEAST2, so it was a good choice for implementation in the prototype. As long as an output power of 10 W is not exceeded the ASIC can provide up to 4 A of continuous output current and is rated for input voltages between 5 and 12 V. The output voltage can be selected between 0.9 and 5 V. A voltage divider is used to set the output voltage. It reduces the output voltage by a fixed factor, which is then compared to a voltage reference (599 mV). To obtain 2.5 V the resistor values in the divider on the hybrid are 316 kΩ between the reference input and ground and 1 MΩ between the reference input and the output voltage. In the pixel detector, the FEAST2 can be disabled during detector operation via a control line. A corresponding line is not foreseen for the Phase-2 OT modules, so the ASIC is configured to be always enabled. The

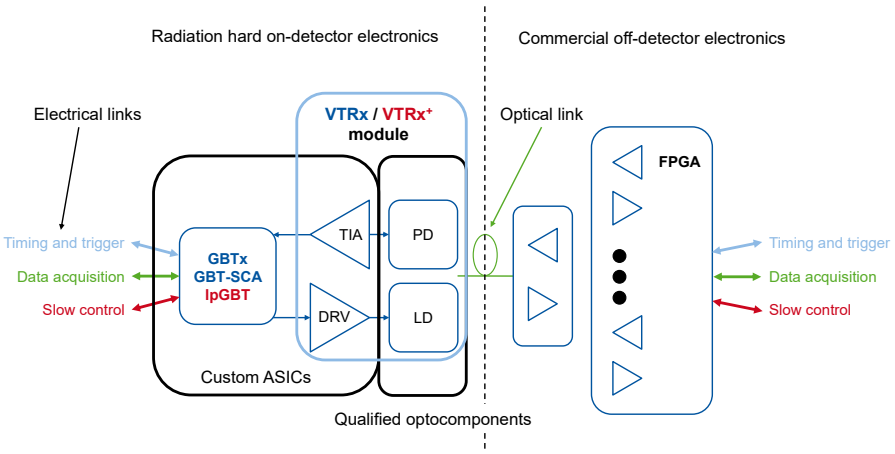

Figure 2.3: General architecture of the Versatile Link. An optical link connects the radiation-tolerant components in the detector with commercial off-detector electronics. Photodiodes (PD) and laser drivers (LD) in combination with trans-impedance amplifier (TIA) and driver (DRV) ASICs perform the opto-electrical conversion. In the VL (respectively VL$^+$) case, the GBTx (lpGBT) extracts the three distinct data paths from the serialized link it receives from the VTRx (VTRx$^+$) module [14].

under-voltage lockout (UVLO) therefore determines the minimal voltage at which the DC-DC converter output is disabled. Due to an intrinsic hysteresis, the UVLO enables the DC-DC converter at 4.8 V and disables it at 4.5 V. A Power Good status signal is output by the DC-DC converter when it is enabled and the output voltage is within $\pm 6.5\%$ of the set value. A soft start procedure prevents large inrush currents. After the DC-DC converter is enabled, the set output voltage is reached after roughly 500 µs. The duration of the soft start depends on the switching frequency, which is set to 2 MHz. The FEAST2 also features an over-current and over-temperature protection. The FEAST2 is packaged in a plastic Quad Flat No-Lead (QFN) package, has 32 pads with a distance of 0.5 mm, and is in total $5.0 \times 5.0 \times 0.9\,\text{mm}^3$ in size. A thermal pad in the center is soldered to the PCB. For more details on the FEAST2 and its usage in the CMS experiment see Reference [35].

LTC3412A

The SEH v3.1 employs the LTC3412A from Analog Devices as the second-stage DC-DC converter [66]. It is a synchronous buck DC-DC converter for industrial applications and not radiation tolerant. The recommended input voltage range is 2.25 V to 5.5 V, which means it operates near its lower edge in the SEH. The DC-DC converter provides up to 3 A of continuous output current. The DC-DC converter's output is enabled for a voltage above 0.5 V on the enable pin. It is connected to

the FEAST2's output, which means that the chip is always enabled if it receives an input voltage. The DC-DC converter features a UVLO for voltages below 2.2 V. An output voltage of 1.6 V is set via a voltage divider. A resistor sets the main switching frequency to 3.05 MHz. A Power Good is output if the DC-DC converter is enabled and the output voltage is within $\pm 7.5\%$ of the set value. The DC-DC converter features a soft start and an over-temperature protection, too. The LTC3412A is also packaged in a QFN package, yet has 16 pads with a distance of 0.65 mm and a size of $4 \times 4\,\mathrm{mm}^2$. In addition, it has a thermal pad in the center that is soldered to the PCB.

Comments on powering

The buck converters need an inductance to store energy. Samples of a custom toroid wound around a ceramic core are used. The same coils with an inductance of 200 nH are chosen for both stages.

The CBCs and CIC require a supply voltage of about 1.25 V. This will also be the supply voltage for the lpGBT and partially for the VTRx$^+$ module. As described, the two DC-DC converters provide 2.5 V and 1.6 V to power the VTRx module and the GBTx and GBT-SCA, respectively. The latter two cannot reliably operate with the lower voltage level required for the CBC and CIC. In consequence, a third voltage is needed for the SEH integration in a prototype module. To preserve the two DC-DC converter approach in the SEH v3.1 design, while still delivering about 1.25 V, two Schottky diodes (NSR20F30NXT5G [67]) are placed between the LTC3412A output and the FEH connectors on the rail carrying the FEH supply voltage. When a current is flowing to the FEH ASICs, the diode lowers the voltage to an acceptable range (480 mV voltage drop at 2 A). Figure 2.4 illustrates the voltage distribution. Like the LTC3412A, the Schottky diodes are not radiation tolerant and consequently the v3.1 SEH cannot be used for irradiation studies.

2.2.3 Summary of the SEH v3.1 prototyping campaign

The I. Physikalisches Institut B of RWTH Aachen University manufactured multiple generations of SEH prototypes. As stated above, the earliest only featured the DC-DC converters and test structures to qualify the data links. They were used to demonstrate the two-step DC-DC conversion powering scheme and its ability to power early versions of the front-ends [59, 60]. The version 3 was the first prototype to also feature the GBTx and GBT-SCA together with the DC-DC converters. It demonstrated the possibility of optically reading out and controlling front-end prototypes with the DC-DC converters as the only power source [61]. The version 3 SEH prototype was only produced in a small number and not distributed in the community. It presented some features that required manual rework of the PCB to make them usable.

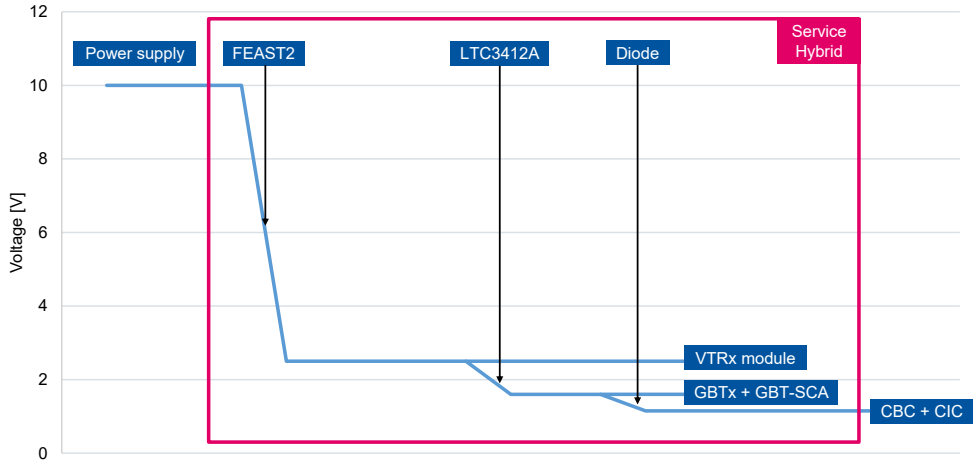

Figure 2.4: Illustration of the voltage distribution scheme on the v3.1 SEH. A power supply provides roughly 10 V of input voltage to the hybrid. In a first step the FEAST2 converts the voltage to 2.5 V which is required by the VTRx module. In a second step, the LTC3412A reduces this further to 1.6 V to power the GBTx and GBT-SCA ASICs. The supply voltage for the CBC and the CIC on a FEH needs to be reduced further by a voltage drop over a diode.

The first SEH version to be produced in larger quantities and intended for distribution in the community and use in module prototyping was version 3.1 (v3.1), of which 34 hybrids were produced in total. Their electrical performance is covered in Section 3.2.

Due to the entirely different ASICs, those hybrids bring only limited insights into the electrical performance of the final SEH. However, they were very valuable for the development of the optical readout software and firmware and for exercising and validating the electrical chain involving the CBC and CIC, which are the front-end ASICs used in the future tracker. The performance of the sensors and front-ends in combination with the two-step DC-DC conversion scheme with the bias voltage filter on the same PCB was extensively studied with little impact observed on the noise performance [68, 69].

The v3.1 also had an impact on the development of the module design and the assembly procedures. Strict precision requirements need to be observed during the manufacturing of a hybrid. However, the final hybrids impose fewer requirements in terms of installation precision than the sensors, bridges, and FEHs. Overall, the chosen connectors were found to suit the application. Earlier work showed good performance of the aluminum shield in combination with toroid inductors. The inductor design was not compatible with an automated assembly and needed to be improved.

The position of the FEH connector in the module was identified as a weakness. The rather short and straight flexible tail made it difficult to plug the SEH to the FEHs and caused stress on the fragile structure. As a result, several connectors were damaged. The proposed solution was to extend the tails into an L-shaped hybrid extension and to place the FEH connector further away from the SEH. The L-shape allows to bend in two axes and compensate for a potential horizontal as well as a vertical offset. It has the added benefit of eliminating the need for a 1.8 and 4.0 mm Service Hybrid variant. The required change in the FEH connector placement can be absorbed on the FEH, which is produced in two design variants anyway. It also helped to level the supply voltage for the CBCs, since the difference in length of the supply path and therefore in voltage drops on a Front-End Hybrid were reduced.

2.3 Final prototypes

In early 2020, the first prototype versions of the common ASICs became available for the CMS Collaboration and other experiments. They were used to assemble for the first time hybrids with the materials and components foreseen by the Technical Design Report [12]. The original design was adapted with the findings from the v3.1 prototyping campaign. In the context of this thesis, the hybrids that were manufactured between 2020 and 2022 by two contractors are referred to as the final Service Hybrid prototypes. They were received and tested in several deliveries.

2.3.1 Goals and limitations

Prototype versions of the custom ASICs CBC and CIC for the OT upgrade were available and tested on hybrid prototypes before 2020 (see previous section). After the successful demonstration of the optical readout and the two-step DC-DC conversion powering scheme for these front-end ASICs, the focus switched to the Service Hybrid chipset foreseen for the final application in the Phase-2 tracker. The common ASICs bPOL12V, bPOL2V5, and lpGBT and the VTRx$^+$ module only became available in 2020. The latter two are significantly smaller than their predecessors GBTx, GBT-SCA, and the VTRx module, thus reducing the required area to fit on the allocated space for the SEH. The DC-DC converter part with the bias voltage filter on the fold-over fulfilled the space requirements in v3.1 and so did the bPOLs' form factors.

The production of the several thousand hybrids that are required for the Phase-2 tracker will be carried out by industrial partners. The CERN procurement rules [70] require an "Invitation to Tender" for all acquisitions costing more than 200,000 CHF that is preceded by a market survey. The order of the OT hybrids surpasses this cost threshold. At the end of the market survey, only two possible suppliers remained.

The assembly and testing of the final prototypes were required to qualify them for the Invitation to Tender and to determine if they could handle the series production.

In addition to the demand for qualifying suitable companies, the prototyping campaign faced several competing interests that need to be taken into account:

- The development of testing hardware as well as software and firmware tools for the readout and data acquisition requires access to functional hardware. They need to be tested and debugged within the application to find problems and gain experience.
- The module prototyping effort has a high demand for functional prototypes. They are required to exercise the module assembly and to read out the front-end. This is vital to qualify the module performance e.g. in terms of noise in a stand-alone operation, on a larger structure, or as a particle detector in a beam test environment. As several different components need to be combined, delays in a single one propagate to the others.
- Prototype hybrids that were installed in modules are significantly harder to debug than bare hybrids. They are difficult or even impossible to remove and the different handling steps introduce additional causes of failure.
- While hybrids might be functional right after assembly or installation in a module, also long-term performance and stress tests e.g. in different temperature environments are required to verify the production quality.
- Issues that are not uncovered during the prototyping, but propagate to the final production, are a risk for the Phase-2 tracker. The delay between e.g. PCB production and large-scale functional testing is quite large and cannot be easily recovered. The same is true for introduced costs and lost materials.

These general considerations were accompanied by several challenges that arose during the campaign:

- Due to their special application, the OT hybrids are fully customized devices, which raises the production costs and increases production time. The prototypes in particular came in small lots and the procurement and manufacturing were lengthy processes.
- The supply of the common ASICs was limited, while they were themselves in the prototyping phase and being prepared for production. Between chip versions, their properties, configuration, and required passive components were subject to change on short notice and documentation was preliminary. It was a challenge to account for all this in the PCB design, the component selection, the software, and the firmware. Features that were encountered during the hybrid testing needed to be disentangled between design, assembly, and component failures, all of which delayed the availability for users.

- The availability of special materials like carbon fiber and passive components in general suffered during the COVID-19 pandemic. This introduced additional delays and costs in the project.

These competing interests and the resources required to overcome the challenges are the reason why the hybrid effort is centrally managed by CERN. The design and campaign outlined below tried to fulfill all the expectations in the final prototypes.

2.3.2 Design

The final SEH prototypes that are designed by CERN have a five-layer flexible polyimide PCB. It is 145 mm long and its copper layers have a nominal thickness of 12 µm – up to 23.1 µm are allowed on outer layers – and are connected by staggered and stacked copper-filled microvias. A 550 µm thick carbon fiber stiffener is glued to the PCB backside, except for the L-shaped tails at both ends. The tails end with the FEH connectors on their bottom side and a 300 µm FR4 stiffener on the top side. The carbon fiber has a CTE of almost zero, while the flexible circuit's CTE is dominated by copper and about 16 ppm/K. To avoid a bow in the hybrid, a CTE compensator made of fiber-glass reinforced polyimide foils is glued to the bottom of the carbon fiber stiffener. Since the compensator is a thermal insulator, it is split into two parts to keep the area free where the hybrid will be glued to and cooled by the AlCF bridges on a module. Figure 2.5 has pictures of both sides of a final SEH prototype. A fold-over section with the sensor bias voltage RC filter is attached to the compensator on the hybrid backside under the DC-DC converter section. The fold-over section also holds the connectors for the sensor pigtails, which provide the bias voltage and the connection to a temperature sensor.

Differential pairs are designed to have an impedance of $90 \pm 10\,\Omega$, except for the 2.56 and 5.12 Gb/s high-speed connections between the VTRx$^+$ module and the lpGBT with an impedance of $100 \pm 10\,\Omega$. To achieve the higher impedance, the design was optimized separately in a 5 mm long area. Underfill is used under the bump-bonded lpGBT and bPOL2V5 ASICs. As in the v3.1 hybrids, the DC-DC converters are covered by an aluminum shield. It is milled from a block to a thickness of 150 µm and plated with nickel and tin to enable its soldering. The milling yields better precision than earlier folded versions. Passive components near the shield are conformally coated. The coating is also applied to the components that make up the bias voltage filter on the fold-over. The main air-core inductors for the two bPOLs are a custom toroid design with a support made of soldering temperature-withstanding Liquid Crystal Polymer. The support provides a guiding pin and well-defined contacts for the assembly and is compatible with reflow soldering. A silicon-free thermal paste is applied in a circle on the top side of the coils to establish thermal contact with the shield. Several tooling holes are present in the hybrid to assist in the lamination of the different components, the folding of the fold-over, to provide guiding holes for the FEH connector, the main inductor coils, and the shield, and to allow the installation

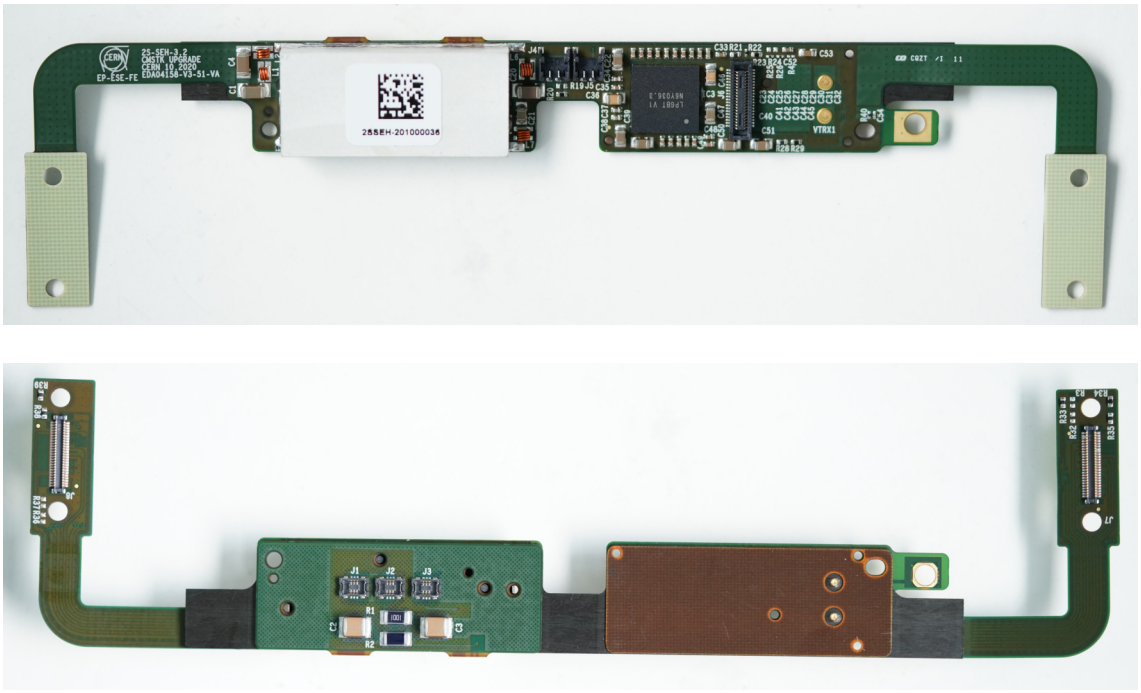

Figure 2.5: Picture of the top and bottom side of a final Service Hybrid prototype.

of two grounding inserts. Some holes are also used in the module assembly. The grounding inserts are press-fit pins that connect the hybrid ground to the carbon fiber stiffener. For the grounding of the AlCF bridges the hybrid has an additional short flap. This is necessary since the glue between the carbon fiber stiffener and the AlCF bridge is also an electrical insulator. The connectors for the hybrid's input voltages (low-voltage supply, ground, and bias voltage) are placed in the middle of the PCB. Figure 2.6 gives an overview of the hybrid stack-up and the placement of major components.

2.3.3 ASICs

For the final SEH prototypes, two different generations of the common ASICs (lpGBT, bPOL12V, and bPOL2V5) are used. A first batch uses prototype ASICs, while the second batch uses production-grade chips. Table 2.1 summarizes the versions used in the assembly of each batch. While the footprints of the different versions are identical and thus the same flex design can be used, some changes exist in the pinout and the functionality. The ASICs' functionalities are documented in the corresponding datasheets. The selection and placement of the passives is driven by the system requirements and the application. The following section describes the relevant features of the chips. In some cases, the prototyping effort foresaw the optimization of passive values. Their functionality and differences between the two batches are discussed. The VTRx$^+$ module is not an intrinsic part of the hybrid, as

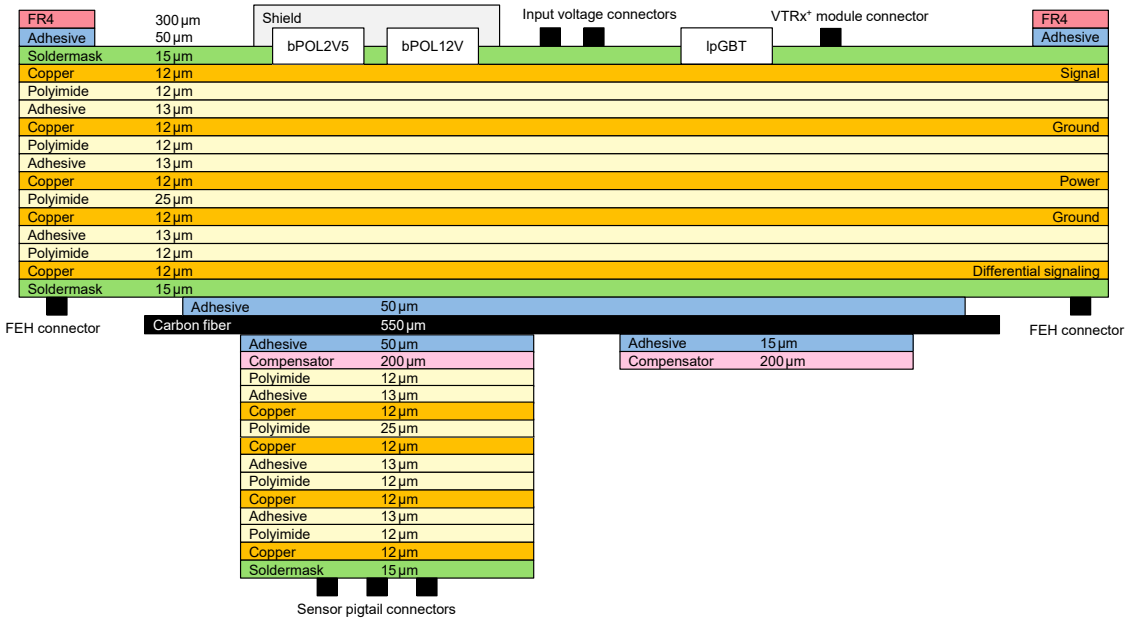

Figure 2.6: Stack-up of a final Service Hybrid prototype. The illustration shows a cross-section along the full length of a hybrid and is not to scale. The ASICs and connectors are also indicated.

it can be replaced easily. The connector pinout is the same on all of those. Depending on their availability, different versions of prototype and production samples were used throughout the developments.

lpGBT

The lpGBT ASIC can fulfill both the GBTx and GBT-SCA ASICs' roles and replaces them in the final SEH prototypes. It combines the three distinct data paths for timing and trigger control, data acquisition, and slow control on a single optical link.

Service Hybrid prototype	Batch 1	Batch 2
ASIC	Version	
lpGBT	v0	v1
bPOL12V	v4	v6
bPOL2V5	v3.2	v3.3
VTRx$^+$ module	All versions compatible with both batches	

Table 2.1: Overview of the ASIC version used for the assembly of the two batches of the final SEH prototypes.

It provides a bidirectional interface with constant latency between the on-module and back-end electronics with high reliability in the harsh radiation environment. To withstand the radiation levels of more than 1 Grad in some applications, "it uses an extensively radiation qualified commercial 65 nm CMOS technology, and special layout techniques" [14]. On the downlink, the lpGBT receives high-speed serial data (fixed to 2.56 Gb/s) and its CDR and phase-locked loop (PLL) circuit recovers the high-speed clock from the data. Subsequently, the data stream is converted to parallel form and decoded (error correction and descrambling). The uplink in the 2S module performs the reverse operations for an uplink speed of 5.12 Gb/s. The data to be sent is DC-balanced via scrambling, encoded with a forward error correction, and serialized. The forward error correction is set to the FEC5 mode, which allows the use of all seven uplink e-link groups. The required user bandwidth in the 2S module would also allow the FEC12 mode, which adds more transmission robustness.

The lpGBT_v0 has three different configuration modes: configuration over I^2C, configuration over the serial control channel, and configuration stored in e-fuses. In the first mode, the lpGBT is the target of an appropriate I^2C controller. It is flexible and easy to implement and does not require an operational high-speed link. In contrast, the second configuration mode utilizes the two IC bits (and/or the EC if not in transceiver mode) per data frame (one frame per 25 ns) of the high-speed link. In general, during normal operation, the user can switch between the I^2C and IC/EC interface by toggling a configuration pin. For the IC/EC to function correctly, the PLL inside the lpGBT has to be locked. This can only be achieved by storing a minimum configuration in the e-fuses and configuring the chip to load them during start-up. For an assembled 2S module, the only available interface is the serial control channel via the optical link. Therefore, any hybrids equipped with the lpGBT_v0 need burning of the e-fuses in a suitable system before assembly into modules. Burning of the fuses requires an I^2C controller and the ability to shortly apply 2.5 V to a dedicated input pin.

The lpGBT_v1 has a default minimal configuration stored in a read-only memory (ROM). When correctly configured, this can be loaded at start-up, enables the PLL to lock, and allows immediate configuration over the serial control channel. For hybrids with these chips, no burning of the fuses is required in order to test or use them in modules. Moreover, the lpGBT developers found the burning of the e-fuses to be unreliable. The OT will therefore use the start-up configuration via the ROM and the serial control channel for the configuration. For the lpGBT_v1, the arbitration between the I^2C interface and the IC/EC is made automatically by the lpGBT and not by a selection pin. The higher priority is given to the I^2C based on activity on the active-low I^2C data lines. In the batch 2 hybrids, the I^2C interface is not required and is unused, and it is ensured that the I^2C pins remain pulled up.

For easier routing, the polarity of the high-speed uplink differential pair is inverted in the PCB routing. The positive part of the lpGBT's differential high-speed output is connected to the $VTRx^+$ module's negative input receiver and vice versa. The

e-fuses can be programmed to invert the polarity internally ("flip" all received data bits). This setting is however not present in the default ROM configuration. When the hybrid is not fused, it is therefore necessary to do the logical inversion on the receiver side in the back-end.

Like the GBTx in the previous hybrid version, the lpGBT is hardwired to operate as a transceiver with bidirectional data communication. E-links connect it to the FEH connectors and the e-port receivers feed the data to the serializer. Transmitter e-ports receive data from the CDR circuit which is forwarded to the FEH. The e-links are set up for the CERN Low Power Signaling (100 Ω receiving end termination, 600 mV common mode voltage, and 1 to 4 mA output current). The settings can be optimized. The clock manager of the lpGBT provides differential clocks needed for the front-end ASICs. The recommended supply voltage is 1.2 V from the second DC-DC converter stage. An active-low reset pin is connected to the input voltage. During start-up, the ASIC's power-up state machine will not proceed before the input voltage reaches 0.9 V.

The lpGBT provides three independent I^2C interfaces, one for the configuration of the VTRx$^+$ module's laser driver and the others for each module side. It also features a 10-bit ADC with eight multiplexed analog inputs, which are all used on the hybrid. Of the 16 GPIO interfaces four are used to provide reset signals for the front-end ASICs and two are used to measure the DC-DC converter Power Good signals.

The lpGBT has a size of $9 \times 9 \times 1.24\,\mathrm{mm}^3$. It is encapsulated in a BGA package with a 17×17 ball matrix with 500 µm pitch.

One known issue with the lpGBT_v0 is of particular importance to the testing of the SEH prototypes: a problem occurred during the physical implementation stage of the I^2C controllers in the lpGBT_v0. Although the resulting weakness in the routing and timing intrinsically affects all three controllers, only failures of controller 0 are observed. Only ASICs that did not show failures on controller 0 during the chip production testing are assembled into hybrids [14].

VTRx$^+$ module

The VTRx$^+$ module is acquired from CERN's VL$^+$ project and used as-is on the hybrid for the opto-electrical conversion. Its development specifically targets the Phase-2 upgrades and the installation in a tracker environment. Compared to the VTRx module it meets more strict requirements in terms of radiation resistance, footprint, and power dissipation. As a result, the VL$^+$ components are also designed with asymmetric data rates and a different amount of channels for the up- and downlink. Functionally the VTRx$^+$ module is connected to the hybrid by an electrical connector, which provides the required supply voltages of 1.2 V and 2.5 V and the interconnection of the differential up- and downlink high-speed electrical links.

The VL and VL$^+$ generations are separated by roughly seven years. Yet, on the downlink, the VTRx$^+$ module utilizes the GBTIA, the same trans-impedance amplifier as the VTRx module. An InGaAs p-in-n diode is selected as the receiver. The VL$^+$ project developed a 4-channel laser driver array ASIC (LDQ10 [71]) that interfaces with a quad VCSEL array to form the uplink part.

All prototypes are used with version v10 of the VTRx$^+$ module design. However, the first received samples use the LDQ10v3, while the latter samples are equipped with the LDQ10v5, which is the final version. Besides a different style in the labeling on the module's backside (visible from the top when plugged into an SEH), the LDQ10v3 samples can be identified by a turquoise fiber protection sleeve. The newer samples feature a dark blue sleeve. In general, a Service Hybrid or module can be used irrespective of the laser driver version. The most notable difference is that the LDQ10v3 has all four laser drivers enabled by default. It therefore consumes more power unless the three unused drivers are explicitly disabled. The additional power consumption contributes to the total input current of the hybrid and needs to be taken into account when comparing measurements with different VTRx$^+$ module versions. The internal register structure and default values also changed for the LDQ10v5 [15]. The registers can be read out via one of the lpGBT's I^{2}C controllers and the two versions are therefore distinguishable by the software.

bPOL12V

As introduced in Chapter 1, the first stage of the two-step DC-DC conversion powering scheme is handled by the bPOL12V. Its design is very similar to the FEAST2. Table 2.2 provides a comparison between the characteristics of the DC-DC converters used in both stages of the 2S module powering scheme.

It is produced in the same 350 nm CMOS technology, but in order to increase the resistance to displacement damage its high-voltage module integrates LDMOS[1] transistors rated for 25 V instead of 60 V. The larger oxide thickness required for higher voltage ratings would result in more threshold shift with TID. The 25 V-rated transistors make it difficult to ensure reliability at input voltages up to or above the eponymous 12 V. Special measures in the chip design are implemented to mitigate the effects at the cost of 1-3% loss of efficiency. It is now rated for an input voltage U_{in} between 5.5 and 12 V and can output 0.63 to 5 V. Whenever possible input voltages below 11 V should be used. The maximal output power is 10 W and the continuous output current is limited to 4 A. The ASIC is specified for a total ionizing dose of up to 150 Mrad.

The bPOL12V output voltage U_{out} is set via a voltage divider connected to the output voltage. The divider output is compared to an internal reference voltage U_{ref} of roughly 600 mV. This voltage is obtained from a small reference voltage generator

[1] laterally-diffused metal-oxide semiconductor

	First stage		Second stage	
	FEAST2	bPOL12V	LTC3412A	bPOL2V5
Total ionizing dose	200 Mrad	150 Mrad	–	100 Mrad
Input voltage range	5-12 V	5.5-11 V	2.25-5.5 V	2.1-2.5 V
Output voltage range	0.9-5 V	0.63-5 V	0.8-5 V	0.6-1.5 V
Output current limit	4 A	4 A	3 A	3 A
Switching frequency	2 MHz	2.5 MHz	3.05 MHz	4 MHz
Under-voltage lockout	Enable 4.8 V Disable 4.5 V	Enable 5.5 V Disable 5.0 V	2.2 V	Enable 1.97 V Disable 1.81 V
Over-current protection	Yes	Yes	No	Yes
Over-temperature protection	Yes	Yes	Yes	Yes
Power Good	Yes	Yes	Yes	Yes
Package	QFN 0.5 mm distance	QFN 0.5 mm distance	QFN 0.65 mm distance	Flip-chip bare die 0.3 mm pitch

Table 2.2: Summary of the DC-DC converter ASIC characteristics. The chips used for the v3.1 prototypes and final prototypes are compared. All values are typical values and can vary between individual samples or due to temperature changes and irradiation.

circuit in a 130 nm CMOS technology that is stacked on the bPOL12V ASIC and contained within the same package. This solution is necessary not to be affected by a drift in the electrical parameters of the 350 nm transistors due to irradiation and annealing in a way strongly influenced by the dose rate and temperature [72]. Based on Figure 1.16, left, the equation to obtain the output voltage is

$$U_{\text{out}} = \frac{R_1 + R_2}{R_2} U_{\text{ref}} \ . \tag{2.1}$$

Note that here the bPOL12V converter output voltage U_{out} is the input voltage of the voltage divider V_{in}. R_1 is fixed to 499 kΩ, while in the first batch R_2 was 158 kΩ, to obtain 2.5 V. For the v6 ASICs U_{ref} was measured to be 630 ± 7 mV and R_2 adjusted to 169 kΩ.

The choice of the switching frequency is a trade-off between reliability, efficiency, and noise emission. For both batches a frequency setting of 2.5 MHz in combination with a 220 nH toroid inductor wound around a plastic core is selected. It is the minimal frequency recommended by the chip designers for this inductor value.

Another important setting is the turn-on threshold of the DC-DC converter and consequently the entire 2S module. As discussed in the detector description, up to twelve modules share the same output channel of a tracker power supply and the return line to the ground. The power supply channel's output voltage is common for the modules, but they can be connected individually to it. Two conditions need

to be met for the DC-DC converter to deliver its output voltage. The input voltage of the DC-DC converter needs to be above 5.5 V and greater than 800 mV at the enable input. The enable/disable level has a hysteresis and the DC-DC converter turns off for smaller input voltages and if less than 500 mV are present at the enable input. Due to the long cables between the power supplies and the 2S modules, voltage drops are unavoidable and no voltage sensing is foreseen to mitigate them. Without using dedicated enabling lines for the DC-DC converter (e.g. solely based on the presence of an input voltage $> 5.5\,\mathrm{V}$) the DC-DC converter could turn on for very low voltages and subsequently draw large currents. These currents will cause a voltage drop on the cables, which might lower the input voltage below the DC-DC converter operation range, which causes it to turn off. This can result in many consecutive turn-ons and turn-offs, referred to as oscillations, which can be harmful and should be avoided. Therefore, a voltage divider from the input voltage sets the voltage level at the enable input. It sets the enable threshold, in terms of the input voltage, to roughly 10 V and the disable threshold (through the hysteresis) to roughly 6.5 V. The enable input is pulled down to ground via an internal resistor R_{pd} parallel to R_2 leading to an effective resistance of R_2^{eff}. The voltage at the enable input U_{en} is given by

$$U_{\mathrm{en}} = \frac{R_2^{\mathrm{eff}}}{R_1 + R_2^{\mathrm{eff}}} U_{\mathrm{in}} \ . \tag{2.2}$$

Since $1/R_2^{\mathrm{eff}}$ is given by

$$\frac{1}{R_2^{\mathrm{eff}}} = \frac{1}{R_2} + \frac{1}{R_{\mathrm{pd}}} \tag{2.3}$$

which is equivalent to

$$R_2^{\mathrm{eff}} = \frac{R_2 R_{\mathrm{pd}}}{R_2 + R_{\mathrm{pd}}} \tag{2.4}$$

this can be solved directly for the required input voltage

$$U_{\mathrm{in}} = \left(R_1 \left(\frac{1}{R_2} + \frac{1}{R_{\mathrm{pd}}} \right) + 1 \right) U_{\mathrm{en}} \ . \tag{2.5}$$

For the v4 of the bPOL12V used for the first batch, the pull-down resistor has a value of 500 kΩ ($R_2 = 9.09\,\mathrm{k\Omega}, R_1 = 100\,\mathrm{k\Omega}$), while in the second batch with v6 ASICs it was reduced to 15 kΩ ($R_2 = 953\,\Omega, R_1 = 10\,\mathrm{k\Omega}$). Figure 2.7 illustrates the impact of the different enable pin thresholds and how the input voltage threshold changes depending on the value of R_2 for the two batches. Due to a temperature dependence of the chip's internal pull-down, which is negative for semiconductors, the threshold decreases with lower temperatures. As measured by the chip designers, the enable threshold on the enable pin will increase with radiation. This translates linearly to the input voltage requirement as illustrated by the figure (here for a 5% increase). Figure 2.8 shows also the disable voltage for the batch 2 configuration

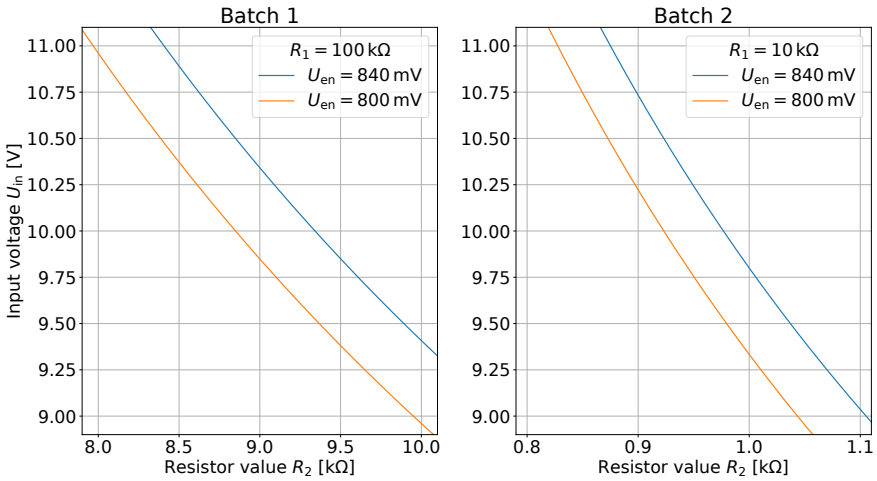

Figure 2.7: Enable threshold of the batch 1 (left) and batch 2 (right) Service Hybrid prototypes in terms of the input voltage versus the resistor R_2 selected in the voltage divider. U_{en} can depend on radiation or vary between samples. The value for R_1 in the voltage divider is fixed and a chip version-dependent internal pull-down resistor is taken into account.

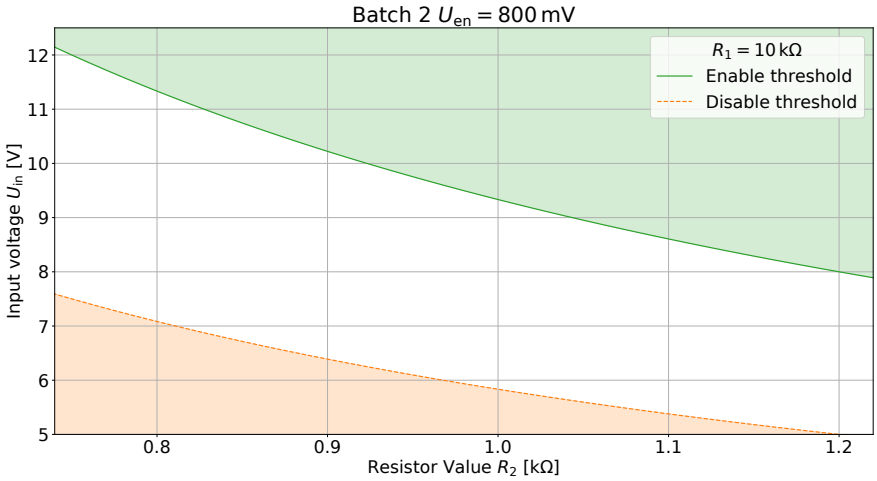

Figure 2.8: Enable and disable threshold of the batch 2 Service Hybrid prototypes in terms of the input voltage versus the resistor R_2 selected in the voltage divider. The value for R_1 in the voltage divider is fixed and the chip version-dependent internal pull-down resistor is taken into account.

as a function of R_2. The absolute hysteresis decreases for lower thresholds and the disable voltage becomes irrelevant when the under-voltage lockout is reached at 5 V.

Chip protection features in addition to the under-voltage lockout are an over-current and an over-temperature protection. A soft start procedure is hardwired to slowly increase the DC-DC converter output voltage, which limits the inrush currents. The chip also provides a Power Good (pulled up to the bPOL2V5 output voltage) during normal operation. It is negated while the chip is disabled, during restart, when a protection feature is active, or e.g. when the output voltage is outside a $\pm 6.5\%$ regulation window. Relative changes in the chip's internal temperature can be monitored from a PTAT voltage. While the absolute value is different for each sample, it changes linearly with the temperature at a slope of 4.85 mV/K.

Like the FEAST2, the bPOL12V is packaged in a plastic QFN32 package with a pitch of 0.5 mm and has a size of $5.0 \times 5.0 \times 0.9\,\mathrm{mm}^3$. The exposed thermal pad in the center is soldered to the PCB [16, 72].

bPOL2V5

The SEH integrates the bPOL2V5 ASIC for its second DC-DC converter stage. The ASIC has an input voltage range of 2.1 to 2.5 V and provides between 0.6 and 1.5 V of output voltage. It is capable to continuously output 3 A load current and tested for a total ionizing dose of up to 100 Mrad.

A voltage divider with $R_1 = 1\,\mathrm{M\Omega}$ and $R_2 = 316\,\mathrm{k\Omega}$ is connected to the DC-DC converter output and the reduced voltage internally compared to a $300\,\mathrm{mV} \pm 1\,\mathrm{mV}$ voltage reference to obtain an output voltage of 1.25 V. In v3.2 ASICs the voltage reference is $291.2\,\mathrm{mV} \pm 6.5\,\mathrm{mV}$, which leads to a higher spread and overall lower output voltage. Therefore, R_2 is adjusted to $287\,\mathrm{k\Omega}$ in batch 1 hybrids. The hybrid enable is connected to the bPOL12V output voltage via a voltage divider (factor 0.91). This effectively always enables the DC-DC converter but limits the maximum voltage at the enable input. The recommended setting for the switching frequency is chosen via a $70\,\mathrm{k\Omega}$ resistor to obtain a typical value of 4 MHz. The higher switching frequency allows the use of a main inductor with a smaller inductance value of 100 nH. To avoid circuit damage, a soft start procedure is implemented that slowly increases the voltage on the output and limits the current to charge the output capacitors. Typically the start-up lasts a few hundred microseconds and includes a small overshoot of the output voltage. The DC-DC converter provides a Power Good signal. It is pulled up to the DC-DC converter output voltage and present when the DC-DC converter is enabled, not in an over-temperature condition, and the output voltage is within $\pm 6.5\%$ of the nominal voltage. The bPOL2V5 also has an analog output signal that is proportional to the on-chip temperature. This PTAT voltage has a slope of 4 mV/K. The offset has a large sample-to-sample variation. Chip

internal protection features safeguard it against a negative and positive over-current condition, over-temperature, or an input under-voltage condition.

The radiation tolerance requirements motivated the design in the 130 nm CMOS technology. The trade-off is a voltage rating for the used devices of 3.3 V, which is close to the voltage specified for the VTRx$^+$ module and thus the bPOL2V5 input voltage. The DC-DC converter will experience over-voltages as a result of the $\mathrm{d}I/\mathrm{d}t$ due to the switching and parasitic input inductors. This requires special consideration in the PCB layout and choice of components which are carried out according to the chip designer recommendations. Compared to wire bonding, the choice to assemble the bPOL2V5 as a flip-chip bare die also minimizes the bonding parasitic inductance. The bPOL2V5 is $3.1 \times 4\,\mathrm{mm}^2$ in size with a 10×13 bump bond matrix (five bumps are missing due to foundry restrictions and for orientation). The SnAg bumps have a pitch of 300 µm, a height of 100 µm and a diameter of 132 µm [17]. This pitch is the smallest of the SEH ASICs.

Notes on the bPOL12V output voltage spread

The output voltage of the bPOL12V is used to power both the bPOL2V5 and the VTRx$^+$ module. The recommended operation ranges of these two devices do not overlap. The input voltage of the bPOL2V5 is stated with 2.1 to 2.5 V and to ensure long-term reliability should be as low as possible [17]. For the VTRx$^+$ module, the general technology limit is 2.25 to 2.75 V ($\equiv \pm 10\%$). However, to guarantee good performance the voltage should be above 2.5 V [15].

The spread of the bPOL12V output voltage U_{out} is determined by three factors. The intrinsic spread of the voltage reference U_{ref}, the influence of the total ionizing dose during radiation, and the variation of the resistors in the voltage divider of the feedback loop R_1 and R_2. In first approximation, the temperature influence is regarded as a simple DC offset. The 1σ variation of U_{ref} is measured to 7 mV. The worst-case deviations with TID are -1.5% to $+4\%$, which we denote with Δ^-_{TID} and Δ^+_{TID}. The observed TID response is not monotone [17]. The resistor values usually have a margin δ of 10%, 1%, or 0.1% and their true values fall within that range. A challenge is that this distribution is in general unknown and all resistors within a lot can potentially vary by the same amount from the nominal value.

Under these assumptions, one can calculate the worst-case spread for 99.7% of hybrids by calculating the lowest and highest DC-DC converter output voltage within a 3σ spread of the reference voltage [73]. The calculation yields

$$\left(1 - |\Delta^-_{\mathrm{TID}}|\right)(U_{\mathrm{ref}} - 3\sigma)\left(1 + \frac{R_1}{R_2}\frac{1-\delta}{1+\delta}\right) < U_{\mathrm{out}} \tag{2.6}$$

and

$$U_{\text{out}} < \left(1 + \Delta^{+}_{\text{TID}}\right)\left(U_{\text{ref}} + 3\sigma\right)\left(1 + \frac{R_1}{R_2}\frac{1+\delta}{1-\delta}\right) \quad . \tag{2.7}$$

If we assume 1% resistors we obtain

$$0.600\left(1 + \frac{0.99R_1}{1.01R_2}\right) < U_{\text{out}} < 0.677\left(1 + \frac{1.01R_1}{0.99R_2}\right) \quad . \tag{2.8}$$

For the choice of $U_{\text{out}} = 2.49\,\text{V}$ with the resistor values $R_1 = 499\,\text{k}\Omega$ and $R_2 = 169\,\text{k}\Omega$ the spread is

$$2.33\,\text{V} < U_{\text{out}} < 2.72\,\text{V} \quad . \tag{2.9}$$

With the bPOL2V5 and VTRx$^+$ module input requirements, it is clear that this spread would be unacceptable. The same estimation with 0.1% resistors yields $2.37\,\text{V} < U_{\text{out}} < 2.67\,\text{V}$ [73]. The only possibility to reduce this spread further is to select bPOL12V chips based on their measured U_{ref}. Here, the Outer Tracker can utilize the fact that similar amounts of bPOL12V ASICs are required for the 2S and PS modules. The spread of U_{ref} per variant can effectively be cut in half by selecting chips with a U_{ref} between 616 and 629 mV (42% of all chips) for the PS module and between 629 and 650 mV (55.5% of all chips with a median of 634 mV) for the 2S module. The median U_{out} can be set for each module variant by adjusting R_2 accordingly. In our example we obtain $2.45\,\text{V} < U_{\text{out}} < 2.68\,\text{V}$ with a median of 2.51 V.

The measurement of the bPOL12V U_{ref} and the binning is performed for all chips that are used for the series production. The series production will also use 0.1% resistors to set the output voltage.

2.3.4 Prototyping campaign

CERN ordered final prototypes from two contractors. In the following, they are referred to as contractor A and contractor B. The suppliers are both consortia of an assembly company and a flex manufacturer. Only the former is a direct contract partner with CERN. In 2020, CERN asked both contractors to produce the bare flex circuits required for prototyping. In a second step, they provided assembled batch 1 hybrids and delivered them in several lots throughout 2021. The first batch contained about 25 hybrids per contractor.

The hybrids were delivered to CERN and underwent a visual inspection. Subsequently, they were functionally tested and qualified by the I. Physikalisches Institut B of RWTH Aachen University. Before they could be distributed further, the lpGBT's e-fuses were programmed, and if necessary the output voltage of the

Arrival at CERN	Contractor	Batch	Quantity	Subsequent actions
05/2021	A	1	13	lpGBT e-fusing Output voltage adjustment Mounting of shield
06/2021	A	1	14	lpGBT e-fusing Mounting of shield
11/2021	B	1	25	lpGBT e-fusing Mounting of shield
05/2022	A	2	22	Shield removal Adjustment of enable voltage divider Mounting of shield

Table 2.3: Hybrid subsets of the final prototyping campaign. The ASIC versions used in each batch are summarized in Table 2.1.

bPOL2V5 was adjusted (by replacing the resistor R_2 in its voltage reference voltage divider) and the shield mounted. The results of the electrical testing are presented in Chapter 5.

After the initial delivery, only one contractor was asked to assemble additional 25 batch 2 hybrids. The flexes produced by contractor B showed cracks in the microvias, which made the hybrids unreliable. The costs and required qualification efforts did not justify a second batch. The batch 2 hybrids are the first with production-grade ASICs[2] and arrived in the middle of 2022. They were inspected visually at CERN, too. The lpGBT did not require an e-fusing, but due to a change in the bPOL12V's internal pull-down resistor at the enable pin, the respective voltage divider needed to be adjusted. It required the removal and later replacement of the shield. A short overview of the subsets of the prototyping campaign is given in Table 2.3.

The prototyping successfully demonstrated the implementation of the two-step DC-DC conversion scheme with the radiation-tolerant DC-DC converter ASICs. It also established the final communication chain between the FEH and the back-end. The material and space constraints for the mechanical parts were met. Hybrids were used in the module prototypes assembly and various system tests for the Outer Tracker [74–78].

[2] Although, for instance, the available lpGBT v1 chips did not undergo the production testing or calibration.

3 Early hardware and prototype testing

This chapter summarizes the developments made at the I. Physikalisches Institut B of RWTH Aachen University that led to the production of PCBs used to test and qualify SEH prototypes. Those PCBs are referred to as test boards. The testable prototype generations reach back to those featuring the GBTx. The development of the test boards spanned several years and iterations. The first section presents the finalized setup state using the test board. The relevant milestones, insights, and adjustments that led to its development are also discussed. The second section documents measurements performed on SEH v3.1 prototypes with the test board. The hybrids were tested before their distribution to the OT community.

3.1 Test board

The final test board (v2.2) shown in Figure 3.1 can be used to qualify and fully test a Service Hybrid. Martin Lipinski and Christian Dziwok compiled the depicted setup to perform active thermal cycles on final Service Hybrid prototypes. The subsequent section summarizes my contributions to the test board developments. The work was required to integrate the test board into the presented setup.

Since the hybrid's destination is inside a tracking detector, it lacks any readily available self-diagnosis features. The test board and several accompanying devices provide the interfaces to assess a hybrid.

The test board is placed on an aluminum base plate that also provides passive cooling. The SEH can be mounted on an aluminum block in a disjoint volume. It can be actively cooled. See Reference [79] for details on the fixture and cooling.

The test board is controlled by a Raspberry Pi miniature computer [80]. It operates several switches and performs various measurements. A 5 V channel of an external power supply powers the test board. An additional daughterboard PCB is powered by the same voltage and enabled by the Raspberry Pi. It provides 1.2, 1.5, 2.5, and 3.3 V from several LDOs and DC-DC converters to active components on the test board. This configuration allows a power cycle of the board without restarting the controlling Raspberry Pi. An external power supply channel provides the supply

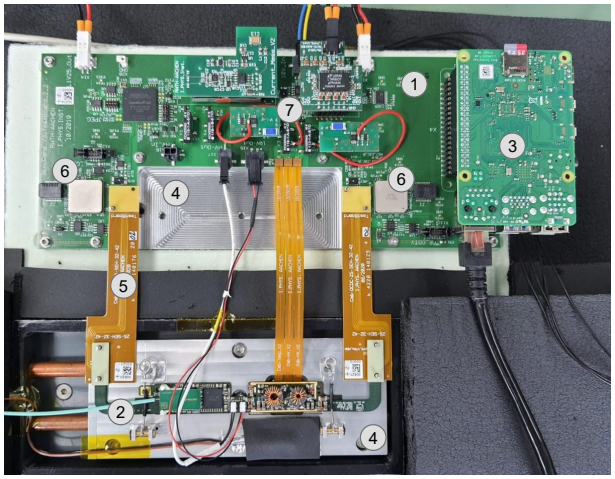

Figure 3.1: Test board setup for active thermal cycles on final Service Hybrid prototypes. Important components are labeled.

voltage for the SEH. The test board measures the voltage and current and a cable forwards it from the test board to the SEH input connector. On the input side, the sensor bias voltage is routed similarly between the test board and the hybrid. Three small and flat flex cables connect the test board and the three connectors on the hybrid fold-over for the upper and lower sensor bias voltage and the thermistor. For details on the custom circuits and daughterboards used to measure the current and voltage of the sensor bias RC filter see Section 3.1.2.

The Service Hybrid's FEH connectors are plugged into flat flex adapter cables. They route the output voltage of the SEH's DC-DC converters to the test board. Here, the output voltage can be connected to an external electrical load, while the voltage and drawn current are measured on the test board.

The adapter cables are also used to connect the additional analog signals and digital data lines between the test board and a Service Hybrid. An overview of these connections between the test board and a Service Hybrid with either an lpGBT or GBTx ASIC is shown in Figure 3.2. An FPGA counts the received transitions of the SEH's 320 MHz differential clock and reset lines and can also set a fixed voltage level for the FEH AMUX lines connected to the SEH's ADC.

To emulate the SEH data stream to and from the FEHs, the test board hosts two GBTx ASICs, one for each side. They are linked to receptacles for VTRx modules and are connected via those with a back-end board. The GBTx ASICs have e-fuses and they are programmed in such a way that they automatically lock themselves to the incoming optical link and are configurable by it. The Service Hybrid e-link receivers are connected to e-link transmitters on the test board GBTxs and vice versa.

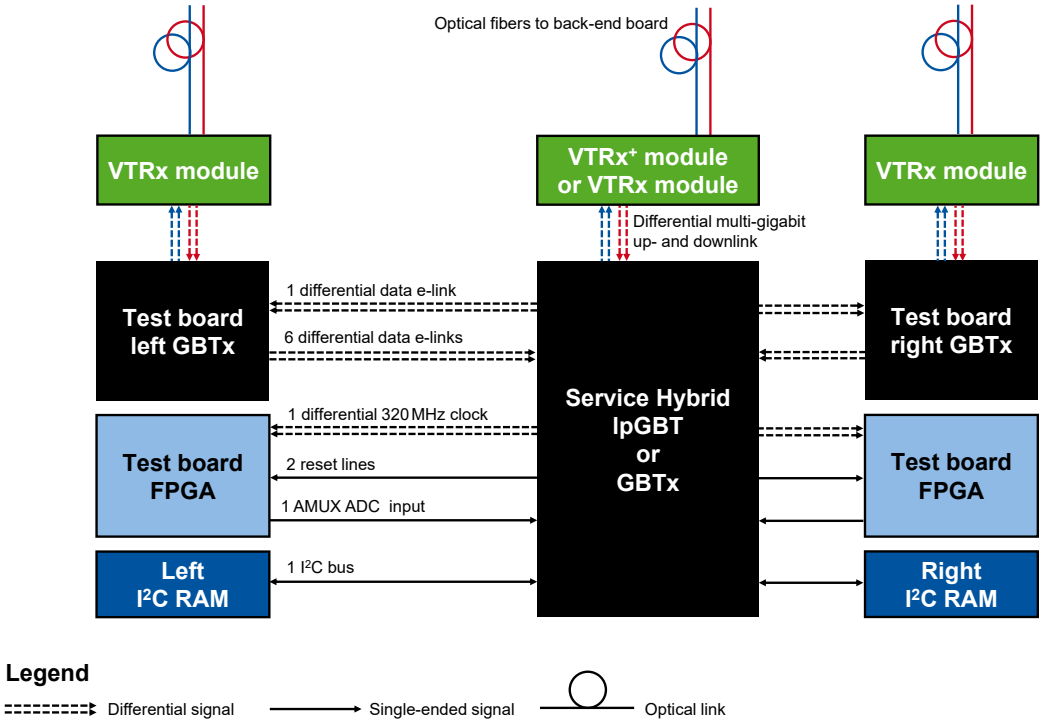

Figure 3.2: Overview of the data connections between components on the test board (left and right) and the Service Hybrid (middle).

To emulate the FEHs' I^2C clients, two I^2C random-access memorys (RAMs) and level translators are present on the test board.

The input and output currents and voltages of the Service Hybrid are measured by dedicated active circuits (Texas Instruments INA226 [81]). The same devices are used to monitor the test board supply currents and voltages. An 8-channel ADC can monitor additional voltages. The Raspberry Pi controls all components on the test board, except the GBTx ASICs and their VTRx modules. The latter two are controlled via the back-end board housed in a µTCA crate connected to a PC. The same PC communicates with the Raspberry Pi in the same network via the Remote Procedure Call protocol [82]. The Raspberry Pi acts as a server in this configuration. An additional back-end board controls the SEH itself via an optical fiber pair. In the presented state of the system, both back-end boards are FC7s.

3.1.1 Summary of the designs and developments

The work on the design of the test board started in 2017. Three different versions were produced: version v2 in 2018, version v2.1 in 2019, and version v2.2 in 2020. The commissioning of the first working version of the test board (v2) is presented in

my Master's thesis [83]. While the low-voltage related circuits, the ADC, I^{2}C RAM, FPGA, and Raspberry Pi in general were found to be functional, several issues prevented a comprehensive SEH testing. For one, neither the leakage current of the sensor bias voltage filter on the SEH nor its connection through the connectors on the fold-over were testable to the full extent and with sufficient accuracy. Secondly, measurements of the number of received clock cycles from the SEH showed missed transactions on the right side. Most importantly, it was impossible to perform a meaningful loop-back test of the optical link. In a loop-back test, data are sent from the back-end board to the GBTx and received back, preferably without data loss. A successful and error-free loop-back is required to perform data transmission checks or bit error rate measurements on the e-links between the SEH and the test board GBTx ASICs since that data are passed through the optical link.

