

Faculteit Industriële
Ingenieurswetenschappen

master in de industriële wetenschappen: informatica

Masterthesis

API and Back-End Development for the AuroraCatcher Network for Aurora Borealis Detection

Stijn Jacobs

Scriptie ingediend tot het behalen van de graad van master in de industriële wetenschappen: informatica

PROMOTOR :

Prof. dr. Kris AERTS

PROMOTOR :

dr. Jaroslav HRUBY

Gezamenlijke opleiding UHasselt en KU Leuven



Universiteit Hasselt | Campus Diepenbeek | Faculteit Industriële Ingenieurswetenschappen | Agoralaan Gebouw H - Gebouw B | BE 3590 Diepenbeek

Universiteit Hasselt | Campus Diepenbeek | Agoralaan Gebouw D | BE 3590 Diepenbeek
Universiteit Hasselt | Campus Hasselt | Martelarenlaan 42 | BE 3500 Hasselt



2024
2025

Faculteit Industriële Ingenieurswetenschappen

master in de industriële wetenschappen: informatica

Masterthesis

***API and Back-End Development for the AuroraCatcher Network for Aurora Borealis
Detection***

Stijn Jacobs

Scriptie ingediend tot het behalen van de graad van master in de industriële wetenschappen: informatica

PROMOTOR :

Prof. dr. Kris AERTS

PROMOTOR :

dr. Jaroslav HRUBY



KU LEUVEN

Preface

This thesis would not have been possible without a lot of support from various institutions and people, and I would like to express my heartfelt gratitude to anyone involved in getting this journey to its completion.

To start, I would like to express my deepest appreciation to Jarda, also known as Prof. Dr. Jaroslav Hruby, whose guidance, support, advice, and basil were irreplaceable in creating this work. Similar feelings are extended to Prof. Dr. Kris Aerts and Dr. Nikolaos Tsiogkas for their feedback and guidance from a technical perspective. Without them, this work would not have been possible.

This project was executed within the OSCAR-team, a group that has ensured that this work managed to fulfill its full potential. This team provided me with the tools I needed to complete this work, and I am grateful for all this support and the collaborative spirit that managed to make this something bigger. Special thanks go out to Simon Wilmots and Yarne Beerden for their advice and feedback, which I may have needed more than I am willing to admit.

Within the scope of this thesis, a prototype of what I developed was tested with the help of Michaël Bessemans and Senne Falagiarda, to whom I would like to express my deepest gratitude. They took care of most of the physical setup of the mobile observatory used during this test. By extension, I would also like to thank the Bachelor's programme in Physics for creating this opportunity.

To create this mobile observatory, a lot of help was received from the OSCAR-team, which I would like to thank again. In particular Rune Rombouts and Senne Bergmans, who created the starting code that was used to read out the detector and helped extensively when I had struggles modifying this to my purpose.

I would also like to thank my parents and sisters for their extensive support over the course of this thesis. Sometimes, style is just a me-problem.

Lastly, I would like to extend my gratitude to my friends and colleagues for all their intellectual, moral and emotional support. Sometimes it is just a hug, a pineapple, a fox, a coffee and some cheese that are needed to overcome some "Technical Difficulties".

Table of contents

Preface	1
List of figures	4
List of tables	7
1 Introduction	15
1.1 Positioning	16
1.2 Thesis Aims	16
2 State of the art.....	17
2.1 Space Weather	17
2.2 Space weather monitoring.....	20
2.3 Current Space Weather Observatory Systems	21
2.3.1 IMAGE	21
2.3.2 INTERMAGNET	22
2.3.3 SuperMAG.....	23
2.3.4 NOAA's Space Weather Prediction Center.....	24
2.3.5 AuroraReach.....	24
2.4 Magnetometry Techniques.....	24
2.5 Quantum diamond magnetometry.....	25
3 Materials & Methods	29
3.1 System design goals	29
3.1.1 The detector	29
3.1.2 System output	30
3.1.3 Other general system goals	30
3.1.4 SuperMAG's data flow	30
3.2 First prototype design	32
3.3 Data processing background	33
3.4 Software tools	35
3.4.1 Cron	35
3.4.2 Docker	35
3.4.3 Apache Web Server	35
3.4.4 InfluxDB.....	35
3.4.5 gRPC.....	36
3.4.6 KNMI's HAPI Timeline Viewer	36
3.4.7 Other non-development tools	36
3.5 Performance Testing methodology	37
3.5.1 Asynchronous gRPC call performance test	37
3.5.2 Unpacked data performance test.....	38
3.5.3 Full pipeline performance test	39
4 Experimental setup.....	41
4.1 System design approach.....	41