To address those issues, a second version of the test board was designed (v2.1). The most notable changes are listed below:

- The leakage current measurement was implemented on the test board. It requires an additional daughterboard PCB. The circuit is described in Section 3.1.2. This change eliminates external parasitic sources of leakage current (e.g. cable isolation).
- The circuit also allows the production of the bias voltage directly on the test board from 5 V. This removes the need for an external high-voltage power supply.
- The quality of the 320 MHz data and clock lines between SEH and the test board was improved. The clock lines were fitted with an external 100 Ω termination. Instead of inputting the differential clock signal directly into the FPGA, a high-speed differential receiver converts it to a single-ended signal that the FPGA can measure.
- The powering tree was reviewed and improved to avoid potential back-powering of the test board both via the SEH and the GBTx USB controller[1].
- The 4.8 Gb/s differential lines between the VTRx module and GBTx were moved to the first layer of the PCB to improve signal quality.
- New supply, load, and bias voltage connectors were implemented to ease handling.

Despite the improved design, the observed error rate during the GBTx loop-back testing remained too high. Studies of various settings of the GBTx and other system variables were performed without yielding positive results. Most notable is the charge pump current of the clock data recovery phase detector. An increase in the

[1] Only the GBTx ASICs on test board v2.2 were fused. Before that, the GBTx was configured as an I^{2}C client via a dedicated USB device.

current reduced the occurrence of transmission errors by the forward error correction and produced a stable phase. However, the failures presented an asymmetry between the two GBTxs. Failures were more pronounced on a GBTx serving the right SEH side. Moreover, equivalent tests were performed on comparable circuits like the SEH prototypes or the VL demo board [84]. The test board was much more sensitive to the various parameters. These observations lead to an investigation of the supply voltage, the start-up scheme, and the reset sequence. Eventually, it was discovered that bypassing the 1.5 V output of the quad DC-DC step-down power module (LTM4643 [85]) led to a reliable operation. The power module provided all voltages deduced from the 5 V supply voltage in the test board v2 and v2.1. For the bypassing, the GBTxs and VTRx modules were powered by an external power supply channel. Further studies demonstrated that the same result can be achieved by replacing the quad power module with an external daughterboard as described above.

Consequently, another version of the test board (v2.2) was designed and produced in 2020. This final series consists of twelve boards. Changes concerning the v2.1 are:

- 1.2, 1.5, 2.5, and 3.3 V are provided from an external add-on PCB.
- A smaller variant replaces the rather large high-voltage relays.
- The values of some passives are updated to reflect the working configuration and to improve some measurement ranges.

Test boards of version v2.1 and v2.2 were used to qualify SEH v3.1 prototypes starting in 2019. Marius Preuten provided the firmware and software for the testing of the GBTxs. The software to control the test board and test the DC-DC converters and the sensor bias voltage RC filter was written by me and I also performed the hybrid testing. Gerhard Pierschel contributed the FPGA firmware.

The v2.2 boards were integrated into the setup to perform active thermal cycles on final Service Hybrid prototypes in Figure 3.1. The system, software, and firmware for it are documented in Reference [79]. It consists of three test boards placed next to each other. Flat flex adaptor cables connect them to Service Hybrids on cooling blocks. The blocks and hybrids are placed in insulated boxes. The cables enter the boxes through cut-outs. This configuration separates the test board from the cold volume reducing also the size of the cold volume. The cooling blocks are connected with copper pipes to two separate cooling loops. One loop is used to cool and the other one to heat. This enables faster thermal cycling since the chillers and most of the coolants do not change temperature. In contrast to the hybrid test system for production, the test board hardware only needs to be qualified for room temperature. Long-term functional testing of multiple hybrids in parallel is possible. During production, the system will also be available for additional measurements on final hybrids.

An early stage of the setup is covered in Oliver Wen's Bachelor's thesis [38]. His work was supervised as part of this thesis. This Bachelor's thesis provides details on the test procedures and commissioning measurements with v3.1 prototypes. He demonstrated that it is possible to cool down the hybrid independently from the attached test board. Improvements were proposed concerning the arrangement of the hardware. Previously, the hybrid was placed in a box above the test board. This vertical separation was changed to a horizontal one. He also experienced freezing of the warm coolant, whose composition was therefore changed. The test board and the testing procedures were found to be functional.

The test board can also be used to test final Service Hybrid prototypes. To make the hardware compatible, only new adapters were required. The FEH receptacle placement on the test board was decided based on the SEH v3.1. The adaptor cables respect the larger distance between both sides.

A major new development of the test board campaign is the circuits for testing the SEH's sensor bias voltage filter. They are also used on the 2S-SEH test card. The following section provides more details on the components and the measurement principle. It also presents performance measurements.

3.1.2 Sensor bias voltage testing hardware

As outlined in Section 1.5.2.1, the SEH serves as the entry point for the negative bias voltage of a 2S module, filters it, and distributes it to the two sensors via flat flex cables. In the test system, the bias voltage filter and distribution are characterized in two different ways. First, the amount of parasitic leakage current of the bias voltage circuit is determined. The current measurement also serves to rule out any possible short circuit. Secondly, the connection is tested. A typical bias voltage for a sensor for a 2S module is $-350\,\mathrm{V}$, but it can reach $-800\,\mathrm{V}$ after irritation. The targeted test voltage for the Service Hybrid RC filter is $-1000\,\mathrm{V}$.

In contrast to the supply and output voltage measurements, the bias voltage and leakage current cannot be measured simultaneously. Any voltage measurement of the order of several 100 V requires a voltage divider and therefore an additional path to ground. In most applications that need a voltage divider, the resulting current is small compared to the overall current. However, in the SEH case, the expected parasitic currents of the filter circuit itself are of the order of some 10 nA. For each 100 V of bias voltage, a feasible voltage divider with $5\,\mathrm{G\Omega}$ resistance leads to a current of 20 nA and therefore affects the measurement. With these small currents, cables from an external power supply need to be well shielded or the path limited in order to obtain a good measurement.

For the test board, these requirements lead to the development of two additional support PCBs to enable the bias voltage evaluation. The components are sketched in Figure 3.3.

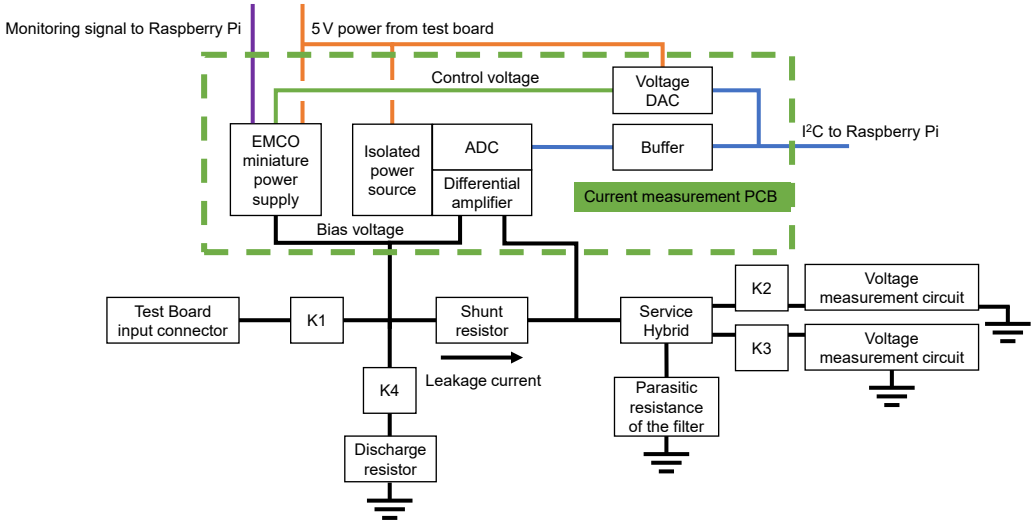

Figure 3.3: Schematic drawing of the circuit to measure leakage current of the sensor bias voltage filter on the SEH. Components in the green dashed box are placed on the measurement PCB (schematic in Figure A.1). All other components are on the test board. Black lines indicate the paths that are presented to the bias voltage potential. Orange lines show the 5 V distribution, blue lines the I^2C connection to the Raspberry Pi, the green line a control voltage, and the purple line a monitoring votlage.

The current measurement PCB features two pin header connectors. One provides the connection for the 5 V power, ground, I^2C, and low-voltage monitoring signals to the test board. The other connector is dedicated to sensing the bias voltage on the power supply and the SEH side of a shunt resistor ($R_{\text{shunt}} = 49.9\,\text{k}\Omega$) on the test board, therefore enabling a differential measurement of the voltage drop and thus a current measurement. The bias voltage can be provided to the test board via an external connector or it can be produced on the PCB itself if a miniature high-voltage power supply (EMCO_P12N [86]) is mounted. The power supply output is directly connected to the power supply side of the current sense connector, which therefore also provides the bias voltage to the test board and hybrid in this configuration. The EMCO_P12N is powered from the 5 V rail. The output voltage level is set proportional (factor 293) to the output of a 10-bit voltage DAC controlled by the test board's Raspberry Pi via I^2C. The voltage drop across the shunt resistor is amplified and subsequently measured by an ADC, which is read out via I^2C from the Raspberry Pi. The amplifier and ADC are separated from the other low-voltage devices by an isolated power source, which consists of a DC-DC converter, two LDOs, and a buffer for the I^2C signals. The EMCO_P12N outputs a monitoring signal proportional to the bias voltage (factor 10^{-3}) to the test board. Four relays and accompanying logic make up four switches named K1 through K4 to configure

Configuration	State K1	State K2	State K3	State K4
Default, no power, and discharge	Open	Open	Open	Closed
Source voltage from EMPO_P12N and measure leakage current	Open	Open	Open	Open
Source voltage from EMPO_P12N and measure filtered voltage	Open	Closed	Closed	Open
Source voltage from external power supply and measure leakage current	Closed	Open	Open	Open
Source voltage from external power supply and measure filtered voltage	Closed	Closed	Closed	Open

Table 3.1: Summary of the allowed configurations of the test board relays that are integrated into the bias voltage measurement.

different states of the measurement on the test board. Table 3.1 summarizes the possible states of the relays and their function. K1 (default state open) allows the connection of the external connector to the rest of the bias voltage circuit. K4 (default state closed) connects the circuit via a resistor to ground. It discharges the system in case of a power cut or at the end of a measurement. K2 and K3 (default state open) connect the bias voltage that is filtered and routed back via the SEH to the test board with the dedicated high-voltage measurement circuits described below. While the leakage current is measured, those relays remain open, whereas they are closed during the connection check.

The measurement range I_{range} of the current measurement is given by

$$I_{\text{range}} = V_{\text{ref}}/f_{\text{amp}}/R_{\text{shunt}} = 5\,\text{V}/201.44/49.9\,\text{k}\Omega = 497.5\,\text{nA}\ . \tag{3.1}$$

Here, V_{ref} is the voltage reference of the ADC, and f_{amp} is the amplification factor of the amplifier. Since this is a differential measurement, positive and negative currents across the shunt resistor between $-248.75\,\text{nA}$ and $+248.75\,\text{nA}$ can be measured.

The EMCO_P12N can provide 0 to $-1,200\,\text{V}$ and a current of up to $2\,\mu\text{A}$. Originally, the voltage level that is provided from the SEH back to the test card should be measured with the ADC on the test board. Thus, a voltage divider with $5\,\text{G}\Omega$ and $5\,\text{M}\Omega$ reduces the voltage by a factor of 1,001 and to a measurable range. It limits the current to ground to the EMCO_P12N's current range but presents too much input resistance for the ADC. Instead, the returned voltage is measured via the measurement circuit (schematics in Figure 3.4). It is attached to the test board by two pin header connectors that were added by hand after the production of the v2.2 test boards. The $5\,\text{V}$ supply voltage is provided by a hand-soldered cable. A shunt voltage reference offsets the voltage divider output by $2.048\,\text{V}$. An additional

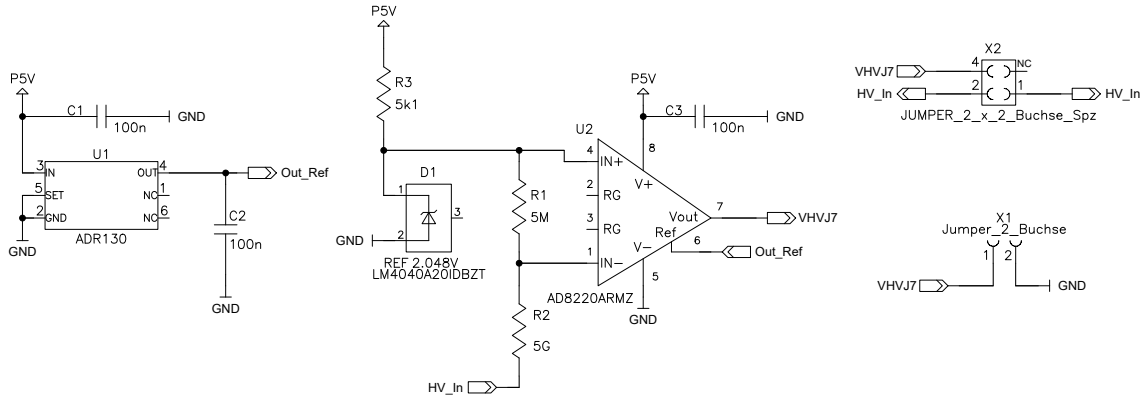

Figure 3.4: Schematic of the circuit to measure the voltage level of the sensor bias voltage that is passed through the SEH's filter.

connection to 5 V via a 5.1 kΩ resistor serves as protection of the reference. The now positive signal is connected to an amplifier. With a gain of 1, it only serves to provide enough current to measure the signal with the test board's ADC.

Measurements with a Service Hybrid are performed to qualify the bias voltage measurement circuit. The EMCO_P12N is programmed to provide 0 to $-1,000$ V in -200 V steps. The monitoring voltage of the EMCO_P12N and the voltage measurement from both SEH bias voltage connectors are measured for five minutes at every voltage step. Figure 3.5 shows the results. All three measured voltages follow the set voltage. Minor undershoots are visible when the voltage changes, which are attributed to the regulation of the power supply [86]. The voltage settles after roughly 10 seconds.

Figure 3.6 shows the three measured voltages versus the set voltage averaged over all values after the settling period. The error bars indicate the standard deviation. For the foreseen connection check, the accuracy of the measurements behind the sensor connectors of ± 5 V is more than sufficient. The output of the EMCO_P12N can be monitored within ± 1 V. No major deviation across different boards is observed.

To qualify the current measurement circuit, a similar measurement is performed. The relays K2 and K3 are now open, so only the leakage current of the SEH is measured. The absolute voltage of the EMCO_P12N is increased in 200 V steps, and the current is measured. After the $-1,000$ V level, the EMCO_P12N's voltage is set to 0 V, and the turn-off is also measured. Figure 3.7 shows the current, the set voltage, and the monitored voltage. The voltage behaves equivalent to Figure 3.5.

Since we apply a negative voltage, a current flow from the power supply to the Service Hybrid is measured as a negative current value across the shunt resistor. The parasitic resistance of the RC filter results in the measured current. We observe

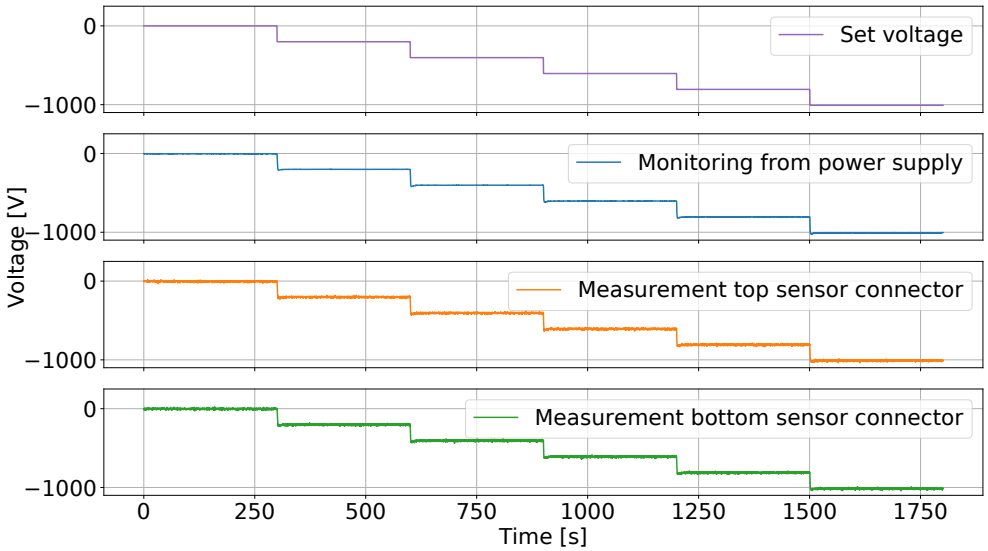

Figure 3.5: Measurement of the applied and measured bias voltage when connected to the test board via a SEH prototype versus time.

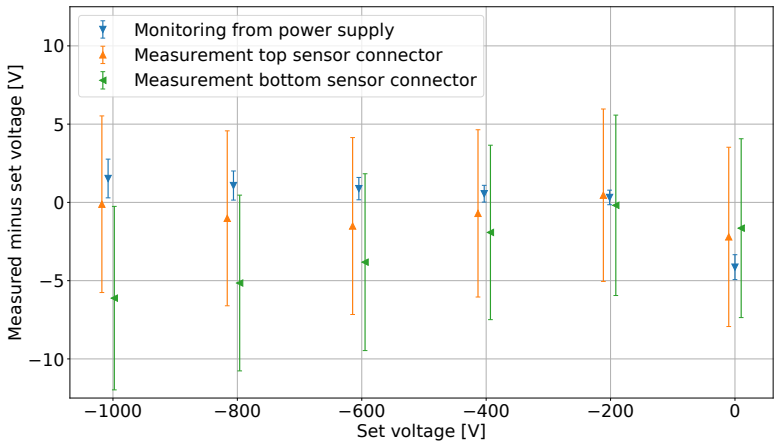

Figure 3.6: Averaged deviation of the measured to the set voltages for six voltage levels. Error bars indicate the standard deviation. The three measurements are offset horizontally to increase readability.

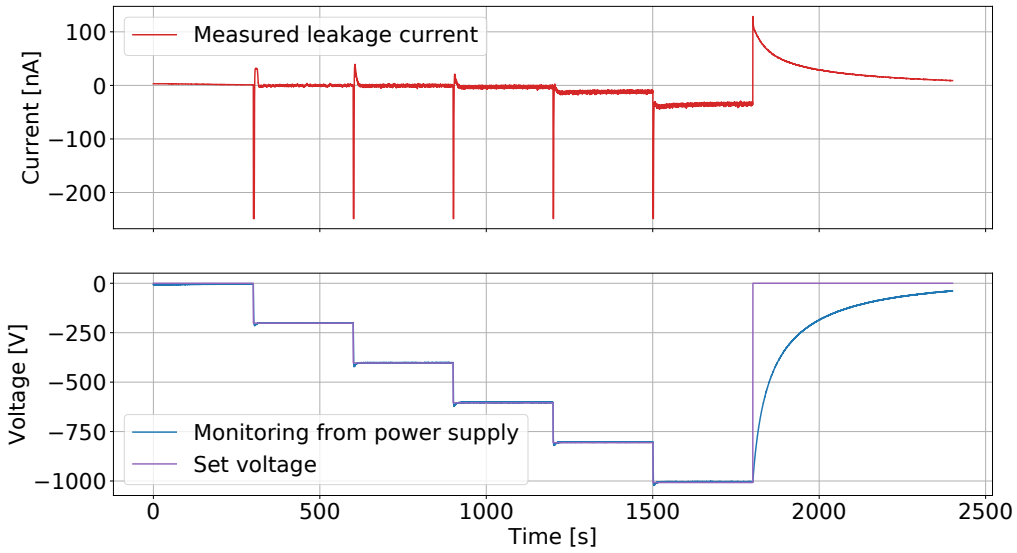

Figure 3.7: Measurement of the leakage current (top) and the set and monitored voltage (bottom) of a SEH's sensor bias voltage filter versus time.

a jump in the measured current for each voltage step. While the voltage is stable, the current settles, but a slight reduction with time remains. This will be discussed in more detail in Chapter 5.2. The current levels do not scale linearly with the voltage, hinting towards non-linear effects on the leakage current. When the voltage changes, additional charge is required to increase the capacitors' voltage. This is visible as spikes in the measurement. The currents during charging are actually larger than the measurement range. For an isolation measurement, these charging effects need to be considered.

After the turn-off, the EMCO_P12N ramps down slowly, since the capacitance of the SEH bias voltage filter can only discharge slowly. This demonstrates the necessity of the discharge resistor. The voltage undershoots also appear as positive values in the current measurement.

The accuracy of the current measurement is assessed by removing the EMCO_P12N from the PCB and providing a negative bias voltage of 1,000 V externally via a Keithley 2410 high-voltage power supply [87]. A comparison between the current measurement of the test board circuit and the power supply is shown in Figure 3.8. In both cases the absolute value of the current drops exponentially after the voltage is turned on and settles at around 30 nA, while decreasing slowly. Both measurements are consistent in value and spread. The measurement was carried out on a test board with a different shunt resistor value so that the measurement range was reduced by a factor of 2 with respect to the final configuration explained above. The power

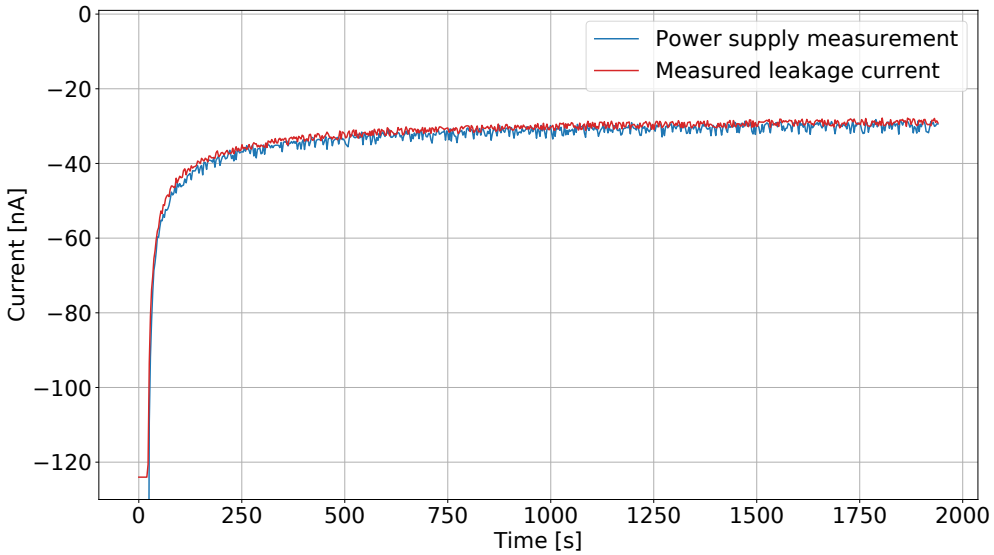

Figure 3.8: Longtime comparison of the leakage current of the SEH sensor bias filter measured by a commercial power supply and measured by the test board.

supply's measurement range is much larger, so it can also measure at the very start of the measurement when the test board circuit is still saturated. Nevertheless, this test demonstrates that the measurement circuit is able to measure the leakage current correctly and at a comparable accuracy to a commercial power supply.

3.2 Measurements of SEH v3.1 prototypes

A total of 34 v3.1 SEH prototypes were evaluated at the I. Physikalisches Institut B of RWTH Aachen University. A first batch of four hybrids arrived in December 2018 and the rest in a second batch in June 2019. For all of these hybrids, the GBTx was fused before assembly. A VL demo board equipped with a socket to hold the unsoldered ASIC was used for the fusing.

Upon arrival, all hybrids were submitted to a basic communication check. The back-end board Gigabit Link Interface Board (GLIB [88]) and custom software were used to read registers of the GBTx and subsequently from the GBT-SCA. Since the hybrid needed to be powered to communicate with the ASICs, this test also verified the basic functionality of the DC-DC converters. The first batch of four hybrids passed the basic check. As a result, the assembly of the second batch took place.

During this basic test of the second batch, several hybrids showed unresponsive GBT-SCAs. The failure was traced back to a GBT-SCA configuration pin that

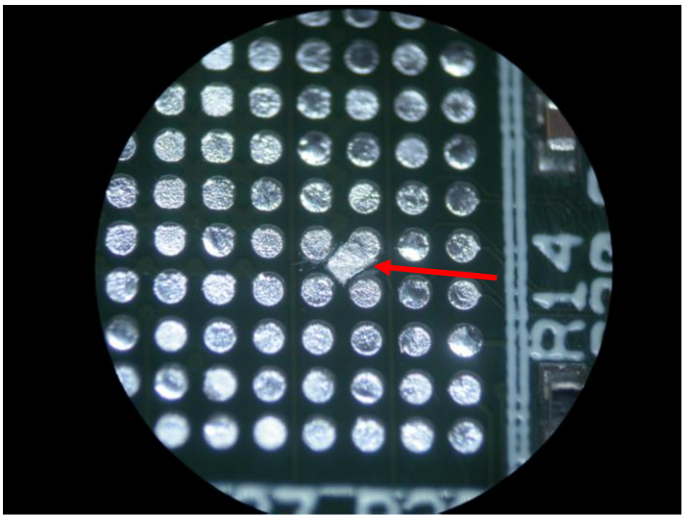

Figure 3.9: Magnified view of the GBT-SCA pad area after the ASIC was removed from the SEH. The enable pad of the auxiliary I^2C client is connected to ground potential by a patch on the PCB level (marked by the red arrow). The patch required removing part of the solder mask and connecting the underlying ground trace to the enable pad.

was left floating in the hybrid's layout. If pulled high, the pin enables an auxiliary I^2C client interface to control the GBT-SCA. An undocumented consequence of the enabled interface is the disabling of the EC interface. The GBTx uses the EC interface in the v3.1 SEH to communicate with the GBT-SCA. As the configuration pin was left floating, an arbitrary voltage level was present and the I^2C interface was active for part of the hybrids. As a consequence, the GBT-SCA ASIC could not be configured. Using probe needles, the differential EC link was contacted and probed with an oscilloscope. The incoming link looked identical on a working and non-working SEH. However, one observed no communication on the outgoing link. This confirmed the disabled EC interface in the GBT-SCA. To recover the hybrids, the GBT-SCA was removed, a patch on the PCB was applied (Figure 3.9), and a new ASIC was assembled.

Of the 30 hybrids of the second batch, communication with the GBTx was possible for 29. Three of these SEHs were damaged or modified during the debugging of the unresponsive GBT-SCAs and their GBT-SCA was not replaced. Of the remaining 26 hybrids, 22 allowed communication with the GBT-SCA after the chip was replaced. They were considered usable for module prototyping. The patching of the hybrid made the bump bonding of the GBT-SCA more difficult and might have resulted in a lower yield. Also, the additional soldering process posed an additional risk of damage.

Before sending the passing hybrids to different module assembly centers, they were tested with the test board at room temperature. The results are summarized below.

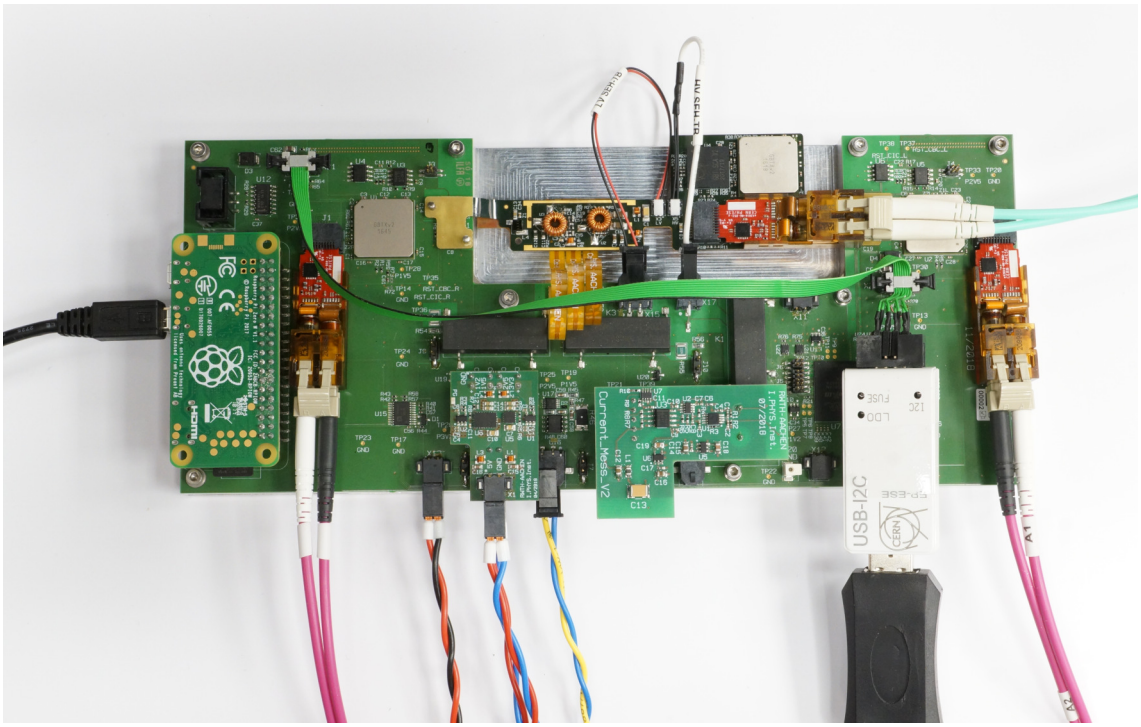

Figure 3.10: Test board setup for the qualification of v3.1 SEH prototypes. A test board v2.1 is used. A hybrid is connected directly to the test board.

Some partially working hybrids were also tested and distributed to support DAQ developments.

The test board setup in its configuration at the time of testing is shown in Figure 3.10. In contrast to the final system shown in Figure 3.1, a test board v2.1 was used. The GBTx ASICs were not fused and were controlled by a USB to I^2C converter. The SEH v3.1 could be plugged directly into the test board without adapter cables. The power daughterboard provided the test board supply voltages from the 5 V input. The measurement of the filtered sensor bias voltage was not yet possible. Only the current could be measured. The back-end board was a GLIB running custom firmware.

3.2.1 Digital readout

To verify the integrity of the optical link, a loop-back test with pseudo-random data sent by the GLIB to the Service Hybrid and back to the GLIB is performed. The hybrid is powered and locks itself to the incoming optical link. The GBTx is programmed with the required configuration to replay the data received on the downlink back on the uplink. Three different kinds of alignment must be performed before a bit error rate can be measured. The gigabit receivers in the GBTx and the

GLIB perform a phase alignment to recover the incoming data correctly. Since the GBTx's high-speed frames are interleaved similarly to the frames used by the lpGBT, the beginning of a frame must be determined by the so-called word alignment. While a fixed pattern is transmitted, the frames received by the GLIB are shifted until they align with the sent frames. Lastly, to account for the latency between sending and receiving the same frame of a pseudo-random pattern in the GLIB, the delay in the GLIB for the frame comparison is varied. This latency alignment as well as the other two alignments are implemented as automated procedures in the software. Subsequently, the firmware within the GLIB compares the sent and received data for five minutes and records mismatches between them. This corresponds to about 10^{12} bits of data.

A total of 23 hybrids are tested in this way. For all of them, the firmware records no errors in the loop-back. According to Poisson statistics, if N bits are transmitted and no transmission errors are observed, the confidence level CL for a specified bit error rate BER_{S} is

$$CL = 1 - e^{-N \times BER_{\mathrm{S}}} \quad . \tag{3.2}$$

The error-free transmission of $N = 10^{12}$ bits corresponds to a bit error rate BER_{S} of less than 6.9×10^{-12} at more than 99.9% confidence level [89].

The data path for this measurement includes the down- and uplink fiber, the receiver in the VTRx module, the GBTx with various data processing steps, the VTRx laser driver, and the receiver in the back-end. The different steps cannot be untangled and the observed rate is only suited to judge the full system but not individual components.

The same test is performed for each GBTx on the test board to also ensure error-free optical communication for the subsequent test of the e-links.

The e-links of the three involved GBTxs are set up to forward or receive data, respectively. The assignment is based on the scheme in Figure 3.2. The phase, word, and latency alignment are performed analogously to the gigabit links for each of the 320 Mb/s e-links. While a phase alignment is required at each transceiver, the word and latency alignment is only required on the receiving end in the back-end board. For five minutes pseudo-random data are sent from the GLIB to the GBTx which injects data on an e-link. Data are received by the second corresponding GBTx, which transmits them to the back-end board, where the data are compared. No wrong bits are recorded in any measurement. For the about 1.2×10^{10} transmitted bits this corresponds to a bit error rate of 7.2×10^{-10} at 99.9% confidence level for the whole system chain.

The two 320 MHz differential clock lines of the SEH are tested using the test board FPGA. The clock drivers are programmed to output their clock signal at the said frequency. Edge counters in the FPGA count the number of edges for 100 µs. The

numbers are compared to the expected amount of 32,000. No deviation of more than ± 1 is observed. Since the phase and frequency of the clock are not aligned with the FPGA, a deviation of ± 1 is expected.

3.2.2 Powering

The SEH needs to provide the supply voltage for the FEH ASICs. To test this functionality, an external one-channel load (BK Precision 8500 [90]) is connected successively to each hybrid sides' output on the test board. The VTRx module is connected to the SEH and the GBTx is configured to receive data on its e-links. The currents and voltages for 0 A, 1 A, 2 A, and for some hybrids 1.7 A are measured. The measurement at 1.7 A became necessary, because of a bug in the test board that prevents an input current measurement for more than 0.82 A at 10 V, which is exceeded by the hybrids at 2 A load. To illustrate the effect, Figure 3.11 shows the output voltage measured at the right side FEH connector output for different loads versus the hybrid's input current. The measurement also includes three hybrids that were tested on a test board v2.2 with a fixed current range. Without a load current flowing across the Schottky diode that separates the second stage DC-DC converter output from the FEH connector output, the hybrids provide 1.6 V. This is the set output voltage of the second-stage DC-DC converter and is needed to power the GBTx and GBT-SCA. The dashed line indicates $1.6 - 0.48 = 1.12$ V, which is the DC-DC converter output voltage minus the voltage drop across the diode (at 2 A and room temperature [67]).

When current is drawn via the FEH connector the voltage drops to about 1.2 V at 1 A load due to the diode. This is the practical working point of the FEH [12]. If the ASICs on the Front-End hybrid require more current, the voltage drop increases further. The input current has to be multiplied by 10 V to obtain the hybrid's power consumption. Since the voltage drop across the diode presents a large penalty on the SEH's output to input power ratio η, it is much worse than foreseen in the final detector. Here, a module power consumption of 5 to 6 W is expected [91]. For the operation of early prototype modules with v3.1 SEH the low η needs to be considered for the module cooling to prevent a shutdown of the DC-DC converters due to the over-temperature protection.

Figure 3.12 shows the output voltage for the left and right side FEH connector output for three different load currents. The 1.7 A measurement is not shown as it overlaps with adjacent data. The same reduction of the output voltage with current is observed. The dashed line shows that the expected voltage drop over the diode of 480 mV at 2 A load is observed. The voltage in general and especially without load current is a bit lower than on the right side due to the longer distance on the SEH to the left side and the resulting voltage drop inside the PCB.

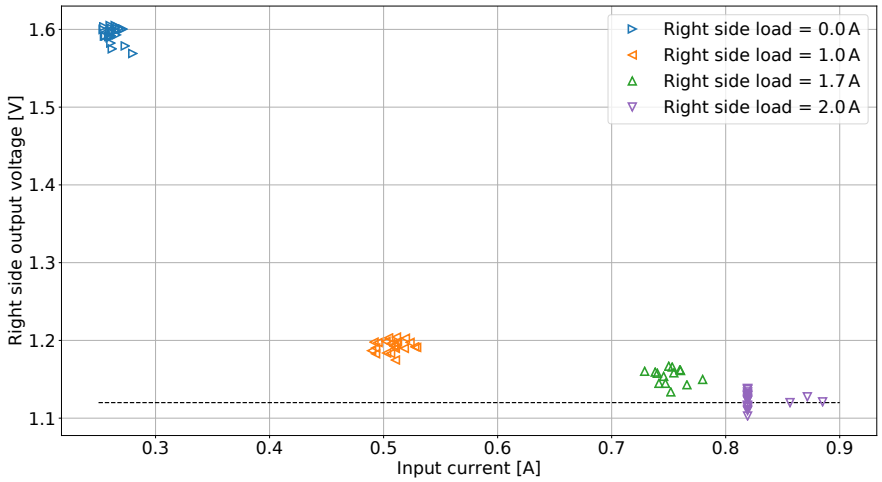

Figure 3.11: Output voltage of a right side FEH connector for different currents drawn at that connector versus the hybrid input current for all tested SEH v3.1. The dashed line corresponds to the nominal DC-DC converter output voltage minus the voltage drop over the output diode.

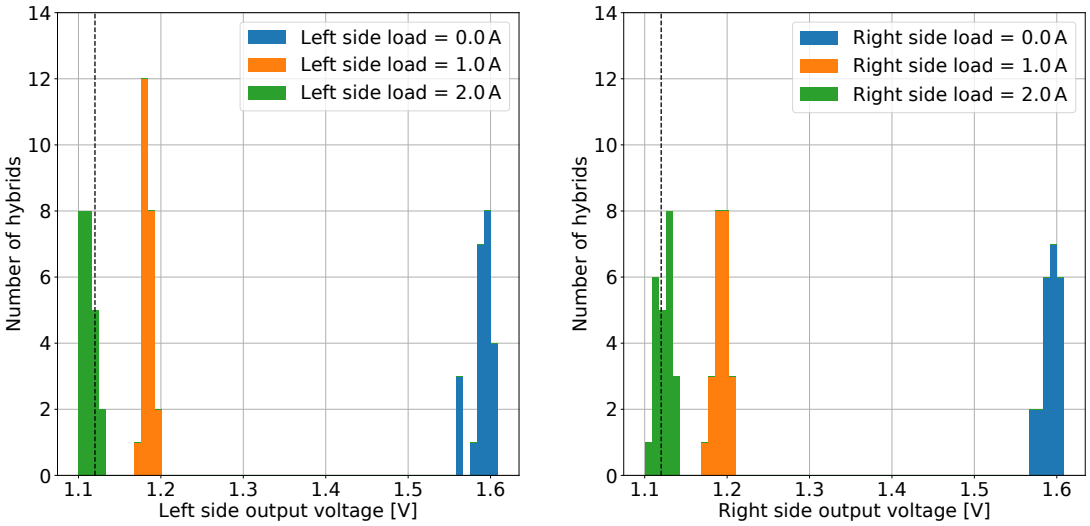

Figure 3.12: Histograms of the output voltages of the left and right side FEH connector for different currents drawn at that connector for all tested SEH v3.1. The dashed line corresponds to the nominal DC-DC converter output voltage minus the voltage drop over the output diode for 2 A.

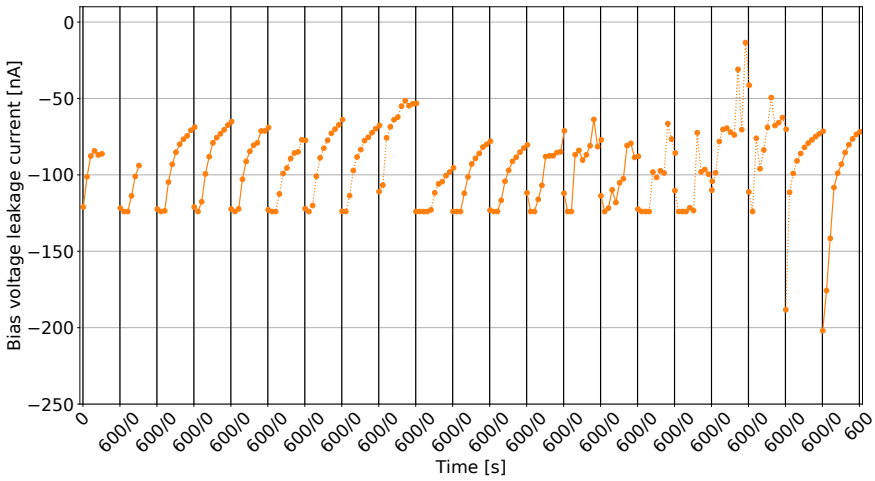

Figure 3.13: Leakage current of the SEH v3.1 bias voltage filter versus measurement time. Each set of data points corresponds to a different hybrid. A negative voltage of 1,000 V is applied. The total measurement time of five minutes was increased to ten minutes after the first two hybrids (left). A different test board with an increased current measurement range was used for the last two hybrids that were tested (right).

The bias voltage isolation of the SEH is tested by applying $-1,000$ V and measuring the resulting leakage current. Each minute ten measurements are taken and the mean is recorded. The result for all tested hybrids is presented in Figure 3.13.

All hybrids show the expected decrease of leakage current with time. After that decrease reached the measurement range, no increase back to the maximum current that would correspond to bad isolation is observed. The monitoring output of the EMCO_P12N power supply always corresponded to the set $-1,000$ V so a voltage breakdown is excluded.

3.2.3 Slow control features

All ASICs on the FEHs are I^2C clients of the GBT-SCA's controllers. The controller for each side is tested by writing a random value once to each register of that side's RAM on the test board. Afterwards, the values are read back from the RAM and compared to the written values. No I^2C transaction errors or mismatches between written and read values are observed for any hybrid or controller.

The GBT-SCA also controls the reset lines of the front-end ASICs. To test those lines, which are connected to the GPIO outputs, the counters of the test board FPGA are used. The test software sends a command to the FPGA to set up the counters. Afterwards, the reset lines are toggled 1,000 times and then the counters

in the FPGA are read. This approach is different from the clock line testing, to allow the synchronization of SEH and test board inside the software. The resets are correctly received for all four reset lines on all tested hybrids.

A number of inputs of the GBT-SCA are also used by the SEH. Two are connected to the right and left side AMUX signals, three monitor the input and output voltages of the DC-DC converters, the received signal strength of the VTRx module is monitored by the V_{RSSI} voltage, and one is operated as a current source and measures the voltage required to drive that current through a fixed resistor on the test board.

These seven voltages and a voltage proportional to the GBT-SCA internal temperature are measured 100 times and the mean is recorded. The AMUX signal is provided by an output of the test board FPGA and its "off" state is also measured.

Figure 3.14 shows the measured voltages for the different channels for the tested SEHs. Where applicable, a dashed line indicates the expected voltage (depending on the voltage levels and voltage dividers used on the hybrid or the sourced current) for a calibrated ADC. In this measurement, the SEH is not calibrated. For all measurements one observes a voltage level that is higher than expected, corresponding to a too-small reference voltage. Across an individual hybrid, the amount of deviation is consistent. Figure 3.15 shows the relative deviation from the mean across all hybrids for the AMUX voltage and the three voltage rail measurements. It proves that the main reason for deviation is the voltage reference. The V_{RSSI} and internal temperature are a system observable of the VTRx module and GBT-SCA, respectively. No common expectation exists for them. Their observed differences across hybrids are considered normal. The 0 V level of the "off" AMUX is correctly measured. The ADC is found to be functional for all Service Hybrids. A calibration is needed for exact measurements on a module.

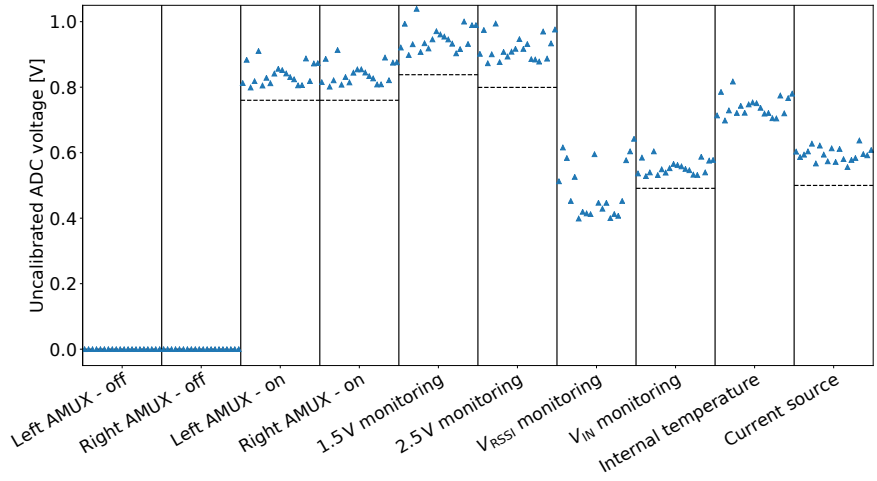

Figure 3.14: Summary of the measurements with the SEH v3.1 GBT-SCA ADC. Each point within a measurement channel belongs to a different hybrid. Dashed lines correspond to the expected level for a calibrated ADC.

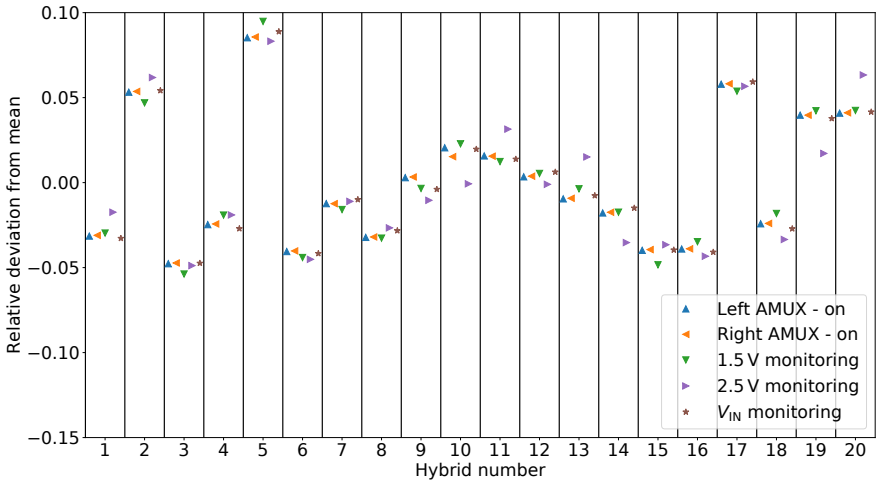

Figure 3.15: Comparison of the relative deviation of selected GBT-SCA ADC measurements per hybrid.

3.2.4 Summary of prototyping experiences

The prototyping campaign of the Service Hybrid v3.1 was successful in delivering functional hybrids for the Phase-2 prototyping effort. After testing and characterization at the I. Physikalisches Institut B of RWTH Aachen University, they were used to exercise the optical readout chain and the powering scheme and to characterize prototype modules and their components [68, 69, 78].

It was possible to integrate the four ASICs and the optical transceiver in a PCB powered by a single 10 V source. Two cascaded DC-DC converters provided the required voltages. Optical communication and data transfer were achieved. The slow-control capabilities required for the use of FEHs were functional. The second stage's output voltage was reduced to a safe level for the FEH ASICs. The sensor bias voltage filter's isolation resistance was found to be sufficient.

The design responsibility for the next prototyping step passed over to CERN [74].

4 Hardware for Service Hybrid production testing

As presented in the previous chapter, the system around the test board greatly contributed to the early development and prototyping phase of the Service Hybrid. After this initial phase, the responsibilities for the different OT hybrid designs were consolidated and centralized at CERN. Around the same time, the preparations for the hybrid production test system as introduced in Section 1.6 started, too. The design began with the backplane and the 2S-FEH hybrid together with its test card. It continued with the PS-FEH and PS-ROH flavors as the required ASICs became available.

This chapter covers the design of the 2S-SEH test card, designed at the I. Physikalisches Institut B of RWTH Aachen University. It also presents performance measurements during the commissioning of the first test card prototypes and the qualification of the final cards for the series production.

4.1 2S-SEH test card design

The design of the hybrids and the test system imposed a high workload on the responsible group at CERN. To support their efforts, the I. Physikalisches Institut B of RWTH Aachen University decided to take over the responsibility for the 2S-SEH test card and the qualification and debugging of hybrid prototypes. This was motivated by the experiences from the design of the test board and the manufacturing of the Service Hybrid v3 and v3.1 prototypes. Before the 2S-SEH test card was developed, both the test board and the PS-ROH test card had been designed. Since both the PS-ROH and the Service Hybrid feature the lpGBT and VTRx$^+$ module, solutions for testing those could be copied.

The presented design covers the second and final 2S-SEH test card version v2. Where relevant, differences to the prototype version are stated.

4.1.1 Design choices

Figure 4.1 shows a photo of the test card and Figure 4.2 the enlarged details of the Service Hybrid fixture. The mechanical constraints that are common for the test

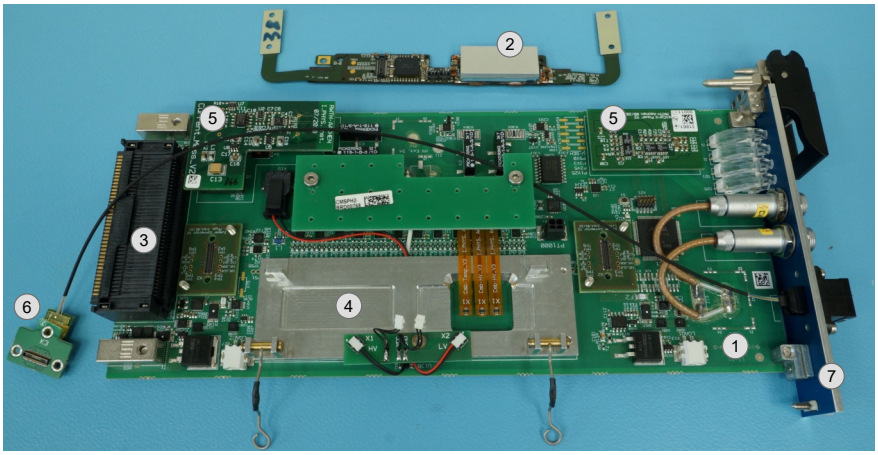

Figure 4.1: Picture of a v2 test card. It is prepared for the mounting of a Service Hybrid (upper part of the picture). The most important components are labeled.

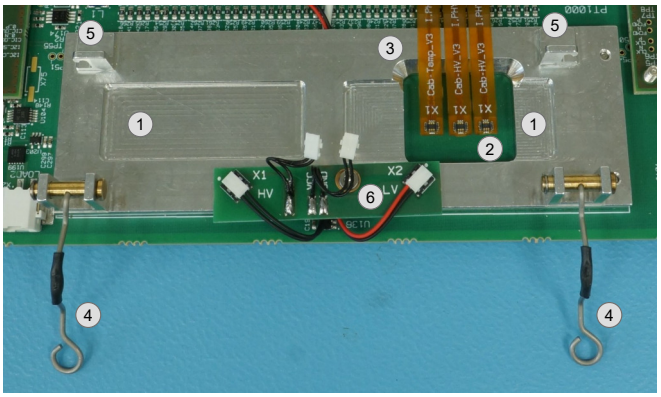

Figure 4.2: Picture of the Service Hybrid fixture on the test card. The most important components are labeled.

cards of all hybrids are summarized in Section 1.6.3.1. The baseline for the design is the PS-ROH test card. The pinout of the main Samtec Searay connector is fixed by the backplane layout. The full schematics of the 2S-SEH test card can be found in Appendix D.

The backplane provides the main supply voltage of positive 3.3 V. A start-up sequence from an active-low input signal (CSEL) is implemented with discrete components. Additional supply voltages (2.5 V and 1.25 V) are provided by voltage regulators on an add-on PCB. It was decided to use this replaceable option since experiences with regulators on the test board showed weaknesses in the GBTx performance due to noise emitted from the power regulators (Section 3.1.1). The add-on PCB option provides more flexibility to address potential issues. The main interface to control the test card is USB. The USB communication is hosted by a CP2130 chip.

This chip is common to all test card flavors and features one-time programmable fuses. The fusible properties include the Product ID that is used to identify the card type and serial number [41]. On the 2S-SEH test card, the CP2130 converts any USB commands to SPI and forwards them to a microcontroller (Microchip ATSAMD51P20A-AF [92]). Subsequently, this microcontroller (µC) controls and reads out all other devices on the test card. If not measured by inputs to the µC itself, voltages and currents are measured with circuits that are also used on the test board (Texas Instruments INA226 [81]) and read out via I^2C.

To perform tests of the bias voltage, a similar circuitry of relays and voltage dividers in combination with the same additional current measurement PCB as on the test board is implemented. For cost reasons, the miniature power supply is not used. Instead, the test bias voltage is applied by an external power supply. The external cable connects to a high-voltage connector on the front panel. A coaxial cable connects the front panel connector to the test card, where its core and shield are soldered directly to the PCB. A second equivalent connection (with the opposite gender type) allows to daisy chain the bias voltage for multiple test cards. A voltage of 5 V is required for the current measurement PCB, which is provided by the USB bus and enabled by the µC.

In order to quantify the SEH's DC-DC converter performance, two programmable transistor loads, based on the Infineon IRL60S216 power MOSFET [93], are integrated on the test card. They can draw up to 2.5 A from either side of the SEH. The µC programming allows the activation of a load current only when the hybrid's output voltage is present. This is necessary to avoid potential harmful oscillations and voltage spikes during the turn-on. For additional studies, e.g. in a bench-top operation, dedicated connectors allow to access the 1.25 V output voltage of the two hybrid sides. Switches enable the different output scenarios.

Three surface-mounted temperature sensors are placed on the test card (one near each load and one next to the hybrid). An external PT1000 sensor can be connected, too.

Level translators for the lpGBT's in- and outgoing e-link, clock, and I^2C controller lines are copied from the design of the PS-ROH test card, as well as the fanout buffers for the incoming e-links. The level translators are required to match the 3.3 V differential transmission from the FC7 with the CERN Low Power Signaling used by the lpGBT. The fanout buffers are necessary since the backplane does not provide enough outputs for all lpGBT inputs. The I^2C between the lpGBT and the test card also undergoes a level-shifting from 1.25 to 3.3 V. As all data transmissions on the backplane are unidirectional and differential, the usually bidirectional I^2C data lines of the right and left side are divided into an in- and outgoing branch on the test card. The I^2C clock lines are unidirectional, driven by the lpGBT, and only converted from single-ended to differential signaling on the test card. A multiplexer

in conjunction with a digital-to-analog converter is used to provide variable voltages to the SEH's AMUX inputs.

In addition, a circuitry to configure the lpGBT via its client I^2C interface and to program the lpGBT's e-fuses is implemented. The programming requires the application of 2.5 V to a lpGBT input for a very brief amount of time. The voltage can either be switched by the µC or by a GPIO output of the lpGBT itself. As a safety measure, a hardware switch usually disables the logic.

An LED array at the front panel provides information about the state of a test card.

Due to the limited available space on the test card, additional connectors for a separate power or USB connection are not present. An additional auxiliary PCB is required to enable the commissioning and evaluation of the boards in a table-top setup without a crate and four-slot backplane. This "mini-backplane" provides the transitions of the signals of the electrical FC7, the supply and test voltages, and the USB connection. A couple of switches allow the selection of a test mode without an FC7.

4.1.2 Mechanics

Counterparts of the Service Hybrid's FEH connectors are placed directly on the test card. However, matching sacrificial adapters are used between the test card and the hybrid. They serve several purposes:

- The connectors are fragile and have a limited number of recommended mating cycles. If a connector on an adapter is damaged the adapter can be replaced.
- Eventual changes in the SEH geometry can be absorbed in a redesign of the adapter.
- They feature readily accessible probe points for debugging.

The test cards will be placed in an upright position in the test crate. The hybrid is connected via a FEH connector on either side. An additional way of fixating the hybrid and avoiding stress on the fragile connectors[1] is needed. In the module, the Service Hybrid is glued to the AlCF sensor bridges. Due to the tight space and material budget constraints, no dedicated features could be implemented on the hybrid for fixation in the test system. In addition, the final outline, dimensions, and component placement might change due to unforeseen circumstances. Any design must be flexible enough to accommodate potential changes.

On the top side of the test card, a space of $35 \times 125\,\mathrm{mm}^2$ is reserved for an aluminum fixture. A detailed view of the fixture is shown in Figure 4.2. It can be attached with

[1] After the inclusion of the L-shaped tails in the v4 prototype design (Section 2.2.3), no more connectors were damaged.

M2 screws via five holes in the test card. Groves for both SEH compensators provide a reference to position the hybrid. The groves are placed so the Front-End Hybrid connectors can be mated without stressing or bending the L-shaped tails. The right grove is extended by a cut-out for the bias voltage filter components on the hybrid backside. The three connectors on the hybrid backside are connected with flat flex cables (similar to the ones used in the module and the test board) to the test card. To avoid damage due to sharp edges, a slope is added on the side where the flex cables are routed. On both sides of the hybrid, single-wire springs press down on the components-free part of the flex with the carbon fiber stiffener and the aluminum surface below. The mechanism to fix the hybrid is identical to the one used in the test board cycling setup. A replaceable patch board is screwed to the fixture and features cables and receptacles for the connection of the hybrid's DF57 connectors.

Cables provide both the low- and bias voltage connection between the test card and the patch board. They are routed from the DF57 connectors patch board in a slit on the backside of the support plate to a connector (a solder pad) for the low-voltage (bias voltage). The originally foreseen connector on the test card for the bias voltage turned out to be too tall for the safe insertion and extraction of the cards in the crate without risking damage to the electrical isolation. The patch board is necessary as the DF57 is only rated for 30 mating cycles and a frequent replacement of the cables on the test card is not feasible due to the routing and soldering. Instead, the patch boards will be replaced.

Each test card has a front panel. It features an extraction handle and labeling for the indicative LED array. Also, the external bias voltage is brought to the test card with the high-voltage connector (male-type) installed in the front panel. The identically installed female-type connector is used to daisy chain the provided bias voltage to additional test cards in the crate. Lastly, the front panel has an MT/MTP adapter (US Conec C18017 [94]) installed. A VTRx$^+$ module is permanently installed on the test card. To store and protect its fiber, an elevated piece of FR4 is fixed on two rods, and the fiber is placed under it. The routing allows the connection and removal of the module while constraining the fiber to the allocated height of the test card. Sacrificial adaptor PCBs are used between the VTRx$^+$ module and the hybrid in order to extend the lifetime of the module's connector.

Exposed pads that present the bias voltage are covered by a conformal coating.

4.1.3 Test descriptions

The test card was designed to provide an interface with each incoming and outgoing connection of the SEH. In combination with the microcontroller and the FC7s' FPGAs, it makes up a highly flexible environment. This section describes how the interfaces, lines, and ASICs are tested on the test card.

In general, tests can be performed for any given input voltage within the operation range of the DC-DC converters. The additional loads can also be adjusted individually for each side. No malfunctions have been observed for any combination of these parameters within realistic bounds of the detector operation conditions. To stay within a reasonable amount of time, testing will be performed at a Service Hybrid input voltage of 10.5 V and a load of 1.2 A per side. The input voltage ensures the hybrid is enabled under all circumstances while leaving some margin to the maximum allowed input voltage. The additional load is the nominal expected one. During prototype testing, an additional test without external load was also performed.

4.1.3.1 Short circuit check

A common failure mode in electronics is a short circuit. Assembly failures or mechanical damage can lead to an abnormal connection between two different voltage nodes. This leads to an electric current flow limited only by the equivalent resistance of the rest of the network. In the test system, all power supply channels are used with a digital fuse that turns off a channel if a set current limit is reached.

After a test card is powered and enabled, the Service Hybrid's test voltage (DC-DC converter input voltage) is not forwarded directly to the hybrid. It is turned on at the start of a test procedure by the software and a voltage measurement on the power supply output channel via the TCP server is performed. If a short circuit is present on the hybrid, the input current will exceed the set current limit and trigger the digital fuse. Consequently, the power supply will turn off the test voltage channel and the subsequent output voltage measurement will be 0 V. A turned-off hybrid test voltage will also disable the test card, so the voltage measurement cannot be performed by the test card itself.

The detection of a short will terminate the test procedure. For the next test, the GUI can turn the power supply channel back on.

4.1.3.2 Service Hybrid I/V behavior

As discussed in Section 2.3.3, the voltages for which the hybrid's DC-DC converters turn on and off are important quantities for the detector operation. In order to measure these voltages a current versus voltage (I/V) scan in the hybrid test voltage is performed. This requires the remote control and synchronization of the power supply with the test software and therefore relies on the TCP server. Starting with the lowest value, the voltage is varied between 5 and 11 V and, after a short delay to allow a stable measurement, the SEH input voltage and current are measured. After the maximal voltage is reached, the voltage is lowered again to record the turn-off process. Close to the expected turn-on/turn-off voltages the voltage steps are reduced to increase resolution. In the analysis, the turn-on can be identified by a large increase in the current and the turn-off by a decrease. A current of roughly

10 mA is drawn by the disabled DC-DC converter. After the enable threshold is passed, the bare hybrid has a power consumption of about 110 mA at 10.5 V. Since the active transistor loads should only draw power while the DC-DC converters provide a stable voltage, the test is performed without enabling the loads. The test also checks the general DC-DC converter functionality across a wider range of input voltages. Failures lead to deviations of the shape of the hysteresis.

After the I/V scan is finished, the input voltage is reset to the starting condition, and the load is applied.

4.1.3.3 Optical link and configuration

While the prototype hybrids with lpGBT_v0 were initially configured and later fused via the test card, hybrids with the lpGBT_v1 are equipped only for configuration through the optical link.

The optical FC7 continuously outputs the downlink frames on all connected fiber pairs. After a hybrid is turned on, its lpGBT will lock itself onto that data stream. In the 2S-SEH case, the lpGBT is configured such that it starts to output its uplink frames autonomously. The software initializes a set-up sequence in the FC7 to recognize the headers of the uplink frames and lock onto the stream. After a successful locking, a configuration is written to the lpGBT via the serial IC interface of the optical link. The configuration is also used in the 2S module and configures the in- and outgoing e-links and clocks with their respective frequency and signal modulation settings and the GPIO direction and level settings.

A failed locking inside the FC7 or errors during the configuration will terminate the test procedure.

4.1.3.4 E-link receivers

Twelve input e-links at 320 Mb/s are used in the SEH. They are connected via the buffers on the test card and backplane to the electrical FC7. Its FPGA can provide arbitrary periodic patterns with a length of 8 bits. Before data can be read out, a phase alignment of the lpGBT inputs is performed. The inputs have 15 different phase settings with a step size of $\pm T/8$, where T is the bit period. The phase alignment relies on the lpGBT internal phase training algorithm. During the training, the FC7 outputs AA_{hex} on all links. Figure 4.3 (left) shows a comparison between a manual phase scan and the outcome of the training. As the phase settings span almost two bit periods, two inefficient regions are visible. In the manual scan, most inefficient settings show low error rates and some e-links present no suboptimal value. This is caused by the coarse phase settings in combination with steep output signal flanks from the test card buffers and the limited statistics of the manual scan. The training results lie well within the error-free region of the manual scan.

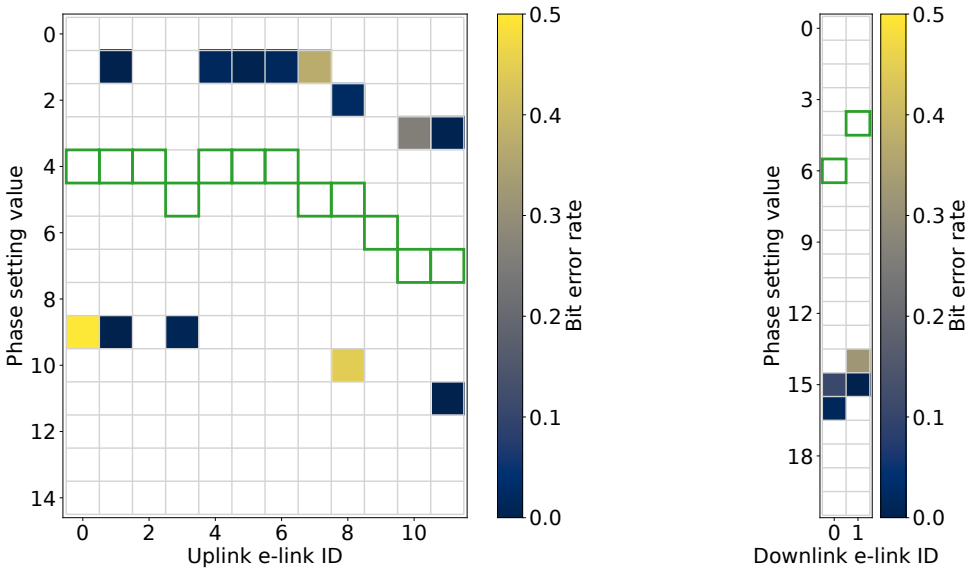

Figure 4.3: Comparison of the automatic phase alignment procedures for the Service Hybrid lpGBT e-links with a manual scan. Green borders mark the results of the automatic procedures. The result of the manual phase scan is color-coded. White fields have a bit error rate of zero. Left: results for the uplink e-link receivers. Right: results for the downlink e-link drivers.

Five different patterns are used for the test of the transmission itself: 00_{hex}, FF_{hex}, CA_{hex}, CC_{hex}, and the phase alignment pattern AA_{hex}. Shorts and opens on the differential lines lead to poor signal quality but low-/high-level transitions might still be correctly received by the lpGBT. The static 00_{hex} and FF_{hex} patterns can identify different types of shorts and opens. For instance, if the negative part of a differential pair is shorted to ground both static patterns will be received as FF_{hex}.

During a test, the firmware extracts four consecutive received words for each e-link from the optical link and provides them to the software. The software performs a word alignment to the data by looking for the best agreement with the broadcasted pattern. The number of non-matching bits per pattern is stored. For a successful test, no wrong bits are allowed for any pattern on any link.

4.1.3.5 E-link drivers

The two outgoing (fast command) e-links are tested in a similar fashion. The lpGBT outputs can be configured to either provide deserialized downstream data or a fixed pattern[2] that is received in the electrical FC7. The fixed pattern AA_{hex} is used for the

[2]Two additional modes are possible, but are not used in the test system: a binary counter counting up and a pseudo-random pattern.

phase alignment in the FC7. A comparison between the automatic and manual phase alignment is shown in Figure 4.3 (right). In the presented case, the automatically chosen value is well separated from the failing region.

For the test, the same fixed patterns as for the receivers are used. From the optical down-stream only idle frames ($C1_{hex}$) are employed. For a successful test, no wrong bits are allowed for any pattern on any driver.

4.1.3.6 Clocks

The SEH provides two 320 MHz clock signals for the ASICs on both FEHs. In the test system, those outputs are routed to the electrical FC7. The firmware features a clock test block that compares the received clock to an internally generated reference one. To do so, signal edge counting is used to count the number of clock edges for 750 µs. The counted amount for the received and reference clock is compared. A failure is reported, if the difference exceeds a tolerance of five counts.

The test is repeated three times. The first run uses the settings employed during module operation. It is not sensitive if one half of the differential clock pair features an open connection or a short to ground or the pair is shortened. The buffers and level translators between the lpGBT and FC7 are the reasons for this weakness. They tolerate a wide range of input waveforms that would normally be insufficient for front-end ASICs. Unlike the e-links, the lpGBT's clocks cannot present a static high or low pattern that easily identifies short and open connections. As a mitigation that avoids a hardware change on the test card, the next run is performed with the minimal clock frequency of 40 MHz and the lowest possible driving strength setting. This configuration can identify a short in a differential pair. To also identify an open connection, the last run is performed at 40 MHz and with the maximal setting of both driving strength and pulse width of the pre-emphasis, which is usually deactivated. The tests need to pass for all three settings.

4.1.3.7 I^2C controllers

The right and left side FEH chips are controlled by I^2C bus lines connected to the controllers 0 and 2 of the lpGBT. On the test card, the two buses are connected to the electrical FC7, which has appropriate I^2C clients implemented. During a test, write transactions to these clients are initiated and the lpGBT internal status register for the I^2C controller is monitored. Any number of transactions can be chosen. However, five consecutive failing transactions will terminate the test to avoid an excessive amount of delay from timeouts. From experience, 10,000 transactions offer a good compromise between testing time and confidence in the link integrity. A failed transaction is considered a bad test.

The I^2C controller 1 is connected to the $VTRx^+$ laser driver ASIC and therefore the same type of client as in the module. It is tested by reading a set of registers and comparing them to their default values. If a transaction fails based on the status register readback, or a value does not match the expected default, the test is considered failed.

4.1.3.8 Reset lines

The four lines for the reset of the front-end chips are connected to the lpGBT's GPIO pins. An ADC in the test card microcontroller can measure their voltage levels. The respective GPIO pins are configured as outputs in the module and the test system. A measurement is performed for the low and high state and compared to the 0 V and the expected level of the second stage DC-DC converter output voltage, which powers the lpGBT and sets the level. A deviation greater than 100 mV for any level or line constitutes a failed test.

In the most recent design for the kick-off SEH (Section 7) an additional reset line for the laser driver of the $VTRx^+$ module was added. It is tested by writing a non-default value to an uncritical register of the laser driver via the I^2C interface, resetting it by releasing and resetting the active-low reset pin connected to the lpGBT GPIO, and rereading the register. If the default value is read back after the reset, the test is passed. Additional register readings are performed before and after the write step to confirm it contains the expected value.

4.1.3.9 ADC lines

Before any ADC measurement is performed, the voltage reference is tuned to 1 V, and the gain and offset are determined. For the tuning, the test card applies a reference voltage from its digital-to-analog converter.

The lpGBT takes ten measurements per external ADC input discussed in Section 1.5.2.3. Coarse cuts for grading are applied to check for obvious failures. The results of the tuning and the measurements are stored for offline analysis and grading.

4.1.3.10 Front-End Hybrid supply voltage

The two DC-DC converters provide the supply voltage of around 1.25 V for the CIC and CBC ASICs. To characterize the SEH in terms of this output voltage, the active transistor loads are used. The test performs sweeps in the amount of current drawn from the bPOL2V5. The number of selected current values depends on the available measurement time. For production testing, we chose current steps of 500 mA with a maximal load of 1.5 A per side. The SEH input voltage and current as well as the second stage DC-DC converter output voltages and currents on both the right and

left side are measured. The first stage output voltage is also monitored from a dedicated signal line that is present in the FEH connector. The line is unconnected on the Front-End Hybrid and only available in the hybrid test system. Each measurement point requires three seconds of settling time for the loads and the measurement circuits. The test procedure is divided into three phases. First, the loads on both sides are set to draw the same current. In the second (third) phase the right (left) side only is stressed. When both sides are loaded, the SEH DC-DC converter output to input power ratio is comparable to the situation of the 2S module. The absolute DC-DC conversion efficiency is not easily quantifiable due to the a-priory unknown power consumption of the active lpGBT and VTRx$^+$ module. Both contribute to the SEH power intake. From the single-sided measurements, the resistance from a hybrid side's output filter to the FEH connector can be estimated.

4.1.3.11 Sensor bias voltage filter

Two aspects of the sensor bias voltage are tested: the leakage current between the bias voltage potential and ground as a measure of the electrical isolation, and the voltage connection to the hybrid, its routing through the RC filter, and the connection back to the test card via flexible Kapton tails.

The leakage current is measured for a negative voltage level of 1,000 V. When the relays are set to the appropriate states and the voltage is applied, the output voltage and current of the power supply and the current on the test card are measured. The duration of the test is limited by the available testing time. During the prototyping campaign, the time was reduced from ten minutes to 150 s. During production testing, the bias voltage will be applied and the current measured while the other tests are running. At the end of the measurement, the voltage is set back to 0 V and the channel is turned off. To pass the test, the measured voltage on the test card must be compatible with the applied voltage, and no large currents are allowed after a charging period.

The measurement of the connection is the last step in testing a hybrid. After the relays are set to the correct state for the connection measurement, the output of the high-voltage power supply is turned on again and the received voltage on the test card is measured. The absolute voltage is increased in steps of 200 V and a delay of 4.5 s before every voltage measurement is added to allow the voltage to settle. The measured voltage must correspond to the set voltage to pass the test. The measurement time can be further reduced by requiring fewer voltage steps. Section 3.1.2 describes the basic measurement principles in more detail.

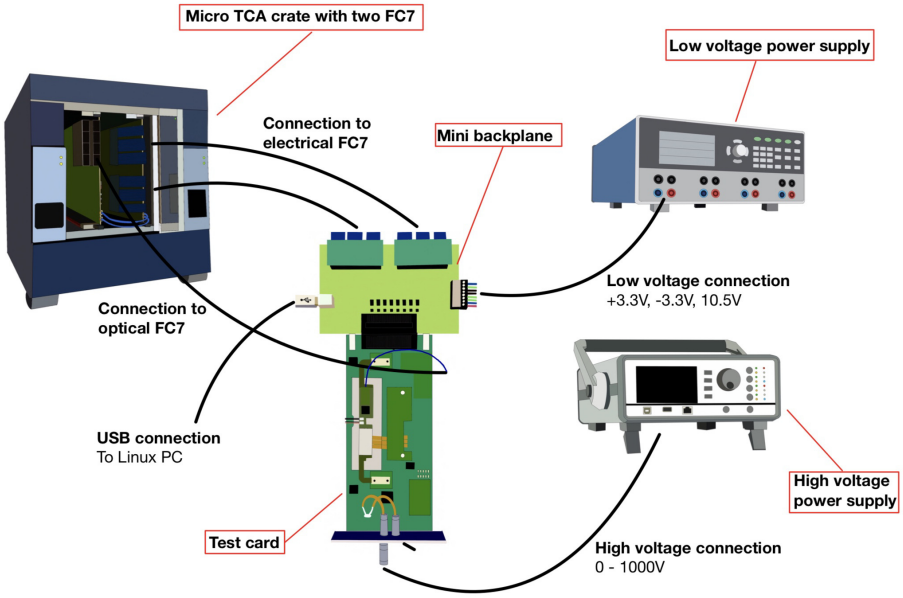

Figure 4.4: Sketch of the table-top setup for the test card commissioning. The most important components are labeled [95].

4.2 Test card commissioning

The general concept of the 2S-SEH test card was compiled within the framework of this thesis. The design and the layout of the board and the auxiliary components were provided by the electrical engineering department of the I. Physikalisches Institut B of RWTH Aachen University under the supervision of Daniel Louis, who also oversaw the production.

A complete evaluation of the performance of a test card is only possible with a fully functional hybrid and vice-versa. The commissioning of the test card and the whole test system was a lengthy and iterative process. The rest of this chapter covers the aspects that are mostly restricted to the test card itself. Three of the next four chapters cover the test results of different hybrid generations that were obtained using those test cards.

A table-top setup is available to connect a 2S-SEH test card to the electrical FC7 and the power supply without a crate. The mini-backplane replaces the backplane and provides the connections in this configuration. Figure 4.4 shows a sketch of the setup.

4.2.1 Software developments

The test routines described in the previous section were implemented in the *Ph2_ACF* software as part of this thesis. The work also included the support of the

2S-SEH test card in the GUI. Routines for the equivalent tests on the PS-ROH test were developed in parallel and with my consultation by Younes Otarid [96].

The abstraction of the test card hardware for the *Ph2_ACF* software is performed in a test card software library. The library is written in *C++* and can also be used stand-alone without the *Ph2_ACF* software. This is useful for hardware validation and debugging since no FC7 interface is required. The 2S-SEH test card interface is implemented as a class. It utilizes a generic USB driver library for Linux-based systems to communicate with the test card's USB-to-SPI bridge CP2130. The software library was created within the framework of this thesis.

When the test card class is initiated, the hardware is identified from the device and vendor ID and claimed. In the case of the production system with up to three connected and active test cards (one per crate), the USB bus can also be specified to claim the correct device.

The CP2130 acts as the SPI master and the µC as the only SPI slave. The software issues a bulk USB transaction for all commands, which prompts a corresponding write-read transaction on the SPI bus.

Daniel Louis implemented all further hardware abstractions like I^2C communication with circuits on the test card in the µC firmware. Measurements on these circuits are performed periodically while the µC is idle. The timeout loop stores new values roughly every 0.8 s. The storing and updating ensure the µC can always ship out a result when a write-read on the SPI bus is performed. The µC interrupts the loop to respond to the requests. If new settings are written to the µC, a flag is raised and the loop processes the values. The new settings need at most one timeout period to take effect. The software needs to account for the delay required to obtain updated values after the test card state or test settings change.

4.2.2 Electrical characterization

In the design section of this chapter, it was discussed that many components were already used successfully in other test systems. The buffers for the e-links, clocks, and I^2C lines were copied from the PS-ROH test card design. To test those without a Service Hybrid a Versatile Link Plus Demo Board ($VLDB^+$) was used [97]. It also hosts an lpGBT ASIC. Jumper cables connected the $VLDB^+$ signals to appropriate test points on the test card. The data path of the e-links and clock lines to the electrical FC7 was verified in this way. Data were recovered correctly by the firmware and displayed by the software. In the opposite direction, data patterns were sent by the electrical FC7. Figure 4.5 shows a measurement of an eye diagram for data from the electrical FC7 measured at test points on the FEH connector. An eye diagram is an oscilloscope display of a repetitively sampled digital signal. The samples are

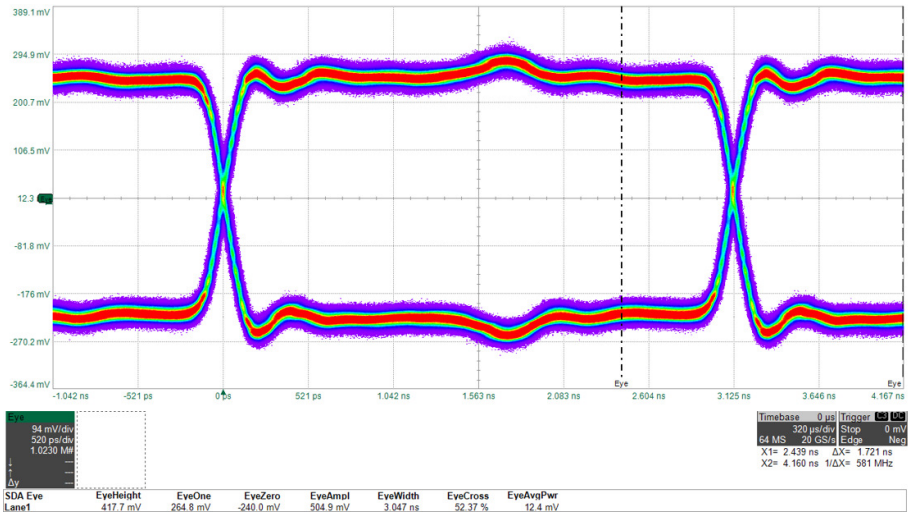

Figure 4.5: Eye diagram of the differential output of a test card buffer. The constant pattern is transmitted by the electrical FC7 and measured on test points on the FEH connector.

triggered with the data rate. The measured eye[3] is wide open with steep flanks and low jitter. It proves a good signal transmission by the buffers and the PCB. A differential oscilloscope and a 100 Ω resistor as termination were used. Oscilloscope measurements of the I^2C transactions and the positive transaction status reported by the $VLDB^+$ verified the I^2C functionality.

The circuits to measure supply or output voltages and currents are a copy of the comparable circuits on the SEH test board. This is also true for the temperature measurements. The PCB and components to measure the sensor bias voltage and isolation were carried over to the test card as well. They have been extensively studied on the test board (Chapter 3). More important was the new aspect of the readout and control via the microcontroller. After the necessary software was implemented, all functionalities were exercised. The measured values were all in agreement with the expectation.

In separate measurements, the card designer verified the general power integrity, the µC programming, and the active transistor loads.

4.2.3 Mechanical compatibility

The size of the test card and the main connector position match the specifications for all cards. No components are placed near the long edges. The card glides safely in the crate guide rails and the main connector mates easily. The retraction handle on the front panel functions properly and a card can be removed easily. Since all

[3]The area enclosed by the sampled waveforms.

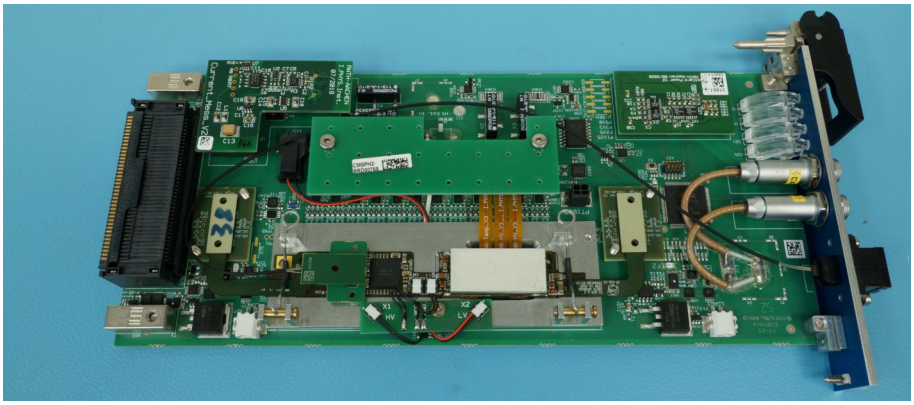

Figure 4.6: Picture of a test card v2 with a mounted Service Hybrid. The card is from the second version and all cables and the VTRx$^+$ module are connected.

components are within the envelope of the front panel, multiple test cards can be installed next to each other without any risk of damaging neighboring cards.

The Service Hybrid compensator positions the hybrid on the cooling block. The single-wire springs keep it in place without any damage. Four screws attach the cooling block to the test card. During its first assembly, the position can be slightly varied to find the correct hybrid placement for which the FEH connector can be mated. The power cable and flexible tail lengths have been optimized and fit the needs. The optimal fiber length of the VTRx$^+$ module was identified and it amounts to 28.5 cm. The fiber describes a gentle 180° turn over most of the card area and does not stick out across the highest parts on the card. Figure 4.6 shows an image of a test card with a connected Service Hybrid. All mechanical constraints are fulfilled.

4.3 2S-SEH Test card production for serial testing

The first test card version was used extensively to qualify the first batches of the final Service Hybrid prototypes. Chapter 5 covers the results. The testing demonstrated all the functionalities of the test cards and the hybrids. The findings led to minor improvements and design corrections. Most notable is the implementation of the daisy-chain connection for the external high-voltage.

To cover the needs during series production, a larger test card quantity is required. The production plans foresee the installation of one full crate (twelve cards) for Service Hybrid testing at the manufacturer and at CERN. Sufficient test cards should be available to unmount and mount hybrids, while others are tested in the climatic chamber. This is necessary to parallelize partially and speed up the testing process. These two requirements give a demand of at least 48 test cards. 90 test cards with the improved second design were produced. This also accounted for losses and provided some spares.

	First stage	Second stage	Third stage
Setup	Table-top	Table-top	Crate
Environment	Room temperature	Room temperature	−35°C and +40°C Climatic chamber
FC7 usage	No	Yes	Yes
SEH usage	No	Yes	Yes
Additional tasks	Visual inspection Microcontroller programming	Cooling block mounting Serial number programming	

Table 4.1: Summary table of the three stages of the test card quality control.

The test cards' functionality was verified during an extensive testing program. The goal was to identify test cards that work reliably. Debugging is not possible at CERN or at the manufacturer. The procedures for the quality control efforts were developed and the work was supervised within the framework of this thesis. The results are compiled in Felix Thurn's Master's thesis [95]. Here, the most important findings will be summarized.

The test card quality control was divided into three stages. Table 4.1 gives an overview of their characteristics. The first stage was performed after the reception of the assembled PCBs. The cards were visually inspected and the auxiliary boards, the front panel, and the cables were mounted. The test cards were also powered for the first time and the microcontroller firmware was programmed. The setup used the table-top configuration but was not connected to an electrical FC7. In total 88 out of 90 cards passed this stage. Two cards could not be powered and programmed.

In the second stage, the test cards were tested with a Service Hybrid prototype. The hybrid was identified as fully functional with the first test card version. Before the hybrid could be mounted, the cooling block was positioned with the hybrid. All tests were performed with a table-top setup and at room temperature. A serial number and card type identifier were permanently stored in the USB-to-SPI bridge. The test script had two parts. First, the supply and output voltages and currents were measured. Different load currents were also applied. This test could spot critical problems before a longer test routine began and high-voltage was applied. Second, a complete hybrid test routine was performed. The results were analyzed and inspected for deviations from the expectations. This stage spotted four conspicuous cards. Wrong resistors were assembled on two cards, one card featured a wrongly orientated chip. These three could be repaired and afterward passed the testing. One card could not be powered when connected to and enabled by the electrical FC7. Using X-ray imaging, the failure was traced back to a missing solder connection in the backplane connector. It could not be repaired. All other cards were inspected with X-rays and open connections were spotted on three cards. The affected pins are redundant connections. After repair, 87 test cards passed the second stage tests.

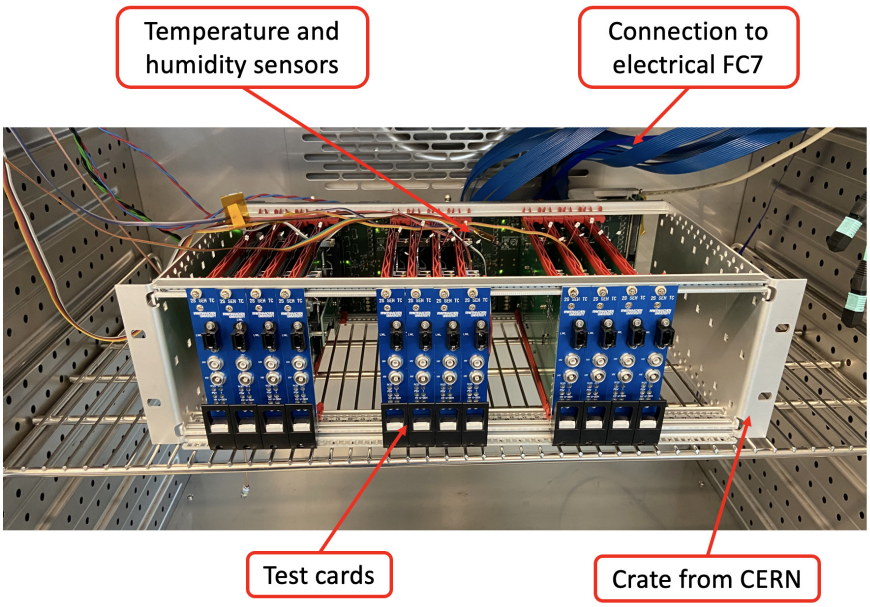

Figure 4.7: Picture of the setup inside the climatic chamber (Binder MK240 [98]) that was used for the third stage of test card quality control. The most important components are labeled. The optical fibers and bias voltage cables are not connected to the front panel [95].

In May 2023, CERN requested 2S-SEH test cards to exercise and finalize the hybrid test system. Thus, 20 test cards were brought to CERN and were not tested in the third stage.

The third and last part of the quality control was carried out with a crate installed in a climatic chamber. The remaining 67 cards were tested at the planned temperatures of +40°C and −35°C. Several thermal cycles with and without hybrids were performed. Testing at cold temperatures was required to ensure full operation of the test system and to avoid hardware-related delays during production. Figure 4.7 shows a picture of the setup inside the climatic chamber.

The tests in the climatic chamber revealed a mistake in the microcontroller firmware that prevented some cards from operating at cold temperatures. Closer investigations revealed that the 5.0 V supply voltage for the current measurement PCB was not switched correctly. The explanation for the issue is illustrated in Figure 4.8. The responsible switch (Semiconductor Components Industries NCP45540 [99]) is powered by the 3.3 V supply voltage rail ("P3V3") and enabled by an active high signal ("Enable") from the microcontroller after its start-up. The left graph shows the time relation of P3V3 and Enable with the original firmware. The microcontroller toggles the enable signal before P3V3 has settled. The switch's manual states that under this condition the enable may not take effect. An additional delay in the firmware

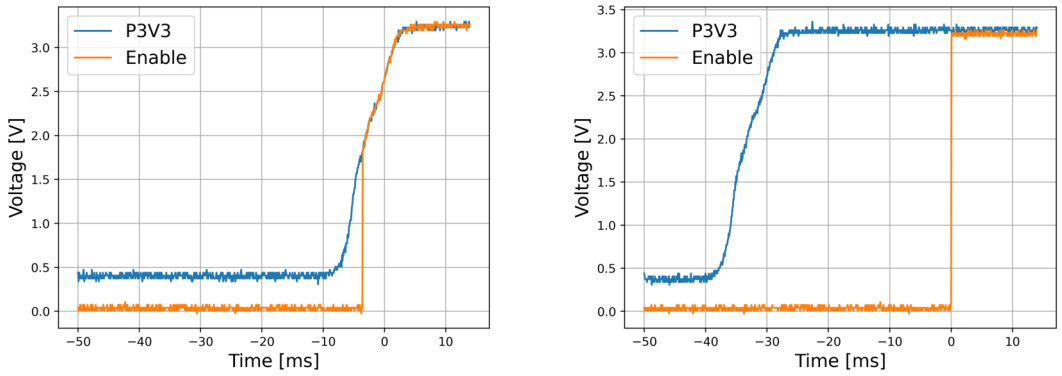

Figure 4.8: Voltage versus time measurements of the 3.3 V supply voltage rail (blue) and the enable signal of the 5.0 V supply voltage switch (orange). Left: first firmware version. Right: adjusted firmware version [95].

(right graph) resolves the issue at cold temperatures and provides a large margin at room temperature.

After the firmware was fixed, the most common observed failure was a missing USB response of the test card. Four of the 67 test cards suffered from this failure mode. CERN reported similar issues for the other test cards. The presumed explanation is an internal weakness in the USB-to-SPI bridge integrated circuit. As mentioned earlier, the CP2130 is common for all test cards.

After all test cards that were repaired, failed in cold, or showed open non-redundant connections in the backplane connector were removed, 58 of the 67 cards still at Aachen passed all tests. They were delivered to CERN for installation at the contractor's side in February 2024.

The testing at room temperature and inside the climatic chamber also served as separate studies of the Service Hybrid. At room temperature, all cards were tested with the same hybrid. For one, this demonstrated the high durability of the connectors and the resilience of the Service Hybrid in terms of handling. The testing also provided a data sample to assess the card-to-card variations of the measured quantities. The variations can be compared with the hybrid-to-hybrid variations obtained during the prototype testing. The prototypes were mainly tested with only a few v1 test cards. The test cards measure voltages with a precision of 0.1%, and currents with an accuracy of 1%. The voltage variations between cards are one order of magnitude lower than those between hybrids. The hybrid results are presented in the next chapters.

The results from the climatic chamber studies provide information on the Service Hybrid behavior at different temperatures. Those results are analyzed in Section 5.2.5.

In addition, the used hybrids underwent a high number of thermal cycles without visible damage.

In conclusion, the quality control process successfully provided a sufficient amount of test cards for the hybrid testing. It also benefited the evaluation of the hybrid performance and characterization.

5 Characterization of final Service Hybrid prototypes

The SEH prototype development, design, and production advanced in parallel with implementing hardware, software, and firmware for the test system. The final Service Hybrid prototypes were the first SEHs produced by CERN. This chapter summarizes the most important findings of the testing campaigns performed for each production lot of hybrids. Since these hybrids were in general urgently expected by module production centers to be used for 2S module prototyping, the testing concentrated on confirming general functionality and finding critical problems. During the span of the prototyping campaigns, the test system underwent several changes and improvements. These changes will be discussed, whenever they affect the consistency of the presented results.

The first batches from contractor A and contractor B required the e-fusing of a minimal configuration inside the lpGBT. The fusing was performed after the hybrid was installed on a test card and before the functional testing of the hybrid. The testing and fusing are performed for the two batches with the first version of the test card. The second batch of contractor A was tested with the second and final test card version.

All hybrids underwent a visual inspection at CERN before their delivery to Aachen.

5.1 Observed failures

The used ASICs of the different batches and their delivery dates are summarized in Table 2.1 and Table 2.3, respectively. Before the quantitative analysis of hybrid properties is discussed, an overview of the outcome of the functional testing is presented. A summary is given in Table 5.1. To better understand and contextualize the failure types they are discussed in detail below.

	Contractor A batch 1	Contractor B batch 1	Contractor A batch 2
Total number of tested hybrids	27	25	22
Functional	16	13	18
Failing	11	12	4
PCB design and manufacturing	0	2*	0
Assembly and handling	0	8	0**
Active components	11	2	4

Table 5.1: Summary of the functional testing results of final Service Hybrid prototypes.
* Does not include failures observed in modules.
** After adjustment of the enable threshold.

5.1.1 Failures attributed to the PCB design and manufacturing

Contractor A batch 1:
The functional testing did not indicate any failures caused by the PCB design or manufacturing.

Contractor B batch 1:
During initial functional testing, one hybrid failed the test of two receiver e-links (the first stub data link and the full module link) routed from the right side FEH connector. During the test, an additional load current was drawn from the DC-DC converters. The test without additional load passed.

Before hybrids were distributed to module assembly centers, they were tested again with the test card. In this test, another hybrid failed the test of the same stub data link, also with additional load. It also failed the test of the left side I^2C controller. Similar failures were later reported for additional hybrids from this batch from tests with the test board setup and from module assembly centers. Investigations of these failures also showed signs of a temperature dependence.

After failures could even be induced by applying mechanical stress to the surface of the hybrid flex in the vicinity of micro-vias, PCB cross-sections were made at Aachen and at CERN [100]. The cross-sections shown in Figure 5.1 revealed low copper thickness in the second inner PCB layer and cracks in the stacked vias. The observed cracks are compatible with the sporadic malfunctioning of the hybrids.

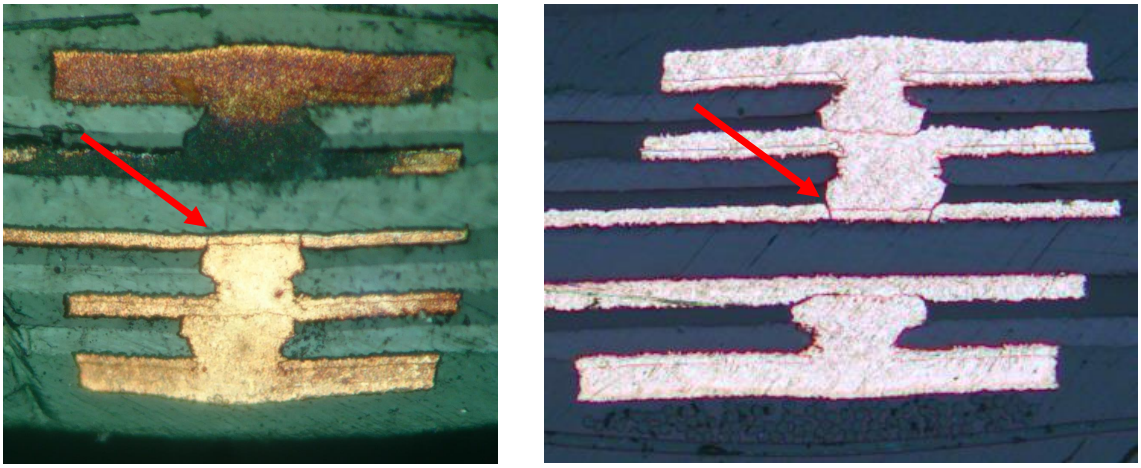

Figure 5.1: Left: cross-section of an assembled contractor B final Service Hybrid. The cross-section was prepared in Aachen. Right: cross-section of an unassembled contractor B final Service Hybrid flex. The cross-section was prepared by CERN [100]. The arrow indicates in both magnified views a stacked via with visible micro cracks and a thin copper layer.

Flexes from the same production lot as the first batch were also intended for the assembly of a second batch with the subsequent ASIC generation. Since these hybrids would likely also suffer from sporadic and unrecoverable bad connections, it was decided not to assemble the second batch with flexes of contractor B.

Contractor A batch 2:
The functional testing did not indicate any failures caused by the PCB design or manufacturing.

5.1.2 Failures attributed to the assembly and handling

Contractor A batch 1:
The functional testing did not indicate any failures caused by the assembly.

Contractor B batch 1:
Two hybrids could not be installed into the test card because connectors were unusable. On one of these hybrids glue from a plastic support for the VTRx$^+$ module had crept into the VTRx$^+$ module connector so that the module could not be plugged. Figure 5.2 (left) shows a picture of the affected connector. The glue reaches up to the pads and discoloration inside the socket is visible. On the second hybrid, a not-identified substance, most likely underfill material, stuck to the left FEH connector. The substances could not be removed, rendering the hybrids unusable. The middle picture in Figure 5.2 shows the residue. Already during the visual inspection, a shortened differential pair (fast com-

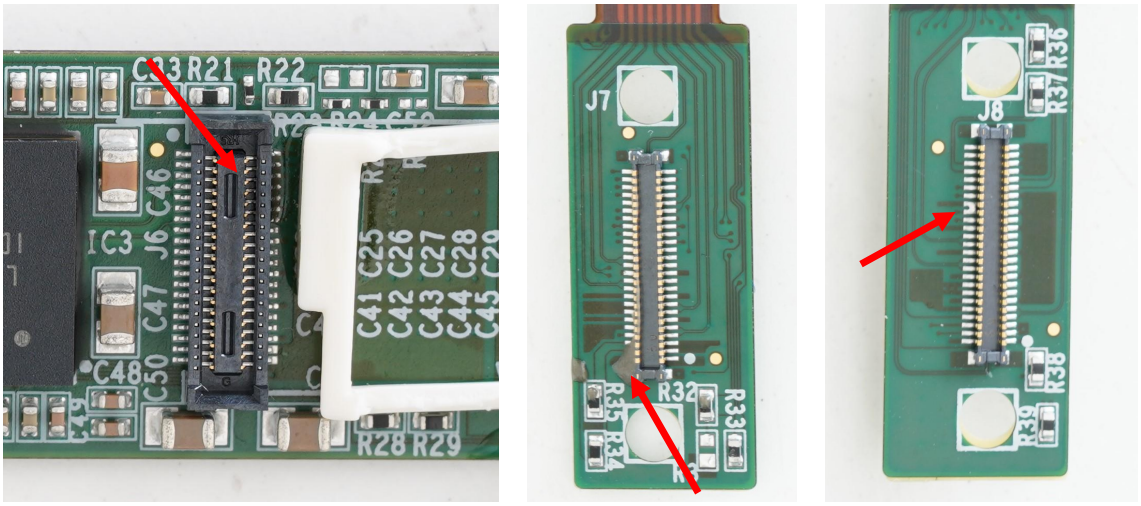

Figure 5.2: Pictures of assembly failures on final Service Hybrid prototypes. The red arrows indicate the locations of the failures. Left: glue contaminates the upper right part of the VTRx$^+$ module connector. Middle: Unknown substance contaminates a FEH connector. Right: a differential pair of a FEH connector is shortened.

mand downlink) was spotted on the right side FEH connector of another hybrid. The functional testing correctly identified this assembly error. Figure 5.2 (right) shows the connector with the short.

One hybrid could not be enabled and provided no output voltage. Inspection of the bPOL12V revealed that several connections were not soldered correctly, one of which was the enable pin. The open connection was also spotted in X-ray images that were part of the quality control at the contractor. A microscope picture of the unconnected pin and the X-ray image with the open connection on the left of the top row is shown in Figure 5.3.

Three more hybrids presented a short of a DC-DC converter output to ground. In two cases it could be traced to the second stage output rail, and in one case to the bPOL12V output rail. Consultation of X-ray images provided by the manufacturer revealed that the former shorts were caused by shortened bump bonds under the lpGBT. Figure 5.4 (left) shows the X-ray of one lpGBT. For the hybrid with a short on the bPOL12V output rail, excessively large bump bonds below the lpGBT are present, but no short. The bPOL2V5 on the same flex is also a bump-bonded chip. The bond quality for that chip may be also poor on this hybrid. However, no X-rays of the bPOL2V5 are available to provide final proof.

Lastly, one hybrid was delivered with a broken flex and stiffener. The missing piece of the PCB contains the 2.5 V line to the VTRx$^+$ module, which is

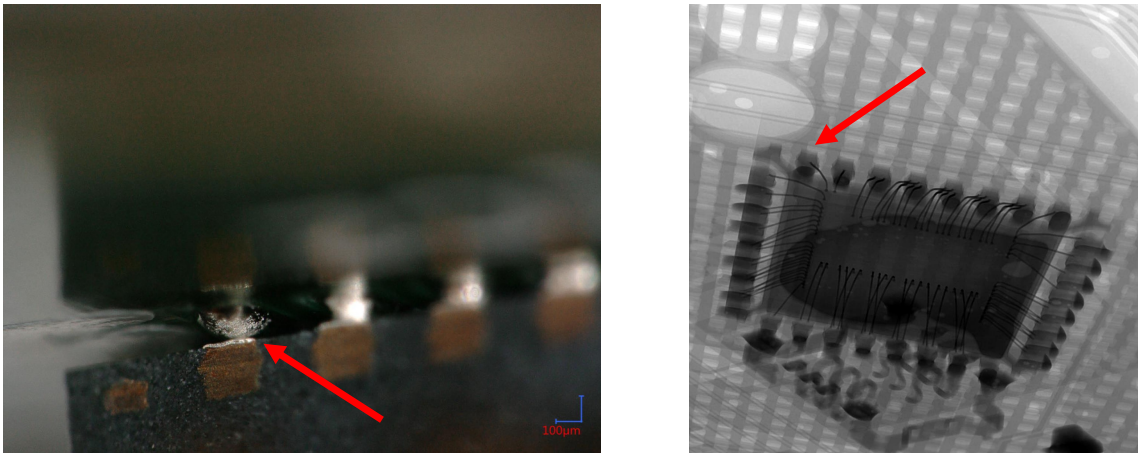

Figure 5.3: Pictures of an assembly failure of the bPOL12V on a final Service Hybrid prototype. Left: the microscope picture shows the unconnected pin. Right: the same open pin is visible in the upper left of the X-ray image. The red arrows indicate the location of the pin.

therefore disconnected. As a result, no optical communication with the hybrid was possible and no functional test was performed. Figure 5.4 (right) shows a picture of the hybrid with the damage close to a tooling hole.

Contractor A batch 2:
Upon first tests of the first samples of this batch, it was discovered that they could not be enabled. The investigations revealed that the bPOL12V's chip internal resistor inside the turn-on voltage divider is different for the bPOL12V_v4 and bPOL12V_v6. This change was not considered in the assembly. This was an error in the information provided by CERN and not by the contractor. To correct it, the shield above the DC-DC converters was removed and the voltage divider was corrected. Afterwards, the shield was assembled again.

5.1.3 Failures attributed to the active components

Contractor A batch 1:
Eleven out of 27 hybrids in this batch failed the test of the I^2C controller 0. Failures were already observed within the first delivered lot and extensively studied. As the general failure mode, the lpGBT reported a time-out during an I^2C transaction and subsequently an erroneous low state of the I^2C data line. Direct measurements with an oscilloscope confirmed the failure. The lpGBT correctly initiates a transaction by addressing the target and indicating a write statement. However, after the eighth clock cycle on the SCL line, the transaction misses the ninth clock cycle for the acknowledgment by the target.

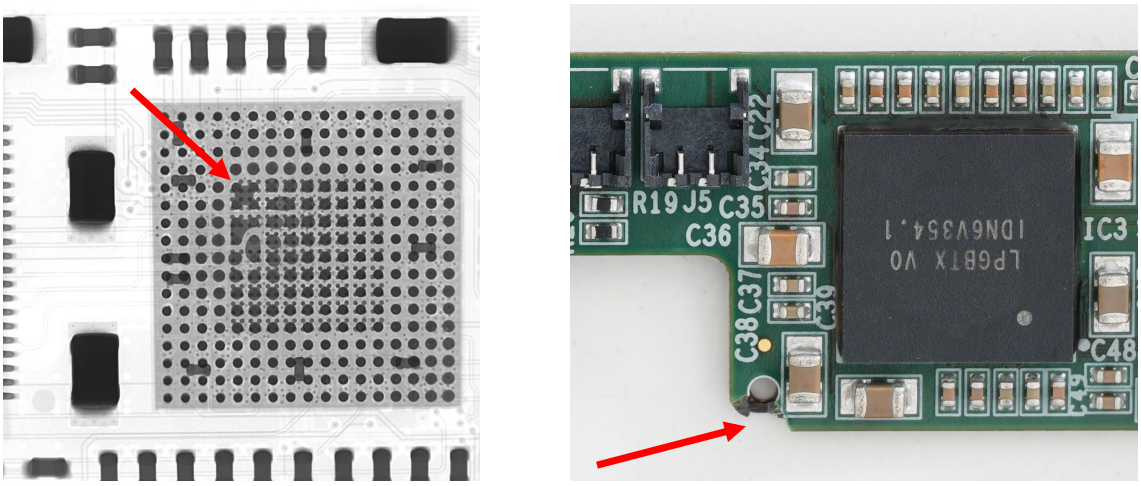

Figure 5.4: Pictures of assembly failures on Final Service Hybrid prototypes. The red arrows indicate the locations of the failures. Left: X-ray image of the lpGBT with shortened bump bonds. Right: damaged flex and stiffener.

Without this clock cycle, the target (implemented in the FC7) indefinitely pulls down the data line. An example of such behavior is shown in Figure 5.5. The recorded transaction therefore fails with a timeout and any subsequent transaction fails due to the low data line. While some controllers failed during the first transactions others only did so after several thousand write attempts. It was also found that some controllers behaved differently when used with the lpGBT's I^2C target interface instead of the IC interface of the optical link. All these failures were attributed to the known issue with the lpGBTv0 discussed in Section 2.3.3. The chips were screened for this failure mode before assembly, but not all proved functional in the application.

The hybrids with this failure were otherwise usable and were also distributed for module prototyping. A patch with small wires allowed the use of I^2C controller 1, which is otherwise used to communicate with the VTRx$^+$ module's laser driver, as a replacement.

Contractor B batch 1:

After batch 1 of contractor A showed problems with the lpGBT's I^2C controller 0 on samples that were considered functional, the lpGBT team retested the chips designated for this batch with a refined test procedure to reduce the number of failing hybrids. Despite this effort, two hybrids failed the I^2C test of controller 0 with the signature described above. This number only includes hybrids that were spotted in the initial testing. Failures could not be excluded for operation at different voltage levels or temperatures e.g. in modules.

Contractor A batch 2:

Failures were observed on four hybrids in this batch that could not be explained

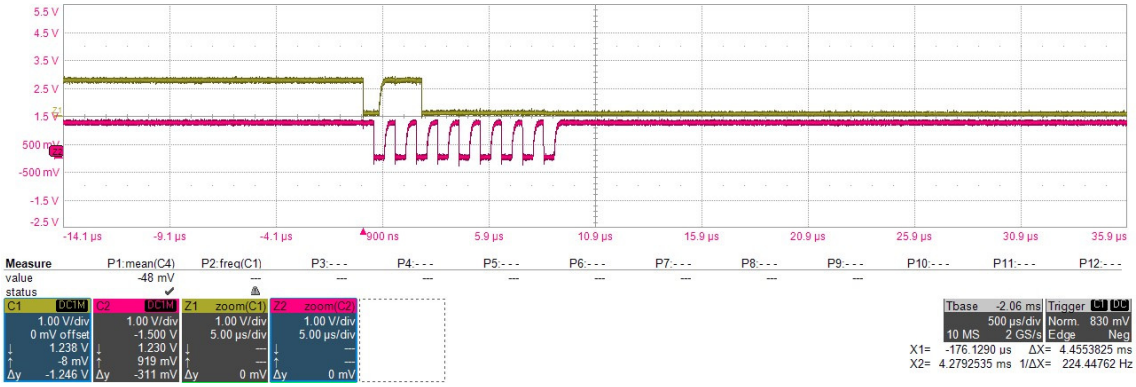

Figure 5.5: Exemplary failure of the lpGBT's I^2C controller 0 captured with an oscilloscope. Only the enlarged traces Z1 and Z2 are shown. Z1 (yellow) is the controller's data line and Z2 (red) is the clock line. The ninth clock cycle is missing and the target (address 0x70) pulls the data line low indefinitely.

by problems with the PCB or assembly. Two hybrids had no output voltage of the bPOL2V5. One hybrid had a bPOL2V5 output voltage significantly lower (1.20 V) than the target of 1.25 V. Another hybrid is missing the Power Good signal that is usually provided by the bPOL2V5. It cannot be excluded that the manipulation during the removal and the additional assembly of the shield damaged the components.

The detailed discussion shows that the hybrids provided by contractor A performed much better than those provided by contractor B. They did not show evidence of manufacturing, assembly, or handling-induced failures. Contractor A's low yield in batch 1 was solely caused by a weakness in the used lpGBT version. The low yield of contractor B is a combination of poor care during assembly and handling, and the flex quality.

5.2 Performance of the different batches

The previous section covered a qualitative look at the Service Hybrids under test. They were only divided into two groups: failing or functional. This was a result of their limited supply and high demand for the assembly of module prototypes. The focus was set on identifying electrically functional hybrids at the expense of not being able to define and apply strict quantitative performance cuts. The latter in particular regards the DC-DC converter performance. During production, a more stringent approach will be required. DC-DC converter modules that are currently used to power the CMS Phase-1 pixel detector were selected based on stringent

objective criteria [35]. Nevertheless, data from the prototype testing can be the foundation for the definition of performance cuts.

This section provides a detailed analysis of the quantitative performance data that were collected during the testing. Hybrids from different contractors and batches and the different ASIC generations will be compared. Especially the second batch of contractor A with the final chips is used to provide input for the final component selection for production.

5.2.1 DC-DC converter output voltages

As discussed in Chapter 2, the output voltages of the DC-DC converters bPOL12V and bPOL2V5 are set via resistor networks. The resulting absolute value also depends on variations and uncertainties of the DC-DC converters' voltage reference and involved passive components. The output voltage plays an important role in the functionality of the ASICs it powers. ASIC designers usually guarantee chip functionality for $\pm 10\%$ of the nominal voltage.

The output voltages of the bPOL12V and bPOL2V5 are extracted from the data of the test card's conversion efficiency measurement routine. The routine includes the setting of zero load current. Per definition, the efficiency is also zero at that point. However, the voltage readings give the output voltage without voltage drops influencing the measurement.

The output voltage of the bPOL12V powers the bPOL2V5 and parts of the VTRx$^+$ module. In the test system, it can be monitored from a single pin in the FEH connector. On the module, it is not output to the FEH, since the counterpart pin is unconnected. The left plot of Figure 5.6 shows the comparison of the output voltage for batch 1 of contractor B, which features the bPOL12V_v4 version of the DC-DC converter, and the second batch of contractor A, with the bPOL12V_v6. A shift to voltages higher than the foreseen 2.5 V and a larger spread is observed for the prototype chips. For the production version, the mean is at the nominal value and the spread is consistent with the information given in the data sheet [16].

The right plot in Figure 5.6 shows the voltage regulation with increasing load applied to the DC-DC converter. The load is indirectly applied through the bPOL2V5. The leftmost data points are measured without output current and correspond to the previously discussed histograms. The error bars represent the distributions' standard deviation. We observe a small decrease in the output voltage, which is similar for both examined batches. Since this voltage is effectively measured in parallel on the test card, the reduction is not caused by voltage drops, except those that are occurring on the ground potential. The observed voltage drop is small compared to the spread within a set of hybrids.

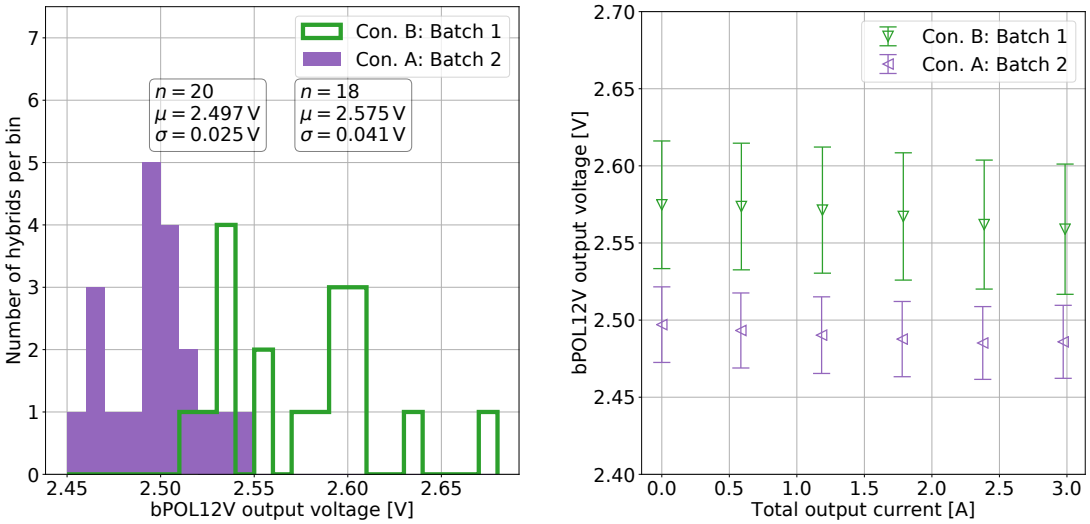

Figure 5.6: Left: output voltage of the bPOL12V on two subsets of final prototype hybrids without additional hybrid load. Right: load regulation of the bPOL12V output voltage on two subsets of final prototype hybrids. The load current is drawn evenly on both sides from the bPOL2V5 in the second DC-DC converter stage. The error bars represent the distributions' standard deviation.