4.1.1	System-wide design approach	41
4.1.2	Detector & observatory requirements approach	44
4.1.3	Data ingest requirements approach.....	44
4.1.4	Data storage requirements approach.....	44
4.1.5	Data processing requirements approach	45
4.1.6	Data output requirements approach.....	45
4.2	First prototype implementation.....	46
4.2.1	Collection server prototype.....	46
4.2.2	Mobile observatory prototype.....	47
4.2.3	Data quality verification	47
4.3	Testing implementation.....	48
5	Results & Discussion	50
5.1	System design & requirements	50
5.1.1	System-wide requirements.....	50
5.1.2	Observatory requirements.....	50
5.1.3	Data ingest requirements	50
5.1.4	Data storage requirements	50
5.1.5	Data processing requirements.....	50
5.2	System architecture	51
5.3	Prototype & data verification results	51
5.3.1	Prototype performance	51
5.3.2	Data processing verification	52
5.4	Performance test results	58
5.4.1	Communication limit test result.....	58
5.4.2	Data structure test result	59
5.4.3	Full pipeline test result	60
5.4.4	Final performance testing discussion.....	62
6	Conclusion & future work	64
	References	65
	Appendices	71
	Appendix A: Requirements list	71
	Appendix B: Protocol buffer used for communication	75

List of figures

Figure	Page
Fig. 1: The Magnetosphere shaped by the solar wind, with its compressed shape on the Sun's side and its tail extending into space	p. 17
Fig. 2: IMAGE magnetometer network	p. 22
Fig. 3: INTERMAGNET's full network, showing both observatories that are operational (red) and those that are currently closed (grey)	p. 23
Fig. 4: SuperMAG's network	p. 24
Fig. 5: Diamond structure with an NV center	p. 26
Fig. 6: Energy diagram of the NV center, showing both possible relaxation paths. It shows both the luminescent relaxation path that emits red light, as well as the alternative relaxation path from $ m_S = \pm 1\rangle$ that does not emit any visible light.	p. 26
Fig. 7: ODMR spectrum with and without Zeeman splitting. The first graph shows the expected ODMR spectrum with no magnetic field present and a single peak in luminescence at 2.87 GHz. The second graph shows how these peaks split once a magnetic field is applied.	p. 27
Fig. 8: OSCAR's NV diamond detector	p. 29
Fig. 9: SuperMAG's data flow, showing the different steps data goes through between a collaborator's submission and its final form. Between these, the data gets averaged, cleaned from errors and rotated to the correct reference frame.	p. 32
Fig. 10: High-level flow of information from the NV diamond to vector magnetic field. The original ODMR-spectrum barely splits so a bias magnet is added to split the peaks and identify the 4 peak pairs. The distance between those peak pairs can then be converted into the vector magnetic field.	p. 33
Fig. 11: The commonly used approach to go from ODMR data to exact peak locations. It starts with averaging and smoothing the data to find rough peak locations, which are then used on specific measurements to find a region around them and fit a third-degree polynomial on its derivative to find the true peak location.	p. 34
Fig. 12: Setup for the first test, which compares the use of gRPC's synchronous and asynchronous calls.	p. 38
Fig. 13: Setup for the second test, which compares the performance difference between sending the detector's byte stream or the processed data through gRPC.	p. 39
Fig. 14: Setup used for the full pipeline performance testing. Here, data is sent out by (mock) detectors, received by the ingest node, stored and then further processed.	p. 40
Fig. 15: A simplified overview of the high-level system architecture, with the backup & cleanup system running regularly at a slower pace and the raw data processing system running regularly at a faster pace. The processing can run in a distributed manner except for the data recombination, which requires data from all detectors in the system.	p. 41
Fig. 16: An example of what a detector spread containing 1483 detectors would look like if the detectors are evenly spread out.	p. 43