The output voltage of the bPOL2V5 is also measured for both sides. The best estimation for the true output voltage is obtained from the right side voltage when no additional load is applied. The left side output also powers the lpGBT and VTRx$^+$ module and is affected by voltage drops. Figure 5.7 shows on the left the bPOL2V5 output voltage results for the first 13 final prototype hybrids. They were delivered as part of contractor A's first batch and feature the prototype ASIC versions. The observed mean value is lower and the spread higher than originally anticipated. However, the results are compatible with the details on the voltage reference given in the bPOL2V5's manual [17]. Voltage drops on the Service Hybrid and FEHs would further decrease the voltage level at the input of the front-end chips, especially for CBCs that have the longest path from the FEH connector. It was therefore decided to increase the voltage setting for the hybrids with prototype bPOL2V5 ASICs to guarantee at least 1.25 V of output voltage. This change required the removal of the DC-DC converter shield and the replacement of a resistor. For future hybrids with the same ASIC version, the bill of materials was updated and the correct value was assembled directly. Therefore no replacement of the shield was required. The right plot of Figure 5.7 shows the output voltage of the first delivery's hybrids after the modification. The spread is maintained, as it depends on the unchanged chip internal reference, but the mean is shifted to a higher value. All hybrids provide an output above 1.25 V. However, the output voltage reaches up to 1.35 V in unloaded cases, which is at the upper limit of operation conditions for the lpGBT and VTRx$^+$

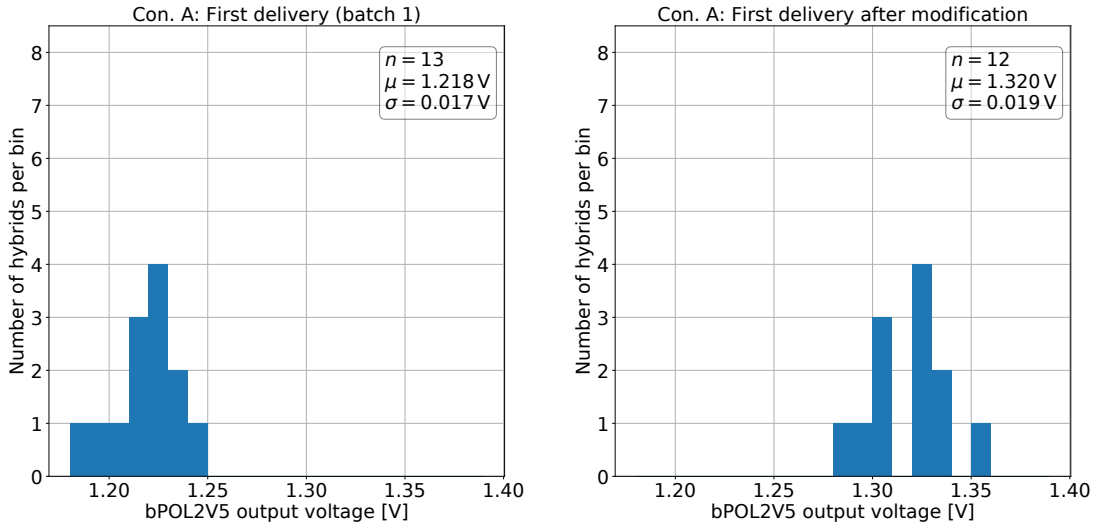

Figure 5.7: Output voltage of the bPOL2V5 on the first delivered subset of hybrids of contractor A's first batch. Left: original configuration. Right: after the modification of the output voltage setting. No additional load is drawn from the DC-DC converters.

module on the Service Hybrid. One hybrid was delivered to a module assembly center before the resistors were modified and is missing in the corresponding distribution.

The data for the second part of batch 1 for contractor A and the entire batch 1 for contractor B are shown in Figure 5.8. They show a behavior similar to the reworked first delivery discussed above. Within the limited statistics, the spread and mean are compatible.

Those three subsets make up all final prototype hybrids with the bPOL2V5_v3.2 chip. Figure 5.9 shows on the left the combined data for those hybrids and on the right the data for the second batch of contractor A. The latter are the only final prototypes with the version bPOL2V5_v3.3 ASIC. The combined statistics of the bPOL2V5_v3.2 hybrids give a broad distribution around 1.312 V. No hybrid has an output voltage lower than the anticipated 1.25 V. The hybrids with a bPOL2V5_v3.3 feature a more narrow distribution, which is centered at a value lower than 1.25 V. For the production, the divider is slightly adjusted and the rating of the resistors is improved to 0.1%. By changing R_2 from 316 kΩ to 309 kΩ the voltage is increased by the required 0.02 V.

When looking at the load regulation of the bPOL2V5's output voltage, the measurement is influenced by voltage drops on the hybrid. In addition, the first version of the test card measured the voltage behind a protective fuse. This card version was used to measure all batch 1 hybrids and its measurements therefore include an additional voltage drop. The second test card version senses the voltage at the level

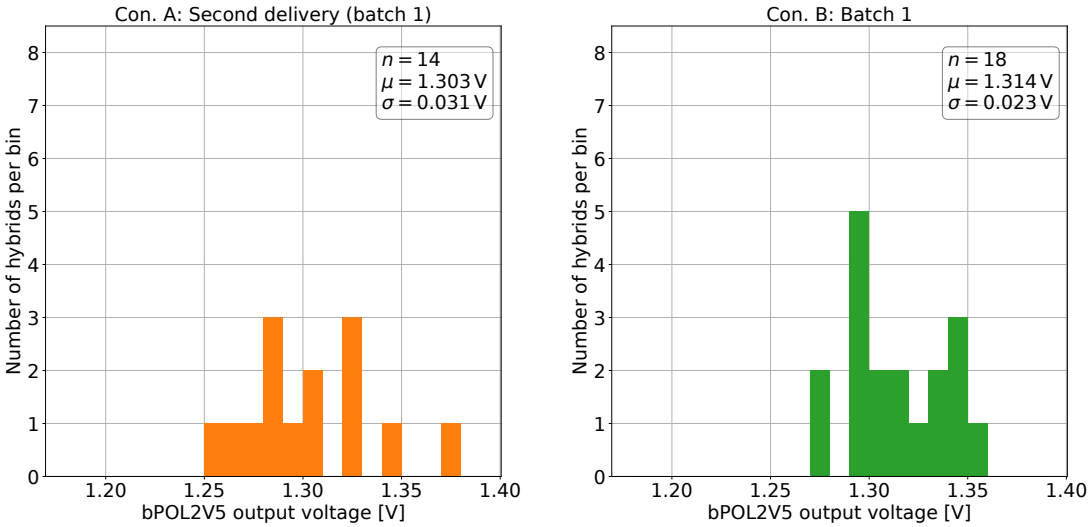

Figure 5.8: Output voltage of the bPOL2V5 on batch 1 hybrids. Left: the second delivery of contractor A. Right: the full first batch of contractor B. No additional load is drawn from the DC-DC converters.

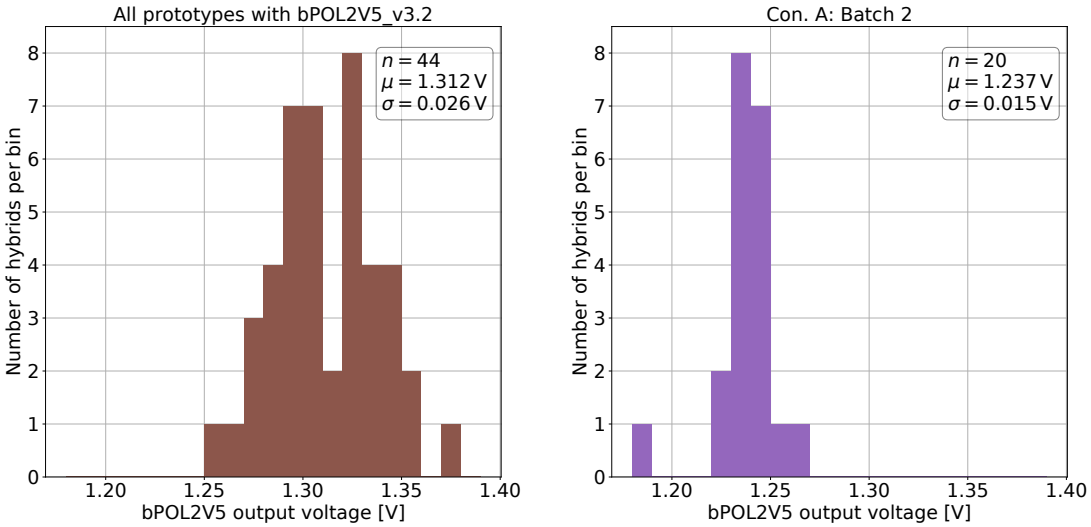

Figure 5.9: Left: output voltage of the bPOL2V5_v3.2 on the entire batch 1 hybrids. Right: output voltage of the bPOL2V5_v3.3 on the hybrids of contractor A batch 2. No additional load is drawn from the DC-DC converters.

Figure 5.10: Output voltage of the bPOL2V5 measured on the left Service Hybrid side versus output current drawn solely on the same side. Different markers represent different subsets of final prototype hybrids. The first delivery of contractor A is shown before and after the output voltage adjustment. The standard deviation within a subset is given by the error bars.

of the FEH connector, which provides a better estimation. For each load configuration of the efficiency measurement, the mean value and standard deviation of the distributions of the different batches are plotted versus the applied load current. The load is only applied to one hybrid side at a time. Figure 5.10 shows the left side output voltage for a load on the left side, Figure 5.11 the right side output voltage for a load on the right side. Data containing batch 1 hybrids behave very similarly with a voltage drop proportional to the output current.

If we assume purely Ohmic losses, we can extract the resistance of the 1.25 V power rail from the slope of the linear voltage decrease. The slope is determined per hybrid measurement and side and the mean per batch is given in Figure 5.12. The error bars represent the standard deviation. A large deviation between the different contractors would indicate differences in the copper layer quality between the different flexes. None are observed. The lower calculated resistance of the second batch is consistent with the 45 mΩ fuse on the test card that does not contribute anymore. The resistance on the left side is higher due to the longer distance between the output of the DC-DC converter and the FEH connector. This leads to additional Ohmic losses in the copper layers.

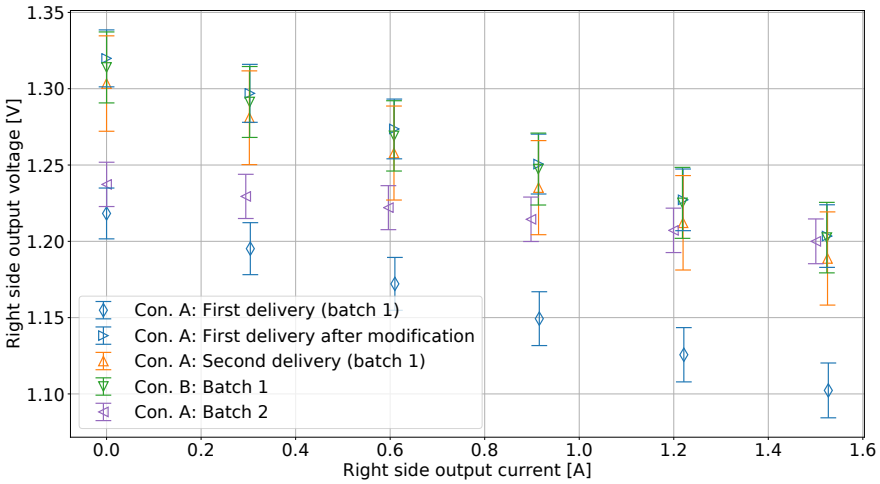

Figure 5.11: Output voltage of the bPOL2V5 measured on the right Service Hybrid side versus output current drawn solely on the same side. Different markers represent different subsets of final prototype hybrids. The first delivery of contractor A is shown before and after the output voltage adjustment. The standard deviation within a subset is given by the error bars.

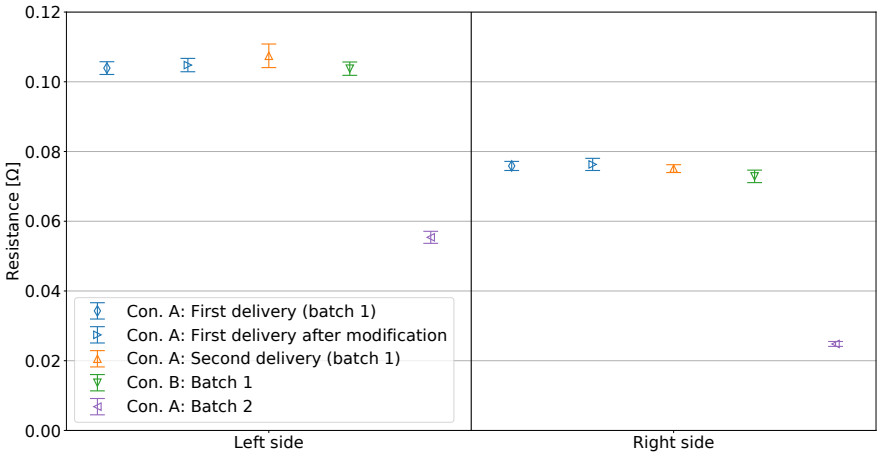

Figure 5.12: Resistance of the 1.25 V rail on the final Service Hybrid prototypes. Different markers represent different subsets of final prototype hybrids. The standard deviation within a subset is given by the error bars. All batch 1 measurements are carried out on version 1 test cards and contain an additional voltage drop of a fuse on the test card. The left and right sets of data points correspond to the left and right side output voltage of the bPOL2V5, respectively.

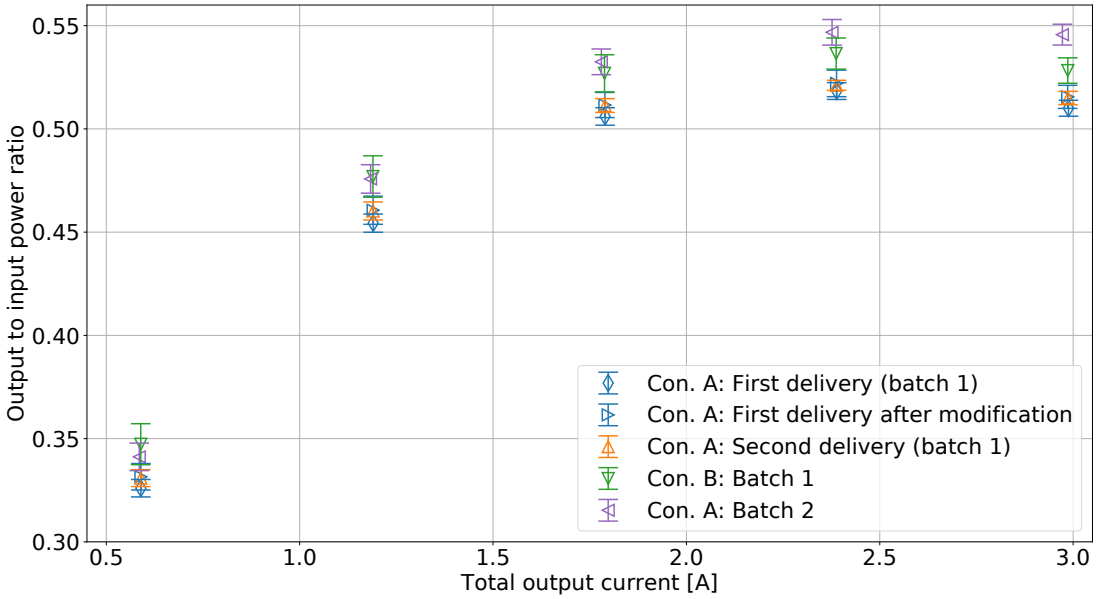

Figure 5.13: Ratio of the output to input power versus load current on the final Service Hybrid prototypes. Different markers represent different subsets of final prototype hybrids. The standard deviation within a subset is given by the error bars. All batch 1 measurements are carried out on version 1 test cards and contain an additional voltage drop of a fuse on the test card.

5.2.2 DC-DC conversion efficiency

So far, we only considered the output of the two DC-DC converter stages. To fully quantify the DC-DC converter performance, the input also needs to be taken into account. By measuring input and output voltages and currents, the ratio of the output to input power η can be obtained. For a single DC-DC converter this quantity is identical to the DC-DC conversion efficiency. Here, the interpretation is more convoluted due to the two-stepped approach and the lpGBT and VTRx$^+$ module present on the Service Hybrid.

Figure 5.13 shows the ratio as a function of output current for the different hybrid batches. The improvement of the ratio with current is typical, as losses inside the DC-DC converter and the optical readout components become less relevant. For large currents, resistive losses and switching losses inside the DC-DC converter limit the efficiency. Contractor A batch 1 shows the lowest η and the change in output voltage does not influence the outcome. The best output to input power ratio is measured for the batch 2 hybrids, especially for large currents. However, those measurements also profit from the reduced voltage drops discussed in the previous section.

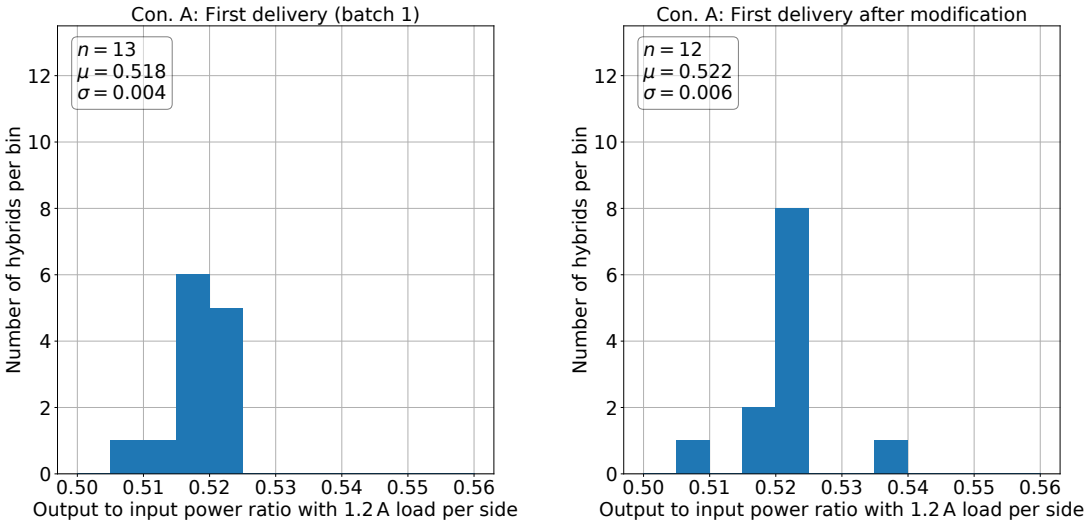

Figure 5.14: Ratio of the output to input power at 1.2 A load current per side on the first delivered subset of hybrids of contractor A's first batch. Left: original configuration. Right: after the modification of the output voltage setting.

Figures 5.14, 5.15, and 5.16 contain the distributions of the ratio for different batches at 1.2 A load per side. This current is a working point roughly equal to the estimated FEH intake. No significant difference between the batches is expected. All distributions have a narrow central population with either high or low outliers. In Figure 5.14 no change due to the adjusted output voltage is observed. The batch 2 hybrids show an improvement in the ratio of 2 percentage points (%pt.) compared to the batch 1 hybrids. A possible explanation for the outliers and the improvement was identified after the measurements: the use of different versions of the VTRx$^+$ module. Depending on the version they feature a different number of active laser drivers and a different power consumption. The VTRx$^+$ module developers state that the power intake of a module with four active laser drivers is about 130 mW higher than for a module with one active driver [64]. This additional power consumption changes the ratio η at 1.2 A load per side by 2-3%pt. The different versions were not consistently used even within a batch and the version information for all measurements except for batch 2 is lost. The central population corresponds to the version mostly used for the testing of a particular batch and the outliers to the alternate version. For the batch 2 measurements, it was confirmed that the low tail was measured for hybrids with more active laser drivers on the VTRx$^+$ module and thus higher power consumption.

To evaluate if the performance of the bPOL12V and the bPOL2V5 in their integrated application in the Service Hybrid is consistent with the single-chip performance, an additional cross-check is performed. The single-chip DC-DC conversion efficiencies

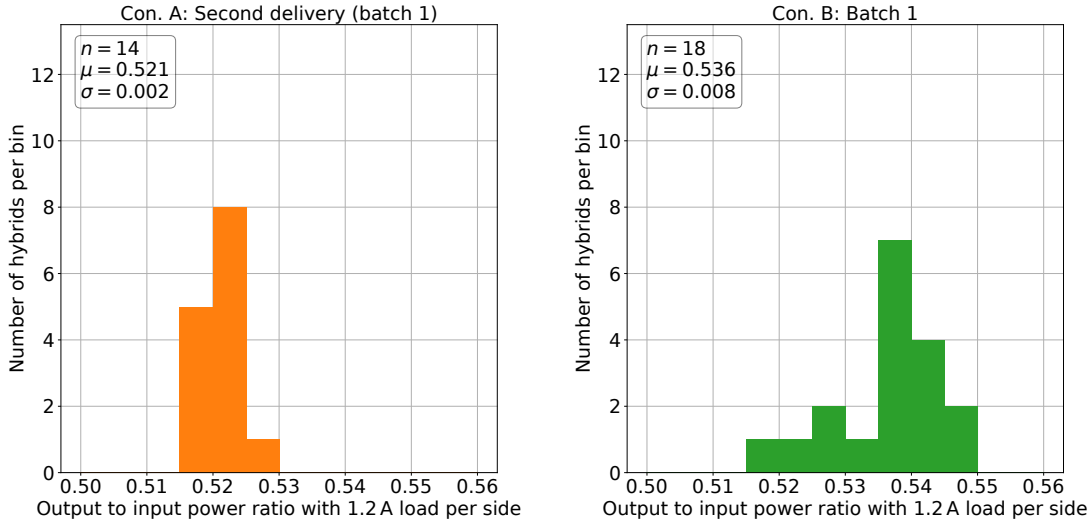

Figure 5.15: Ratio of the output to input power at 1.2 A load current per side on batch 1 hybrids. Left: the second delivery of contractor A. Right: the full first batch of contractor B.

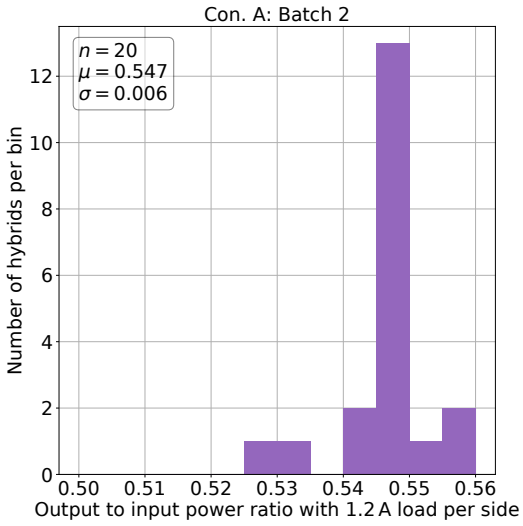

Figure 5.16: Ratio of the output to input power at 1.2 A load current per side on contractor A batch 2 hybrids.

as a function of load current are given in the data sheets [16, 17]. They are obtained using optimized boards with little voltage drops. To allow a fair comparison, the SEH's input voltage and DC-DC converter output voltages are measured directly at the filter capacitors and in the following used as inputs to a current-based model. Additional inputs to the model are the Service Hybrid load current, the singe-chip efficiencies, and the current consumption of the lpGBT and VTRx$^+$ module. The latter was measured for the configuration of a configured lpGBT and VTRx$^+$ module with all four laser drivers active to be 90 mA on the 2.5 V power rail and 193 mA on the 1.25 V power rail [96].

From the datasheet, we read off the single-chip efficiencies of 77% for the bPOL12V and 88% for the bPOL2V5 and vary in an iterative process the only free parameter, the output current, until the model's predicted output current of the bPOL12V reaches 1 A. At this point, the assumed efficiencies are valid (Scenario A). The model predicts the input current of the Service Hybrid. Table 5.2 summarizes the content of the model. The columns $U(x)$, $I(x)$, and $P(x)$ hold the voltage, current, and power at a specific input or output on the Service Hybrid. η_{DC} holds the single-chip efficiency relevant at that stage and δP_{DC} the power lost inside the DC-DC converter. I_{FE} and δP_{FE} stand for the current and power consumed by the lpGBT and VTRx$^+$ module. A Service Hybrid output current of 1.4 A is required for Scenario A and the model predicts an SEH input current 329 mA. For the measurement the 1.4 A load current is applied to a batch 2 hybrid. The measured input current is 337 mA.

Scenario B is based on the expected FEH current consumption of about 2.05 A [12]. The single-chip efficiencies do not change considering the scale of the read-off uncertainty. The prediction for the input current is 452 mA and the measurement is 459 mA.

It can be concluded that the DC-DC converters on the Service Hybrid behave quite similarly to their stand-alone performance. The observed differences between the model and the measurement on the hybrid in this cross-check are small. It is plausible that uncertainties on the model inputs like the single-chip efficiency measurements and I_{FE} are large enough to cover the deviation.

As the converters behave as expected and the observed voltage drops are reasonable, the obtained input to output power ratio agrees with the expectation.

5.2.3 bPOL12V turn-on and turn-off voltage

The procedure for the measurement of the I/V behavior was implemented after the delivery and testing of the batch 1 hybrids. As discussed in Section 2.3.3, the enable circuit of bPOL12V_v6 used in batch 2 has different properties compared to the batch 1 bPOL12V_v4. Therefore, only batch 2 provides relevant results for the production.

Scenario A	$U(x)$ [V]	$I(x)$ [mA]	$P(x)$ [W]	η_{DC} [%]	δP_{DC} [W]	I_{FE} [mA]	δP_{FE} [W]
Service Hybrid input	9.80	329	3.23	-	-	-	-
bPOL12V input	9.80	329	3.23	77	0.74	-	-
bPOL12V output	2.47	1006	2.48	-	-	90	0.22
bPOL2V5 input	2.47	916	2.26	88	0.27	-	-
bPOL2V5 output	1.25	1593	1.99	-	-	193	0.24
Service Hybrid output	1.25	1400	1.75	-	-	-	-
Combined DC-DC converters	-	-	-	69	1.01	-	-
Scenario B	$U(x)$ [V]	$I(x)$ [mA]	$P(x)$ [W]	η_{DC} [%]	δP_{DC} [W]	I_{FE} [mA]	δP_{FE} [W]
Service Hybrid input	9.80	452	4.43	-	-	-	-
bPOL12V input	9.80	452	4.43	77	1.02	-	-
bPOL12V output	2.47	1380	3.41	-	-	90	0.22
bPOL2V5 input	2.47	1290	3.19	88	0.38	-	-
bPOL2V5 output	1.25	2243	2.80	-	-	193	0.24
Service Hybrid output	1.25	2050	2.56	-	-	-	-
Combined DC-DC converters	-	-	-	68	1.40	-	-

Table 5.2: Current based model to cross-check the Service Hybrid output to input power ratios with the single-chip DC-DC converter power conversion efficiencies. Voltage drops are not considered. Fixed inputs to the model are marked in yellow, the variable input of the SEH load current is marked in pink, and the SEH input current predicted by the model is marked in green. Values in white cells are the results of calculations.

Figure 5.17 shows the I/V response measurement of a batch 2 final Service Hybrid prototype. The result was obtained using a test card in table-top mode. The reason is that a test card cannot be enabled for test voltages lower than 5 V in the crate. In the final test system the part of the curve below 5 V will be missing.

When the voltage is first ramped up from 0.0 V, a small bump in the measured current is present between 1.5 and 3.0 V. It is likely caused by parasitic resistive paths inside the bPOL12V that are present when the chip logic is not powered sufficiently. At about 3.0 V parts of the controlling logic become active and the current disappears. At about 5.0 V the voltage regulators which provide 3.3 V for the drivers of the power transistors are enabled [16]. This is an important part of the internal chip logic and causes a 10 mA current step. Until the enable voltage, which is deduced from the input voltage with a voltage divider, reaches the enable threshold,

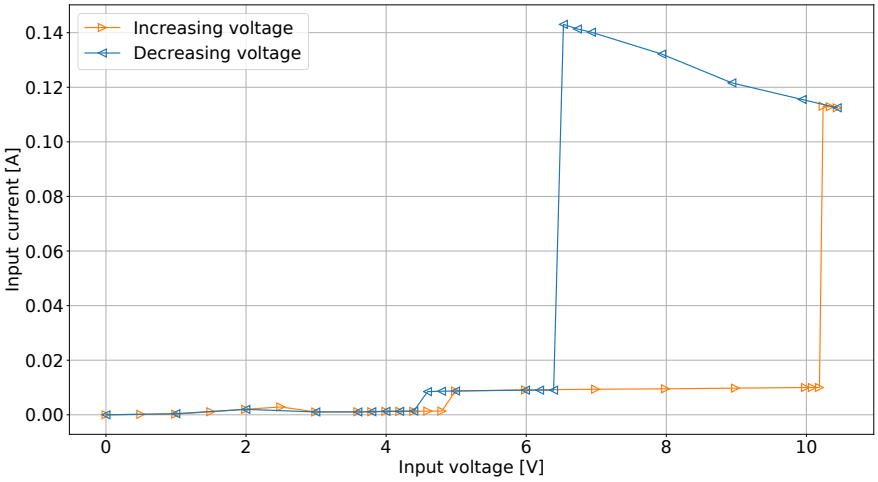

Figure 5.17: I/V response of a final Service Hybrid prototype. The orientation of the markers indicates the direction of the voltage change.

the current remains constant. Above the enable threshold the DC-DC converter becomes operational and provides current to the VTRx$^+$ module and the bPOL2V5 and the subsequent lpGBT. The current increases up to the SEH idle current of about 117 mA. The maximal tested voltage is 10.5 V. From there, the voltage is reduced. As the power demand from the bPOL12V remains constant, the current increases until the voltage at the enable node drops below the disable threshold. Here, the current decreases to the level where the 3.3 V regulators are active. The logic enabling the regulators has a hysteresis and disables them at roughly 4.5 V.

We define the turn-on and turn-off voltage as the voltage after the current step to and from the idle current. The finite voltage step size (0.1 V at the turn-on and 0.2 V at the turn-off point) limits the precision of the measurement. It also introduces a systematic overestimation (underestimation) of the turn-on (turn-off) voltage. The current-induced voltage drop on the supply cable is included in the measured voltage and leads to a systematic underestimation in both cases. A finer voltage step size would increase the measurement duration and the algorithm for the turn-on detection requires a larger step size than the current-induced voltage drop. Figure 5.18 shows the binned results for the batch 2 hybrids. The measured turn-on voltages are binned and plotted so that the actual turn-on voltage can be inferred from the distribution. This is not possible for the turn-off voltage due to the systematic underestimation of the measurement. For the given voltage divider and an assumed threshold of 800 mV, a turn-on voltage of 10.0 V is expected. The obtained result is closer to 10.2 V and suggests a threshold of 820 ± 10 mV. The observed spread is about twice the size of the measurement uncertainty and can be caused by variations of the enable threshold or the internal pull-down resistance.

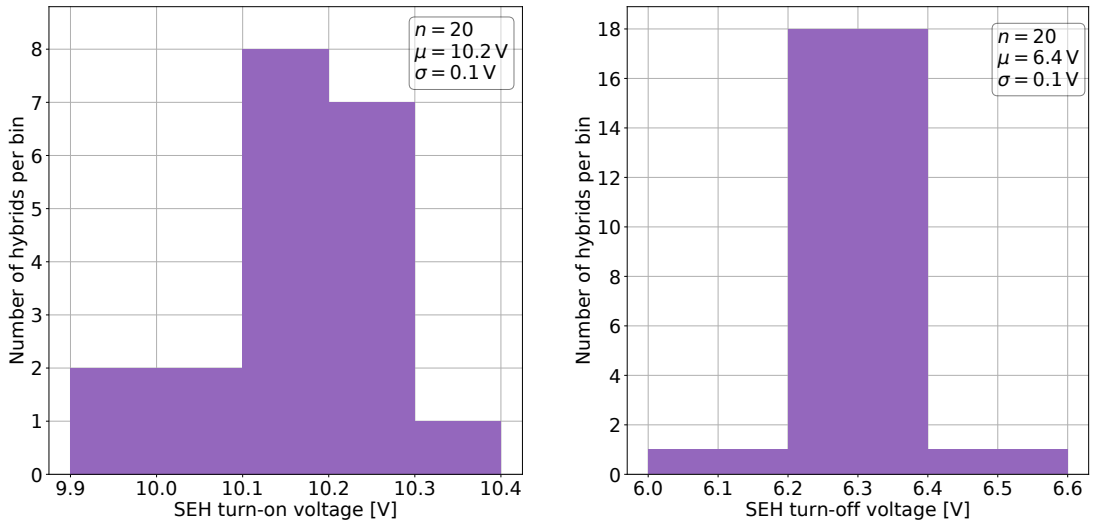

Figure 5.18: Left: turn-on voltage of contractor A batch 2 final Service Hybrid prototypes. Right: turn-off voltage of those hybrids.

The measurement of the turn-off voltage has a lower precision and the obtained turn-off threshold is $510 \pm 10\,\mathrm{mV}$.

5.2.4 Leakage current of the sensor bias voltage filter

The bias voltage filter design is the same in both batches. The capacitors of the RC filter use the X7R dielectric [101]. The leakage current of the sensor bias voltage filter is measured for all hybrids at $-1,000\,\mathrm{V}$. For the first batch, the current is measured for ten minutes. The time is reduced to 150 s for batch 2. During the test, no additional load current is drawn from the DC-DC converters.

Figure 5.19 shows five measurements of the leakage current versus time for the same hybrid (contractor A batch 2). In all cases, the negative voltage was applied for 150 s. A first measurement at room temperature shows an exponential decrease of the (absolute) current after it reaches the measurement range at 30 s. In a second measurement five days later, a similar exponential behavior is present, but the current is lower at all times. Later room temperature measurements that were performed after the hybrid was subjected to the bias voltage for an extended amount of time in the test board setup are similar. The temperature dependence is discussed in the next section, but for comparison, Figure 5.19 also contains a measurement at $-40°\mathrm{C}$, where the current is lower.

The data show that the history of voltage exposure of a hybrid has a major impact on the measured leakage current. The various prototypes were subject to various debugging and testing activities in different setups and with different amounts of

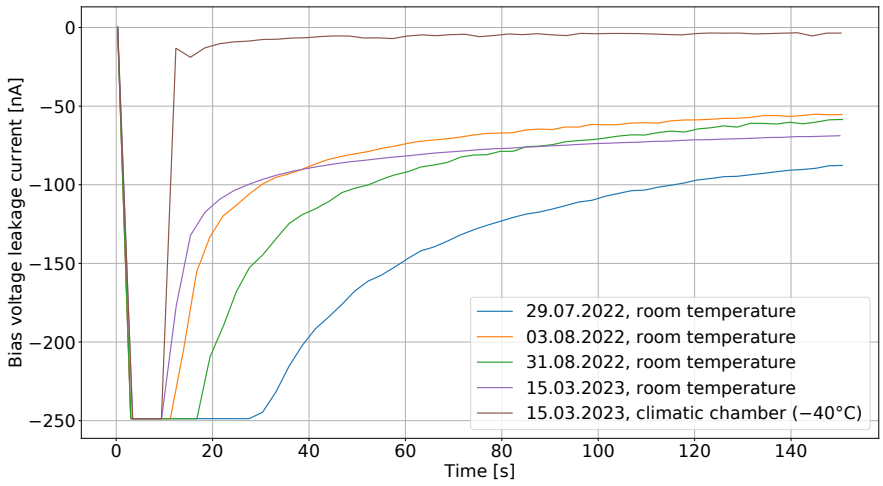

Figure 5.19: Measurements of the leakage current of the sensor bias voltage filter versus time performed for the same hybrid. Refer to the text for details on the measurement conditions. The applied negative voltage is 1,000 V. The current measurement circuit on the test card is designed to measure ± 250 nA.

bias voltage application. The largest and most comparable currents can only be obtained for the first measurement. Figure 5.20 shows the data of the first bias voltage exposure for all hybrids of the first batch of contractor B (left) and the second batch of contractor A (right). For the second batch, the leakage current is measured on the test card and by the power supply simultaneously, which yields consistent results and extends the range to higher currents. However, the measurement time is shorter. All hybrids present a similar exponential behavior, which is also illustrated by Figure 5.21 that shows both data sets on the same time scale.

The output voltage of the power supply remained at $-1,000$ V during all tests. No increase in current that would indicate an isolation failure was observed.

After the voltage is applied, the resulting time dependence of the current is dominated by two factors. The first factor is the current to charge the filter capacitors during the ramp-up of the power supply. The time constant of an exponential charging of the capacitors through the resistors of the bias voltage RC filter is too fast to be resolved by the setup ($\tau = RC = 8.2\,\mathrm{k\Omega} \times 54\,\mathrm{nF} = 443\,\mu\mathrm{s}$). The power supply is set to a ramp-up speed of 200 V/s. Assuming a linear behavior, the capacitance is charged with 10.8 µA for 5 seconds. The current measurement circuit is designed to measure ± 250 nA and will be saturated during the charging. Second, leakage currents due to parasitic resistances in the PCB and especially the capacitors will be present. These currents approach the true DC leakage currents only after hours of applied voltage. Currents caused by polarization and temperature fluctuations dominate before that [102].

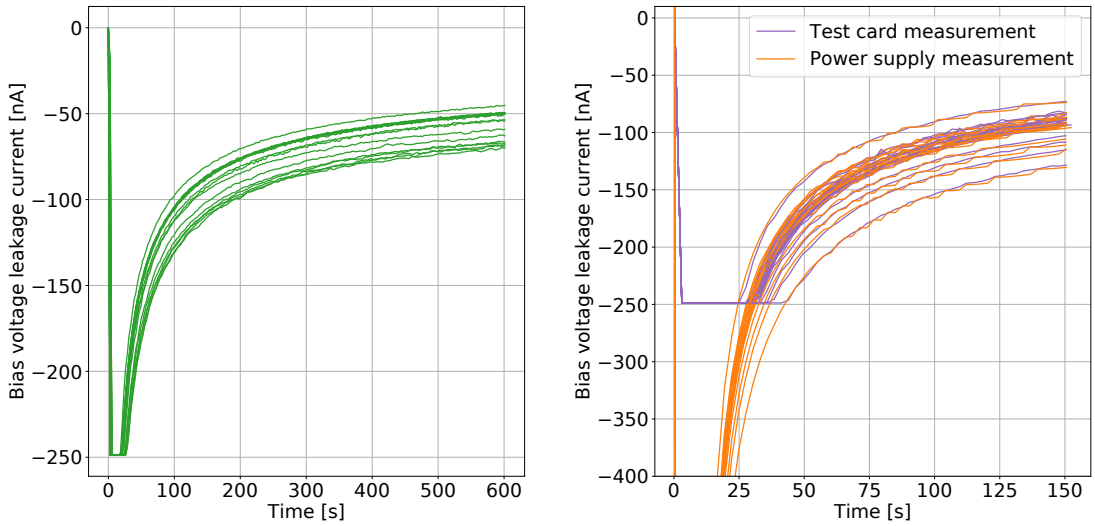

Figure 5.20: Initial measurement of the leakage current of the sensor bias voltage filter on final Service Hybrid prototypes versus time. Left: all hybrids from contractor B batch 1. Right: all hybrids from contractor A batch 2. Here, the leakage current measurement of the power supply and test card are shown. The applied negative voltage is 1,000 V. The current measurement circuit on the test card is designed to measure ± 250 nA.

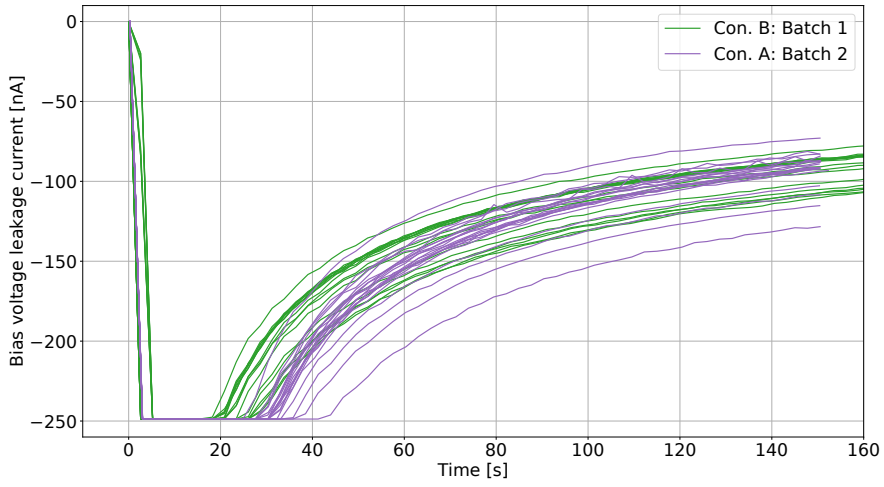

Figure 5.21: Comparison of the initial measurement of the leakage current of the sensor bias voltage filter versus time. The applied negative voltage is 1,000 V. The plot shows all final prototype hybrids from contractor A batch 1 and batch 2.

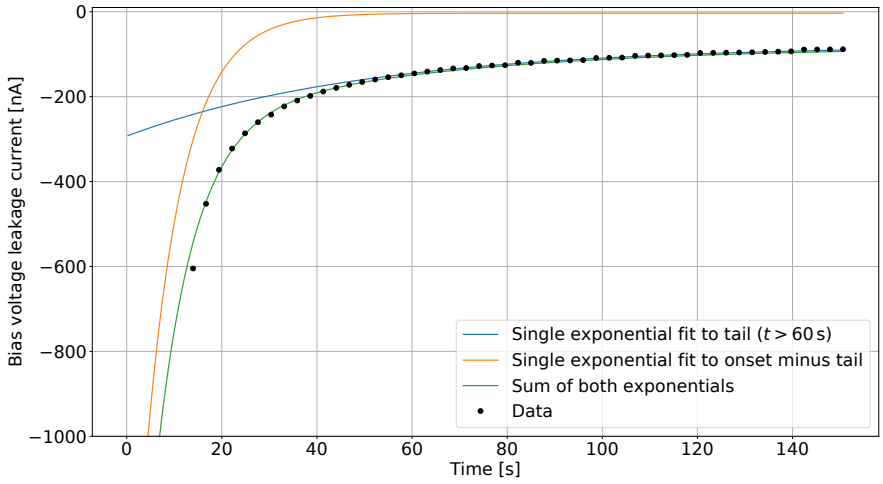

Figure 5.22: Fits to describe the time evolution of the leakage current of the sensor bias voltage filter. Black points are data. The blue curve is obtained by an exponential fit for the tail $t > 60\,\mathrm{s}$. The orange curve is the exponential that describes the onset. The green curve is the sum of both exponentials. The same data as in the first measurement in Figure 5.19 are used.

After the charging of the capacitors, the data can be described by the sum of two exponential functions:

$$I(t) = c_1 \cdot e^{-c_2 \cdot t + c_3} + c_4 \cdot e^{-c_5 \cdot t + c_6} + c_7 \quad . \tag{5.1}$$

Figure 5.22 illustrates this for a current measurement with the power supply for the first bias voltage exposure of a contractor A batch 2 hybrid. Only the data for $t > 15\,\mathrm{s}$ are considered. The tail for $t > 60\,\mathrm{s}$ can be described by a single exponential function plus offset, which is fitted first. To also model the onset, a faster decaying term needs to be added. It is obtained by fitting Equation 5.1 for $15\,\mathrm{s} < t < 60\,\mathrm{s}$ with c_4, c_5, c_6, and c_7 fixed to the values obtained in the tail. For the example in Figure 5.22 time constants $1/c_2 = 8.33\,\mathrm{s} \pm 0.21\,\mathrm{s}$ and $1/c_5 = 51.1\,\mathrm{s} \pm 3.7\,\mathrm{s}$ and an offset corresponding to the true DC leakage current of $c_7 = -79 \pm 2\,\mathrm{nA}$ are obtained. The data in Figure 5.19 suggest that the decay towards the true DC leakage current is happening faster as the hybrid is subjected more often to the bias voltage.

5.2.5 Expected temperature influence

As discussed in Section 4.3 the test cards underwent a quality assurance program that involved thermal cycles and tests at various temperatures. Since the final pro-

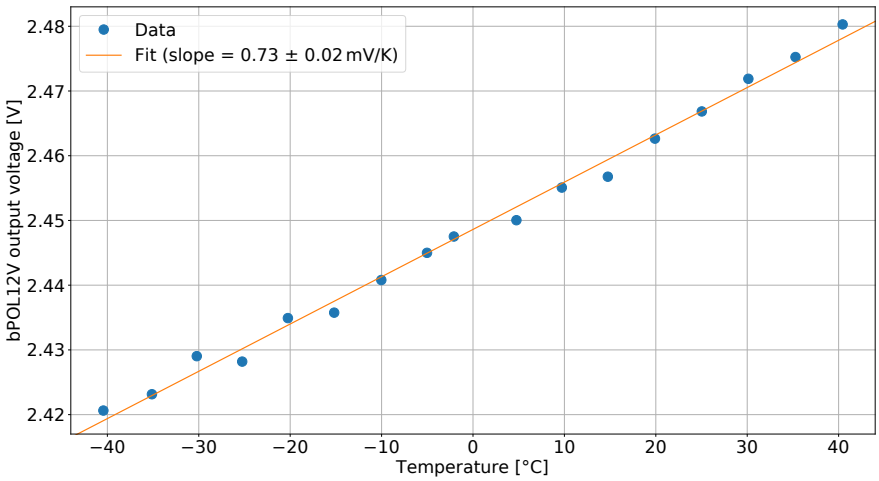

Figure 5.23: Output voltage of the bPOL12V versus temperature of the climatic chamber, measured for a batch 2 final Service Hybrid prototype.

totype hybrids were used as devices under test, the results also provide insights into the impact of temperature variations on the studied parameters.

The measurements during the thermal cycles are performed at −35°C and +40°C. They are carried out with two batch 2 hybrids and repeated on 15 different test cards each. In addition, a scan in the temperature range between −40°C and +40°C is conducted for two pairs of a hybrid and a test card. The temperature data for the presented results are measured by a sensor on the test card at the start of the test routine.

For the output voltage of the bPOL12V, a change of 51 mV between the two temperature levels is observed for one of the hybrids. For a second hybrid, a change of just 17 mV is measured. The results are the average of the measurements on 15 different test cards. The card-to-card variation of the (absolute) output voltage measurement is 7 mV. However, multiple measurements on the same card yield an uncertainty at the level of only ± 1 mV [95]. The bPOL12V data sheet does not state an expected value for the output voltage drift. More dedicated measurements are necessary to fully quantify the effect. Corresponding results are presented in Section 8.1.1.

Figure 5.23 shows the measured bPOL12V output voltage versus temperature for the first hybrid. A linear rise of the output voltage with temperature is observed. The slope is consistent with the mean change measured on multiple test cards.

A measurement of the temperature dependence of the voltage reference (300 mV) of the bPOL2V5 exists. It gives an average positive change with temperature of about 1 mV per 12 K or 0.028%/K. The results obtained from the two hybrids are 0.036%/K and 0.020%/K. Figure 5.24 presents the measurement for one of the hybrids. It clearly shows a linear dependence.

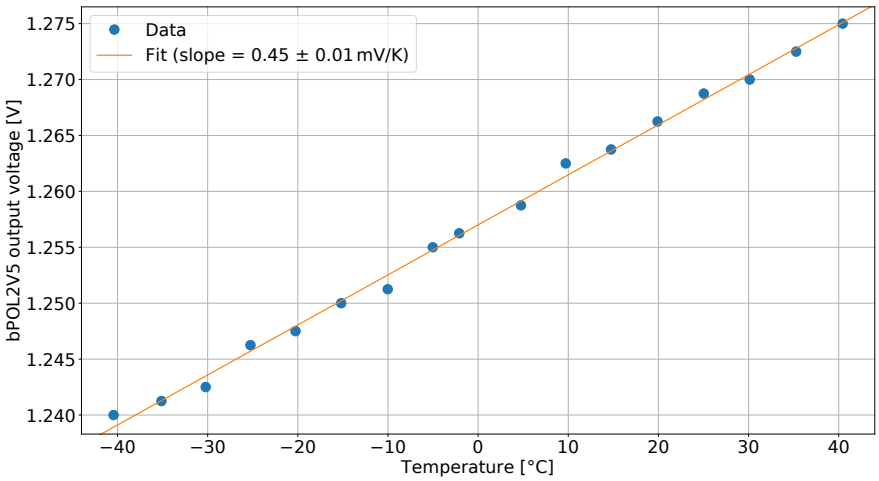

Figure 5.24: Output voltage of the bPOL2V5 versus temperature of the climatic chamber, measured for a batch 2 final Service Hybrid prototype.

The efficiency of a DC-DC converter improves at lower temperatures. Figure 5.25 shows that for a given output current of 1.25 A per side the ratio of output to input power presents a linear relationship with temperature. Felix Thurn showed in his Master's thesis that a polynomial of grade five can be used to describe the efficiency data and extrapolate to different currents. The order five polynomial $P_5(I)$ that describes the ratio η can be adapted by a current-dependent correction factor in order to describe the temperature influence [95].

He obtained

$$\eta(I,T) = P_5(I) - (T - 293\,\mathrm{K}) \cdot \left(2 \cdot 10^{-4} \frac{1}{\mathrm{AK}} \cdot I + 1.11 \cdot 10^{-4} \frac{1}{\mathrm{K}}\right) \quad . \tag{5.2}$$

The order of magnitude of the effect is consistent with the information in the bPOL2V5 datasheet, which gives an improvement of $6.7 \cdot 10^{-5} \frac{1}{\mathrm{AK}} \cdot I + 1.4 \cdot 10^{-4} \frac{1}{\mathrm{K}}$. They also find an increase in the effect with load current. No information is given for the bPOL12V. In combination with the unknown behavior of the lpGBT and VTRx$^+$ module, it is not possible to give a more accurate prediction.

The relationship in Equation 5.2 is also valid for the second hybrid ($\Delta\eta = 4.4$ percentage points).

The result of the temperature scan for the turn-on voltage is given in Figure 5.26. The obtained voltage reduces in steps for sinking temperatures. The steps result from the discrete test voltage settings used in the I/V scan. The set and observed step size is 100 mV. A 50 mV difference between the output voltage of the power supply and the measured voltage on the test card is caused by current-induced

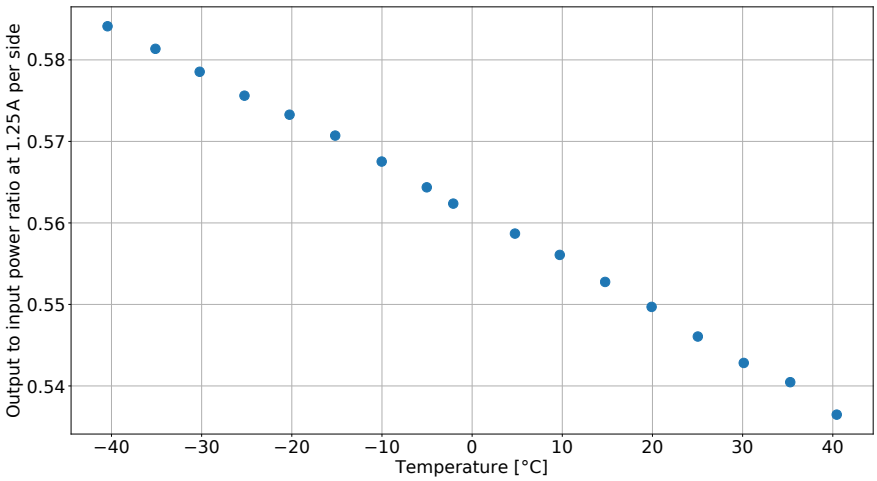

Figure 5.25: Output to input power ratio versus temperature of the climatic chamber, measured for a batch 2 final Service Hybrid prototype.

voltage drops. This offset can be considered constant and does not have an effect on the general assessment of the drift with temperature. The best estimator for the turn-on versus temperature measurement is the last point before an increasing step. For lower temperatures the turn-on voltage is lower and the set voltage at the power supply after the turn-on is too high. At the step, the set voltage is matching the required voltage. The slope of the linear function fitted to the data is therefore anchored on three data points. The observed change of the turn-on voltage and the corresponding voltage reference is 0.035%/K and similar to the drifts observed in the output voltages. The second hybrid displays a shift of 20 mV between the two extremes of the temperature levels, which is compatible.

To assess the temperature influence on the bias voltage leakage current, both thermal cycles and temperature ramps are used. In Figure 5.27 measurements at −40°C and +40°C are shown for 13 cycles. In agreement with the findings of the previous section, the measurement becomes more reproducible after the first cycles. The behavior is quite different for the two temperature levels. In cold, the current drops fast to a low and stable level.

The results for a temperature scan between −40°C and +40°C in 5 K steps are shown in Figure 5.28. One hybrid per contractor is tested. The leakage current is extracted as the mean of the last ten data points in the time-dependent measurement. The error bars indicate the standard deviation of those data. An alternative approach of fitting an exponential yields compatible results. The data are shown on a logarithmic scale and a clear exponential dependence is visible. This is consistent with the assumption that the current is mainly attributed to the X7R dielectric capacitors, which feature this behavior [101]. For small temperatures and currents, the exponential behavior is not a good description. However, the current scale is small and

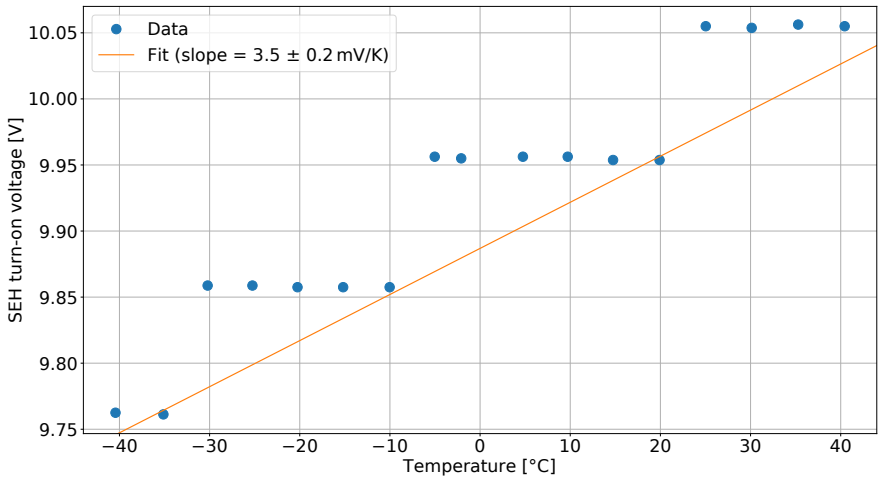

Figure 5.26: Turn-on voltage of the bPOL12V versus temperature of the climatic chamber, measured for a batch 2 final Service Hybrid prototype.

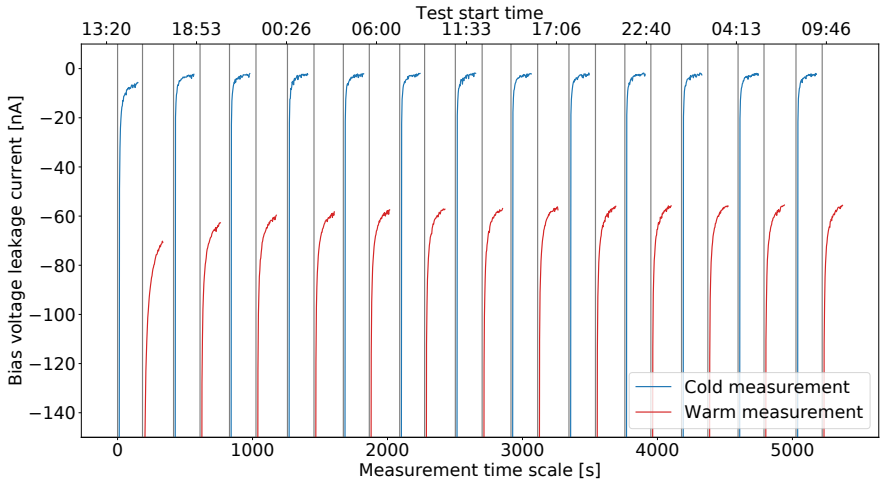

Figure 5.27: Measurements of the leakage current of the sensor bias voltage filter versus time performed for one hybrid during thermal cycling. Blue curves are obtained at −40°C, red curves at +40°C. The lower x-axis is the time scale of the leakage current measurement. The grey vertical lines indicate the start time of a measurement and refer to the upper x-axis. The applied negative voltage is 1,000 V.

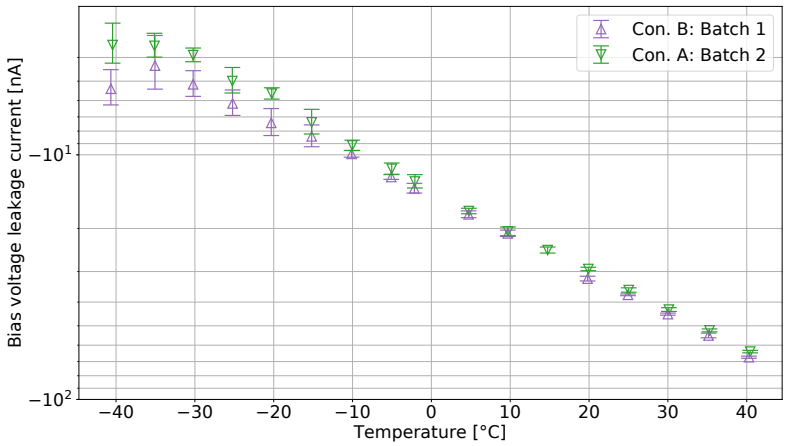

Figure 5.28: Leakage current of final Service Hybrid prototypes at the end of a 150 s long measurement versus temperature. One hybrid per batch is tested. The applied negative voltage is 1,000 V.

these measurements at low temperatures are performed first and suffer most from the early effects discussed above. All results indicate that the highest currents are expected at warm temperatures during the very first test and grading needs to be based on that.

5.3 Input for production

In addition to the quantitative assessments above, no issue concerning the digital communication of the lpGBT and VTRx$^+$ module were observed in the batch 2 hybrids. In earlier studies, the hybrids of batch 1 either suffered from the instabilities of the I^{2}C controller 0 or the weaknesses of the electrical vias in the flexes of contractor B. The good performance of hybrids featuring the production ASICs marked an important finding of the studies and gave confidence in the Service Hybrid performance for the final application. Only reliably operating prototypes enabled extensive testing and validation of modules and confidence in obtained results. Those tests in conjunction with the results of the hybrid test system provided the approval for the use of the hybrid design as part of the CMS detector.

The results from the second batch of contractor A were used to evaluate the behavior of the final ASICs.

The observed spread of the bPOL2V5 output voltage was consistent with the uncertainty of the internal voltage reference and the resistors in the voltage divider. The absolute value was slightly below the targeted 1.25 V at room temperature. Since

the decrease of the output voltage with temperature also needed to be considered, the resistor values were adjusted to increase it.

In the bPOL12V case, the stated variations of the internal voltage reference were too large to use random samples in the Service Hybrids. The final prototype measurements confirmed the datasheet information. The remeasurement and binning of the bPOL12V chips for the production addressed the issue. The voltage divider was adjusted (R_2 reduced to 165 kΩ) to account for the higher mean voltage reference in the binned sample for the 2S module.

The measurements showed that the ratio of output to input power is compatible with the assumed single-chip DC-DC conversion efficiencies and single-chip power consumption of the lpGBT and VTRx$^+$ module. The obtained results of the PCB resistance or voltage drops and the temperature influence served as input for the power model of the whole detector. Since the statistics of the temperature studies were limited, more data from additional tests were required. Corresponding studies are presented in Chapter 8.

The testing collected valuable information on the behavior of the enable and disable threshold of the bPOL12V. A mean of 10.2 V, a spread of $\sigma = 0.1$ V, and a variation with temperature was observed. The absolute value of the turn-on voltage needed to be reduced to account for an increase in the threshold with radiation. R_2 in the enable voltage divider was increased from 953 Ω to 1070 Ω.

The design and the components of the bias voltage filter were not definite in the final prototypes. They were adjusted based on the availability of parts and the experience of the studies of the module noise discussed in Chapter 7. The corresponding results obtained in this chapter can therefore not be transferred to the production Service Hybrids.

6 System tests with the M1 magnet

Of the many challenging operation conditions the new Phase-2 tracker will face the 3.8 T magnetic field is of particular interest. Even special iron alloys are fully magnetically saturated at 2 T [103]. As a consequence, only air core inductors can be used for the DC-DC converters. The designers of the common ASICs are aware of the conditions inside the particle physics experiments and design and test their chips accordingly. Naturally, this is only done on a single unit level. Since the OT hybrids and modules are composed of several different ASICs, PCBs, and passive components, it is vital to also test the full objects. The required magnetic field strength and instrument volume necessitate special facilities for such a test. Besides dedicated high-energy physics facilities, MRI magnets provided by medical facilities have been used in the past [104]. This chapter presents a measurement campaign with a Service Hybrid prototype at the M1 magnet at CERN. The hybrid is placed in the magnetic field and powered. The performance of the DC-DC converters and the communication with the lpGBT are studied.

6.1 Facility

The evaluation of the prototype performance under the influence of a substantial magnetic field was carried out at the CERN North Area Multi-Purpose Superconducting Magnet Facility in July 2021. The facility is located in building 887 on the CERN Prévessin Site. With the M1 and Morpurgo H8 it features two multi-purpose superconducting magnets and utilizes the Super Proton Synchrotron beam. The beam momentum ranges from 10 to 400 GeV/c [105].

The M1 magnet (max. 3 T central field) offers a higher magnetic field than the Murpurgo H8 (1.6 T). It also has a large free and warm opening, which makes it a good choice for prototype testing.

The magnet was manufactured in the late 70s as part of the European Hybrid Spectrometer experiment. The central part of the experiment was a rapid cycling bubble chamber, which was housed inside the magnet [106].

The magnet itself consists of two separated coils, similar to a pair of Helmholtz coils, with a horizontal axis. The windings have an outer (inner) diameter of 2.4 m

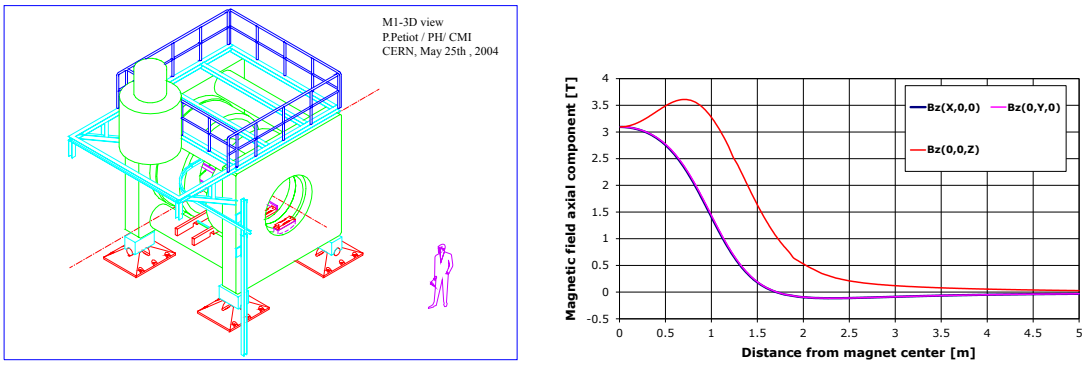

Figure 6.1: Left: schematic drawing of the M1 magnet with a person for scale. The red dashed lines indicate possible beam directions. Right: axial component of the M1 magnet's field versus distance from its center. The plot shows the dependency on the distances from the magnet center for three different directions [107].

(1.65 m) and are separated by 1.13 m. The warm free distance is 0.82 m horizontally along the axis, and the warm bore diameter is 1.4 m. This offers an acceptance angle from the central plane of $\pm 18°$. Figure 6.1 (left) shows a sketch of the magnet and the surrounding mechanical structure with a person for scale. The free volume is easily accessible.

The two NbTi coils are each submerged in a liquid helium bath at 4.2 K. The nominal current is 4,000 A per turn, and the total stored energy 55 MJ. The resulting magnetic field has a strength of about 3 T at the center [105] and has a cylindrical symmetry. Figure 6.1 (right) shows the axial magnetic field strength as a function of distance from the center for the two radial and one axial direction [107].

The entire magnet can be rotated by 90° to realize different field configurations with respect to the beam axis. For this first test of the OT components in a strong magnetic field no beam was present. Therefore it was sufficient to rotate the tested devices in order to achieve two orientations. The two orientations are "barrel" and "endcap" and refer to the tracker substructures' placement in the CMS detector. For all 2S barrel modules, the field is parallel to the sensors, for the 2S endcap modules it is perpendicular.

The magnet has no return yoke, which leads to significant fringe fields outside the magnet. Along the central axis, the fringe field has a strength of 125 mT at 1 m in front of the magnet and 15 mT at 4 m distance. In a radial distance from the outer magnet edge of 1 m the fringe field is 78 mT, while it is negligible after 4 m [107].

Measured voltage	Picoscope channel	DC range	AC range
Second stage output left side	A	2 V	20 mV
SEH input	B	20 V	50 mV
Second stage output right side	C	2 V	20 mV
First stage output	D	5 V	20 mV

Table 6.1: Summary of the settings used for the Picoscope that measures the input and output voltages of the SEH.

6.2 Setup

During the test campaign in July 2021, three different OT devices were tested inside the M1 magnet, a 2S module prototype, a PS-POH, and a Service Hybrid. This thesis only covers the measurements with the SEH.

Figure 6.2 shows the setup to test a SEH. It was designed to be installed without dedicated fixtures and without ferromagnetic materials. The hybrid rests on a 3D-printed polycarbonate plate, which is fixed with cable ties to an aluminum table. The table top is positioned at the center of the magnet. Cable ties fix the hybrid and a break-out PCB to the polycarbonate plate. The hybrid is not actively cooled, which needs to the respected during the testing. The break-out PCB is required as the interface to the hybrid voltages. The FEH connectors are plugged directly into the board. The input voltage is provided by a HMP4040, which also measures the input current. A custom-made three-channel active programmable load box (schematic in Appendix B.1) is used to draw current from the left and right side voltage output of the 1.25 V rail. The load box also measures the current and voltage of both sides. Voltage sensing wires are installed for the input power supply and load box connections to the break-out PCB to eliminate the influence of voltage drops on the supply cables on the measurements. A short cable routes the input voltage from the break-out PCB to in SEH input connector. The input voltage, both sides' 1.25 V outputs, and the 2.5 V output voltage are connected to a USB oscilloscope (Picoscope 5444B [108]). Screw terminals are used for the connections to the power supply and the load, BNC connectors and cables for the oscilloscope connections. The SEH's 2.5 V output is only intended for voltage sensing and no load can be applied. The communication with the lpGBT is realized by a VTRx$^+$ module and a 30 m long optical fiber pair, which is connected to an FC7 in the control room. A near field probe (ETS-Lindgren Model 7405-903 [109]) is placed on the DC-DC shield and connected to a second Picoscope. Table 6.1 summarizes the settings used by the Picoscope that measures the hybrid voltages. Part of the common infrastructure is a set of environmental sensors to monitor the temperature and humidity. One environmental sensor is placed next to the SEH.

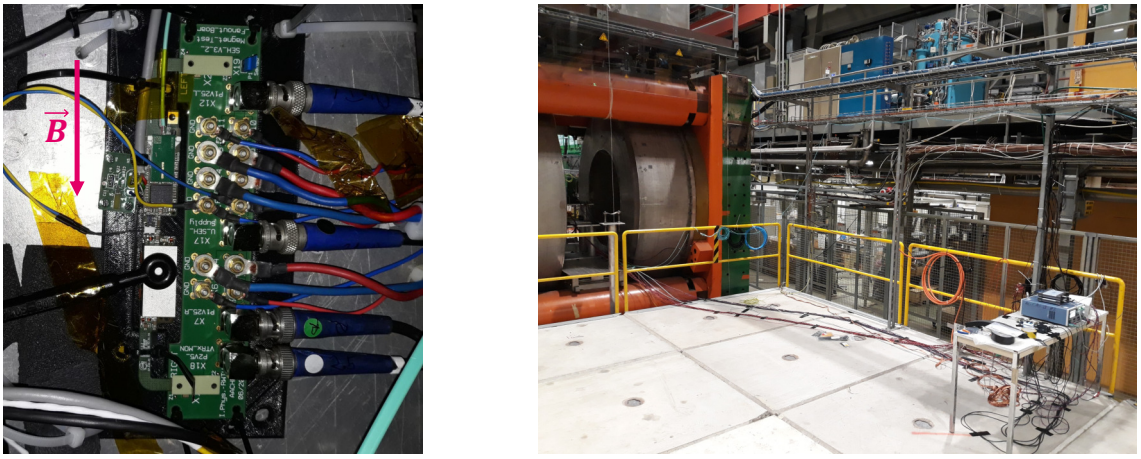

Figure 6.2: Left: SEH and break-out PCB after the installation inside the M1 magnet in the barrel configuration. Right: placement of the table with power supply, load, and oscilloscopes relative to the M1 magnet in the background.

To avoid problems caused by the magnetic fringe field, only the devices under test, environmental sensors, and passive connections are placed inside the magnet volume. The oscilloscopes, powers supply, and the load are placed on a table on a platform that is 4 m of radial distance away from the magnet and offset in axial direction by 1 m. The BNC cables and the power and sensing wires are roughly 10 m long. The oscilloscopes and load are read out and controlled via 30 m-long USB connections to a PC. The power supply has an Ethernet interface. The devices are integrated into a *Labview* program to carry out the measurements. An adapted branch of the module readout software is used from a PC console to perform the readout of the lpGBT.

The magnetic field strength is deduced from the current inside the magnet's coils. Measurements are performed at three different field strengths. A first measurement with no current in the coils and thus no magnetic field serves as a reference. The last measurement is performed at the full field strength of 3.0 T. As an intermediate step, a measurement at the field strength of 1.5 T is included. It also allows us to assess the conditions before the full magnetization of iron.

For the barrel configuration shown in Figure 6.2, the hybrid is placed flat on the table. The endcap configuration is achieved by setting the plate upright on the short edge. Figure 6.3 shows how the endcap configuration is realized for the three tested devices (left) and in a more detailed view for the SEH (right).

Due to time and resource restrictions, a single SEH prototype is used for all measurements. It was manufactured by contractor A in the first batch. Therefore, it features a bPOL12V_v4, bPOL2V5_v3.2, and lpGBT_v0.

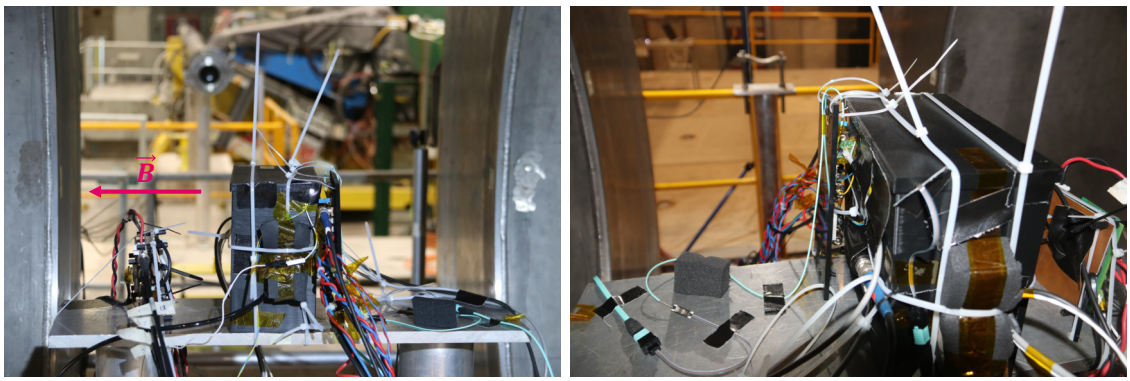

Figure 6.3: Left: the three tested Outer Tracker devices after installation on the table inside the M1 magnet. From left to right: a PS-POH, a 2S module in a light-tight box, and a SEH. They are all in endcap configuration. Right: close-up of the SEH installation in endcap configuration.

6.3 Results

The measurements cover five different aspects of the Service Hybrid. The functionality of the DC-DC converters is tested by scanning the phase space of input voltage and output load and recording the resulting current and voltages. We also study the dynamic behavior of the hybrid during a turn-on and turn-off sequence. Moreover, a measurement of the output voltage ripple is performed. The near-field probe is used to check the shielding and electromagnetic emission of the DC-DC converters. Lastly, the communication with the lpGBT via the VTRx$^+$ module and some slow control features are tested.

All parameters are also accessible in standard lab-based measurements. Those measurements do not require long cables, which introduce losses, damp signals, or pick up additional noise, and allow a more controlled temperature setting. They are to be preferred for quantitative studies of the hybrid performance. The tests inside the magnet focus on relative comparisons between the different configurations and field strengths.

The hybrid was operational for all field strengths and orientations. The magnetic field had no negative mechanical impact on the PCB or components. In spite of the missing direct cooling, no downtime due to the converters' over-temperature protection was required.

6.3.1 DC-DC conversion performance

The DC-DC conversion performance is measured by the power supply and the load. The input voltage is varied between 11.0 V and 6.5 V in 0.5 V steps. The maximum voltage that is considered safe for the long-term operation of a bPOL12V is 11 V.

The hybrid turns off for voltages lower than 6.5 V. It was planned to test with an applied load of up to 1500 mA per side, but the voltage drop on the cables to the load box is too high for the load circuit to operate in that case. The maximal achievable current was 1350 mA in a single-sided measurement and about 1200 mA when a load is applied to both sides. The applied load is varied in fixed steps in terms of DAC units. However, the response of the load is not linear and the current steps decrease for larger currents. Three different scenarios are measured: First, the load is applied only for the left, then for the right side, and in the end simultaneous on both. After a working point has been set and a settling delay of one second has been respected, the currents and voltages on the input and output are measured. The lpGBT configuration stored in the e-fuses is used. The VTRx$^+$ module is connected, and downlink frames are present. The measurement is performed three times per field configuration.

We calculate the ratio of output to input power, η. It combines all observables in a single quantity. Due to the sensing, it is independent of cable losses. Figure 6.4 shows η versus the total output load for a measurement in the 3 T barrel configuration for all different input voltage settings. The hybrid presents the expected ordered curves with η rising for smaller input voltages (smaller conversion ratio) [34]. As expected, η is increasing with the load current. Notably, both the lpGBT and VTRx$^+$ module consume power, which is accounted for only in the input power but not in the output power. The quantity η is therefore only related to, but not identical to, the efficiency of the two-stepped DC-DC conversion scheme. Since no particular influence of the input voltage is observed for the highest magnetic field, we compare the different configurations for a fixed input voltage of 10.5 V. Figure 6.5 shows the average of all measurements in a configuration. As the differences are quite small, the plots only show the highest output current region. The right selection also contains the errors based on the standard deviation of the three measurements. It illustrates that the measurement of η is consistent across all configurations. The small influence of the magnetic field on the load current is most likely a setup effect.

When a load is only applied to one side, the measurement of the output voltage load regulation is possible. We study the right side, as it includes less voltage drop in the PCB and does not supply the lpGBT and VTRx$^+$ module. The right and left side output voltage rails are separated after the bPOL2V5 and have different LC filters. In Figure 6.6, the bPOL2V5 output voltage measured on the right side is plotted versus the load applied to that side. The voltage falls linearly with the load current. The observed slope of roughly $|25\,\mathrm{m\Omega}|$ is compatible with a resistance measurement on the PCB. No significant differences between the configurations are observed.

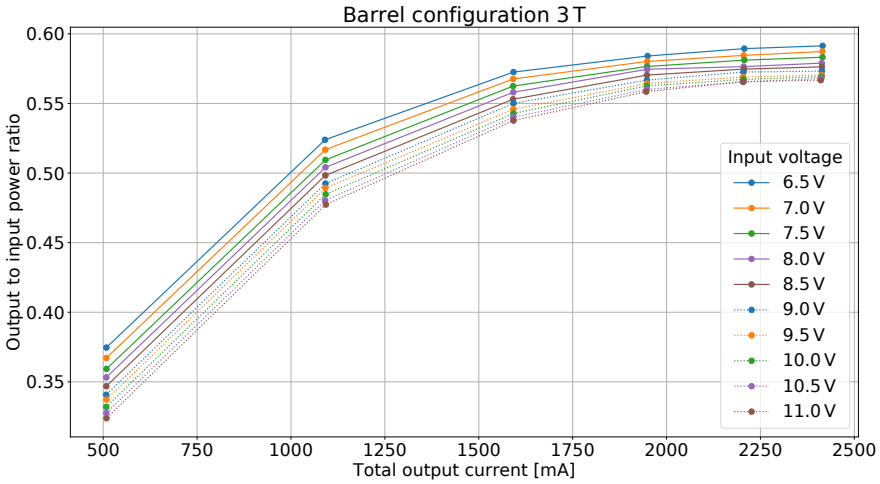

Figure 6.4: Ratio of output to input power versus total load current of the SEH in the 3 T barrel configuration. The various curves represent different input voltage settings.

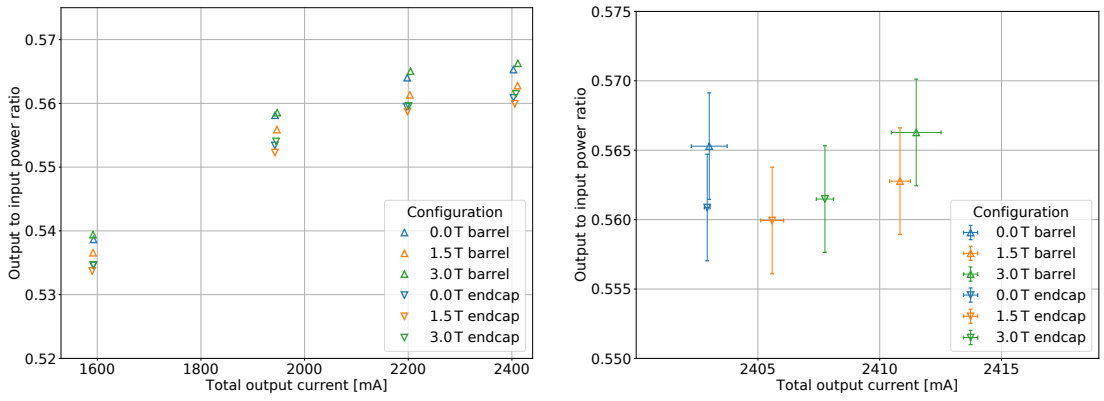

Figure 6.5: Comparison of the ratio of output to input power versus total load current of the SEH for the different magnetic field strengths and configurations. Only the zoomed view in the right figure shows the error bars obtained from the standard deviation of three measurements.

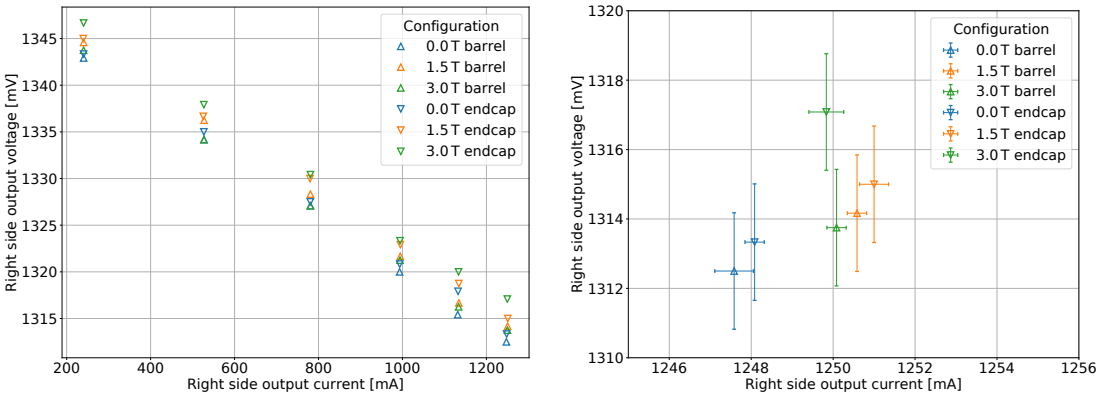

Figure 6.6: Left: comparison of the right side hybrid output voltage versus output current for the different magnetic field strengths and configurations. Right: zoomed view of the left plot at maximal output current. It shows the error bars obtained from the standard deviation of three measurements.

6.3.2 Voltage ripple

To measure the voltage ripple on the SEH output that is caused by the DC-DC converters, the electrical load and the oscilloscope are used. 1,250,000 samples are recorded with a random trigger and an 8 ns time resolution. The AC voltage signals at each field configuration are measured for three different load current settings (0.0 A, 0.8 A, and 1.2 A per side) and three times per set.

The electrical load box introduces large oscillations in the measured waveforms. The frequency of these oscillations changes with the regulated current but is always much smaller than the switching frequency of the DC-DC converters. The time-based data is therefore transformed into the frequency domain to allow a quantitative assessment. First, a Hanning window is applied to the data. Second, a fast Fourier transform (FFT) is performed [110]. The frequency and amplitude are determined following the prescription in Reference [111]. Figure 6.7 shows the resulting spectra for the three SEH output voltages for a measurement in the 3 T endcap configuration. The left spectrum was recorded with 0.0 A load, and the right one with 1.2 A load per side. While there are differences in the spectra for frequencies below 2 MHz which are attributed to the electrical load, the spectra are consistent for higher frequencies, where the DC-DC converters contribute. The largest contribution from the bPOL12V is visible in the first stage output voltage spectra with frequencies around 2.55 MHz. The first higher harmonics are also clearly visible. Corresponding peaks can be recognized in the spectra of the second stage bPOL2V5, too. The largest contribution of the bPOL2V5's switching frequency of about 3.55 MHz is expected in its second stage output. The right-side output shows higher contributions due to the shorter distance between the DC-DC converter output and the FEH connector

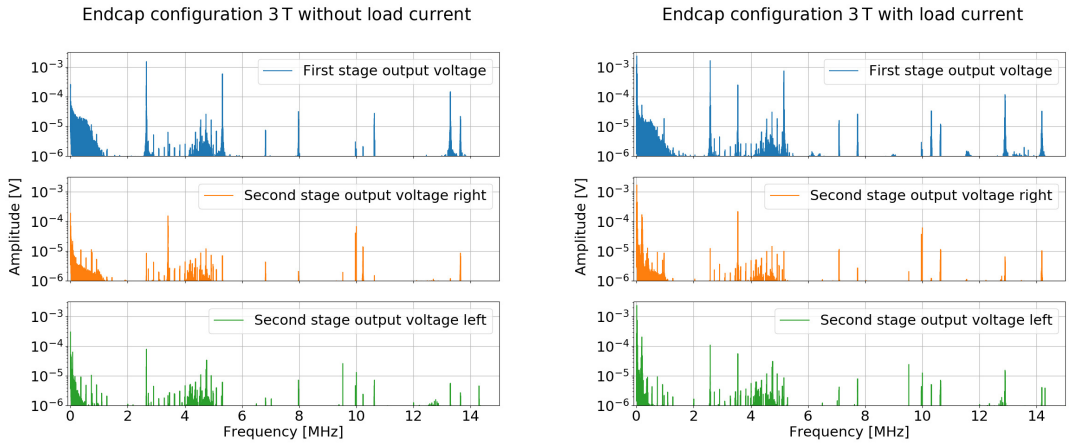

Figure 6.7: Spectra of the voltage ripple measured for the three hybrid output voltages in the 3 T endcap configuration. Left: spectrum without additional load current. Right: spectrum with an additional load of 1.2 A load current.

than on the left side. The left graph in Figure 6.8 shows the same measurement at 1.2 A load per side for the 0 T endcap configuration. No significant deviations caused by the magnetic field are present. For a quantitative comparison, the integrals of the bPOL12V switching peak in the first stage output and the bPOL2V5 switching peak on the right side output are determined. The right graph in Figure 6.8 shows a close-up of a bPOL12V switching peak on a linear scale. The dashed lines border the integration window of ± 4 kHz, which was optimized to cover the entire peak area.

Figure 6.9 shows the combined results for all field configurations as a function of the total applied load current. The full markers belong to the bPOL12V switching peak integral, and the open markers to the bPOL2V5 integral. The error bars represent the standard deviation of the three measurements per setting. Systematic effects are expected to cancel for relative comparisons. No dependency on the output current is observed. The results for the bPOL12V peak are consistent within their uncertainties across all configurations. This is also true for the 1.5 and 3.0 T bPOL2V5 peak integrals. The 0.0 T measurements are in themselves consistent but show a higher voltage ripple than with magnetic field. We exclude a negative effect from the operation inside a magnetic field on the voltage ripple.

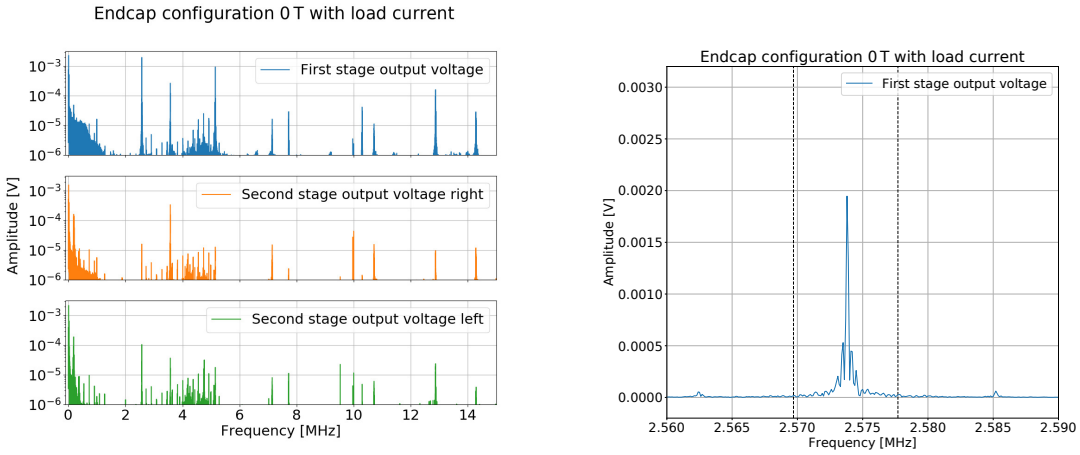

Figure 6.8: Spectra of the voltage ripple measured for the three hybrid output voltages in the 0 T endcap configuration with additional load current. Left: spectra for all three voltages for comparison with Figure 6.7. Right: zoom of the bPOL12V output peak for the first stage output voltage. The dashed vertical lines indicate the integration window (see text for details).

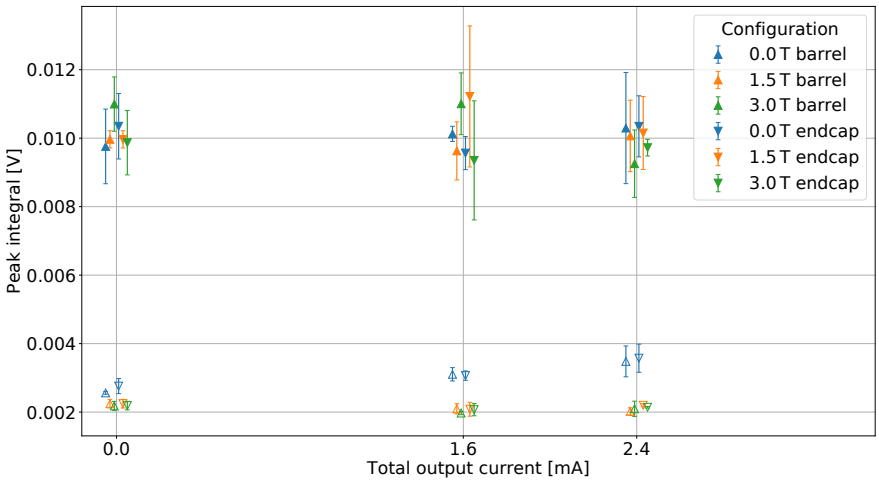

Figure 6.9: Comparison of the voltage ripple versus total load current of the SEH for the different magnetic field strengths and configurations. The error bars are obtained from the standard deviation across three measurements. Filled markers correspond to the bPOL12V switching peak and unfilled markers to the bPOL2V5. The measurements at a current setting are horizontally offset to improve visibility.

6.3.3 Radiated emission and shielding

The near field probe that is placed on top of the DC-DC converter shield can pick up electromagnetic emissions from the converters' main inductors that pass through the aluminum shielding. This RF process should not be influenced by the static magnetic field. The second Picoscope is used to record the voltage signal induced in the probe. It records 1,250,00 samples with a random trigger and 8 ns time resolution. The channel range is $\pm$ 10 mV AC coupled. During the measurement, the SEH is powered with 10.5 V, the lpGBT is configured, and a load of 1.2 A per side is applied. It is the maximal achievable load in the system. In general, the oscillating current inside a DC-DC converter's main inductor and thus the emitted radiation is independent of the provided current [34]. An approach analogous to the extraction of the voltage ripple is used to obtain a spectrum of the induced voltage that is a measure of the electromagnetic emission. Figure 6.10 shows on the left a spectrum obtained in the 3 T endcap configuration. Besides the DC part, the bPOL12V switching peak at 2.55 MHz is the highest contribution. Its higher harmonics are also distinguishable. In this hybrid, the bPOL2V5 has a switching frequency of about 3.5 MHz. The peak sits close to a broad noise contribution. It is still possible to extract the peak integral for both switching peaks. The result is shown in the right part of Figure 6.10. Since the whole setup was moved between the barrel and endcap measurements, the position of the coil was different for both conditions. The results are therefore not directly comparable. The results in the endcap configuration for both switching peaks show no change for the different field strengths. The ratio between the two DC-DC converter peaks depends heavily on the exact placement of the coil on the shield. This might explain the discrepancy between the barrel results with and without a magnetic field since the setup was touched before the magnet ramp-up. After the ramp-up, both peaks have a similar size, while the no-field result shows a difference in magnitude. Still, it is close to the endcap result. We conclude from the measurement that the shielding performs well in the presence of the magnetic field and the radiated electromagnetic emission of the DC-DC converter does not increase significantly or to any critical level.

6.3.4 Turn-on and turn-off sequence

The turn-on and turn-off behavior of the bPOL chips is discussed in Section 2.3.3. The largest current changes occur during the turn-on and turn-off so it is important to study it, especially in the presence of a large magnetic field. The measurements are performed for a single input voltage setting and the power supply channel is remotely switched on or off. The load box is set to the maximal achievable current of 1.2 A per side. In this way the largest current on the input is reached, too. The Picoscope measures the SEH input and output voltages with 8 ns time resolution. For the turn-on (turn-off) 1,250,000 (2,625,000) samples are recorded with a rising

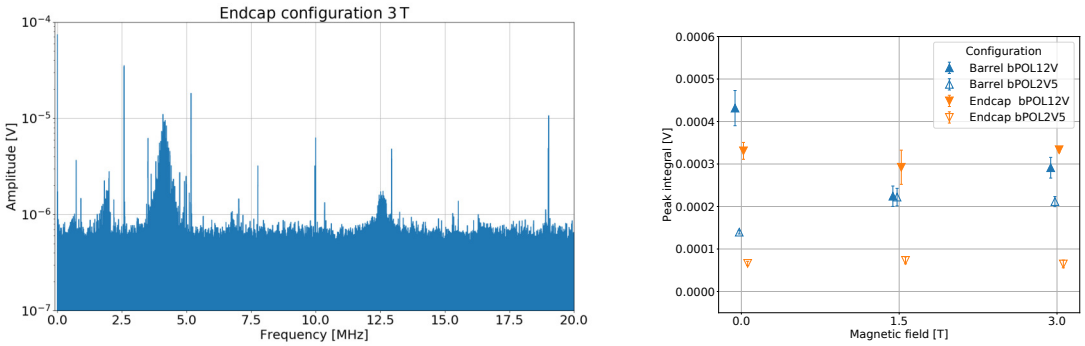

Figure 6.10: Measurement of the SEH's emitted electromagnetic radiation. Left: spectrum of the voltage induced by the emission in the 3 T endcap configuration. Right: comparison of the switching peak integral for the two DC-DC converter stages versus the magnetic field for the two different configurations. The error bars are obtained from the standard deviation of three measurements. The measurements at a field strength setting are horizontally offset to improve visibility.

(falling) edge trigger of 9 V on the input voltage. We conduct five measurements per configuration.

Figure 6.11 shows a turn-on sequence for the 3 T barrel configuration. After the power supply channel is enabled, the input voltage rises to 10.5 V within 6 ms. Due to the power supply regulation, the increase is not linear. It is the fastest ramp-up possible with this setup. At about 10 V the bPOL12V is enabled via a voltage divider on the input voltage and starts its soft start procedure. The bPOL12V increases its output voltage linearly and settles at the set voltage of 2.5 V. In a very similar way the bPOL2V5 is enabled via a voltage divider on the first stage output voltage at about 2.12 V. The bPOL2V5 also conducts its soft start procedure with a linear increase, before the output voltage settles at 1.3 V. The small overshoots are intrinsic to the soft starts. The observed right and left side output voltage and behavior are the same.

The waveforms for all measurements in the different configurations are inspected for obvious differences. As none are observed, we move on to compare the configurations based on their characteristic measures. We measure the turn-on voltage of the enabling supply voltage for each stage. For the bPOL12V the input voltage from the power supply is measured at the time when the first stage output voltage reaches 10 mV. We also define this time as the beginning of the soft start. The uncertainty of the turn-on time determination has a negligible impact on the turn-on voltage measurement compared to the Picoscope's noise level. The turn-on time uncertainty is limited by the steep increase of the bPOL12V output voltage, and the SEH input voltage measurement has a noise level of about 20 mV due to the large required DC

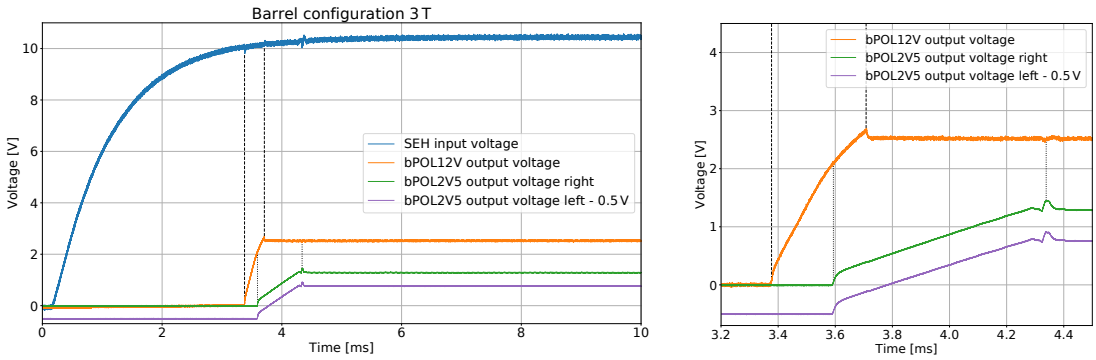

Figure 6.11: Measurement of the SEH turn-on in the 3 T barrel configuration. The right plot shows an enlarged section of the left one. The blue curve shows the SEH input voltage from the power supply. The orange curve shows the bPOL12V output voltage and the green and purple curves the output voltage of the bPOL2V5 measured on the right and left side, respectively. The bPOL2V5 left side output voltage is offset for visibility. The grey dashed (dotted) lines illustrate characteristic measures of the bPOL12V (bPOL2V5) turn-on.

range of Picoscope. The definition introduces a systematic underestimation of about 5 µs for the soft start duration. We define the maximum of the regulator overshoots as the end of the soft start and measure the maximum voltage. As the DC range is smaller, the noise level on the first stage output voltage measurement is only about 5 mV. The enable voltage, the duration of the soft start, and the soft start overshoot voltage are measured analog for the bPOL2V5.

Figure 6.12 summarizes the measured turn-on voltages. The uncertainties are based on the standard deviation of the five measurements and are consistent with the noise on the voltage waveforms. The turn-on voltages are consistent within their uncertainties. The soft start overshoot voltages are obtained in the same way and shown in Figure 6.13. Neither for the bPOL12V nor for the bPOL2V5 any significant effects of the magnetic field are observed.

The duration of the soft start is calculated from the difference in time between the start of the turn-on and the maximum of the regulation peak. The results are shown in Figure 6.14. For each orientation, the durations are consistent within their uncertainties except for the 3 T endcap measurement for the bPOL2V5. The deviation is small with ten nanoseconds and consistent with the barrel results.

The turn-off sequence is analyzed similarly. A waveform example is the 3 T barrel measurement in Figure 6.15. When the power supply channel is disabled, the input voltage drops linearly until the disable voltage of the bPOL12V is reached. The output voltage of the bPOL12V starts to decrease until the bPOL2V5 is disabled. Already during this phase, the bPOL2V5 output voltage drops (Figure 6.15, right).

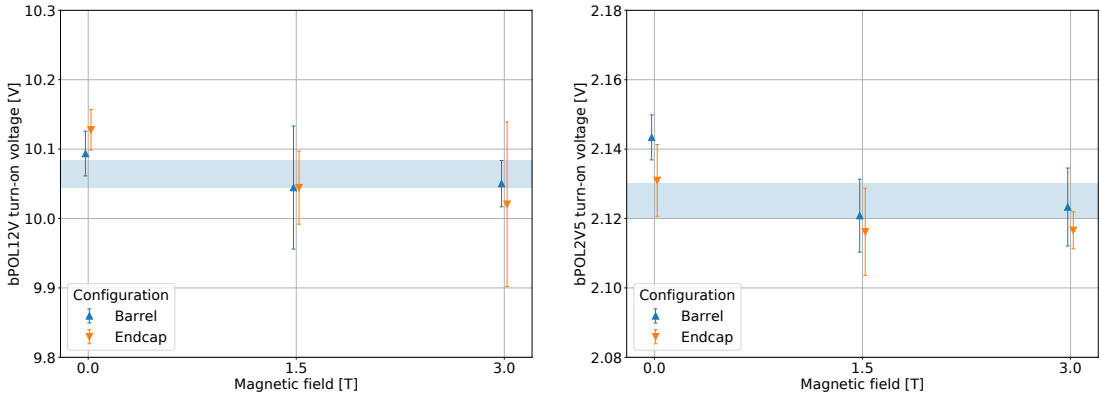

Figure 6.12: Left: turn-on voltages of the bPOL12V for the different orientations versus the magnetic field strength. Right: measurement of the bPOL2V5 turn-on voltage. The blue band represents the mean value with its width representing the noise of the Picoscope waveform. The measurements at a field strength setting are horizontally offset to improve visibility. The error bars are obtained from the standard deviation of three measurements.

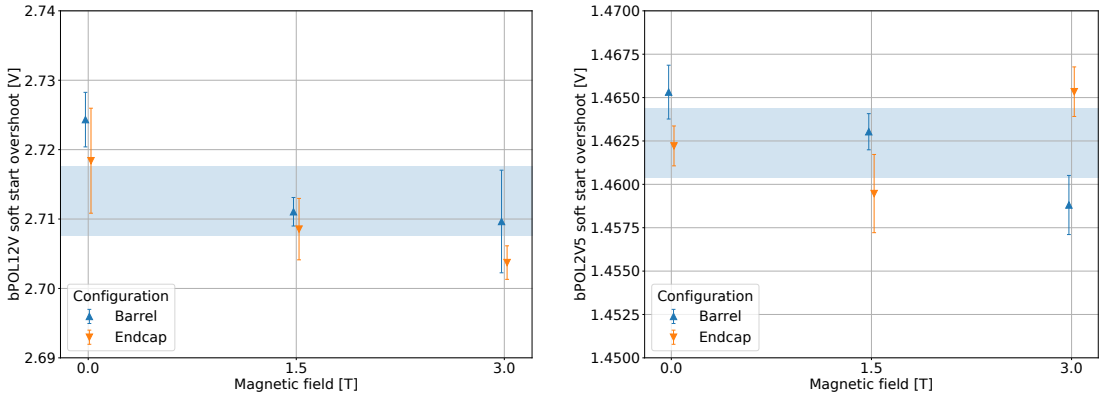

Figure 6.13: Left: soft start overshoot voltage of the bPOL12V for the different orientations versus the magnetic field strength. Right: measurement of the bPOL2V5 soft start overshoot voltage. The blue band represents the mean value with its width representing the noise of the Picoscope waveform. The measurements at a field strength setting are horizontally offset to improve visibility. The error bars are obtained from the standard deviation of three measurements.

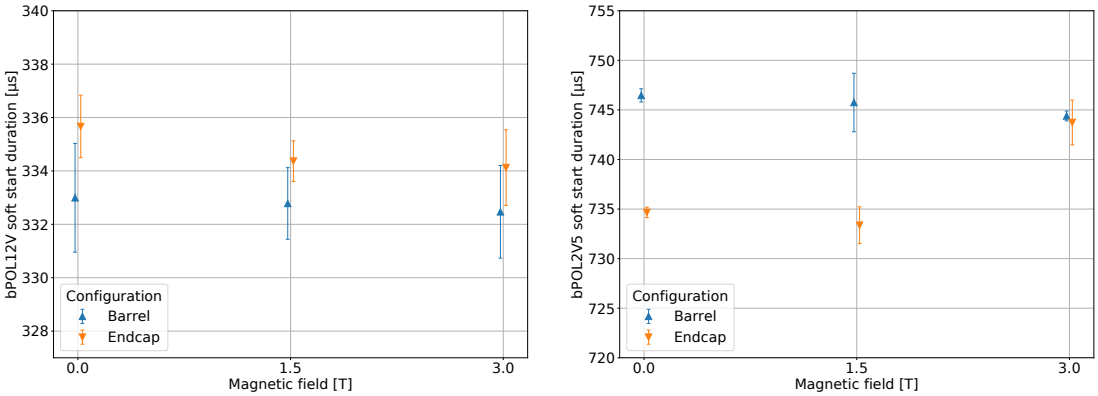

Figure 6.14: Left: soft start duration of the bPOL12V for the different orientations versus the magnetic field strength. Right: measurement of the bPOL2V5 soft start duration. The measurements at a field strength setting are horizontally offset to improve visibility. The error bars are obtained from the standard deviation of three measurements.

After it is disabled, the output voltage decreases exponentially within a short time. Since both DC-DC converter inputs feature a high impedance after the DC-DC converters are disabled, the input voltage and bPOL12V output voltage decrease only slowly at this point. Upon inspection of all waveforms, we do not observe apparent differences.

Two measures are used to characterize the turn-off sequence: the turn-off voltage of the bPOL12V and the duration of the voltage decrease after the turn-off of the second-stage DC-DC converter. Only the prior is an intrinsic property of the DC-DC converters, as the latter strongly depends on the system impedance. We define the turn-off voltage as the SEH input voltage at the point where the first stage output voltage decreases below 2.4 V for the first time. The turn-off voltage results are shown in the left plot of Figure 6.16. Uncertainties are estimated from the standard deviations of the five measurements per configuration. The data are consistent within their uncertainties. The duration of the voltage decrease is measured as the time between the decrease of the second stage output voltage below 1.247 V and the decrease to 0.2 V. As shown in the right part of Figure 6.16 this time has a small dependence on the magnetic field strength. Closer inspection of the waveforms yields that a small negative offset (-10 to -50 mV) is present as the voltage level approaches zero. This offset increases with the magnetic field. The time of the 0.2 V level crossing is sensitive to the offset. The offset might be caused by the large currents inside the magnet coils, which influence the ground levels of the Service Hybrid power supply and the Picoscope. The size of the effect is not considered critical. In summary, no harmful effects of the magnetic field are observed for the turn-on or turn-off sequence.

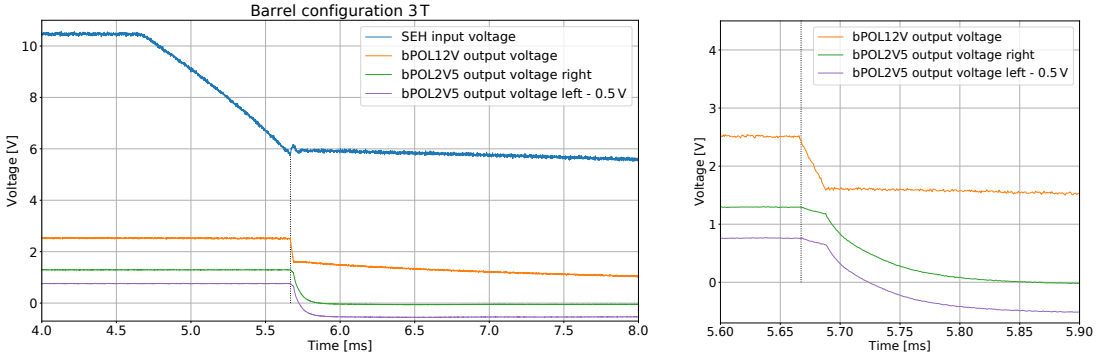

Figure 6.15: Measurement of the SEH turn-off in the 3 T barrel configuration. The right plot shows an enlarged section of the left one. The blue curve shows the SEH input voltage from the power supply. The orange curve shows the bPOL12V output voltage and the green and purple curves the output voltage of the bPOL2V5 measured on the right and left side, respectively. The bPOL2V5 left side output voltage is offset for visibility. The grey dotted line is placed at the turn-off point.

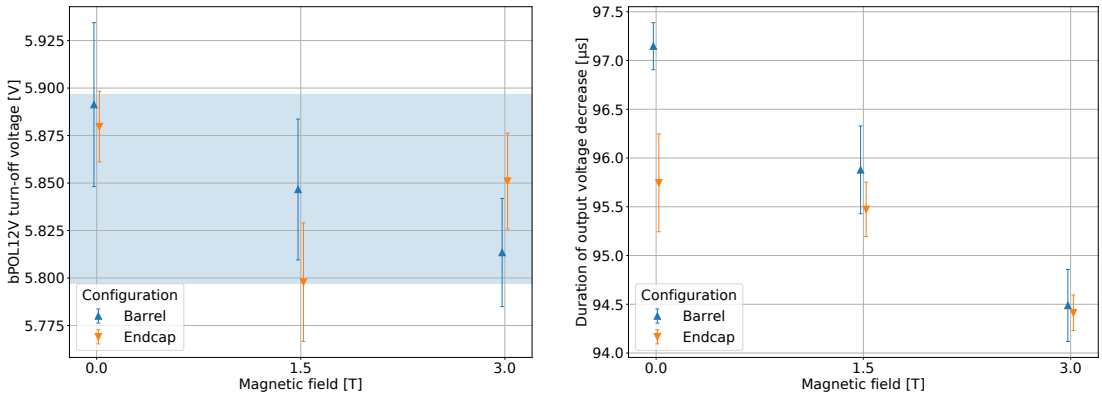

Figure 6.16: Left: turn-off voltage of the bPOL12V for the different orientations versus the magnetic field strength. The blue band represents the mean value with its width representing the noise of the Picoscope waveform. Right: measurement of the duration of the voltage decrease after the turn-off of the second-stage DC-DC converter. The measurements at a field strength setting are horizontally offset to improve visibility. The error bars are obtained from the standard deviation of three measurements.

6.3.5 lpGBT communication and slow control

After the SEH is powered up at 10.5 V, communication with the lpGBT is possible through the optical link. A set of communication tests is performed five times at each configuration, while it is repeated for a setting of 0.0 A additional load current per side, and 1.2 A per side (2.4 A total output current). Most tests only record pass or fail. They are summarized below with their outcome:

- The lpGBT's ability to lock onto the optical link is checked. A default SEH configuration is applied by programming a set of registers with the IC interface. No problems are observed.
- The lpGBT's internal pattern generator is used to inject fixed patterns on the uplink. They are correctly recovered by the back-end.
- We use the I^2C controller 1 to communicate with the VTRx$^+$ module's laser driver. The default values of its registers are successfully read back, which confirms the functionality of this interface.
- The Power Good signals of the two DC-DC converter ASICs are measured with the lpGBT GPIO interface. Their signals are always present.

In addition, the ADC interface is used to monitor a set of voltages. For the lpGBT_v0 that is used in the tested hybrid, no calibration for the ADC's voltage reference was performed. Therefore, the measurements are carried out with an untrimmed voltage reference. This leads to a systematic upward deviation of roughly 10% for all measured voltages. A single measurement is composed of 100 consecutive ADC conversions of the same channel. As the standard deviation is found to be smaller than 1‰, only the mean value is subsequently used.

There are four values of interest accessible on a bare Service Hybrid. The input monitoring voltage, the bPOL2V5 output monitoring voltage, the bPOL2V5's PTAT voltage, and the VTRx$^+$ module's RSSI voltage. Unfortunately, due to the untrimmed voltage reference and the discrete voltage divider on that SEH version, the dynamic range of the latter is not sufficient for any meaningful analysis.

The input monitoring voltage is analyzed first. It is derived from the input voltage of the SEH via a voltage divider. Since the input power supply uses sensing, it should be independent of the current consumption. Figure 6.17 shows the measured voltage for each configuration in the order in which they were obtained in the barrel (left) and endcap (right) orientation. We observe only a small deviation ($< 1\%$) of the absolute value. However, a clear trend is visible. The measured voltage drops with the measurement number. This effect is less pronounced in the measurements without load current, which also tends to have higher starting values. According to the calibration of the lpGBT_v1, the voltage reference has a temperature dependence ($\mathcal{O}(-0.5‰/\mathrm{K})$) [14]. It directly propagates to the measured voltage. We therefore conclude that the hybrid is heated up between each measurement, which is performed

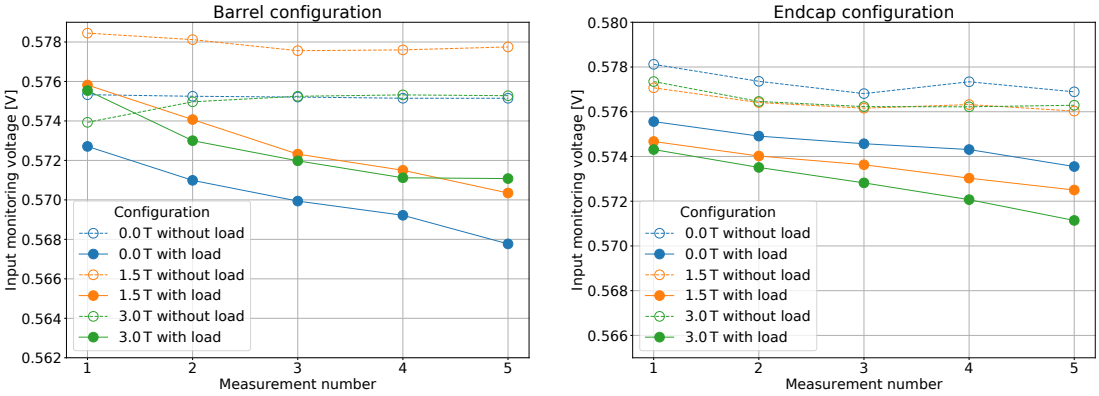

Figure 6.17: Measurements of the SEH input monitoring voltage with the uncalibrated lpGBT ADC for different field strength and output load conditions. Left: in barrel orientation. Right: in endcap orientation.

roughly 30 to 60 seconds after another. This can also explain why the effect is more pronounced in the loaded case, as power losses in the DC-DC converters also increase.

The increase in temperature is also measured by the bPOL2V5 PTAT voltage. The left plot in Figure 6.18 shows the measured PTAT voltage in the endcap configuration versus the temperature of the nearby positioned environmental sensor. The temperatures and their uncertainties are estimated from the timestamp of the measurement and the logged environmental data. The uncertainty of the PTAT voltage measurement can be estimated from the standard deviation of the 100 ADC conversions and is found to be negligible. We only consider the endcap data, since the environmental sensor is placed in the vicinity of the hybrid for that configuration and in better thermal contact in a more confined placement. The measurement confirms the rising temperature of the PCB in between measurements and also the higher level in the loaded case. The dashed line represents the expected increase in PTAT voltage if the rise is solely attributed to the warmer environment. Since the increase is much larger than this, we observe a significant self-heating of the PCB.

The right plot shows the measurement of the bPOL2V5 output monitoring voltage versus the total load current. All configurations feature two distinct voltage levels for the different load currents. This load regulation and the Ohmic losses are also observed in the direct measurement (Figure 6.6). The deviation for the 3.0 T endcap measurement at 0 A load is present in both measurements. The uncertainty is derived from the standard deviation of the five measurements and is larger in the loaded case. This is consistent with the temperature change.

None of the presented measurements show significant effects that can be attributed to the magnetic field. We conclude that the lpGBT communication and the slow control capabilities are functional in the presence of a strong magnetic field.

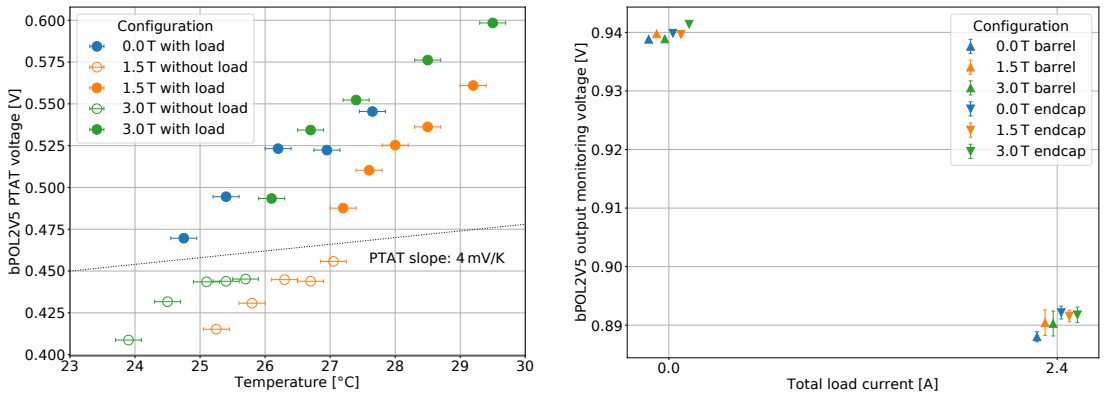

Figure 6.18: Left: measurements of the bPOL2V5 PTAT voltage with the uncalibrated lpGBT ADC versus air temperature for different field strength and output load conditions in the endcap orientation. The grey dashed curve indicates the expected PTAT slope. The error bars indicate the uncertainty of the temperature determination. Right: comparison of the measured bPOL2V5 output monitoring voltage with the uncalibrated lpGBT ADC versus total output load current for the different configurations. The measurements at a current setting are horizontally offset to improve visibility.

6.4 Summary of the magnetic field system tests

The facility and the setup that was used to characterize a batch 1 Service Hybrid prototype inside a magnetic field of up to 3 T was presented. The tests were carried out in July 2021 and included tests of the DC-DC converters, measurements of the output voltage ripple and radiated emission, validation of the turn-on and turn-off sequence, and communication and control tests of the lpGBT. The campaign demonstrated the successful operation of all SEH components in the magnetic field. Only one observable showed a small dependence on the magnetic field. The effect is not caused by the active components, but instead by the system impedance and by the influence of the magnetic field on the measurement devices. All other observables show no dependence on the magnetic field.

The ASIC technology is identical between the prototype ASIC versions used in this test and the production versions. Therefore, the experience of this test is transferable to production pieces. The same holds for the PCB design, passive components, and other materials. In the future, more tests on whole modules and with a particle beam will be carried out in a magnetic field. The quantities studied in this dedicated setup will not be accessible there. So, this study provides valuable insights into the hybrid performance inside the magnetic field.

7 Hybrid production kick-off

The decision on the manufacturer of the hybrid flexes and the assembly company was settled after the conclusion of the prototyping phase. It was originally planned to have also the design and the component selection finalized at that point and proceed with the production of detector-grade hybrids.

However, an average noise of about 1220 electrons (e^-) was observed on prototype 2S modules constructed with final prototype Service and Front-End Hybrids. In comparison, earlier modules featured a noise of 940 e^-. Both quoted noise levels are the values measured at room temperature and decrease with temperature. For the targeted hit efficiency of $> 99.5\%$, the noise of the final system should be below 1000 e^- at the detector's operating temperature [12, 112].

The earlier modules used the same sensors, readout ASICs, and a very similar FEH design. However, as outlined in Chapter 2 the SEH components and design changed significantly from the v3.1 prototype to the final prototype.

The collaboration took the decision to pursue multiple approaches to understand and mitigate the noise source. The activities included, among others, the evaluation of different grounding and powering schemes, exploratory measurements with near field probes, direct measurements on the power nets, and injection of time-varying magnetic fields [75, 112]. One of the studies, a direct comparison of AC ground potential differences for different Service Hybrids, is presented in Section 7.2.

The studies identified two factors that made a 2S module with final prototype Service Hybrids more susceptible to noise:

- The bias voltage flex tails and the sensors with their connected front-ends form a loop with a small impedance. Time-varying magnetic fields from the DC-DC converters' main inductors are not fully shielded and generate large currents in the loop. These currents lead to noise in the analog front-end. The v3.1 prototypes feature a stiffener with 70 µm copper on both sides, and thus the bias voltage circuit on the backside is better shielded.
- The DC-DC converters and their generated switching currents induce an AC voltage difference on the ground potential between the left- and right-side Front-End Hybrid. The voltage difference creates an AC flow through the sensors and directly injects noise into the analog front-ends. The AC noise is present on a side's power and ground rail and cannot be measured in a

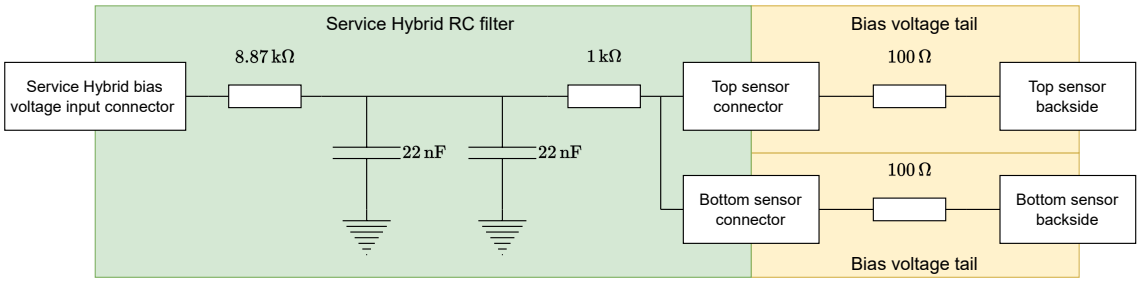

Figure 7.1: Simplified circuit diagram of the bias voltage RC filter on the Service Hybrid backside and the resistors on the bias voltage tails. Wire bonds connect the tails electrically to the sensor backsides.

Figure 7.2: Illustration of the ground-balancing PCB assembled into a 2S module and connected to the Front-End Hybrids. The rest of the module is grayed out.

differential measurement between those potentials on one hybrid side. In other words, the common mode voltage is different for both Front-End Hybrids.

In response to the first observations, a 100 Ω resistor was added to the design of the bias voltage tails. The resistor limits the current flow between the sensors and greatly improves the noise performance. The simplified circuit diagram of the bias voltage RC filter on the Service Hybrid backside and the resistors on the bias voltage tails is shown in Figure 7.1. To address the second finding, a ground-balancing PCB was added to the module design. As shown in Figure 7.2 it is a passive PCB on the side of the module opposite the SEH that provides a low-impedance ground connection between the two FEHs. This alternative path reduces the current flow on the sensors. In single-module measurements, the grounding of the FEHs to the module carrier has a similar effect.

As this only mitigates the effects of the AC voltage difference and requires a design change of the FEHs, a third approach was pursued that also required an additional SEH design and order. The resulting hybrids are referred to as “kick-off” hybrids as the order starts the transition from prototyping to production. The single remaining contractor produced the hybrids.

In the Service Hybrid case, flexes with two different designs were produced. The first one, called the "regular" design, is in principle very similar to the design of the final prototypes. The changes are summarized in the following:

- All vias are now implemented as staggered vias. More redundancy of vias is added. The routing of power and ground planes is improved and their traces are enlarged.
- The ground plane on the top layer under the shield now covers a larger area and the connection and solderability of the shield are improved.
- The output voltage of the bPOL12V is routed on the third layer to its output LC filters to allow some improvements on that rail, too.

The second version, called the "split-plane" design, includes all previously listed changes. In addition, it differs from the regular design only in one major aspect, illustrated in Figure 7.3 and Figure 7.4. In both hybrid variants, the ground potential is present on parts of the top layer L1 and most of the inner layers L2 and L4. For the split-plane variant shown in Figure 7.4, the only connection between the two inner ground layers within the hybrid area that contains the DC-DC converters is at the input voltage filter. The ground layer L2 is split behind the bPOL2V5 output filter so that only the common reference ground is provided to the right FEH. The intention is to contain conducted common mode (CM) noise on the layers L1 and L2 and obtain a clean ground on L4. The inner layers are interconnected at the FEH connectors (not shown in Figures 7.3 and 7.4). The regular design in Figure 7.3 has a continuous L2 and additional vias that connect L2 and L4 under the DC-DC converters.

The contractor provided assembled flexes of both variants in late September 2023.

7.1 Experiences from the kick-off Service Hybrid testing

For the first time, the functional testing of the 40 kick-off SEHs was not conducted at the I. Physikalisches Institut B of RWTH Aachen University. They were tested at and by CERN using the test cards and procedures provided by the I. Physikalisches Institut B of RWTH Aachen University and described in Section 4. The results are thus not included in this thesis. Due to a mistake in the initial assembly, the hybrids provided an incorrect bPOL2V5 output voltage (2.5 V instead of 1.25 V). A resistor with the wrong value had been assembled in the voltage divider for the bPOL2V5's feedback loop. To obtain usable hybrids, the DC-DC converter shield was removed by the contractor and replaced after the error was corrected. The shield replacement is a severe intervention and as is discussed below the yield of the hybrids with the regular design was rather low afterwards (about 50%). After the replacement, 39

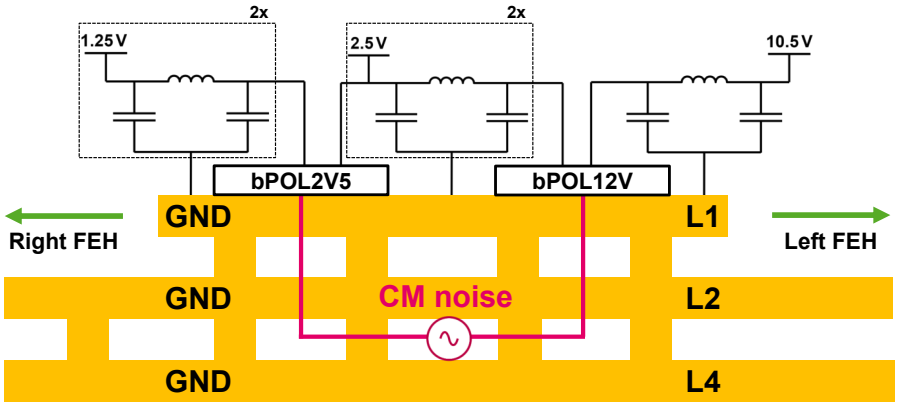

Figure 7.3: Schematic illustration of the concept for the regular design. The ground (GND) layers are interconnected with vias. Each DC-DC converter has two output filters: the bPOL12V for the VTRx$^+$ module and the input of the bPOL2V5; and the bPOL2V5 for the left and the right FEH.

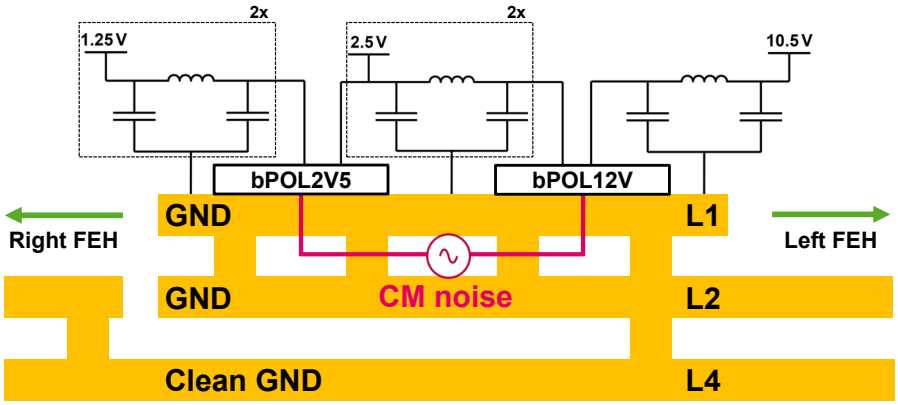

Figure 7.4: Schematic illustration of the concept for the split-plane design. A common reference ground (GND) for both hybrid sides exists at the input filter. Each DC-DC converter has two output filters: the bPOL12V for the VTRx$^+$ module and the input of the bPOL2V5; and the bPOL2V5 for the left and the right FEH.

hybrids were tested. Only 20 hybrids passed the test. A total of ten hybrids failed the testing and in eight cases the testing was not complete. It was the first time the 2S-SEH test card was used outside the developing institute and technical difficulties affected the stability of the test procedures.

The I. Physikalisches Institut B of RWTH Aachen University supported the ensuing debugging investigations and failure analysis of the 18 hybrids that did not pass the functional testing (10 split-plane design hybrids and 8 regular design hybrids). All issues discussed below were identified in Aachen. The author retested the hybrids in a more stable table-top 2S-SEH test card setup and all split-plane hybrids passed the test, while all regular design hybrids failed it. In the malfunctioning hybrids, the bPOL2V5 either provided a wrong (too high) output voltage or did not provide any output voltage. Both issues were sensitive to changes in temperature or mechanical stress on the hybrid. Most likely, the failures can be traced back to the bump bonds of the bPOL2V5. The left image of Figure 7.5 shows a cross-section of a bump bond from the first row for a malfunctioning hybrid prepared in Aachen. In contrast to the usual ball shape, the bump appears deformed and compressed. A large void is visible where the bump connects to the ASIC pad. The first row bumps are responsible for the enabling of the DC-DC converter as well as the feedback loop of the output voltage. An open connection can explain the observed failures. Also, the sensitivity to temperature and mechanical stress is consistent with a weak bump bond. Fresh bPOL2V5 ASIC samples were also inspected and showed good bumps. All assembled hybrids had their shield replaced. Unfortunately, it cannot be known for certain if the bumps were soldered correctly and if they were damaged during the repair. The failure was only observed for hybrids with the regular design. No design-related reason for the failure was identified.

Closer inspection of the hybrids also revealed additional yet less critical assembly errors, which prevent the full monitoring via the lpGBT. In addition, the right picture in Figure 7.5 shows a design error in the bottom side's solder mask that could potentially result in a short and needs to be covered with isolating tape.

The hybrids were also designed to use the FEC12 mode for the lpGBT's forward error correction. The 2S module uses one 320 Mb/s e-link less per side compared to what is available on the lpGBT, and the additional bandwidth on the uplink could be used to utilize the stronger forward error correction. This change requires new firmware for the optical readout. However, one stub data e-link was not rerouted to the required lpGBT receiver group for this FEC change and prevented the readout of stub data from the right FEH. The forward error correction mode is hardwired and inaccessible under the bump-bonded lpGBT. The routing mistake was discovered in Aachen. Due to the required bandwidth and number of e-links, the PS module needs to be operated in the FEC5 forward error correction mode. By design, the PS modules are placed closer to the beam interaction point and thus experience higher dose and flux rates than the 2S modules. As the FEC5 mode is considered safe for the operation of the PS modules, the collaboration decided to revert to this mode

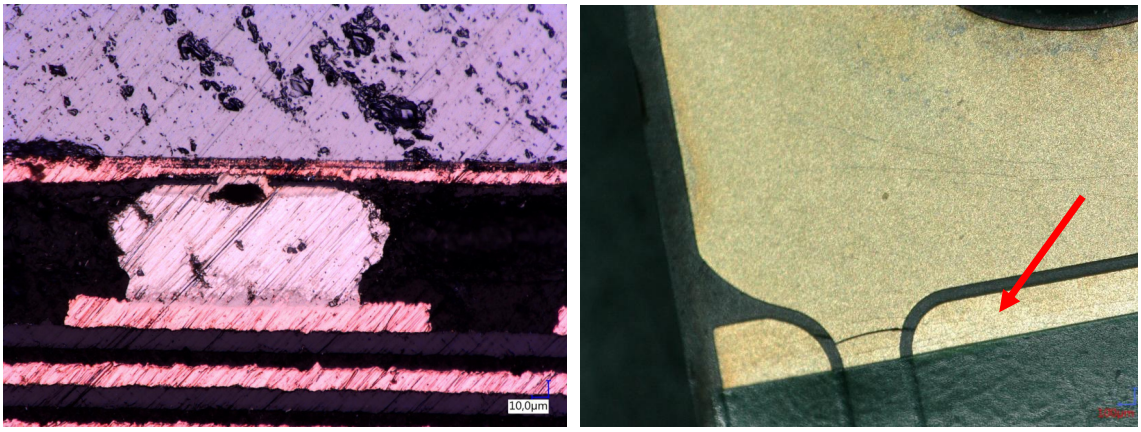

Figure 7.5: Left: cross-section of the bump bond required to enable the bPOL2V5 ASIC. The picture belongs to a malfunctioning hybrid. The bump bond is deformed and a large void is visible. Right: detailed view of the bottom side of a kick-off SEH. The solder mask in the grounding flap region erroneously exposes a 1.25 V plane. A (unrelated) crack in the adjacent ground plane copper is also visible. The red arrow indicates the locations of the 1.25 V plane

for the 2S production. The change requires fewer modifications of the Service Hybrid design and was well-tested with the final prototypes.

The measurement of the module noise is possible without the stub data and enough regular and split-plane design hybrids were deemed usable to compare the noise performance of the two layouts on kick-off modules. Figure 7.6 shows the summary of the studies performed by the CMS Collaboration [113]. Modules with a ground-balancing PCB and a regular design SEH yielded the lowest noise levels. Unexpectedly, the split-plane design showed significantly higher noise.

The testing at CERN and the debugging at Aachen were performed under time pressure. The noise assessment was only possible after the hybrids were mounted in modules. Due to the low yield caused by the problems described above, the kick-off Service Hybrids were not well suited to assess the expected performance of series-production hybrids.

The shortcomings of the kick-off hybrid production were addressed in an additional assembly, which is presented in the next chapter. Those hybrids were required to validate the regular Service Hybrid design with a sufficient number of modules.

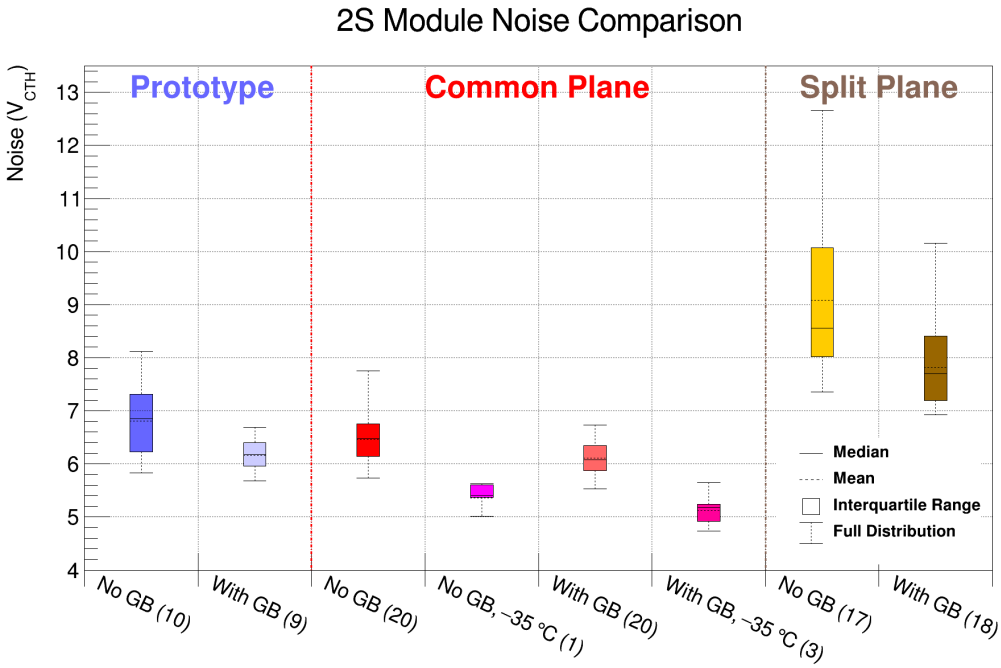

Figure 7.6: Noise measurements with prototype and kick-off 2S modules in different configurations performed at −350 V bias voltage. The measurements were performed either at room temperature or at −35°C, where indicated. The boxplot representations of the obtained strip noise levels per configuration are given in terms of the CBC threshold unit, V_{CTH}. The distributions contain all individual channel noise measurements and include measurements from several modules; quantities are indicated in the parenthesis. The median and mean are indicated by a solid and dashed horizontal line, respectively. The box encloses the central half of a distribution and the whiskers show its full extent. The two right data sets refer to the split-plane design ("Split Plane"). The others refer to the final prototypes ("Prototype") and the regular design hybrids ("Common Plane"). Measurements with and without the ground balancing PCB (GB) are shown [113].

7.2 Comparison of the AC voltage difference between ground potentials

As a module assembly center, Aachen contributes to the investigations of the noise behavior of the 2S modules. The ground of either silicon sensor is per design connected to the ground potential of both sides' FEHs. Since the beginning, noise studies with prototype modules were found to be very sensitive to the grounding of the FEHs to the module carriers. Exploratory oscilloscope measurements hinted at the presence of an AC voltage unbalance between the left and right hybrid ground potential. A voltage difference would result in a current flow. In a grounding scheme where both FEHs feature a low-impedance ground connection to the carrier, most of the current takes that path. Otherwise, depending on the grounding, some or all current would flow across the silicon sensors, which can lead to additional noise in the front-end. Those exploratory studies were performed on the desk with module prototypes or skeletons (hybrids without silicon sensors) and suffered from large background noise or pick-up.

A dedicated setup was put together to perform a clean measurement and objectively compare different hybrids. Figure 7.7 shows the basic principle. The Service Hybrid is connected to a custom PCB for this measurement. Its schematic is presented in Figure C.1 in the Appendix. The PCB receives the ground potential from both FEH connectors and routes them in a symmetric way to a 1 kΩ resistor that connects both sides. The voltage difference across the resistor is measured with an active differential probe (Teledyne LeCroy AP033 with AC coupler [114]) and a compatible oscilloscope (Teledyne LeCroy Waverunner 9254M-MS [115]). The hybrid, the PCB, and the probe are placed inside a copper box to shield them from external fields. The 10.5 V input voltage of the SEH is provided with a power supply (Keysight E3633A [116]). The supply cable is routed in a metallic mesh that is connected to the laboratory ground at the power supply. The mesh is also connected to the copper box.

Pictures of the setup are shown in Figure 7.8. To obtain a short and simple connection, the SEH supply cable and the mesh are connected to screw terminals at the back of the power supply. The copper box is isolated from the electrostatic discharge protection mat on the desk. Inside the box, the PCB is also isolated from the box. The PCB features different connectors and measurement points for both the v3.1 and all later hybrids with L-shaped tails.

For the actual measurement, the hybrid under test is placed on the PCB in the closed box and powered with 10.5 V. No VTRx or VTRx$^+$ module is connected. All tested Service Hybrids are assembled with a shield to avoid radiated pick-up in the probe. The ground difference is measured with AC coupling by the probe and transferred to the oscilloscope, where a FFT is performed. A total of 128 million sample points are

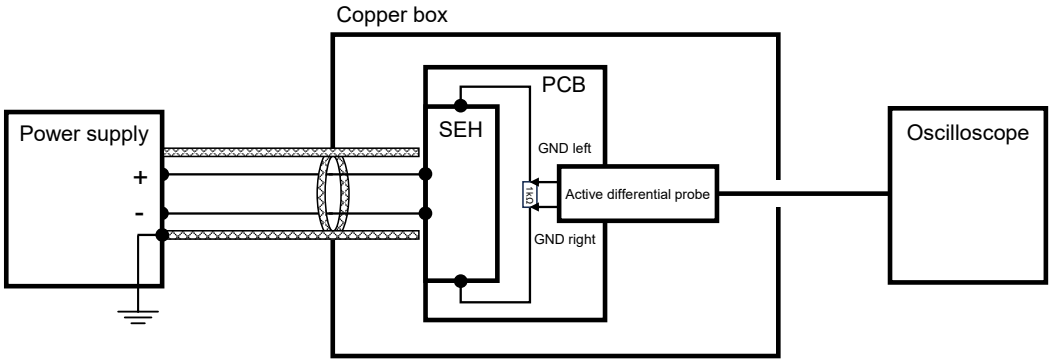

Figure 7.7: Schematic illustration of the concept for the differential measurement of the Service Hybrid grounds on both sides.

collected at 40 GS/s. The FFT is calculated and recorded for frequencies between 0 and 500 MHz and with a step size of 0.6 kHz.

Data are collected for a v3.1 prototype hybrid, which hosts a FEAST2 and a LTC3412A ASIC in the first and second DC-DC converter stage, and also for a final prototype hybrid with the latest bPOL chip versions bPOL12V_v6 and bPOL2V5_v3.3. A regular and a split-plane design hybrid from the kick-off batch, which host the same bPOL versions, are measured as well.

Figure 7.9 shows a background measurement that is obtained without a powered hybrid. The amplitude of the observed voltage difference for each frequency interval is given in µV. To reduce the impact of the width of individual peaks on the interpretation, the amplitude is summed for twenty neighboring bins and plotted at the center of the bins. Compared to the measurements with a hybrid, the background measurement is very clean. Broadcasting from local radio stations gives some peaks in the higher frequency range between 80 and 110 MHz, while broader excesses are present below 15 MHz. Studies with time-varying magnetic fields indicate that the module is most sensitive to frequencies between 10 and 30 MHz [75]. The DC-DC converter switching frequencies lie between 2 and 4 MHz, and their higher harmonics will be present mostly below 35 MHz. This frequency range is shown in the right part of Figure 7.9.

Figures 7.10 and 7.11 compare the measurements for a v3.1 prototype hybrid and a final prototype hybrid up to 150 MHz and 35 MHz, respectively. The two observed spectra are quite different. Based on the DC-DC converter activity, a large peak for the FEAST2 at about 2.0 MHz and a smaller one for the commercial DC-DC converter at 2.4 MHz is present in the case of the v3.1 prototype hybrid. The higher harmonics are also visible but decrease in size. Prominent features that appear not related to the DC-DC converters are large peaks at multiples of about 33.3 MHz.

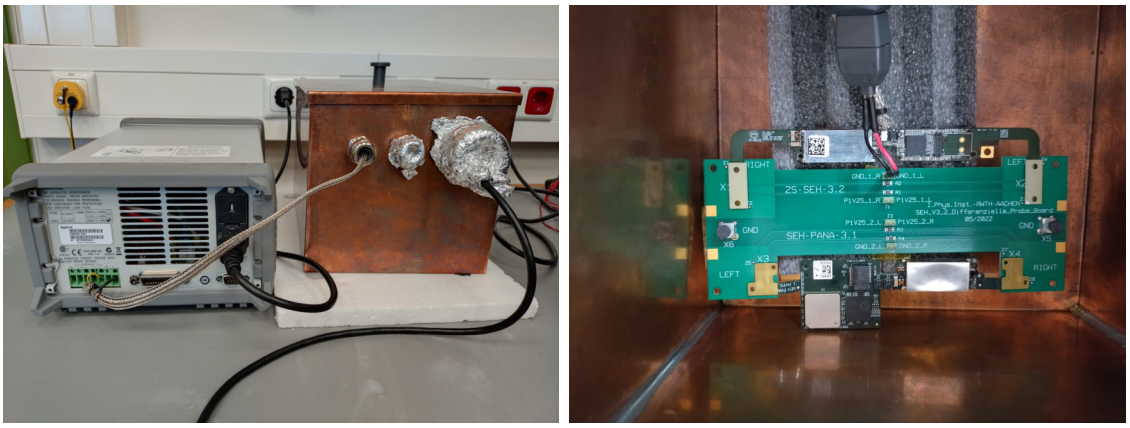

Figure 7.8: Pictures of the setup to measure the difference in ground potential of both Service Hybrid sides. Left: backside of the power supply with the connection to the SEH supply voltage. It also shows the copper box and interface openings. Right: look inside the copper box with the PCB, plugged Service Hybrids, and the differential probe. Only the upper hybrid is connected to the power supply and the probe.

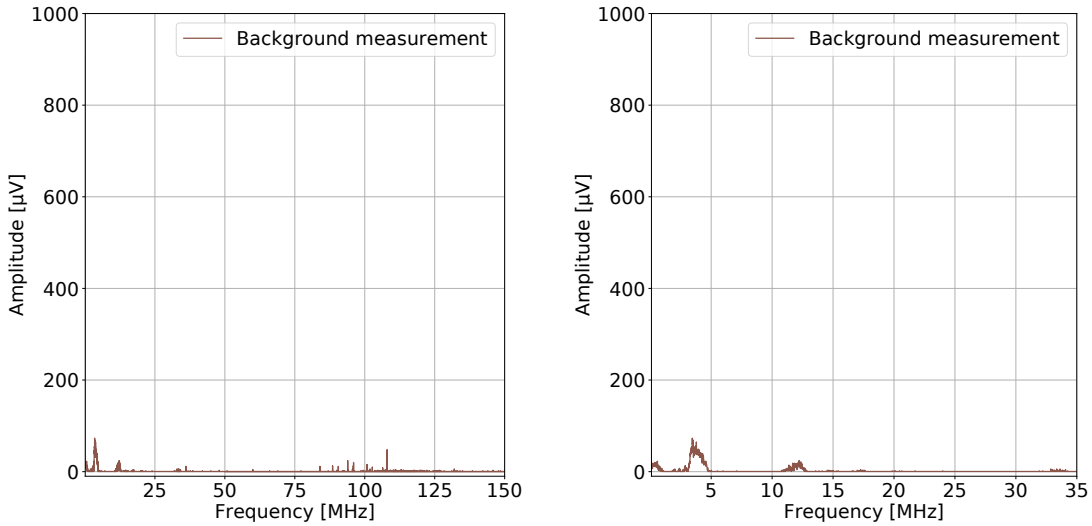

Figure 7.9: Background measurement for the study of the difference in ground potential of both Service Hybrid sides. The plot on the right shows a zoom to the area of the first DC-DC converter harmonics and a sensitive range of the front-ends.

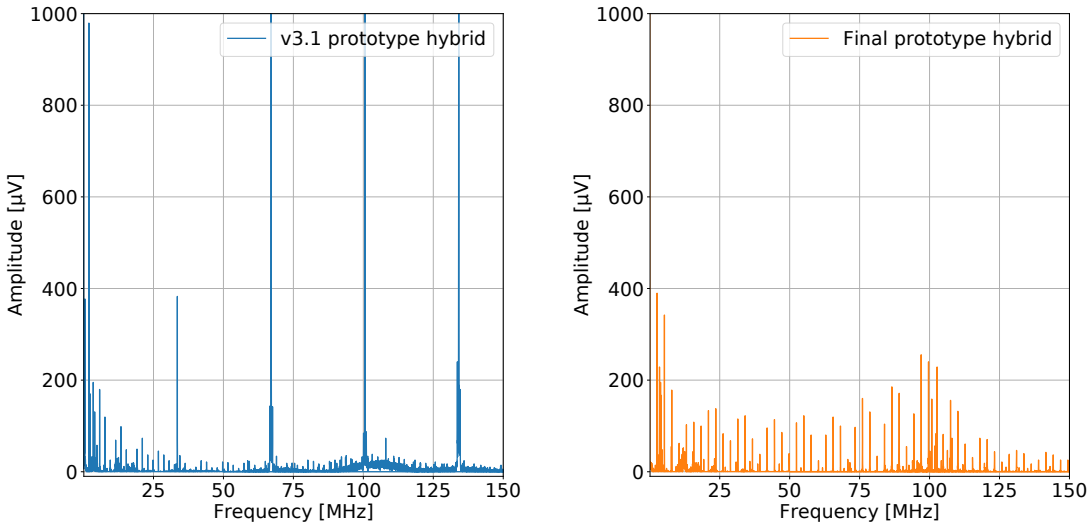

Figure 7.10: Measurement of the difference in ground potential of both sides for two different Service Hybrid generations. Left: v3.1 prototype hybrid. To improve the readability, the top of the peaks at about 66, 100, and 133 MHz with their height of about 3680, 1120, and 2060 µV are not shown. Right: final prototype hybrid.

It is possible that these are due to the unlocked GBTx and GBT-SCA ASIC on those hybrids. The main switching peaks for the final prototype have initially a smaller size than those of the v3.1 prototype. However, beyond the second or third higher harmonics they are higher and remain clearly visible up to > 100 MHz. In the presumably sensitive frequency range, the measured amplitudes for the final prototype are higher, which could constitute a higher module noise.

The comparison of the data for the two different kick-off variants is shown in Figure 7.12 and on the zoomed frequency scale in Figure 7.13. The regular design shows a significant improvement compared to the final prototype design. The main switching peaks have a smaller but comparable size. While still visible, the higher harmonics are well suppressed. Out of the measured hybrids, the split-plane design presents the spectrum with the most noise contributions. The bPOL12V's switching peak is comparable in height to the FEAST2's and all overtones are clearly visible beyond 100 MHz. The overtones' heights are an order of magnitude higher than with the other hybrids. These observations align well with measured noise levels on the kick-off modules, where the regular design performs significantly better than the split-plane design [117]. In the split-plane design, both the right- and left-side output filters of the bPOL2V5 are connected to the noisy ground as shown in Figure 7.4. This creates an asymmetric path for the 1.25 V rail and the ground connection, especially on the right hybrid side. Radiated noise can also couple into layer 4. In the regular design in Figure 7.4, the induction loop is closed by vias between layer 2 and

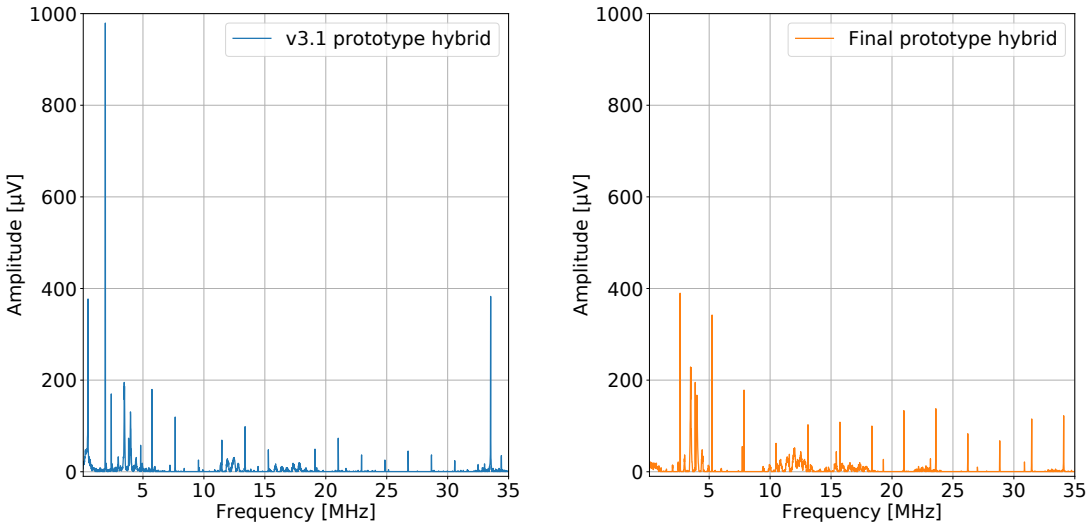

Figure 7.11: Measurement of the difference in ground potential of both sides for two different Service Hybrid generations. Left: v3.1 prototype hybrid. Right: final prototype hybrid. This zoomed view of the Figure 7.10 data shows the frequency region where the 2S module is susceptible to injected magnetic fields.

layer 4 under the bPOLs. The split-plane design forms a large loop, which includes the whole sensors and the Front-End Hybrids.

7.3 Production kick-off conclusions

Two different PCB variants of the Service Hybrid were manufactured as part of the hybrid production kick-off. The functional testing was performed at CERN and for the first time exclusively with the serial crate-based system presented in Section 1.6. The test system encountered some technical difficulties when testing multiple hybrids in series that affected the stability and delayed the qualification. Missing features that improve the presentation and interpretation of the results were identified. The hybrids were not tested at different temperatures in the climatic chamber. The lessons from the kick-off testing were used to improve the test system further and prepare it for the pre-production hybrids.

The Service Hybrid yield suffered most likely from an assembly error that necessitated rework. The yield of the regular design hybrids was 53%, while no split-plane design hybrids suffered damage from the rework. Without the rework, the hybrid would not have been usable. In addition, a design error in the PCB design and incorrectly mounted components made it difficult to judge the design and hybrid

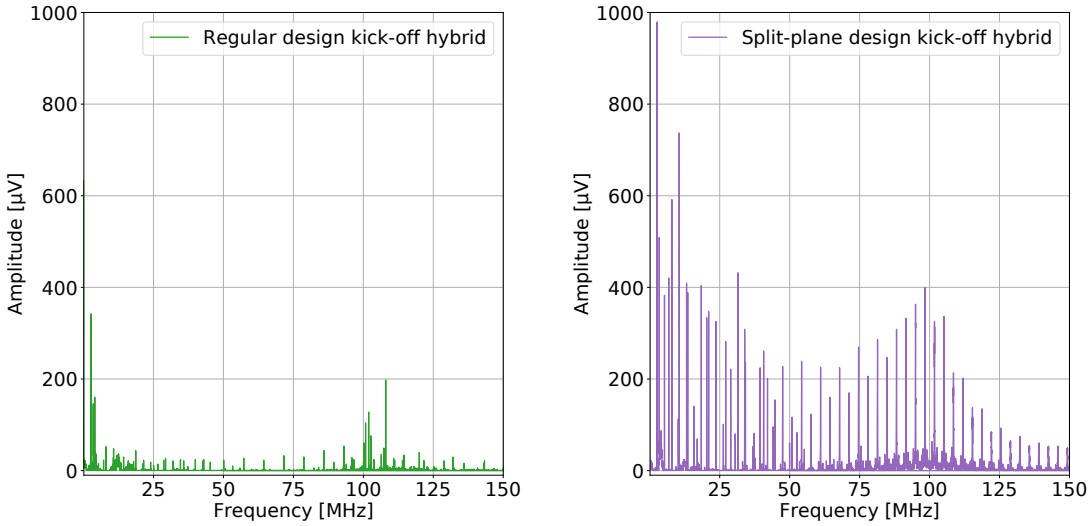

Figure 7.12: Measurement of the difference in ground potential of both sides for two different kick-off Service Hybrids. Left: a hybrid with the regular design. Right: a hybrid with the split-plane design.

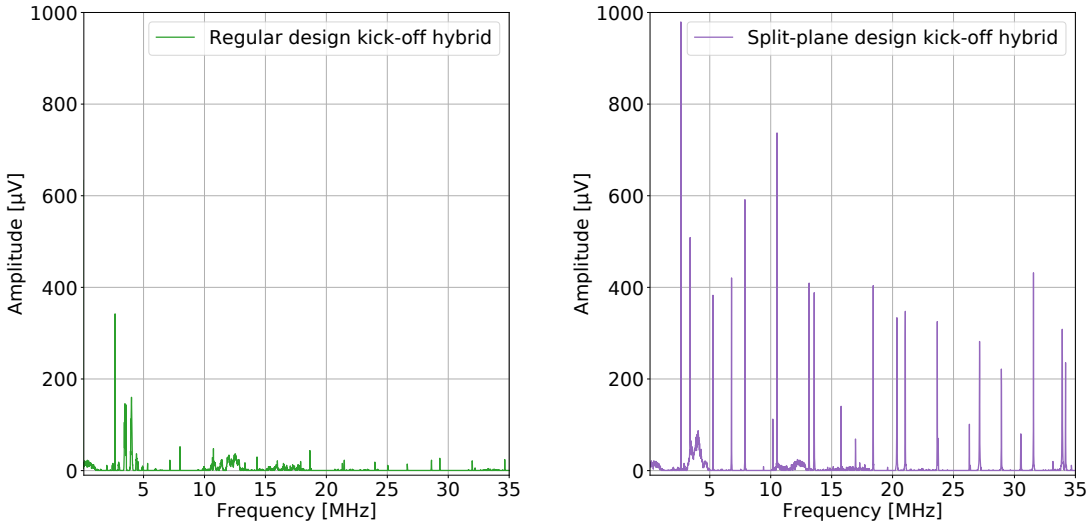

Figure 7.13: Measurement of the difference in ground potential of both sides for two different kick-off Service Hybrids. Left: a hybrid with the regular design. Right: a hybrid with the split-plane design. This zoomed view of the Figure 7.12 data shows the frequency region where the 2S module is susceptible to injected magnetic fields.

quality in its entirety. Only a small number of hybrids of the regular design version could be used for module prototyping.

However, the regular design version performed significantly better in terms of module noise. Direct measurements also favored the regular design, which became the baseline for the production. Additional hybrids with this design were required before the series production. These hybrids and their performance are discussed in the next chapter.

8 Pre-series Service Hybrids

After the production of the kick-off hybrids, the collaboration faced a couple of tasks. The regular SEH design proved favorable in terms of noise performance and was selected for the series production of the hybrid flexes. However, the yield after assembly and un- and resoldering of DC-DC converter shields was unsatisfactory and a problem with the general bPOL2V5 assembly could not be fully excluded. Lastly, because they were set to the FEC12 forward error correction mode, the SEHs could not be used to read out stubs on a full module, which made the kick-off modules unusable as detector stand-ins.

Instead of waiting for the first series of production flexes and their assembly, left-over kick-off flexes are used to evaluate the measures that address those challenges. It is possible to correct the forward error correction mode on the level of the PCB before the lpGBT is assembled. The contractor was tasked to perform this patch before the assembly.

Based on the experience of the kick-off batch, only regular design flexes are used. The assembly is performed in the same way and with the same parts as foreseen for the series production. The resulting hybrids are therefore considered part of the hybrid production's "pre-series" and will be referred to as such. They, in particular, use bPOL12V chips that were binned in U_{ref}. The wrongly placed solder mask on the grounding flap is of course still present with those flexes, but can be patched by the user.

A total of 20 assembled hybrids were delivered to CERN in March 2024. All hybrids passed the visual inspection without noteworthy observations. Following the experiences of the kick-off hybrid testing, the crate-based test system and its stability had been improved. As the developer of the test routines, I supported the testing in order to streamline the operation and make the process ready for the series production. The results are presented in the following section.

8.1 Characterization of first pre-series hybrids

For the functional testing at CERN, all hybrids are mounted on 2S-SEH test cards and installed in a crate. The crate is positioned inside a climatic chamber (Climats EXCAL^2 4025 [55]). Figure 8.1 (left) shows twelve test cards with mounted hybrids inside the crate. The individual optical fibers for each card and the bias voltage cable

Figure 8.1: Crate-based Service Hybrid test setup at CERN. Left: fully installed crate with 2S-SEH test cards inside the climatic chamber. Right: rack adjacent to the climatic chamber with auxiliary devices.

between the power supply and the left-most card are connected. The bias voltage is distributed via a daisy chain from test card to test card. Figure 8.1 (right) shows the rack adjacent to the climatic chamber. It holds the necessary µTCA crate with the optical and electrical FC7s, the power supplies for the supply and test voltage, and the power supply for the bias voltage. The setup in the climatic chamber is designed for three crates with the top one being dedicated to the testing of 2S-SEHs and PS-ROHs. The tests are carried out in series for all twelve hybrids in a crate by the hybrid testing GUI using the most recent firmware and software versions at +40°C and −35°C. The temperatures are set manually. The bias voltage isolation is tested in parallel to the rest of the testing routines except for the I/V scan and the bias voltage connection check. Table 8.1 summarizes the tested aspects. The table lists the aspects in the order their procedures are executed and provides details on the applied acceptance criteria for the pre-series hybrids and the duration of the procedure. A full test of a hybrid takes about four and a half minutes. The procedures were discussed in detail in Section 4.1.3.

Before they are installed, test cards and hybrids are associated with one another by scanning their attached bar codes. The initial grading after a test is performed based on cuts acquired during the characterization of earlier generations. The cuts target to spot severe assembly errors. Results are uploaded to the central CMS construction database. The data are also stored for in-depth offline analysis. All hybrids are tested at least once at both temperature levels. Figure 8.2 shows temperature and dew point measurements of three typical climatic chamber cycles [56].

All 20 tested hybrids are found to be fully functional at +40°C. Some failed due to tight cuts in the initial grading, which is discussed in detail below. After review, the hybrids are considered to have successfully passed.

Tested aspect	Acceptance criteria	Duration
Short circuit check	No short detected	3 s
Service Hybrid I/V behavior	Detection of turn-on and turn-off but no cuts applied for pre-series	47 s
Optical link and configuration	Successful configuration	11 s
E-link receivers	No detected errors for any line and pattern	25 s
Reset lines	Measured GPIO output level within ±100 mV of nominal voltage for all lines and levels	4 s
I^2C controllers	No transaction errors for all controllers	30 s
ADC lines	No cuts applied for pre-series	36 s
Clocks	No detected errors for any line and setting	1 s
E-link drivers	No detected errors for any line and pattern	2 s
Front-End Hybrid supply voltage	$1.1\,\mathrm{V} < U_{\mathrm{out}}^{\mathrm{bPOL2V5}} < 1.4\,V$, $2.4\,\mathrm{V} < U_{\mathrm{out}}^{\mathrm{bPOL12V}} < 2.7\,V$, and DC-DC conversion efficiency within acceptance range (Section 8.1.2)	48 s
Sensor bias voltage filter leakage current	Leakage current within acceptance range (Section 8.1.5)	160 s
Sensor bias voltage connection	Measurement of applied voltage on test card	50 s

Table 8.1: Overview of the tested aspects, their acceptance criteria, and the required time for the qualification of the pre-series Service Hybrids in the order they are executed. The measurement of the sensor bias voltage leakage current is performed in parallel to other tests.

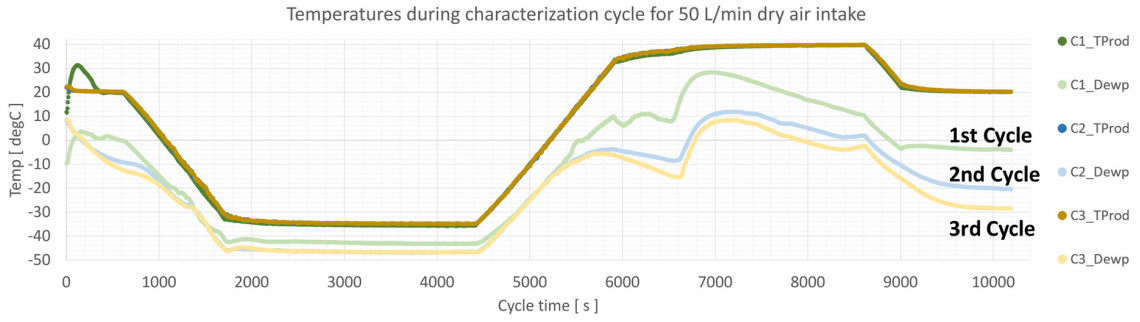

Figure 8.2: Example of three sequential thermal cycles (C1-C3) executed with the same dry air intake configuration. Measurements of the dew point (Dewp) and temperature (TProd) are done by sensors within the climate chamber [56].

A total of 17 hybrids passed all tests at −35°C. Three hybrids failed at least once during the locking phase of the optical link, which prevented further testing. Within the available time for testing, two out of these three hybrids were successfully retested at cold. The failed locking of the optical link can have many different causes and many of them are not related to the hybrids themselves. It was also the first cold run for the VTRx$^+$ modules and other optical components inside the chamber. The hybrids in question are retained at CERN, where it is planned to perform further studies of the test system performance with a focus on the optical components. The preceding I/V curves and the successful test runs show no hint of a hybrid-related problem.

In light of the occasionally low yield for previous hybrid generations, a successful test rate of 100% at +40°C and of at least 85% at −35°C with no signs of assembly-related failures is an encouraging result.

The test system and software showed a stable performance. Minor adjustments were made to improve the safety and usability by non-expert users. The importance of the quality control outlined in Section 4.3 was underlined by the unresponsiveness of two used test cards at −35°C. They were brought to CERN in an initial delivery and did not undergo the procedure.

In the following, the most important measured quantities are presented and compared for both temperature levels. The results can be used to judge the expected behavior for the later production.

8.1.1 DC-DC converter output voltages

Figure 8.3 shows the measurement of the bPOL12V output voltage for all measured hybrids at a given temperature. As with the prototypes, this measurement is performed without any additional load applied to the hybrids. The one hybrid that could not be tested at −35°C is missing for the corresponding measurement shown in blue.

The pre-series hybrids display an almost flat output voltage distribution with a mean of 2.516 V and 2.554 V at −35°C and +40°C, respectively. The standard deviation is similar in both cases (0.008 V and 0.009 V). Compared to the prototype hybrids the spread is significantly smaller (down from 0.024 V). This shows the impact of the binned U_{ref} and the benefit of using 0.1% resistors. The expectation of $\overline{U_{\mathrm{out}}} = 2.555\,\mathrm{V}$ for the selected $R_2 = 165\,\mathrm{k}\Omega$ at +40°C is nicely met. The measurement also reaffirms the assumption that the temperature influence presents itself as a DC offset. A shift of about −0.02%/K is observed. If a comparable distribution is achieved in the production, the bPOL2V5 and the VTRx$^+$ module can be operated safely in the 2S module.

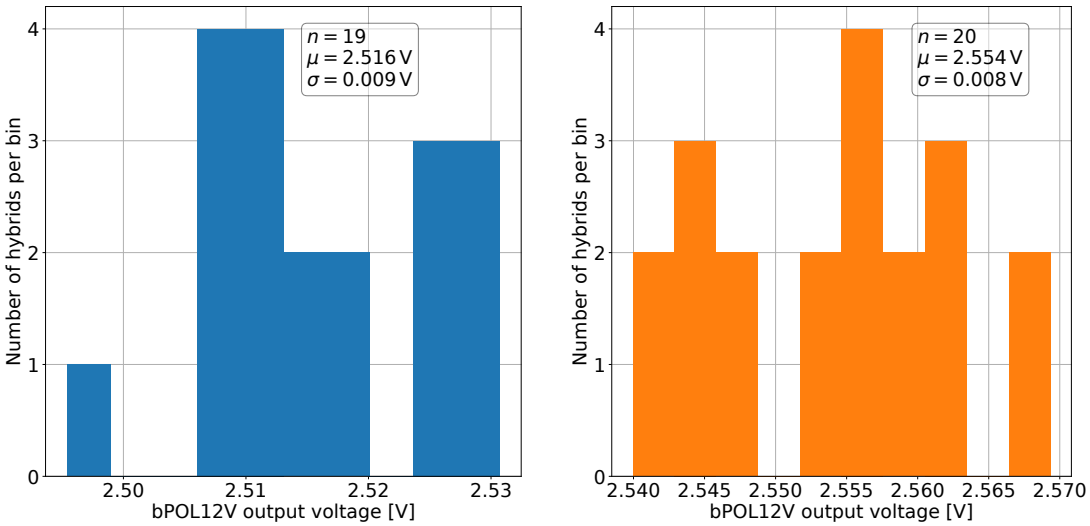

Figure 8.3: Output voltage of the bPOL12V measured on pre-series Service Hybrids. Left: histogram for all hybrids tested at −35°C. Right: histogram for all hybrids tested at +40°C.

Figure 8.4 (left) shows all measurements of the bPOL12V output voltage per pre-series hybrid. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test. Some hybrids have multiple measurements at the same climatic chamber set point. Additional reference runs and cross-checks are performed at other temperatures and are also included in the overview. The plot illustrates a certain spread in the output voltage temperature shift across different hybrids. The measurement-to-measurement variations for the same hybrid at a given temperature are small (measurements for the same hybrid are performed with the same test card).

The analogous results for the bPOL2V5 output voltage are given in Figure 8.5 and Figure 8.4 (right). The measurement yields a narrow distribution at both temperatures with a standard deviation of 4 to 5 mV. The temperature drift of the mean output voltage from 1.268 to 1.243 V is equivalent to −0.026%/K, which is compatible with experiences from the prototypes. The right plot in Figure 8.4 illustrates that the drift is not constant across all hybrids, which matches the observation of the study with two prototypes, too.

In terms of a safe module operation, the prototype measurements of the bPOL2V5 output voltage raised no concerns. The pre-series hybrid results back up this conclusion. The pre-production testing uses the coarse limits of $1.1\,\mathrm{V} < U_{\mathrm{out}} < 1.4\,\mathrm{V}$ for the bPOL2V5 and $2.4\,\mathrm{V} < U_{\mathrm{out}} < 2.7\,\mathrm{V}$ for the bPOL12V output voltage. Those cuts are only suited to spot severe problems. An additional offline analysis and grading are required for the series production.

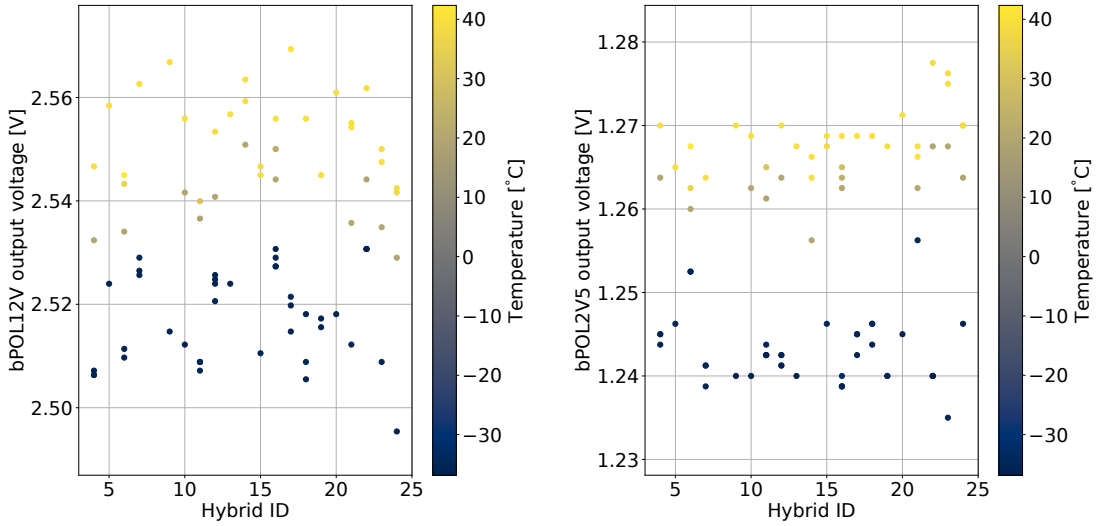

Figure 8.4: Left: output voltage of the bPOL12V measured per pre-series Service Hybrid. Right: output voltage of the bPOL2V5 measured per pre-series Service Hybrids. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

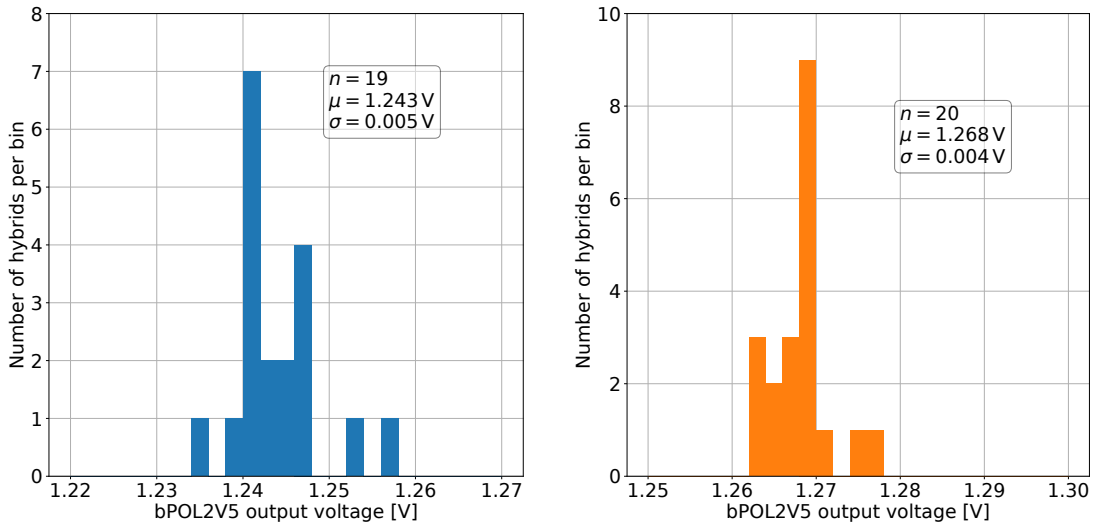

Figure 8.5: Output voltage of the bPOL2V5 measured on pre-series Service Hybrids. Left: histogram for all hybrids tested at −35°C. Right: histogram for all hybrids tested at +40°C.

So far, we considered only the case without an additional load current. The available data also allow quantification of the output voltage that the left and right FEHs receive based on their current consumption. It is mainly influenced by the DC-DC converter intrinsic regulation and the resistive PCB properties. Both effects cannot be disentangled in the measurement. If the variation is solely caused by the intrinsic regulation, it can be quantified by the load regulation R, which is given by the voltage difference at minimal load I_{min} and with full load I_{max} divided by the voltage at full load given in percent:

$$R[\%] = \frac{U(I_{\mathrm{min}}) - U(I_{\mathrm{max}})}{U(I_{\mathrm{max}})} \cdot 100 \quad . \tag{8.1}$$

The full load of the bPOL2V5 is 3 A. The SEH system features the complication that current can be drawn from both sides. The worst case is considered, which means the voltages are measured on each side with load current drawn from both sides symmetrically. This case also covers the use case in the module, where two very similar FEHs are the consumers. The results for the left and right side at −35°C are shown in Figure 8.6. The average observed load regulation for the 1.25 V output voltage rail of the bPOL2V5 at 100% load is 6.1% for the left side and 3.5% for the right side. The difference between the sides is due to the higher resistance along the longer current path.

The temperature influence on this quantity is shown in Figure 8.7. The two sides behave differently with changes of 2 and 1%pt. for the whole temperature range. If only the DC-DC converter ASICs' load regulation would be temperature dependent, we would not expect a difference. However, the lpGBT's and VTRx$^+$ module's change in current consumption and changing resistances play a role, too.

When a load is only applied to a single side we can interpret the resulting voltage drop as a resistance and can compare it to the results obtained with the prototypes. Table 8.2 summarizes the average results of the pre-series Service Hybrids. For the batch 2 prototype hybrids, we found 55 and 25 mΩ at room temperature. The new result of the pre-series hybrids at +40°C is fairly consistent with the old results if we take the different flexes and temperatures into account. The resistance $R(t)$ of copper changes with temperature according to

$$R(t) = R_0(1 + \alpha t) \quad , \tag{8.2}$$

where R_0 is the resistance at 0°C and α is the temperature coefficient. For copper, the value of α is 0.00393 K^{-1} [118]. We obtain a resistance change by a factor of 1.34 when the temperature is increased from −35°C to +40°C. The factor is sufficient to explain the observed changes on both sides and the intrinsic regulation can be disregarded. This is compatible with the load regulation measured by the chip designers, which is smaller than 1‰ [17].

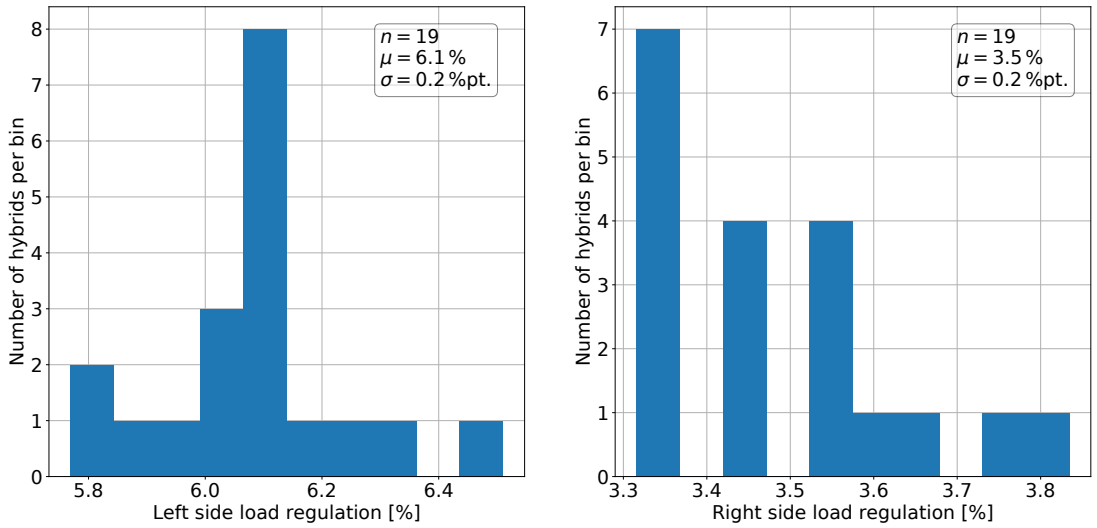

Figure 8.6: Load regulation of the bPOL2V5 output voltage measured on pre-series Service Hybrids. The numbers are obtained at −35°C and for 100% load (1.5 A per side). The left and right plots are for the left and right output, respectively.

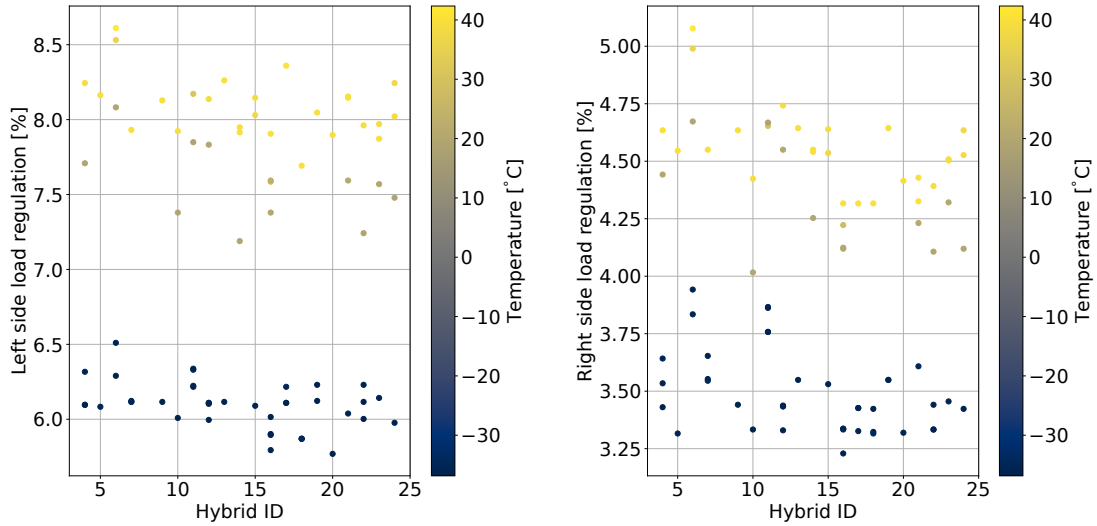

Figure 8.7: Load regulation of the bPOL2V5 output voltage at 100% load (1.5 A per side) measured per pre-series Service Hybrid. The left and right plots are for the left and right output, respectively. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

	Resistance [mΩ]	
Temperature	Left side	Right side
+40°C	52	26
−35°C	39	19
Relative change	1.33	1.37
Resistance change copper	1.34	

Table 8.2: Reinterpretation of the load regulation as a resistance. The table shows the average values across the pre-series Service Hybrids at the two temperature levels.

8.1.2 DC-DC conversion efficiency

The ratio of the output to input power is measured in a similar fashion as for the prototype hybrids. However, the current step size is adjusted to 0.5 A to cover the module working points better. This is done in sync with the PS-POH testing routine.

Figure 8.8 shows the two temperature sets of results for all tested hybrids. It also shows in purple the old limit of $\pm\,(\eta(I,T)\cdot 2\% + 1\,\%\text{pt.})$ based on Equation 5.2. A few hybrids narrowly fail this strict criterion at low currents. As the limits are supposed to spot critical assembly errors, it is decided to use the more relaxed limit of $\pm\,(\eta(I,T)\cdot 5\% + 1\,\%\text{pt.})$. After the functional testing, offline grading can be performed. All measurements in a set have the same basic trend and are fairly close together. The shape of the ratio curves was discussed in detail in Section 5.2.2. The ratio improves at lower temperatures, as illustrated again by Figure 8.9. Only the module working point of 1 A current consumption per side is shown. Due to differences in the routing of the power in the SEH PCB layout, higher voltage drops are present and the ratio is 2%pt. lower compared to what was measured for the final prototypes. This was already observed during the kick-off testing and accounted for in the limits.

8.1.3 bPOL12V turn-on and turn-off voltage

The SEH response during an I/V scan was discussed in detail in Section 5.2.3. In the crate-based system, a minimal voltage of 5 V is used, as the test cards are deselected for lower test voltages. The scanned voltage points are optimized based on the expected result range, the desired measurement accuracy, and the available time. Most hybrids are initially tested with the voltage set [..., 8.0 V, 9.0 V, 9.1 V, ...]. As can be seen in Figure 8.10 (left) this choice causes a loss of sensitivity for

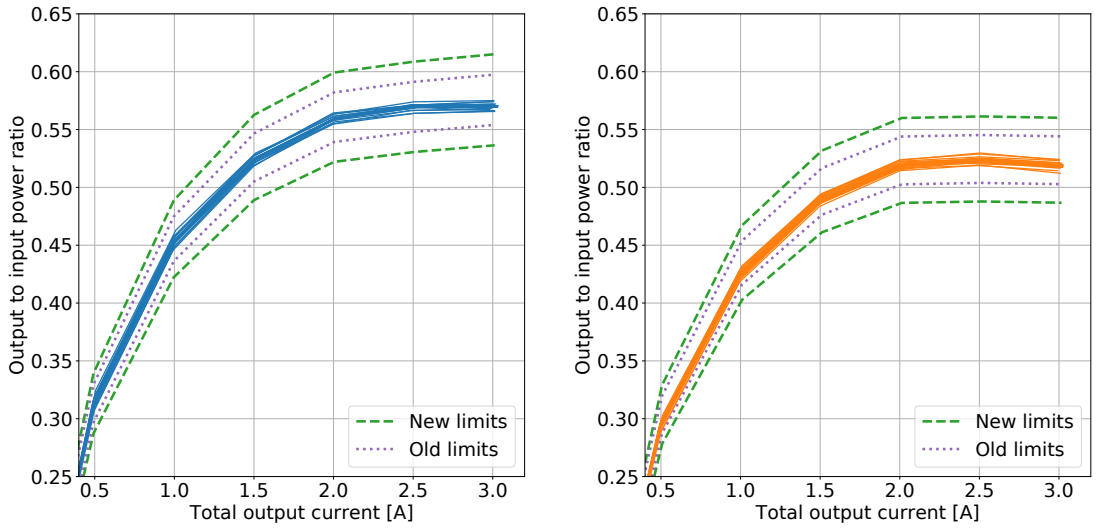

Figure 8.8: Ratio of the output to input power versus load current on pre-series Service Hybrids. Left: all hybrids measured at −35°C. Right: all hybrids measured at +40°C. Two different limits for the performance are shown. See text for details.

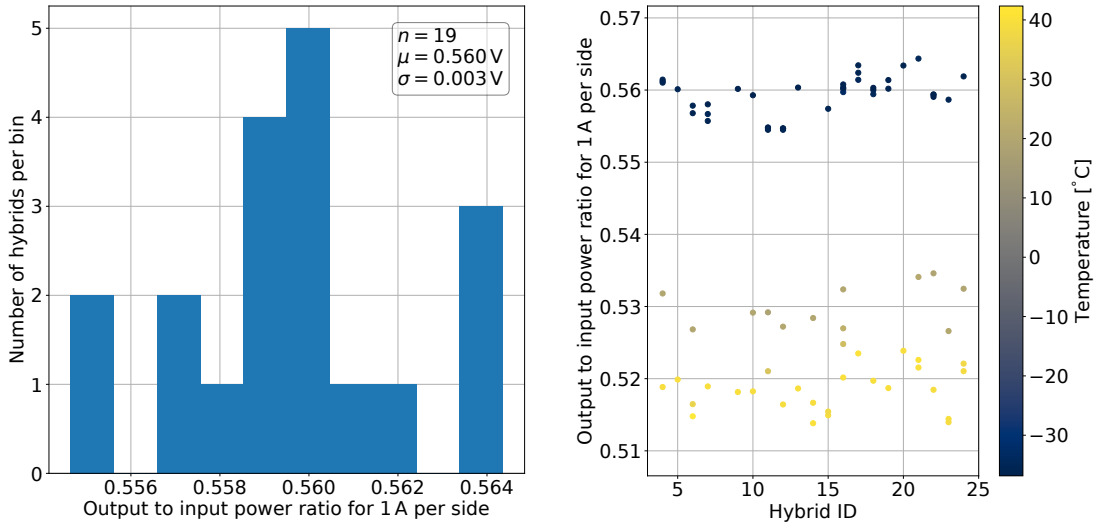

Figure 8.9: Ratio of the output to input power measured on pre-series Service Hybrids for a load of 1 A per side. Left: histogram for all hybrids tested at −35°C. Right: measurements per hybrid. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

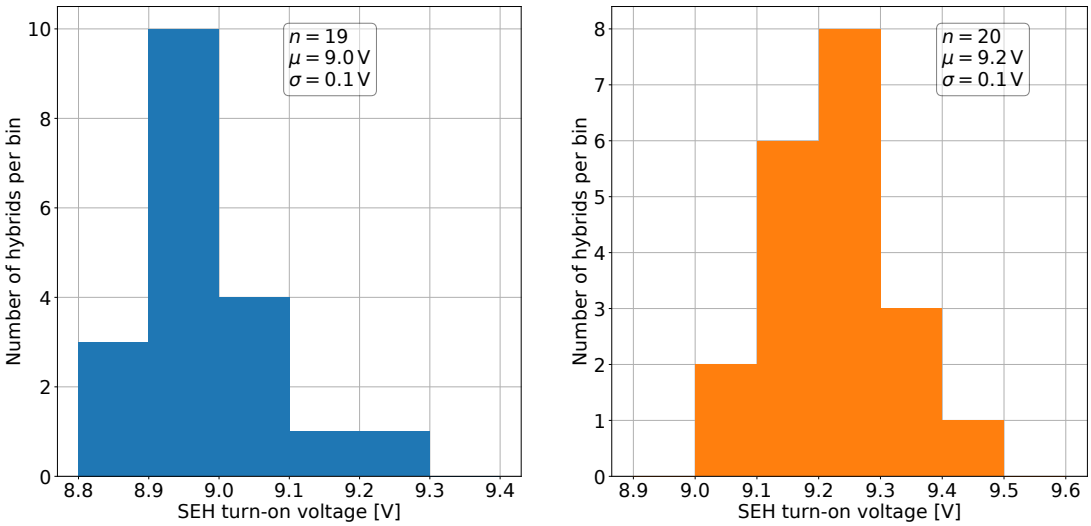

Figure 8.10: Turn-on voltage of the bPOL12V measured on pre-series Service Hybrids. Left: histogram for all hybrids tested at −35°C. See text for additional discussion. Right: histogram for all hybrids tested at +40°C.

turn-on voltages below 9.0 V[1], which are observed for the −35°C measurements. For the measurements at +40°C no turn-on below 9 V is observed, so it is not affected[2]. The voltage vector is corrected to include [8.7 V, 8.8 V, 8.9 V]. Some hybrids are tested after the change and some are retested, as can be seen in Figure 8.11 (left). Only the latest measurement per hybrid is included in the histogram in Figure 8.10.

The mean turn-on voltage shifts by 200 mV and is 9.0 V at −35°C and 9.2 V at +40°C. The temperature-dependent shift, the Gaussian distribution, and the spread in the hybrid distribution are consistent with the observations from testing the final prototype hybrids. The overall lower turn-on setting by choosing $R_2 = 1070\,\Omega$ corrects the too-high set-point of the prototypes and accounts for the expected threshold increase with radiation. For the distribution in cold, one has to account for the bias for $U_{\text{in}} < 9.0$ V, which distorts the shape.

The turn-off voltage is measured with a step size of 0.2 V. Figure 8.11 shows the correlation between the measured turn-on and turn-off voltage at +40°C as a two-dimensional histogram. We observe a clear correlation. Most hybrids turn off below 5.8 V. For −35°C we observe a down shift of the threshold by 200 mV, too.

[1]For a measurement with a binning of [..., 8.0 V, 9.0 V, 9.1 V, ...] a measured voltage of about 9.0 V corresponds to a turn-on voltage between 8.0 and 9.0 V. With a binning of [..., 8.9 V, 9.0 V, 9.1 V, ...] it corresponds to a turn-on between 8.9 and 9.0 V.

[2]The lowest observed turn-on occurs during a voltage increase from 9.0 to 9.1 V

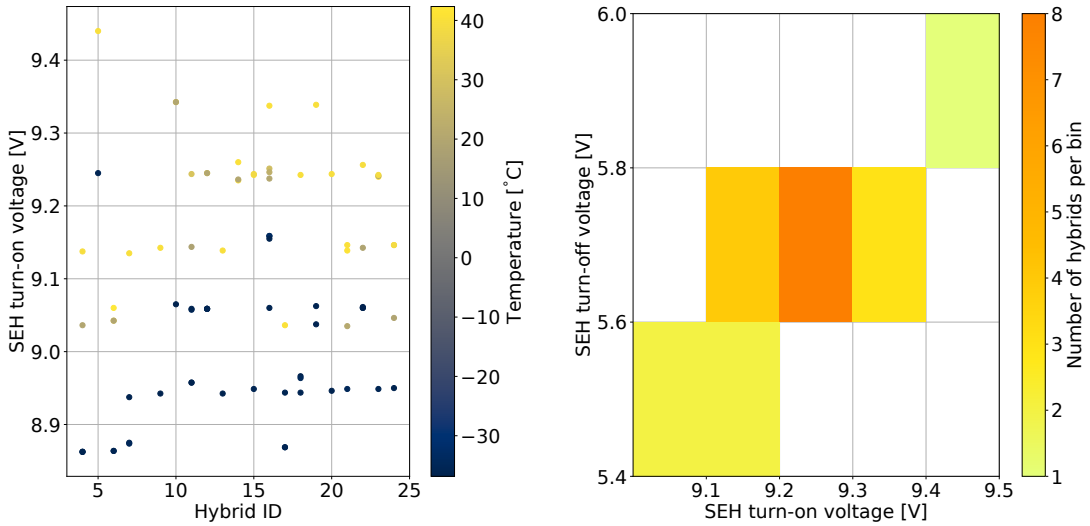

Figure 8.11: Left: turn-on voltage of the bPOL12V measured per hybrid. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test. Right: correlation between the turn-on and turn-off voltage of the bPOL12V for pre-series Service Hybrids measured at +40°C.

8.1.4 Slow control monitoring

The digital interfaces (high-speed optical link, e-links, clock, I^2C) of the lpGBT_v1 have been well demonstrated in the hybrid test procedures of the final Service Hybrid prototypes, the kick-off hybrids, and the various ROH versions. The optical readout of modules has also been demonstrated successfully. Moreover, no failures related to the lpGBT have been observed in the pre-series delivery, other than the failed link locking described above. The PCB level patch back to the FEC5 forward error correction mode is successful and the affected stub link is fully recovered.

For various reasons, the full foreseen slow control features of the Service Hybrid via the lpGBT (Section 1.5.2.3) have not been demonstrated before. This is addressed in the pre-series.

During assembly, a lpGBT is not associated with a hybrid. However, it can be identified by its ID that is stored redundantly in the e-fuses [14]. Figure 8.12 (left) demonstrates that this ID can consistently be retrieved via the optical readout. With this information, a lpGBT can be linked to an identifiable Service Hybrid or module.

The identification of the lpGBT is relevant to utilize calibration data for the lpGBT's ADC. They are provided in an open format for all production-grade ASICs. In the hybrid test system, an additional tuning of the reference voltage based on a test voltage provided by the test card is performed. The results are shown in Figure 8.12

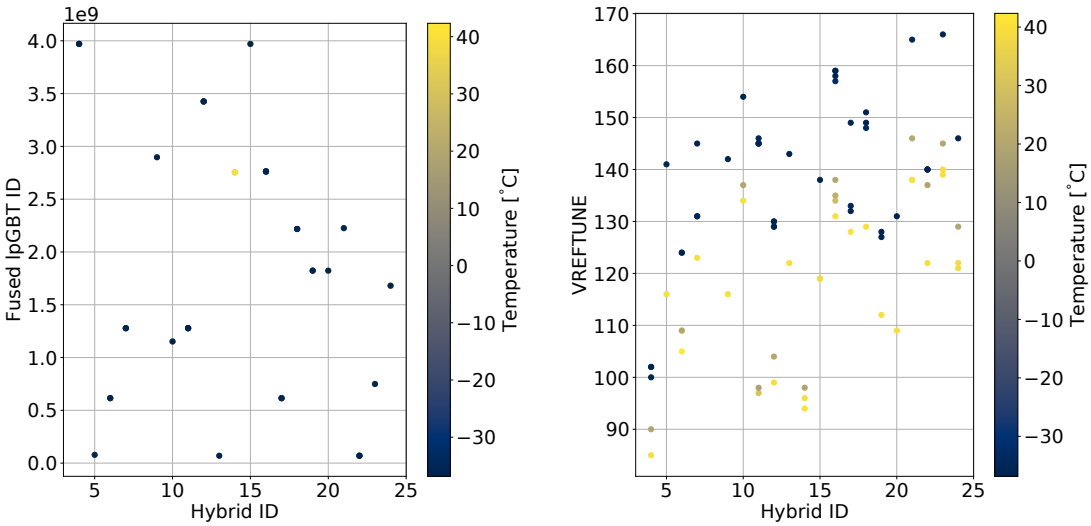

Figure 8.12: Left: lpGBT ID read from the e-fuses per pre-series Service Hybrid. Right: result of the lpGBT's voltage reference tuning per pre-series Service Hybrid. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

(right) in terms of the value that needs to be programmed to a certain lpGBT register. The temperature dependence that necessitates the tuning is clearly visible.

After tuning, the ADC can be used to monitor the input voltage of the hybrid. Figure 8.13 (left) shows that the voltage is in general within $\pm 1\%$ of the expected value (0.516 V for a power supply output of 10.5 V, not accounting for voltage drops). One hybrid is an outlier at all temperatures but still within tolerances of the voltage divider resistances.

Figure 8.13 (right) shows the measured RSSI voltage used to monitor the VTRx$^+$ module's received signal strength. The resistors within the measurement network have been updated compared to the final prototypes to better cover the sensitive range of the ADC. The plot also shows the conversion to the indicative current on its secondary vertical axis. The conversion is done according to Equation 1.5. An increasing current corresponds to a decreasing monitoring voltage. The responsivity across VTRx$^+$ modules usually varies by $\pm 10\%$. This is consistent with the observed spread in data. Most hybrids also present an increasing current for lower temperatures.

Temperature changes of the DC-DC converters can be monitored via their PTAT voltages. Figure 8.14 shows the obtained measurements for the bPOL12V and bPOL2V5 PTAT. For both signals a clear separation between the two temperature

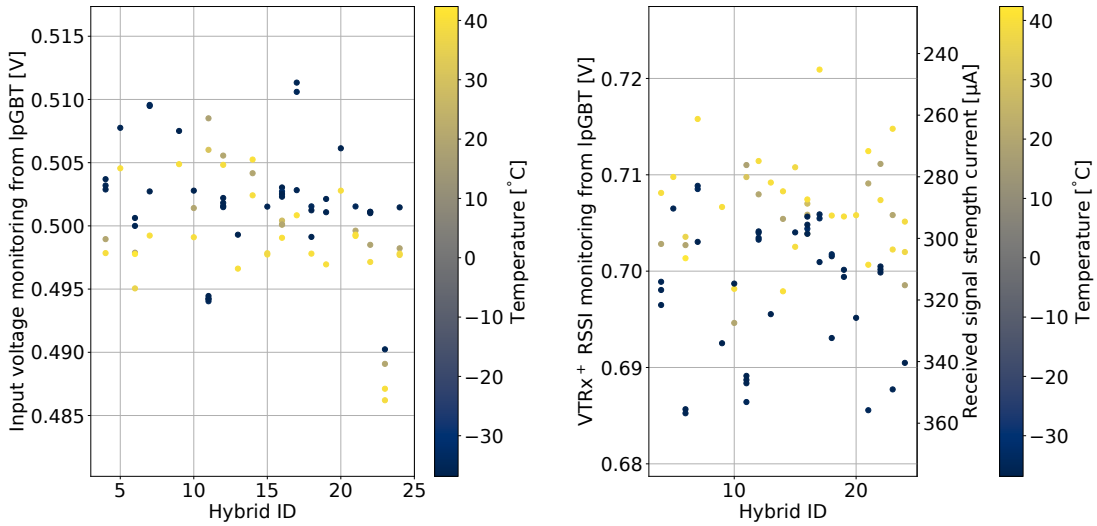

Figure 8.13: Left: input voltage monitoring of the lpGBT per pre-series Service Hybrid. Right: VTRx$^+$ RSSI measurement per pre-series Service Hybrid. The left axis shows the voltage measurement from the lpGBT, and the right axis the conversion into the received signal strength current. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

levels is visible. Figure 8.15 shows the same data but with the value of the −35°C measurement subtracted. In this way, the data can be compared to the expectation. From the ASIC datasheets, we expect a voltage increase of about 3 mV/K for the bPOL12V PTAT. The required voltage divider is already considered. Therefore, an increase of 225 mV is expected for the 75 K temperature difference. This is well represented by the data. Notably, two hybrids show a significantly lower voltage difference than expected. For the bPOL2V5 the corresponding numbers are a voltage increase of 4 mV/K or in total 300 mV. The overall sightly higher observed PTAT voltage at +40°C correlates with higher conversion losses inside the chip, which can explain the observation.

8.1.5 Leakage current of the sensor bias voltage filter

The leakage current of the RC filter for the sensor bias voltage is measured in parallel to the testing of the lpGBT and the DC-DC converters. Only the I/V scan and the bias voltage connection check are decoupled. The parallelization is necessary to limit the test time. In the prototyping phase, the hybrid isolation was tested serially for 150 s. Since the parallel tests take about as long, the obtained results are comparable.

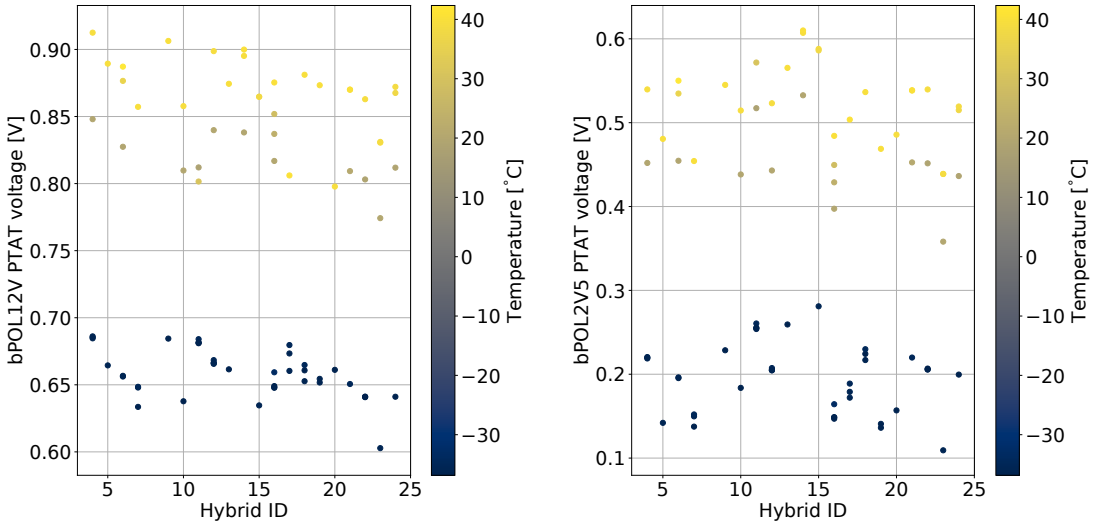

Figure 8.14: PTAT voltages for temperature monitoring of the DC-DC converters per pre-series Service Hybrid. Left: data for the bPOL12V. Right: data for the bPOL2V5. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

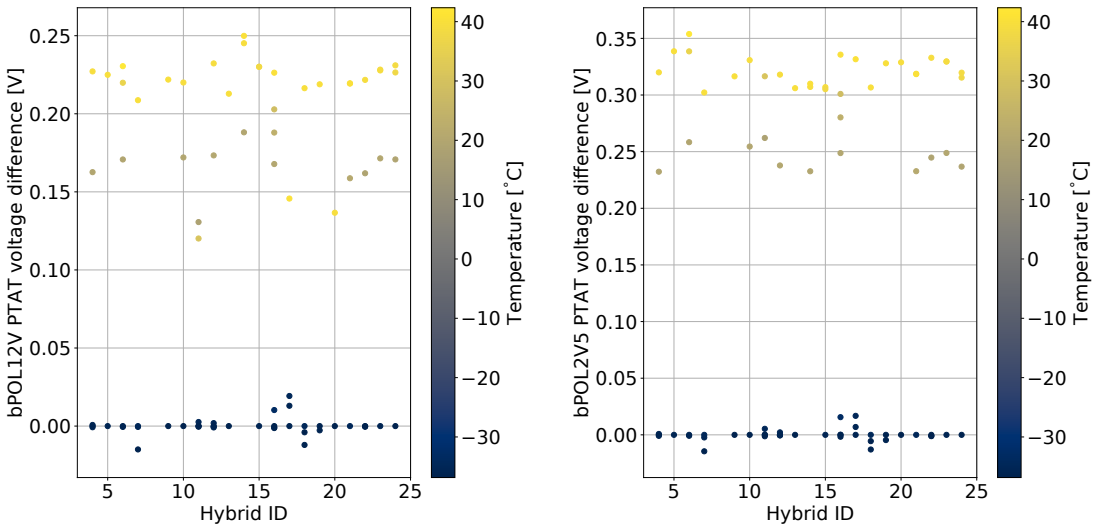

Figure 8.15: PTAT voltages for temperature monitoring of the DC-DC converters per pre-series Service Hybrid. The data are offset for each hybrid by a measurement at −35°C. Left: data for the bPOL12V. Right: data for the bPOL2V5. Each data point represents an individual measurement and the color indicates the temperature measured on the test card at the beginning of a test.

We need to consider the effects of the temperature and the voltage exposure history discussed in Section 5.2.4.

Most hybrids are exposed to the bias voltage for the first time at +40°C. Because of a lack of initial bias voltage exposure results at this temperature with the prototypes, they are initially graded by a criterion based on room temperature measurements. The results are presented in Figure 8.16 (left). Most (17 of 20) show the same behavior. Between 20 and 25 s the current reaches the measurable range and reduces exponentially. Distortions in the exponential waveform are caused by the changes in the load current during the measurement of the output to input power ratio. The original limit is shown as a purple dotted line and is failed by those 17 hybrids. Two hybrids are first measured at room temperature. At +40°C they show an almost flat curve below 100 nA after a few seconds. A new limit is proposed based on the ensemble and is shown as a dashed green line. A single hybrid has a higher leakage current and would also fail the new criterion. The same hybrid is included in a set of measurements at room temperature shown in Figure 8.16 (right). Here, it also has a higher leakage current than the rest. After the testing, the hybrid was inspected. No noticeable problems with the bias filter or its assembly were found. The quantity will be monitored during production. If required, additional actions will be taken. The room temperature plot also contains a measurement of one hybrid, whose first bias voltage exposure was during that measurement. The obtained curve with the exponential decay matches the results from the first exposure of the prototypes (Figure 5.21). The other curves correspond to hybrids with their first bias voltage exposure before the measurement at room temperature.

The results of all hybrids measured at −35°C are summarized in Figure 8.17. As expected from the prototype testing, the current is small (< 30 nA) throughout the measurement duration. No voltage breakdown at the power supply or sudden current increase is observed in any measurement.

8.2 Summary of the first pre-series experience

The first 20 Service Hybrids of the pre-series were assembled. The used flexes are taken from the regular design kick-off production. Pre-series ASICs and components populate the hybrids. The forward error correction mode and assembly errors in the previous versions were corrected.

The contractor successfully delivered hybrids, which passed the visual inspection without any noteworthy observations. The crate-based test system for the hybrid series production was used to test the hybrids. A separate 2S-SEH test card was used for each hybrid and the testing was performed serially with a full crate in a climatic chamber. All hybrids passed the functional testing at +40°C. Few failures at −35°C require additional investigations. The overall yield was 85%. The system's capability to qualify a production lot was demonstrated.

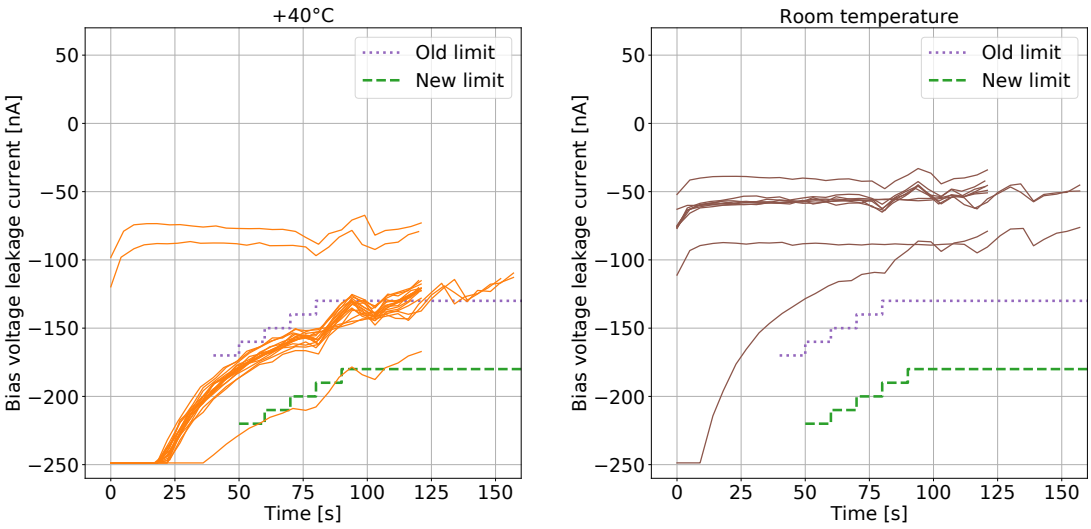

Figure 8.16: Measurement of the leakage current of the sensor bias voltage filter on pre-series Service Hybrids versus time. Left: first measurement of every tested hybrid at +40°C. Right: room temperature measurements for a set of hybrids. The old limit for the leakage current based on room temperature measurements and the new limit based on the +40°C data are plotted.

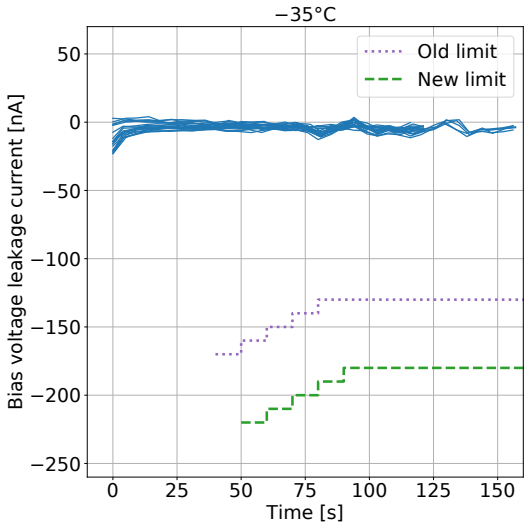

Figure 8.17: Measurement of the leakage current of the sensor bias voltage filter on pre-series Service Hybrids versus time. The first measurement of every tested hybrid at −35°C, the old limit for the leakage current based on room temperature measurements, and the new limit based on the +40°C data are plotted.

The testing data provide valuable input to predict the detector performance. The component selection and their respective values are validated.

The research and development phase and the design of the Outer Tracker hybrids concluded with the production of the presented hybrids. The studies performed in the scope of this thesis show that the design can be used in the detector. The CMS Collaboration will continue with the production of additional pre-production hybrids. Those hybrids use newly produced flexes and can be used for the construction of the Outer Tracker. The number of produced hybrids per batch will increase significantly ($\mathcal{O}(100)$) and a successful production relies on the test system.

9 Summary and outlook

Starting in 2026, the LHC will be upgraded to achieve an instantaneous luminosity of up to $5 \times 10^{34}\,\mathrm{cm^{-2}s^{-1}}$. The CMS detector requires a series of upgrades to maintain and improve its performance in collecting LHC collision data, collectively known as the Phase-2 upgrades. A major aspect of these upgrades is the complete replacement of the CMS strip tracker with the Outer Tracker. This thesis covers the development of the Service Hybrid for the 2S module of the Phase-2 Outer Tracker.

The Outer Tracker consists of modules with two closely spaced sensors. The 2S module has two strip sensors and is used in the outer parts of the Outer Tracker. The correlation of hits on both sensors enables a coarse p_T discrimination on module level. High-p_T track information is forwarded to the CMS L1 trigger at the bunch crossing frequency of 40 MHz. The required readout bandwidth of more than 3.84 Gb/s necessitates the implementation of a multi-gigabit transceiver ASIC (lpGBT) and opto-electrical converter (VTRx$^+$ module) on every Outer Tracker module.

Compared to the current CMS strip tracker the front-end power consumption of the Outer Tracker will be a factor of two higher. A direct powering scheme would lead to unacceptable losses in the cables. Like the CMS Phase-1 pixel detector, the Outer Tracker will use a DC-DC converter powering scheme. A voltage of about 10 V will be fed to the detector and DC-DC converters on every module will provide the required voltages for the front-ends. In the 2S module case, two converter ASICs are necessary that deliver about 2.5 V (bPOL12V) and 1.25 V (bPOL2V5). The conversion is performed in two steps, the 2.5 V output voltage of the bPOL12V is the input voltage for the bPOL2V5. The Service Hybrid also distributes and filters the bias voltage for the silicon strip sensors.

All PCBs used on the Outer Tracker modules are referred to as hybrids. In the 2S module, the transceiver, the opto-electrical converter, and the DC-DC converters are placed on the same PCB, the Service Hybrid. An early prototype version of the Service Hybrid was designed and produced in 2018 by the I. Physikalisches Institut B of RWTH Aachen University. Since the required ASICs were not yet available, it utilized the GBTx, the VTRx module, the FEAST2, and the LTC3412A instead.

A custom test board was designed to qualify the prototypes. Different versions were produced between the years 2018 and 2020. The version 2.1 was used to test hybrids before their use in prototype modules. The v2.2 boards were installed in a setup for active thermal cycling of hybrids. Circuits to measure the bias voltage and the

leakage current of the corresponding RC filter were a noteworthy new development. The results obtained in the context of this thesis demonstrated the functionality of the two-stepped DC-DC converter powering scheme in combination with the optical readout components. The Service Hybrids were found to be functional for their usage in prototype 2S modules.

The implementation of the lpGBT, the VTRx$^+$ module, the bPOL12V, and the bPOL2V5 in the Service Hybrid design was done by CERN, which was responsible for all hybrid designs of the final prototyping phase. CERN also developed infrastructure for production-scale testing of hybrids. It is based on hybrid-specific test cards. Up to twelve cards are installed in a crate with a common backplane and tested in series. In 2020, the I. Physikalisches Institut B of RWTH Aachen University designed the test card for the Service Hybrid (2S-SEH test card). It combined features of the Service Hybrid test board and previously designed test cards but also required new electrical and mechanical solutions.

The 2S-SEH test card was commissioned and integrated into the framework for hybrid testing in the scope of this thesis. It was also used to qualify the final Service Hybrid prototypes. Two contractors manufactured hybrids for CERN and delivered them in 2021. To be useable for module prototyping an e-fusing of the lpGBT was required. A total of 74 hybrids were tested. The DC-DC converters and the two-stepped powering scheme were characterized. Failures due to the manufacturing, assembly, and functional components were identified. The design was found to be functional for use in the 2S module. CERN selected only one contractor for the series production. The results obtained in this thesis were valuable input for the decision.

The performance of the Service Hybrid in the presence of a strong magnetic field was evaluated with the M1 magnet at CERN in July 2021. A dedicated setup and an interface PCB were prepared for the test. The test setup was placed outside the magnet and connected with several-meter-long cables to the hybrid. No change in the DC-DC converters' behavior or the optical readout was observed for 0 T, 1.5 T, or 3.0 T magnetic field strength in two different field orientations.

To fulfill the 2S-SEH test card needs for the hybrid series production 90 of them were produced in 2022. An extensive testing campaign ensued. Since the cards will be operated in a climatic chamber, the campaign included thermal cycles between +40°C and −35°C. The measurements were also used to characterize the Service Hybrid prototypes in terms of their response to temperature change and to evaluate the uncertainties of the measured quantities on the test cards.

A total of 78 2S-SEH test cards were delivered to CERN in 2023 and 2024. They will be used for functional hybrid testing at CERN and the remaining contractor. The setup at CERN was used to test two production kick-off designs of the Service Hybrid. The I. Physikalisches Institut B of RWTH Aachen University identified a routing error in the designs and a problem with the bump bonds connecting the bPOL2V5 to the hybrid during further tests and investigations of failures. The

bump bonds caused a low yield for one design. Besides the routing error, the hybrids were used to compare the designs' influence on the 2S module strip noise. A direct measurement of one noise source was also performed in the scope of this thesis. All measurements favor the so-called regular design.

The last set of Service Hybrids that is covered in this thesis is the hybrid pre-series. It provided fully functional hybrids with the regular design and was tested by the author at +40°C and −35°C at CERN in early 2024. The characterization at the two temperatures provided valuable information for the Service Hybrid operation in the Outer Tracker.

Beyond the scope of this thesis is the installation of the hybrid test system at the contractor. Over two years, 8,000 Service Hybrids will be produced and tested. The production relies on a functional test system. The size of production lots increases from a few dozen to several hundred. The results of the in-depth studies of the earlier hybrid versions are important to assess the quality of the coming hybrid production.

This work made important contributions to the development of the 2S module for the Phase-2 Outer Tracker and the production test system. The contributions provided functional Service Hybrid prototypes for vital activities of the upgrade project, confidence in the Service Hybrid performance in the future detector, and an understanding of encountered failures. The test board and test card systems will be used throughout the hybrid production to continue this effort.

Appendix

The following section contains circuit diagrams of PCBs developed by the I. Physikalisches Institut B of RWTH Aachen University that were used in the scope of this thesis.

A Leakage current measurement daughterboard

Figure A.1 shows the schematics of the leakage current measurement daughterboard. It is used with the EMCO_P12N on the Service Hybrid test board and without it on the 2S-SEH test card.

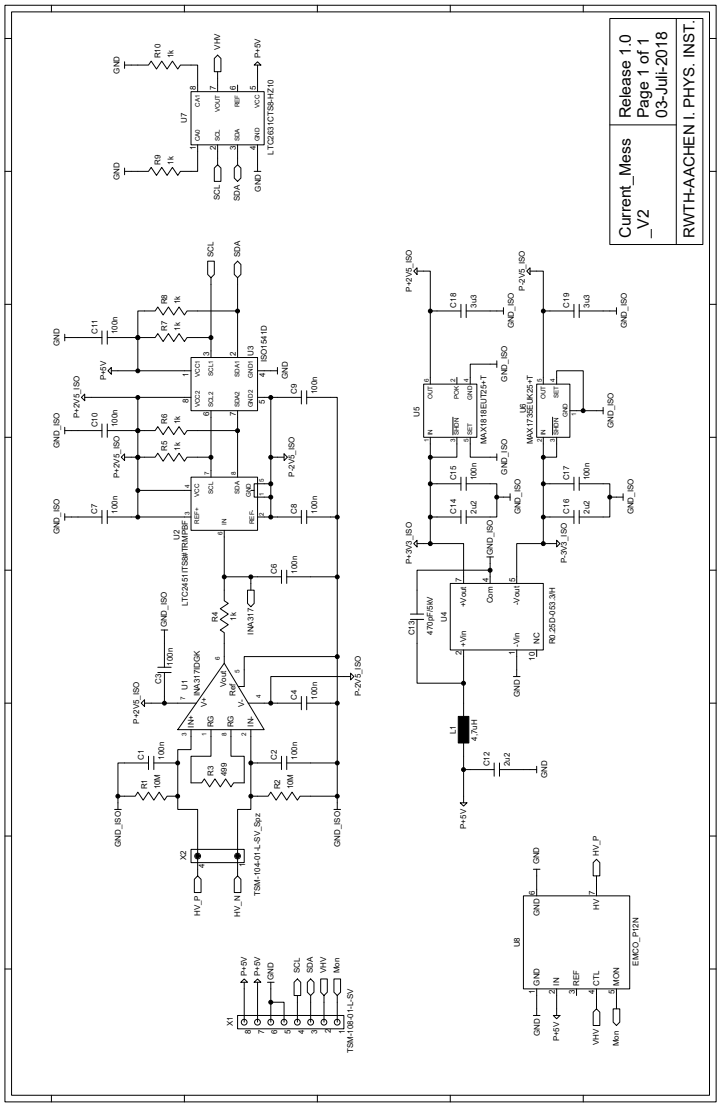

Figure A.1: Schematic of the leakage current measurement daughterboard.

B Load Box

Figures B.1, B.2, and B.3 show the schematics of the PCB used inside the load box. It has three different input channels that can draw a constant load current. It was used during the testing of a final Service Hybrid prototype in a magnetic field.

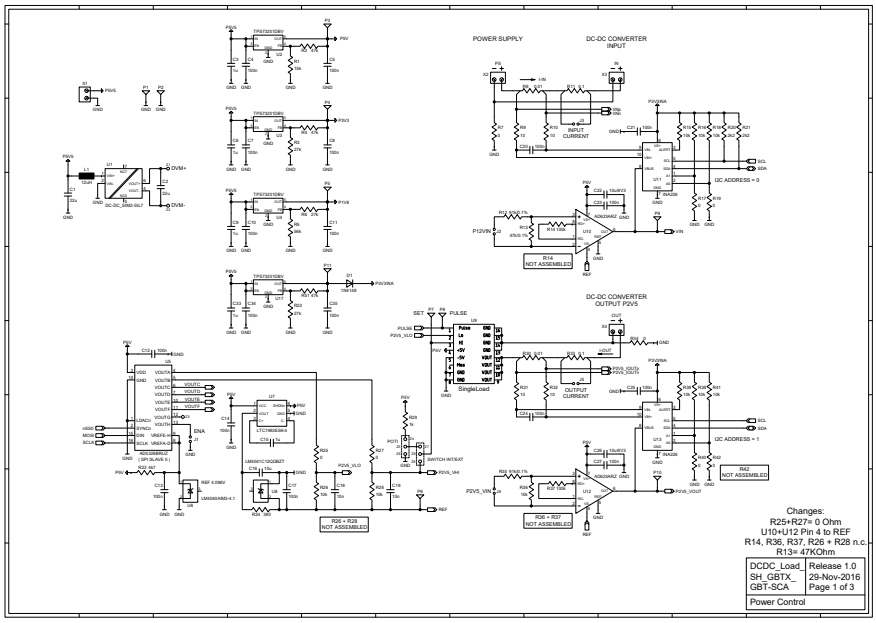

Figure B.1: Schematic of the load box PCB. Page 1 of 3.

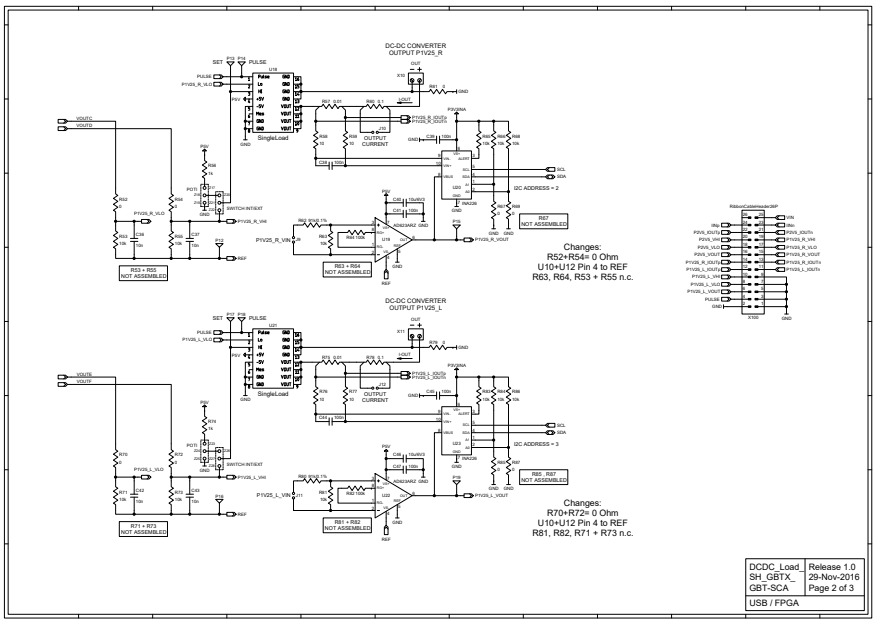

Figure B.2: Schematic of the load box PCB. Page 2 of 3.

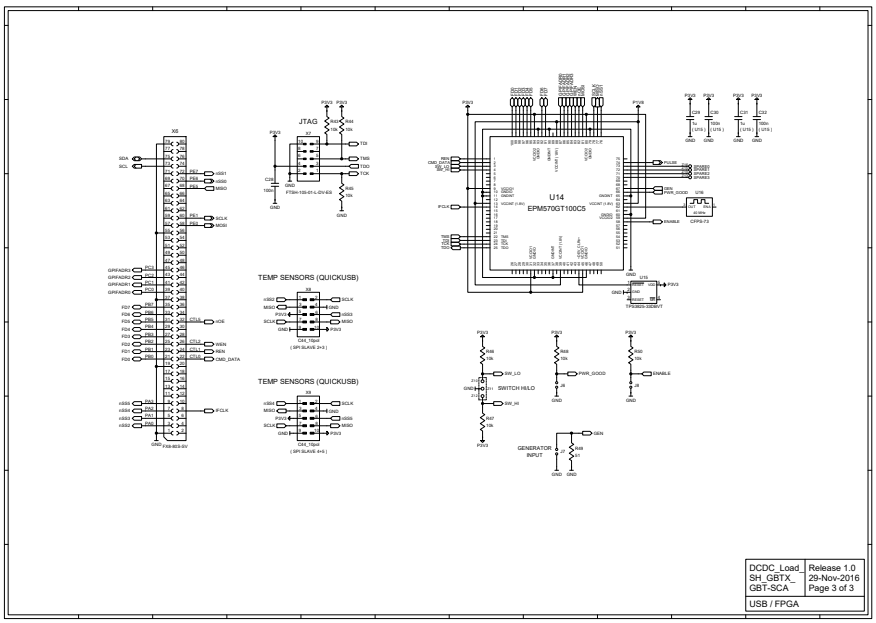

Figure B.3: Schematic of the load box PCB. Page 3 of 3.

C Ground difference measurement PCB

Figure C.1 shows the schematics of the PCB that was used to compare the ground potential between the Service Hybrid sides. It features connectors for the v3.1 Service Hybrids and the hybrids that feature the L-shaped tails and were developed by CERN.

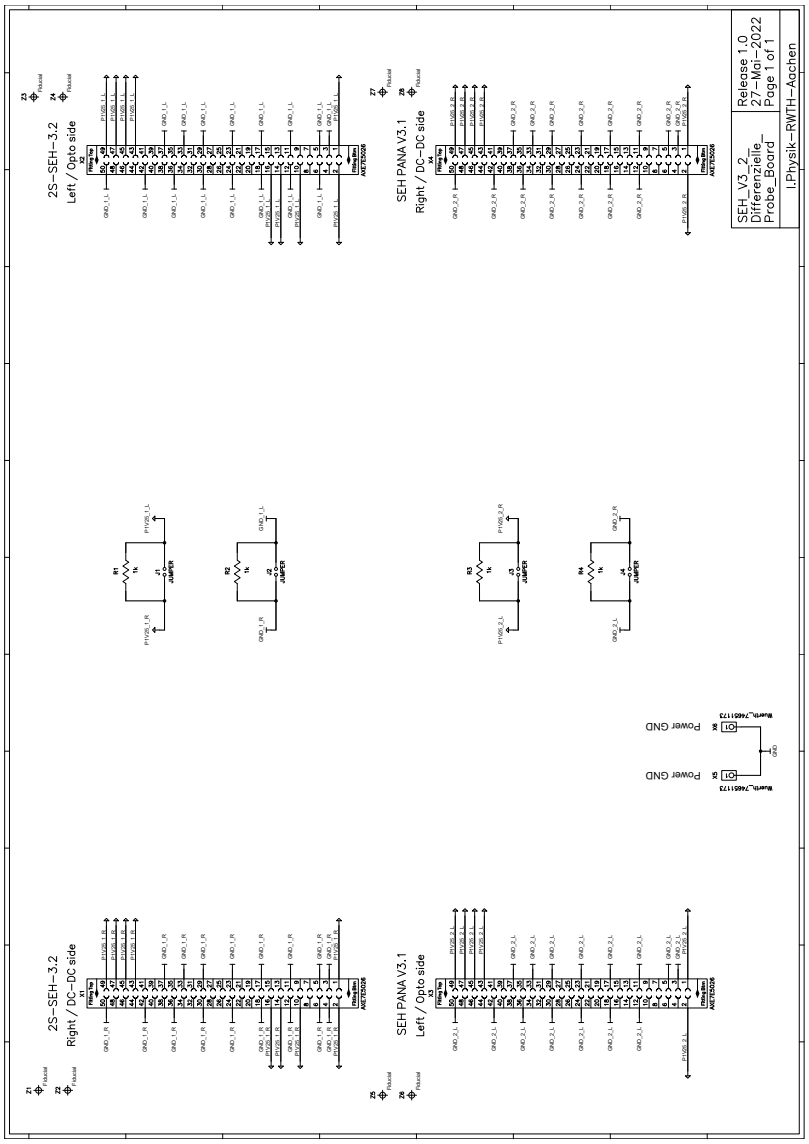

Figure C.1: Schematic of the PCB that is used to measure the AC voltage difference on the ground potential of the Service Hybrid.

D 2S-SEH test card

Figures D.1 to D.20 show the schematics of the 2S-SEH test card. The board was designed by Daniel Louis (I. Physikalisches Institut B of RWTH Aachen University). The schematics correspond to the version that is used for the testing at CERN and the contractor.

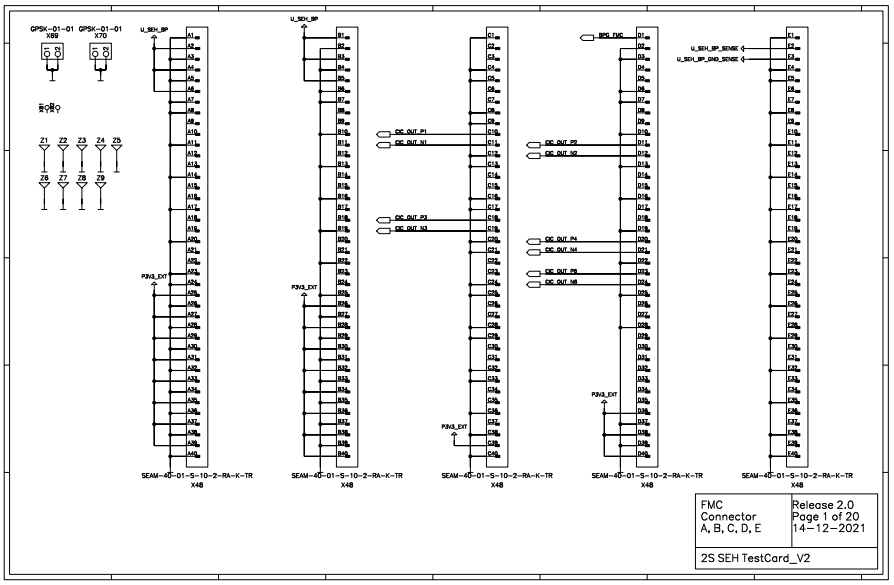

Figure D.1: Schematic of the 2S-SEH test card. Page 1 of 20.

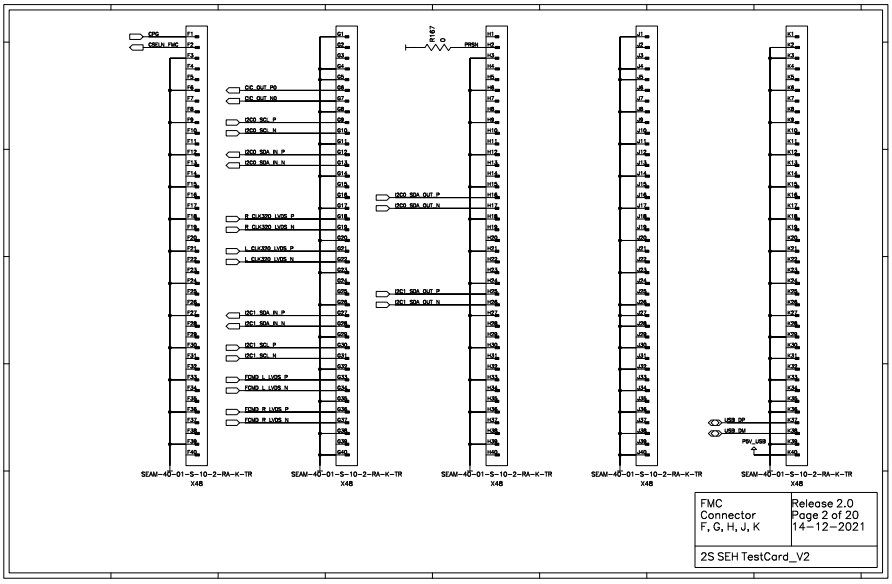

Figure D.2: Schematic of the 2S-SEH test card. Page 2 of 20.

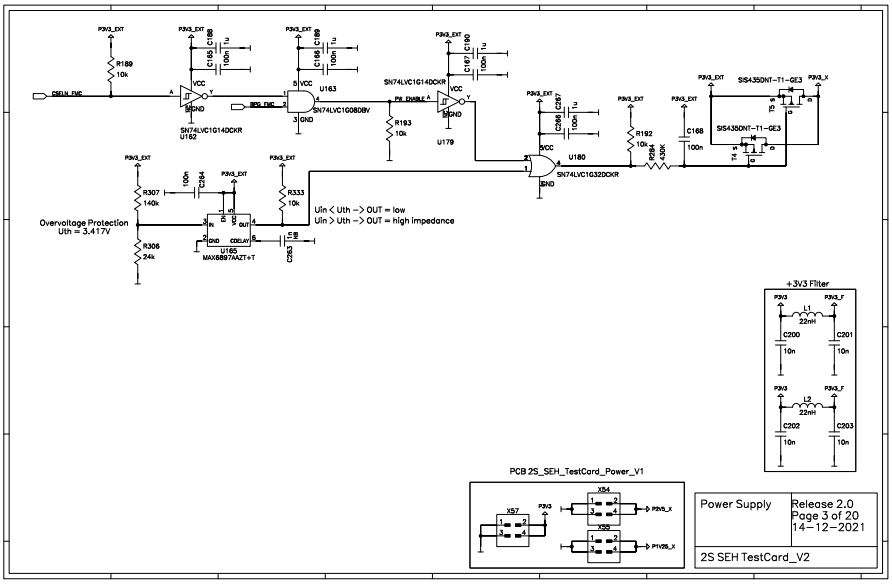

Figure D.3: Schematic of the 2S-SEH test card. Page 3 of 20.

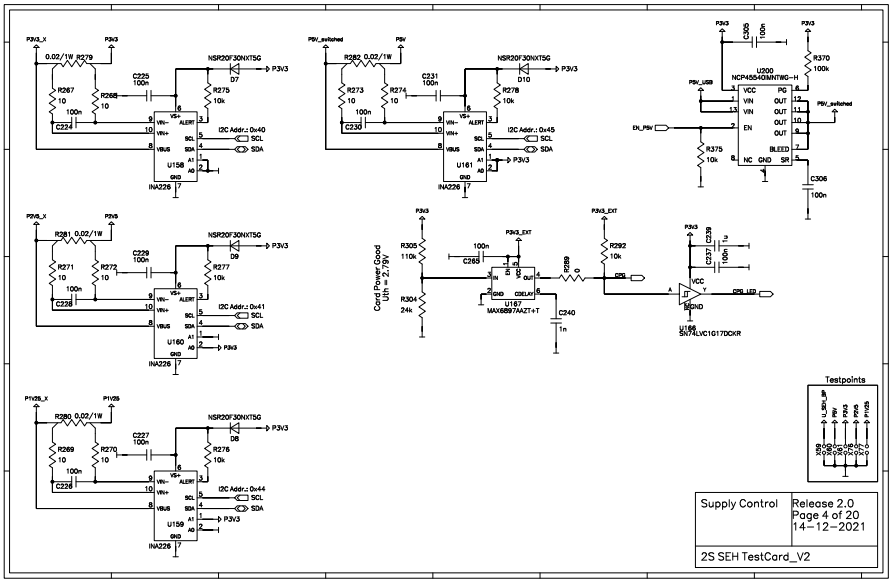

Figure D.4: Schematic of the 2S-SEH test card. Page 4 of 20.

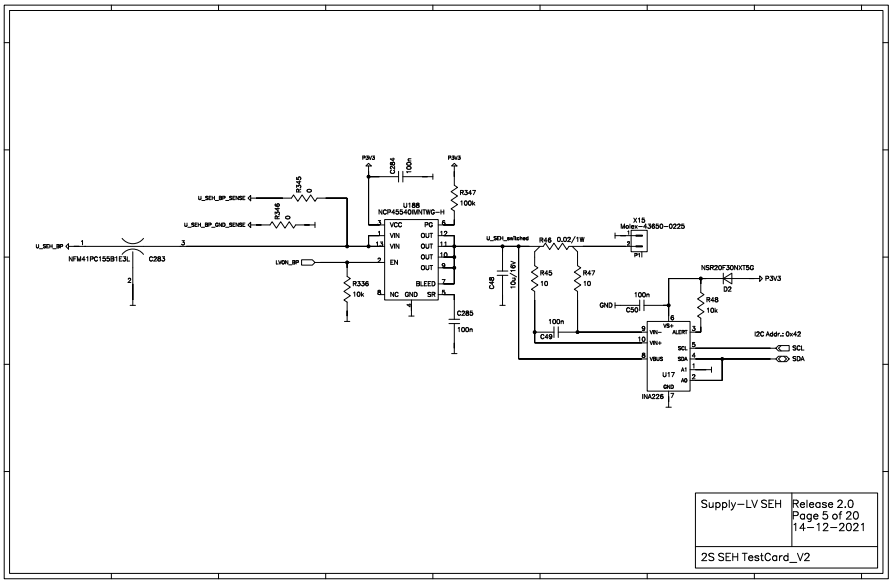

Figure D.5: Schematic of the 2S-SEH test card. Page 5 of 20.

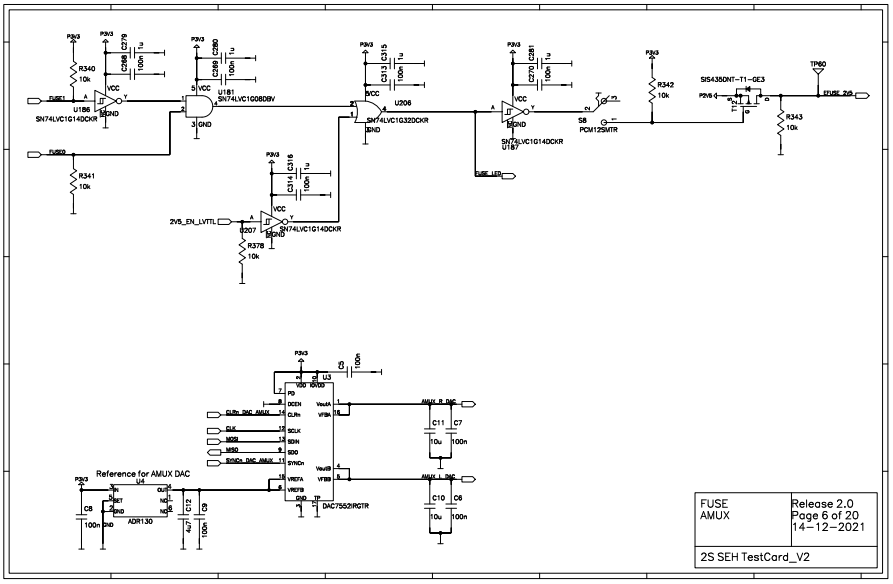

Figure D.6: Schematic of the 2S-SEH test card. Page 6 of 20.

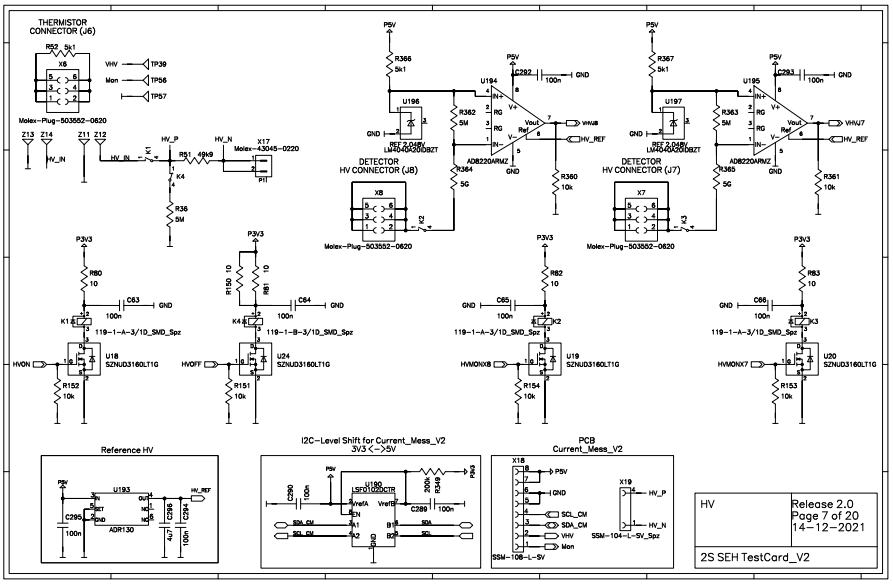

Figure D.7: Schematic of the 2S-SEH test card. Page 7 of 20.

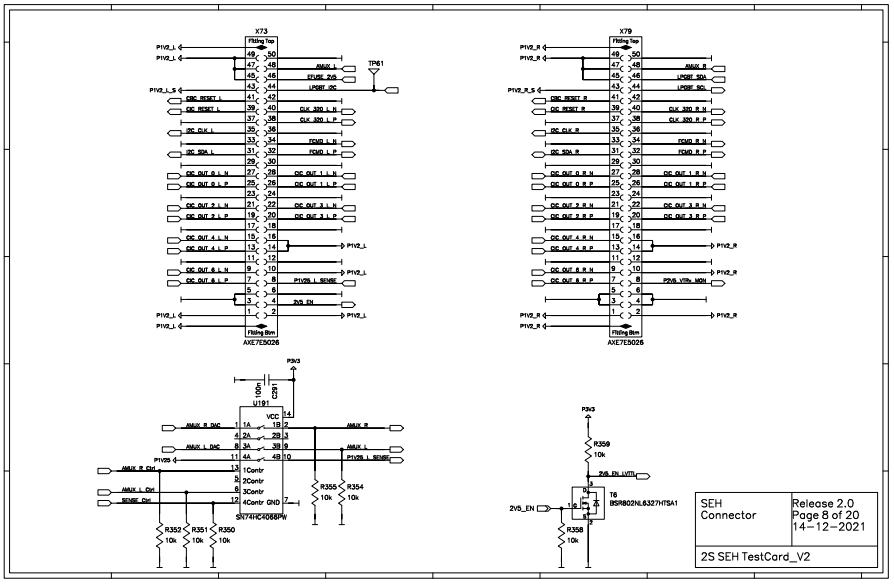

Figure D.8: Schematic of the 2S-SEH test card. Page 8 of 20.

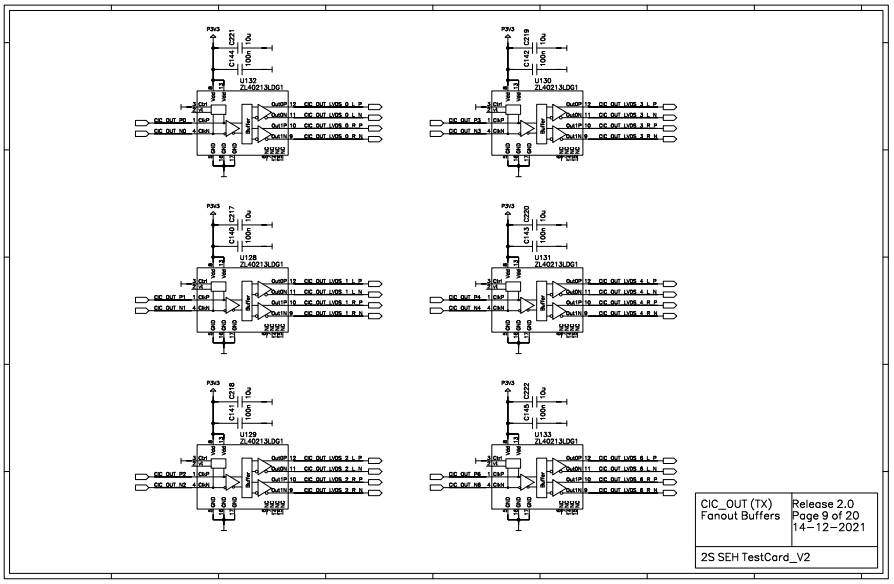

Figure D.9: Schematic of the 2S-SEH test card. Page 9 of 20.

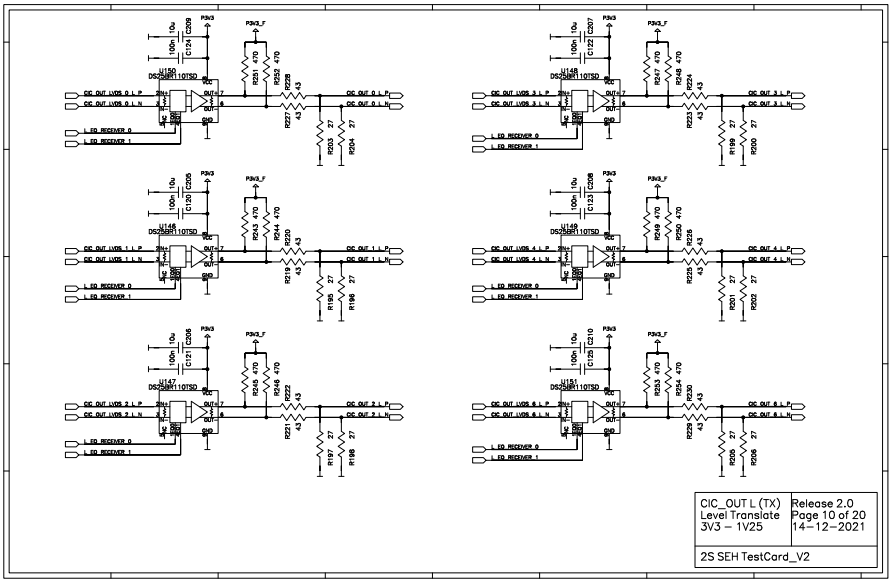

Figure D.10: Schematic of the 2S-SEH test card. Page 10 of 20.

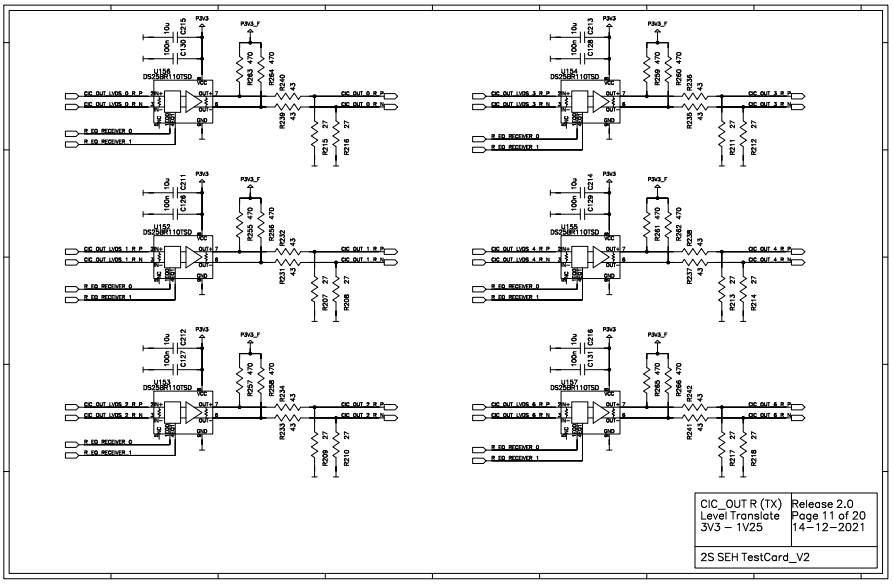

Figure D.11: Schematic of the 2S-SEH test card. Page 11 of 20.

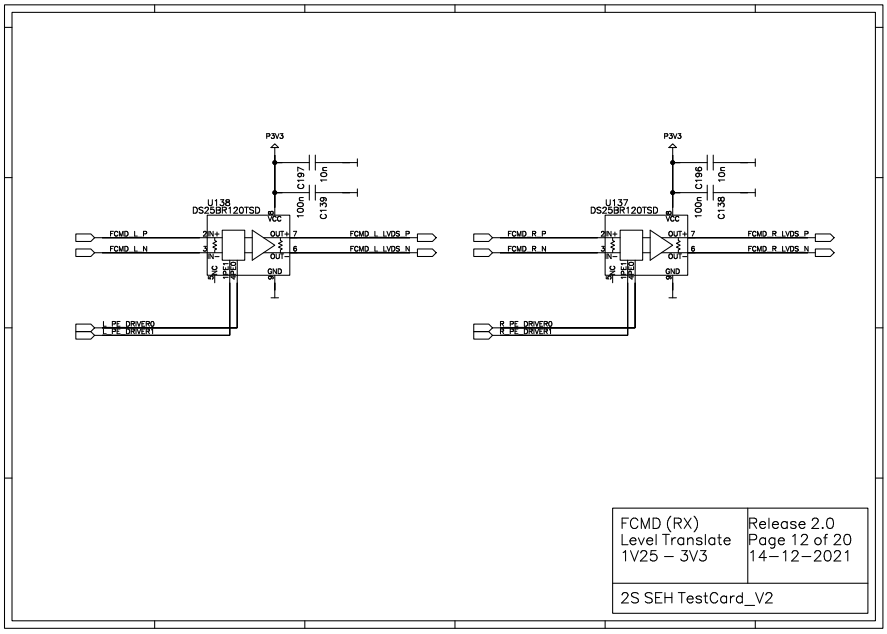

Figure D.12: Schematic of the 2S-SEH test card. Page 12 of 20.

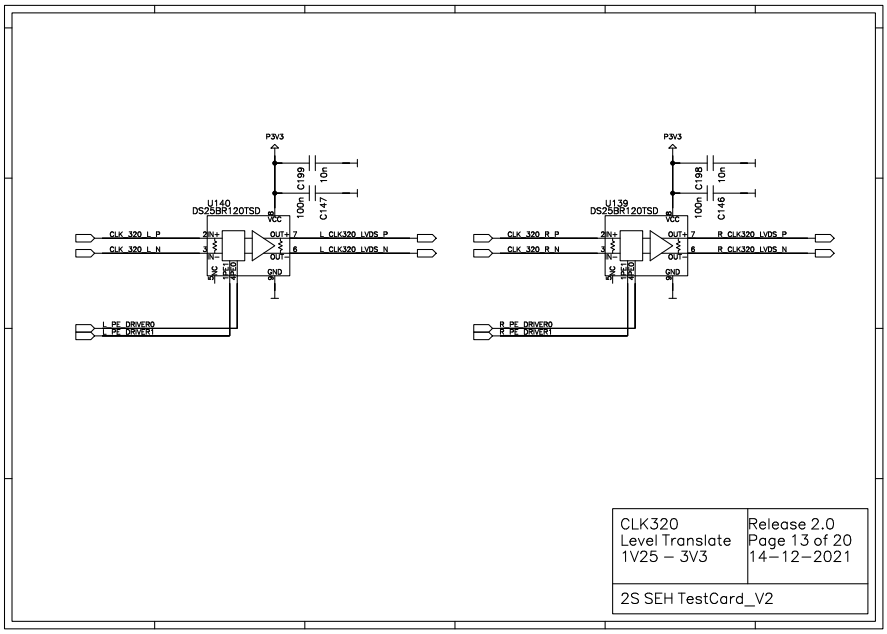

Figure D.13: Schematic of the 2S-SEH test card. Page 13 of 20.

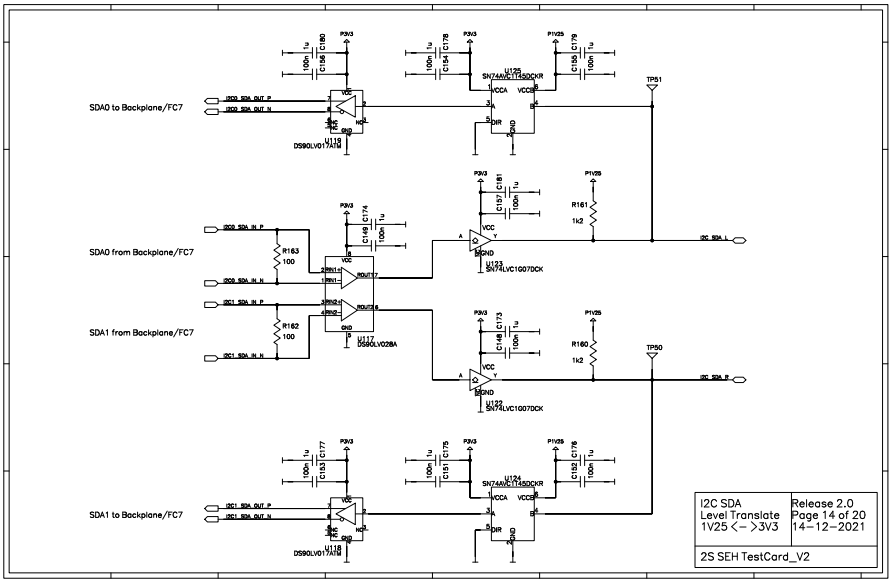

Figure D.14: Schematic of the 2S-SEH test card. Page 14 of 20.

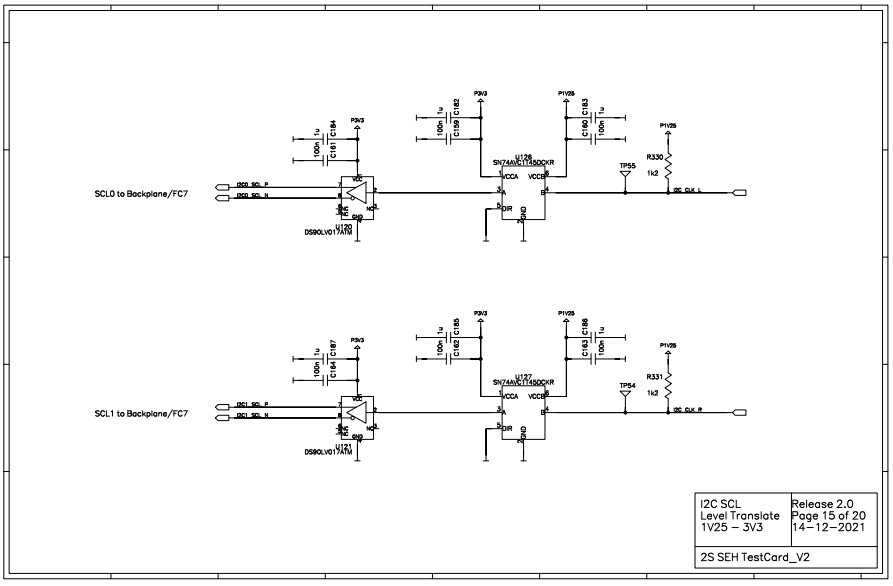

Figure D.15: Schematic of the 2S-SEH test card. Page 15 of 20.

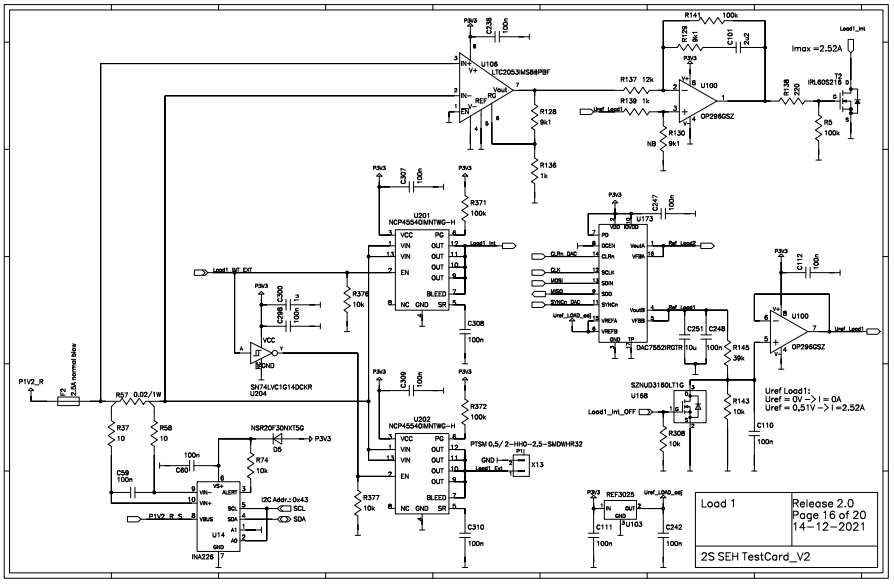

Figure D.16: Schematic of the 2S-SEH test card. Page 16 of 20.

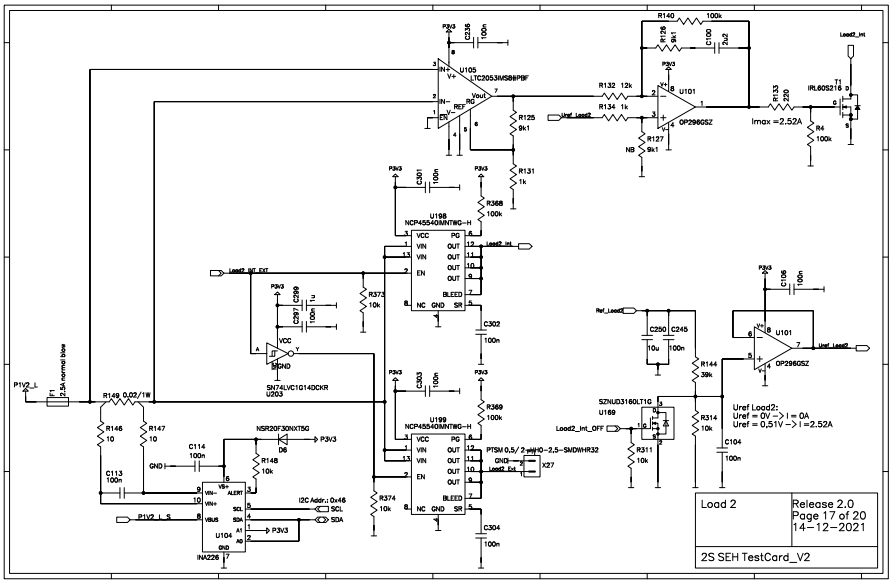

Figure D.17: Schematic of the 2S-SEH test card. Page 17 of 20.

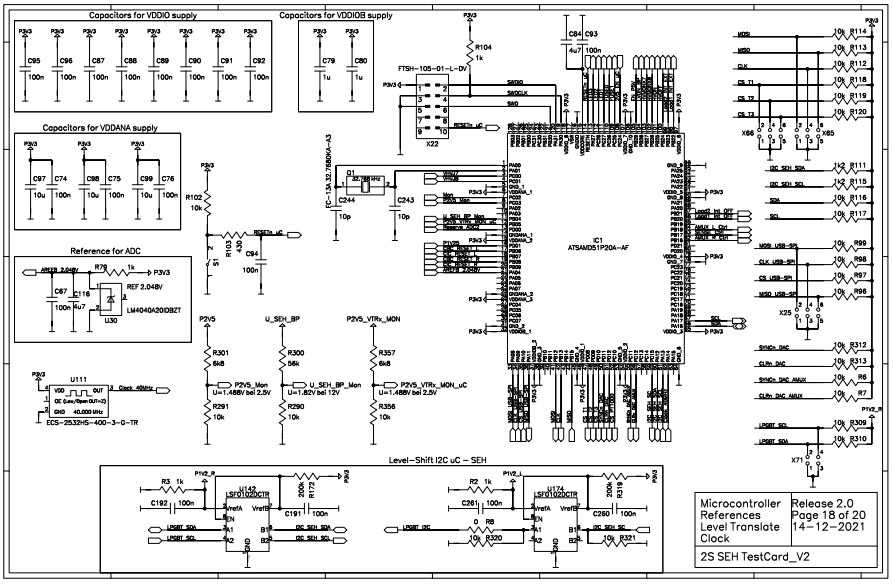

Figure D.18: Schematic of the 2S-SEH test card. Page 18 of 20.

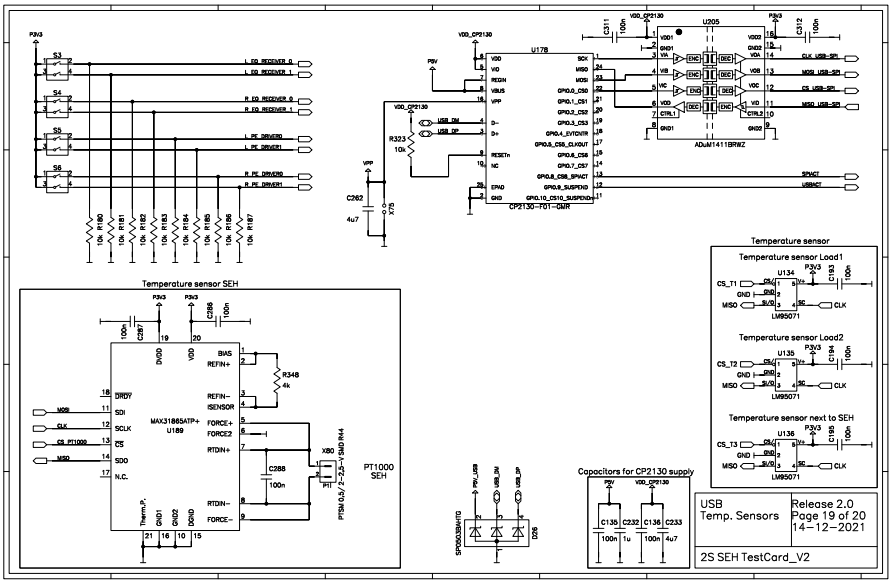

Figure D.19: Schematic of the 2S-SEH test card. Page 19 of 20.

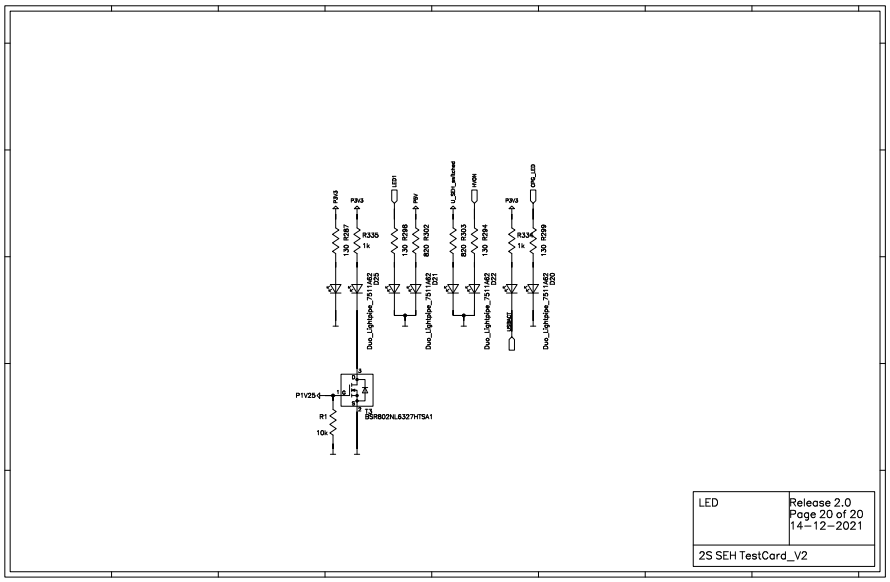

Figure D.20: Schematic of the 2S-SEH test card. Page 20 of 20.

Acronyms

%pt. percentage point

2S strip and strip

AC alternating current

ADC analog-to-digital converter

AlCF aluminum carbon fiber

ALICE A Large Ion Collider Experiment

AMC Advanced Mezzanine Card

AMUX analog multiplexer

ASIC application-specific integrated circuit

ATLAS A Toroidal LHC ApparatuS

BGA ball grid array

bPOL buck point-of-load converter

BRIL Beam Radiation, Instrumentation, and Luminosity

BTL barrel timing layer

CBC CMS Binary Chip

CDR clock and data recovery

CERN Conseil européen pour la recherche nucléaire

CIC Concentrator Integrated Circuit

CM common mode

CMOS complementary metal–oxide–semiconductor

CMS Compact Muon Solenoid

CPU central processing unit

CSC cathode strip chamber

CTE coefficient of thermal expansion

DAC digital-to-analog converter

DAQ data acquisition

DC direct current

DC-DC direct current to direct current

DT drift tube

e-link electrical link

EC External Control

ECAL electromagnetic calorimeter

ETL endcap timing layer

FEC forward error correction

FEH Front-End Hybrid

FFT fast Fourier transform

FMC FPGA Mezzanine Card

FPGA field-programmable gate array

GBLD Gigabit Laser Driver

GBT-SCA Gigabit Transceiver Slow Control Adapter

GBTIA Gigabit Trans-Impedance Amplifier

GBTx Gigabit Transceiver ASIC

GEM gas electron multiplier

GLIB Gigabit Link Interface Board

GPIO general purpose input/output

GUI graphical user interface

HCAL hadron calorimeter

HL-LHC High Luminosity LHC

HLT high-level trigger

I^{2}C inter-integrated circuit

I/V current versus voltage

IC Internal Control

ID identifier

IT Inner Tracker

L1 Level-1

LDO low-dropout regulator

LHC Large Hadron Collider

LHCb Large Hadron Collider beauty

linPOL linear point-of-load converter

lpGBT low-power Gigabit Transceiver

LS Long Shutdown

LVDS Low Voltage Differential Signaling

MCH MicroTCA Carrier Hub

MPA Macro-Pixel ASIC

NTC negative temperature coefficient

OT Outer Tracker

PCB printed circuit board

PG Power Good

PLL phase-locked loop

POH Power Hybrid

PS macro-pixel and strip

PTAT Proportional To Absolute Temperature

QA/QC quality assurance and quality control

QFN Quad Flat No-Lead

RAM random-access memory

RF radio frequency

ROC readout chip

ROH Readout Hybrid

ROM read-only memory

RPC resistive-plate chamber

RSSI received signal strength current

SEH Service Hybrid

SFP Small Form-factor Pluggable

SiPM silicon photomultiplier

SM standard model

SSA Strip Sensor ASIC

TB2S Tracker Barrel with 2S modules

TBPS Tracker Barrel with PS modules

TEDD Tracker Endcap Double Disk

TID total ionizing dose

UVLO under-voltage lockout

VL Versatile Link

VL$^+$ Versatile Link plus

VLDB$^+$ Versatile Link Plus Demo Board

VTRx Versatile Transceiver

VTRx$^+$ Versatile Transceiver plus

µC microcontroller

µTCA Micro Telecommunications Computing Architecture

Bibliography

[1] L. Evans and B. Bryant (editors), "LHC Machine", 2008 *JINST* **3** S08001 (2008), DOI: 10.1088/1748-0221/3/08/S08001.

[2] CMS Collaboration, *Public CMS Luminosity Information*, URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults (visited on June 7, 2024).

[3] CERN, *CERN's accelerator complex*, URL: https://home.cern/science/accelerators/accelerator-complex (visited on May 17, 2024).

[4] CERN, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN Yellow Report CERN-2020-010, 2020, DOI: 10.23731/CYRM-2020-0010.

[5] CMS Collaboration, "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC", *Phys. Lett. B* **716** (2012) 30–61, DOI: 10.1016/j.physletb.2012.08.021.

[6] CMS Collaboration, *CMS Publications*, URL: https://cms-results.web.cern.ch/cms-results/public-results/publications/ (visited on May 17, 2024).

[7] CMS Collaboration, "The CMS experiment at the CERN LHC", 2008 *JINST* **3** S08004 (2008), DOI: 10.1088/1748-0221/3/08/S08004.

[8] T. Sakuma and T. McCauley, "Detector and Event Visualization with SketchUp at the CMS Experiment", *J. Phys. Conf. Ser.* **513** 022032 (2014), DOI: 10.1088/1742-6596/513/2/022032.

[9] CMS Collaboration, "The CMS trigger system", 2017 *JINST* **12** P01020 (2017), DOI: 10.1088/1748-0221/12/01/P01020.

[10] CMS Collaboration, "Development of the CMS detector for the CERN LHC Run 3", 2024 *JINST* **19** P05064 (2024), DOI: 10.1088/1748-0221/19/05/P05064.

[11] M. Zerlauth and O. Brüning, "Status and prospects of the HL-LHC project", *PoS* (**EPS-HEP2023**) 615 (2024), DOI: 10.22323/1.449.0615.

[12] CMS Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, Technical Design Report CMS-TDR-014, 2017, DOI: 10.17181/CERN.QZ28.FLHW.

[13] CERN, *Extension of RD53*, 2018, URL: https://cds.cern.ch/record/2637453.

[14] P. Moreira et al., *lpGBT documentation: release*, 2022, URL: https://cds.cern.ch/record/2809058.

[15] F. Vasey et al., *Versatile Link+ Technical Specification, part 2.1 Versatile Link+ Transceiver (VTRx+)*, URL: https://edms.cern.ch/file/1719329/1/VTRxPlus_spec.pdf (visited on Apr. 11, 2024).

[16] G. Ripamonti et al., *bPOL12V_V6 datasheet – V1.6*, URL: https://power-distribution.web.cern.ch/assets/datasheets/bPOL12V_V6%20datasheet%20V1.9.pdf (visited on July 15, 2024).

[17] G. Ripamonti et al., *bPOL2V5_V3.3 datasheet – rev3.0*, URL: https://power-distribution.web.cern.ch/assets/datasheets/bPOL2V5_V3.3%20datasheet%20rev3.pdf (visited on July 15, 2024).

[18] A. Rose, "Serenity - An ATCA prototyping platform for CMS Phase-2", *PoS* (**TWEPP2018**) 115 (2019), DOI: 10.22323/1.343.0115.

[19] CMS Collaboration, *A MIP Timing Detector for the CMS Phase-2 Upgrade*, Technical Design Report CMS-TDR-020, 2019, URL: https://cds.cern.ch/record/2667167.

[20] A. Caratelli et al., "The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments", 2015 *JINST* **10** C03034 (2015), DOI: 10.1088/1748-0221/10/03/C03034.

[21] CMS Collaboration, *The Phase-2 Upgrade of the CMS Barrel Calorimeters*, Technical Design Report CMS-TDR-015, 2017, URL: https://cds.cern.ch/record/2283187.

[22] G. Pásztor, "The Phase-2 Upgrade of the CMS Detector", *PoS* (**LHCP2022**) 045 (2023), DOI: 10.22323/1.422.0045.

[23] CMS Collaboration, *The Phase-2 Upgrade of the CMS Endcap Calorimeter*, Technical Design Report CMS-TDR-019, 2017, DOI: 10.17181/CERN.IV8M.1JY2.

[24] G. Ripamonti et al., *linPOL12V datasheet – rev3.4*, URL: https://power-distribution.web.cern.ch/assets/datasheets/linPOL12V%20datasheetV3.4.pdf (visited on July 30, 2024).

[25] CMS Collaboration, *The Phase-2 Upgrade of the CMS Muon Detectors*, Technical Design Report CMS-TDR-016, 2017, URL: https://cds.cern.ch/record/2283189.

[26] M. I. Kovacs et al., on behalf of the CMS Collaboration, "Front-end hybrid designs for the CMS Phase-2 Upgrade towards the production phase", 2022 *JINST* **17** C04018, (2022), DOI: 10.1088/1748-0221/17/04/C04018.

[27] A. Caratelli et al., "Short-Strip ASIC (SSA): A 65nm silicon-strip readout ASIC for the Pixel-Strip (PS) module of the CMS Outer Tracker detector upgrade at HL-LHC", *PoS* (**TWEPP2017**) *031* (2018), DOI: 10.22323/1.313.0031.

[28] D. Ceresa et al., *Macro Pixel ASIC (MPA): The readout ASIC for the pixel-strip (PS) module of the CMS outer tracker at HL-LHC*, 2014, DOI: 10.1088/1748-0221/9/11/C11012.

[29] N. Röwert, on behalf of the CMS Tracker Group, "Design and construction of the CMS Outer Tracker for the HL-LHC Upgrade", *Nucl. Instrum. Meth. A* **1064** (2024) 169317, DOI: 10.1016/j.nima.2024.169317.
[30] Dupont, *Polyimide Films*, URL: https://www.dupont.com/electronics-industrial/polyimide-films.html (visited on May 17, 2024).
[31] Panasonic, *A35S SERIES (0.35MM PITCH)*, URL: https://na.industrial.panasonic.com/products/connectors/board-board-connectors/lineup/narrow-fine-pitch-connectors/series/82630 (visited on May 17, 2024).
[32] Hirose, *DF57 Series*, URL: https://www.hirose.com/en/product/series/DF57 (visited on May 17, 2024).
[33] S. Paoletti, on behalf of the CMS Collaboration, "The powering systems of the CMS tracker for HL-LHC", 2022 *JINST* **17** C12015 (2022), DOI: 10.1088/1748-0221/17/12/C12015.
[34] S. Roberts, *DC/DC Book of Knowledge: Practical Tips for the User*, 2015, URL: https://recom-power.com/en/support/technical-resources/book-of-knowledge/book-of-knowledge.html?3.
[35] M. Lipinski, "System tests and qualification of pixel modules and DC-DC converters for the phase-1 upgrade of the CMS pixel detector", PhD thesis, RWTH Aachen University, 2019, DOI: 10.18154/RWTH-2019-01452.
[36] G. Bergamin et al., *CIC2/2.1 technical specification*, URL: https://espace.cern.ch/Tracker-Upgrade/Electronics/CIC/Shared%20Documents/Specifications/CIC2_specs_v1p4.pdf (visited on May 6, 2024).
[37] C. Soós et al., "Versatile Link+ transceiver production", 2023 *JINST* **18** C03003 (2023), DOI: 10.1088/1748-0221/18/03/C03003.
[38] O. Wen, "Thermal Dependence of the Performance of the Service Hybrids for the CMS Phase-2 Tracker Upgrade", BA thesis, RWTH Aachen University, 2019, URL: https://wiki.physik.rwth-aachen.de/pub/Feld/FeldAbschlussarbeiten/Bachelorarbeit_Oliver_Wen.pdf.
[39] Samtec, *SEAM-40-01-L-10-2-RA-GP-TR*, URL: https://www.samtec.com/products/seam-40-01-l-10-2-ra-gp-tr (visited on May 29, 2024).
[40] N. Raŝević, *PS Readout Hybrid Test Card*, private communication, 2019.
[41] Silicon Labs, *CP2130 - Single-Chip USB-to-SPI Bridge*, URL: https://www.digikey.de/htmldatasheets/production/1427193/0/0/1/cp2130-datasheet.html (visited on May 29, 2024).
[42] M. I. Kovacs et al., "A High Throughput Production Scale Front-End Hybrid Test System for the CMS Phase-2 Tracker Upgrade", *PoS* (**TWEPP2019**) 082 (2020), DOI: 10.22323/1.370.0082.
[43] PICMG, *MTCA.0R1.0 Specification*, URL: https://www.picmg.org/openstandards/microtca/ (visited on Jan. 4, 2024).
[44] PICMG, *AMC.0R2.0 Specification*, URL: https://www.picmg.org/openstandards/advanced-mezzanine-card/ (visited on Jan. 4, 2024).

[45] E. Hazen et al., "The AMC13XG: a new generation clock/timing/DAQ module for CMS MicroTCA", 2013 *JINST* **8** C12036 (2013), DOI: 10.1088/1748-0221/8/12/C12036.

[46] M. Pesaresi et al., "The FC7 AMC for generic DAQ & control applications in CMS", 2015 *JINST* **10** C03036 (2015), DOI: 10.1088/1748-0221/10/03/C03036.

[47] C. Ghabrous Larrea et al., "IPbus: a flexible Ethernet-based control system for xTCA hardware", 2015 *JINST* **10** C02019 (2015), DOI: 10.1088/1748-0221/10/02/C02019.

[48] M. Haranko, "Development of a test DAQ system for the CMS Phase-2 outer tracker upgrade", PhD thesis, Hamburg University, 2020, DOI: 10.3204/PUBDB-2020-00885.

[49] FS, *MTP/MPO vs Traditional LC Fiber Cables: Which One to Choose?*, URL: https://community.fs.com/article/mtpmpo-vs-traditional-lc-fiber-cables-which-one-to-choose.html (visited on May 21, 2024).

[50] USConec, *MTP Connectors*, URL: https://www.usconec.com/connectors/mtp-connectors (visited on May 21, 2024).

[51] I. Mateos Domínguez et al., "Software tools for hybrid quality control for the CMS Outer Tracker Phase-2 Upgrade", 2023 *JINST* **18** C01048 (2023), DOI: 10.1088/1748-0221/18/01/C01048.

[52] Rohde&Schwarz, *R&S HMP4000 Power supply series*, URL: https://www.rohde-schwarz.com/uk/products/test-and-measurement/dc-power-supplies/rs-hmp4000-power-supply-series_63493-47360.html?change_c=true (visited on May 21, 2024).

[53] iseg, *SWITCHABLE HIGH END HIGH PRECISION AC/DC DESKTOP HV SOURCE-MEASURE-UNIT*, URL: https://iseg-hv.com/en/products/detail/SHR (visited on May 21, 2024).

[54] W. Eddy, "Transmission Control Protocol (TCP)", *RFC* **9293** (2022), DOI: 10.17487/RFC9293.

[55] Climats, *EXCAL^2 range*, URL: https://www.climats-tec.com/en/products/our-product-lines/temperature-environmental-test-chamber/excal2-range-3.html (visited on June 25, 2024).

[56] P. Szydlik et al., on behalf of the CMS Collaboration, "Commissioning of the test system for PS and 2S hybrids for the Phase-2 Upgrade of the CMS Outer Tracker", 2024 *JINST* **19** C04040 (2024), DOI: 10.1088/1748-0221/19/04/C04040.

[57] CMS Collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, Technical Design Report CMS-TDR-015-02, 2015, DOI: 10.17181/CERN.VU8I.D59J.

[58] F. Faccio et al., "FEAST2: A Radiation and Magnetic Field Tolerant Point-of-Load Buck DC/DC Converter", 2014 *IEEE Radiation Effects Data Workshop*, 2014, DOI: 10.1109/REDW.2014.7004569.

[59] L. Feld et al., "First Implementation of a two-stage DC-DC conversion powering scheme for the CMS Phase-2 outer tracker", 2017 *JINST* **12** C03090 (2017), DOI: 10.1088/1748-0221/12/03/C03090.

[60] M. Preuten et al., "Functional tests of 2S modules for the CMS Phase-2 Tracker Upgrade with a MicroTCA-based readout system", *PoS* (**TWEPP2017**) 069 (2017), DOI: 10.22323/1.313.0069.

[61] K. Klein et al., on behalf of the CMS Tracker Group, "Service hybrids for the silicon strip modules of the CMS Phase-2 Outer Tracker upgrade", *PoS* (**TWEPP2018**) 127 (2019), DOI: 10.22323/1.343.0127.

[62] P. Moreira et al., *GBTX manual; V0.18 draft*, 2021, URL: https://cds.cern.ch/record/2809057.

[63] J. Troska et al., "Radiation tolerant optoelectronics for high energy physics", *Nucl. Instrum. Meth. A* **1052** (2023) 168208, DOI: 10.1016/j.nima.2023.168208.

[64] F. Vasey et al., *The Versatile Link Application Note V2.7*, URL: https://edms.cern.ch/file/1309737/1/Application_Note_V2_7.pdf (visited on Dec. 4, 2023).

[65] F. Faccio and S. Michelis, *FEAST2.1 datasheet – rev1.0*, URL: https://power-distribution.web.cern.ch/assets/datasheets/FEAST2.1%20datasheet.pdf (visited on July 16, 2024).

[66] ANALOG DEVICES, *LTC3412A: 3A, 4MHz, Monolithic Synchronous Step-Down Regulator Data Sheet (Rev. G)*, URL: https://www.analog.com/media/en/technical-documentation/data-sheets/ltc3412a.pdf (visited on Dec. 4, 2023).

[67] LLC Semiconductor Components Industries, *Schottky Barrier Diode NSR20F30 NXT5G*, URL: https://www.onsemi.com/download/data-sheet/pdf/nsr20f30-d.pdf (visited on Nov. 17, 2023).

[68] A. Droll, "Performance Studies on 2S Module Prototypes for the Phase-2 Upgrade of the CMS Outer Tracker", PhD thesis, Karlsruhe Institute of Technology, 2021, URL: https://publish.etp.kit.edu/record/22088.

[69] A. La Rosa, "The CMS Outer Tracker for the High Luminosity LHC upgrade", 2020 *JINST* **15** C2029 (2020), DOI: 10.1088/1748-0221/15/02/C02029.

[70] CERN, *Procurement and Industrial Services Group*, URL: https://procurement.web.cern.ch (visited on Dec. 13, 2023).

[71] Z. Zeng et al., "LDQ10: a compact ultra low-power radiation-hard 4 × 10 Gb/s driver array", 2017 *JINST* **12** P02020 (2017), DOI: 10.1088/1748-0221/12/02/P02020.

[72] F. Faccio et al., "The bPOL12V DCDC converter for HL-LHC trackers: towards production readiness", *PoS* (**TWEPP2019**) 070 (2020), DOI: 10.22323/1.370.0070.

[73] A. Zografos et al., on behalf of the CMS Tracker Collaboration, "CMS Outer Tracker Phase-2 Upgrade on-module powering", 2024 *JINST* **19** C02031 (2024), DOI: 10.1088/1748-0221/19/02/C02031.

[74] A. Zografos et al., on behalf of the CMS Tracker Group, "Power, readout and service hybrids for the CMS phase-2 upgrade", 2022 *JINST* **17** C03034 (2022), DOI: 10.1088/1748-0221/17/03/C03034.

[75] N. Röwert et al., on behalf of the CMS Tracker Group, "Influence of high-frequency magnetic fields on the noise behavior of CMS 2S module prototypes", 2023 *JINST* **18** C01013 (2023), DOI: 10.1088/1748-0221/18/01/C01013.

[76] F. Wittig et al., on behalf of the Tracker Group of the CMS Collaboration, "Beam test studies with silicon sensor module prototypes for the CMS Phase-2 Outer Tracker", 2023 *JINST* **18** C01005 (2023), DOI: 10.1088/1748-0221/18/01/C01005.

[77] T. Ziemons, "Assembly and Test of Prototype 2S Modules for the Phase-2 Upgrade of the CMS Outer Tracker", PhD thesis, RWTH Aachen University, 2022, URL: https://cds.cern.ch/record/2836729.

[78] The Tracker Group of the CMS Collaboration, "Beam Test Performance Studies of CMS Phase-2 Outer Tracker Module Prototypes" (2024), arXiv: 2404.08794v1, URL: https://cds.cern.ch/record/2895766.

[79] C. Dziwok, "Test systems for the Phase-2 Upgrade of the CMS Outer Tracker", PhD thesis in preparation, RWTH Aachen University, 2024.

[80] Raspberry Pi Ltd, *Raspberry Pi Documentation*, URL: https://www.raspberrypi.com/documentation/ (visited on Nov. 17, 2023).

[81] Texas Instruments, *INA226*, URL: https://www.ti.com/product/INA226 (visited on May 29, 2024).

[82] R. Thurlow, *RPC: Remote Procedure Call Protocol Specification Version 2*, RFC **5531**, 2009, DOI: 10.17487/RFC5531.

[83] A. Pauls, "Optimization Studies and Development of Test Procedures for Service Hybrids of CMS Silicon Modules at HL-LHC", MA thesis, RWTH Aachen University, 2019, URL: https://cds.cern.ch/record/2657609.

[84] R. Lesma et al., "The Versatile Link Demo Board (VLDB)", 2017 *JINST* **12** C02020 (2017), DOI: 10.1088/1748-0221/12/02/C02020.

[85] INC ANALOG DEVICES, *LTM4643 Ultrathin Quad µModule Regulator with Configurable 3A Output Array*, URL: https://www.analog.com/media/en/technical-documentation/data-sheets/LTM4643.pdf (visited on Nov. 17, 2023).

[86] XP Power, *P Series DC-HVDC CONVERTER*, URL: https://www.xppower.com/portals/0/pdfs/SF_P_Series.pdf (visited on Nov. 17, 2023).

[87] Keithley Instruments, *Model 2410 High-Voltage SourceMeter*, URL: https://www.tek.com/en/products/keithley/source-measure-units/2400-standard-series-sourcemeter (visited on Nov. 17, 2023).

[88] P. Vichoudis et al., "The Gigabit Link Interface Board (GLIB), a flexible system for the evaluation and use of GBT-based optical links", 2010 *JINST* **5** C11007 (2010), DOI: 10.1088/1748-0221/5/11/C11007.

[89] J. Redd, "Explaining those BER testing mysteries", *Lightwave* (2004), URL: https://www.lightwaveonline.com/home/article/16647704/explaining-those-ber-testing-mysteries.

[90] BK Precision, *USER MANUAL 8500 SERIES DC ELECTRONIC LOADS*, URL: https://bkpmedia.s3.amazonaws.com/downloads/manuals/en-us/85xx_manual.pdf (visited on Nov. 22, 2023).

[91] A. Cassese on behalf of the CMS Tracker Group, "Powering of the CMS Phase-2 Upgraded Tracker", *PoS* (**TWEPP2019**) 066 (2020), DOI: 10.22323/1.370.0066.

[92] Microchip, *ATSAMD51P20A*, URL: https://www.microchip.com/en-us/product/atsamd51p20a (visited on May 29, 2024).

[93] Infineon Technologies, *IRL60S216*, URL: https://www.infineon.com/cms/de/product/power/mosfet/n-channel/irl60s216/ (visited on July 9, 2024).

[94] USConec, *18017, Adapter, MT/MTP, SC Footprint, Full Flange, Black, 20N, No Dust Plug*, URL: https://www.usconec.com/products/adapter-mtmtp-sc-footprint-full-flange-black-20n-no-dust-plug (visited on May 29, 2024).

[95] F. Thurn, "Quality control for service hybrid test cards and thermal long term tests of electronics components for the CMS Phase-2 tracker upgrade", MA thesis, RWTH Aachen University, 2023, URL: https://wiki.physik.rwth-aachen.de/pub/Feld/FeldAbschlussarbeiten/Masterarbeit_Felix_Thurn.pdf.

[96] Y. Otarid, "Pixel-Strip Modules for the CMS Tracker Phase-2 Upgrade: From DAQ test system development to module assembly and qualification", PhD thesis, Hamburg University, 2023, URL: https://cds.cern.ch/record/2876414.

[97] D. H. Montesinos et al., "The Versatile Link+ Demo Board (VLDB+)", 2022 *JINST* **17** C03032 (2022), DOI: 10.1088/1748-0221/17/03/C03032.

[98] Binder, *Modell MK 240*, URL: https://www.binder-world.com/de-de/produkt/mk-240 (visited on July 10, 2024).

[99] Semiconductor Components Industries, *Protected Power Switches NCP45540*, URL: https://www.onsemi.com/products/power-management/protected-power-switches/NCP45540 (visited on July 10, 2024).

[100] M. I. Kovacs et al., on behalf of the CMS Collaboration, "Hybrid designs and kick-off batch production experience for the CMS Phase-2 Upgrade", 2024 *JINST* **19** C04052 (2024), DOI: 10.1088/1748-0221/19/04/C04052.

[101] KYOCERA AVX, *X7R Dielectric, KGM Series General Specifications*, URL: https://datasheets.kyocera-avx.com/KGM_X7R.pdf (visited on Mar. 15, 2024).

[102] H. Lee et al., "Leakage Currents in Multilayer Ceramic Capacitors", *IEEE Transactions on Components, Hybrids, and Manufacturing Technology* **7** (1984), DOI: 10.1109/TCHMT.1984.1136375.

[103] G. Bertotti and F. Fiorillo, *7.1.1.4.1 Magnetization curve, hysteresis loop, losses: Datasheet from Landolt-Börnstein - Group III Condensed Matter · Volume 19I1: "Magnetic Alloys for Technical Applications. Soft Magnetic Alloys, Invar and Elinvar Alloys" in SpringerMaterials*, ed. by H. P. J. Wijn, 1994, DOI: 10.1007/10065028_7.

[104] M. Rauch, "Untersuchung der Magnetfeldtoleranz von DC-DC-Konvertern für das Upgrade des CMS-Pixeldetektors", BA thesis, RWTH Aachen University, 2013, URL: https://wiki.physik.rwth-aachen.de/pub/Feld/FeldAbschlussarbeiten/Bachelorarbeit_Max_Rauch.pdf.

[105] S. Singh et al., "CERN North Area Multi-Purpose Superconducting Magnet Facility", *IEEE Transactions on Applied Superconductivity* **32** (2022) 6, DOI: 10.1109/TASC.2022.3146809.

[106] H. Desportes et al., "Superconducting magnet for EHS", *IEEE Transactions on Magnetics* **17** (1981) 1, DOI: 10.1109/TMAG.1981.1060951.

[107] D. Lazic, *M1 Magnet in H2 Test Beam Area*, URL: https://magnet-m1.web.cern.ch/Index.html (visited on Nov. 20, 2023).

[108] Pico Technology, *PicoScope 5000 Series FlexRes Oscilloscopes*, URL: https://www.picotech.com/oscilloscope/5000/flexible-resolution-oscilloscope (visited on Nov. 22, 2023).

[109] ETS-Lindgren, *Model 7405-703*, URL: https://www.ets-lindgren.com/sites/etsauthor/ProductsManuals/Probes_Monitors/7405.pdf (visited on Nov. 22, 2023).

[110] National Instruments Corporation, *Understanding FFTs and Windowing*, URL: https://www.ni.com/en/shop/data-acquisition/measurement-fundamentals-main-page/analog-fundamentals/understanding-ffts-and-windowing.html (visited on Nov. 21, 2023).

[111] A. Bayen et al., *Python programming and numerical methods: a guide for engineers and scientists*, 2021, DOI: 10.1016/C2018-0-04165-1.

[112] A. Zografos et al., "CMS Outer Tracker Phase-2 Upgrade module noise and mitigation", 2022 *JINST* **17** C12008 (2022), DOI: 10.1088/1748-0221/17/12/C12008.

[113] S. Maier et al., on behalf of the CMS Tracker Group, *2S Module Noise Comparision*, private communication, 2024.

[114] Teledyne LeCroy, *AP033*, URL: https://www.teledynelecroy.com/probes/differential-probes-1500-mhz/ap033 (visited on May 31, 2024).

[115] Teledyne LeCroy, *WaveRunner 9254M-MS*, URL: https://www.teledynelecroy.com/oscilloscope/waverunner-9000-oscilloscopes/waverunner-9254m-ms (visited on July 15, 2024).

[116] Keysight Technologies, *E3633A 200W Power Supply*, URL: https://www.keysight.com/gb/en/support/E3633A/200w-power-supply-8v-20a-20v-10a.html (visited on July 15, 2024).
[117] N. Röwert, "Assembly, Electrical, and Thermal Performance of 2S Modules for the CMS Phase-2 Tracker Upgrade at the LHC", PhD thesis in preparation, RWTH Aachen University, 2024.
[118] J. M. Eargle, "Resistance Change with Temperature for Copper", *Electroacoustical Reference Data*, 2002, DOI: 10.1007/978-1-4615-2027-6_52.

Acknowledgement

I would like to thank everyone who supported me over the last few years and helped me in the process of writing this thesis. It is impossible to name everyone but some people deserve a special thanks.

First, I want to thank Prof. Dr. Lutz Feld for the chance to obtain my PhD in his group. He closely followed all steps and provided guidance in all aspects. Furthermore, I would like to thank Prof. Dr. Jörg Pretz for being the second assessor of this thesis. Many thanks to Dr. Katja Klein for her valuable support in all phases of my PhD.

This work would not have been possible without the current and former members of the electronics and the mechanical departments of the I. Physikalisches Institut B of RWTH Aachen University, especially Michael Wlochal, Clara Ebisch, Irfan Özen, Jan Andres, Stefan Schmitz, Waclaw Karpinski, Karl-Heinz Dreger, Franz-Josef Neus, and Gerhard Pierschel. Daniel Louis deserves special thanks for his work on the test card.

The staff, colleagues, and students of the institute made this a very enjoyable workplace. Many thanks to Tanja Bingler, Natalie Drießen, Dr. Georg Schwering, and all members of the working group.

A special thanks goes to my fellow PhD students Danilo Meuser and Nicolas Röwert, who started this journey with me when we were all Bachelor students. From all the people who shared our sometimes crowded office, I would like to thank especially Dr. Martin Lipinski for his support. Thanks to Felix Thurn for this work on testing the test cards. My thoughts are with Alexander Bogner and his family whose fate reminds us that life can be cruel and unfair and we should cherish what we are given.

At CERN, I got to experience the special spirit that this place and its collection of exceptional people have. I would like to name and thank in particular Patryk Szydlik, Angelos Zografos, and Irene Mateos Dominguez.

Many thanks to my family for their trust and patience. Without the support of my partner Hanna this work would not have been possible.

Eidesstattliche Erklärung

Alexander Pauls

erklärt hiermit, dass diese Dissertation und die darin dargelegten Inhalte die eigenen sind und selbstständig, als Ergebnis der eigenen originären Forschung, generiert wurden.

Hiermit erkläre ich an Eides statt

1. Diese Arbeit wurde vollständig oder größtenteils in der Phase als Doktorand dieser Fakultät und Universität angefertigt;

2. Sofern irgendein Bestandteil dieser Dissertation zuvor für einen akademischen Abschluss oder eine andere Qualifikation an dieser oder einer anderen Institution verwendet wurde, wurde dies klar angezeigt;

3. Wenn immer andere eigene- oder Veröffentlichungen Dritter herangezogen wurden, wurden diese klar benannt;

4. Wenn aus anderen eigenen- oder Veröffentlichungen Dritter zitiert wurde, wurde stets die Quelle hierfür angegeben. Diese Dissertation ist vollständig meine eigene Arbeit, mit der Ausnahme solcher Zitate;

5. Alle wesentlichen Quellen von Unterstützung wurden benannt;

6. Wenn immer ein Teil dieser Dissertation auf der Zusammenarbeit mit anderen basiert, wurde von mir klar gekennzeichnet, was von anderen und was von mir selbst erarbeitet wurde;

7. Kein Teil dieser Arbeit wurde vor deren Einreichung veröffentlicht.

Datum

Unterschrift