Fig. 17: An example of what a detector spread containing 122 detectors would look like if the detectors are evenly spread out.	p.43
Fig. 18: High-level architecture of the prototype used to prove the viability of the system. It uses a single database for all data from detector to visualiser.	p. 46
Fig. 19: High-level structure of the code used for the test, with the technology-specific elements split off to a separate file to both ensure equal conditions for all tests and enable easier testing with different technologies in the future.	p. 48
Fig. 20: The full system architecture, showing individual processing & storage steps	p. 51
Fig. 21: Directional and total magnetic field data recorded by the NV magnetometer during the polar circle excursion as visualised by the web visualiser.	p. 52
Fig. 22: The magnetic field values measured by the NV-based magnetometer before filtering over the course of a day during the polar circle excursion. An arbitrary unit was used, as the purpose of this figure is to show similarities in trends, and the bias magnet makes absolute measurements misleading.	p. 53
Fig. 23: The magnetic field values measured by the reference magnetometer before filtering over the course of a day during the polar circle excursion. An arbitrary unit was used, as the purpose of this figure is to show similarities in trends, and the bias magnet makes absolute measurements misleading.	p. 53
Fig. 24: The magnetic field values measured predicted by the CHAOS model for the same points as those displayed in Fig. 22 and Fig. 23 in nanotesla.	p. 54
Fig. 25: All 8 peak locations from data collected during the polar region excursion, and the boundaries used for each to identify and filter out outliers.	p. 55
Fig. 26: All 8 peak locations from data collected during the polar region excursion after filtering out outliers.	p. 55
Fig. 27: Directional magnetic field measured during the polar region excursion as measured by the NV-detector and the reference magnetometer after filtering outliers during peak detection.	p. 56
Fig. 28: All 8 peak locations from the Helmholtz coil-generated data, showing the response to the artificially applied external magnetic field along all 3 axes.	p. 57
Fig. 29: Comparison of the directional and total magnetic field between the NV-based magnetic field measurements and the reference magnetometer's measurement under the same conditions generated by the Helmholtz coil.	p. 58
Fig. 30: Data retention rate (%) of the data ingest system when using synchronous calls from a single device, asynchronous calls from a single device and synchronous calls from multiple devices. The multi-device tests only start from the point where data loss started to show for the single-device tests to ensure the results are representative of the receiver's limits. This test has a target retention rate of 95%.	p. 59
Fig. 31: Data retention rate (%) of the data ingest module when using either an unpacked or packed detector payload for various sending rates. This has a target retention rate of 95%.	p. 60
Fig. 32: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 64Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.	p. 61

Fig. 33: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 80Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.

p. 61

Fig. 34: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 100Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.

p. 62

List of tables

Table	Page
Table 1: NOAA's Geomagnetic storms scale	p. 14

Abstract (English)

The Sun's influence on the Earth's magnetic field can interfere with satellites, power grids and telecommunication networks. These disruptions can be anticipated and mitigated by observing the Earth's magnetic field with networks of ground-based observatories using fluxgate magnetometers: a technology that requires regular manual recalibration, which disincentivises remote observatories and incurs operational costs.

A novel quantum technology based on Nitrogen Vacancy diamonds' unchanging lattice structure promises to measure absolute magnetic field without recalibration. This technology creates ODMR spectra, which require processing to determine the magnetic field. This thesis explores the design of a global system based on this technology, the ability to manage the expected volume of data, and a prototype of this system and its data processing steps.

It first shows a requirements list and architectural design for this system focussed on scalability. Then, a subset of both was implemented in a prototype, focussing on the feasibility of the required data flow in a real-world environment, testing its robustness and speed. Additional performance testing showed that the pipeline would be capable of handling 64 observatories at 1 Hz, which can be scaled up by distributing the data collection and processing service.

In summary, this thesis shows the feasibility of an NV-based earth magnetic field observatory system and paves the way for this technology to be used on a larger scale.

Abstract (Nederlands)

De Zon stoot continu geladen deeltjes uit die het aardmagnetisch veld beïnvloeden, wat satellieten, telecommunicatie en elektronica kan verstoren. Deze storingen voorspellen en hun impact beperken kan door het aardmagnetisch veld te meten met fluxgate magnetometers in grondgebonden observatoria. Deze technologie vereist echter manuele herkalibratie, wat operationele kosten veroorzaakt en afgelegen observatoria ontmoedigt.

Een nieuwe kwantumtechnologie gebaseerd op de onveranderlijke structuur van een stikstofgat (NV) in diamant kan het aardmagnetisch veld meten zonder deze herkalibratie. Deze technologie creëert ODMR-spectra, die extra verwerking vereisen om het magnetisch veld te bepalen. Deze thesis onderzoekt een ontwerp voor een globaal systeem op basis van deze technologie, verkent hoe dit het verwachte datavolume kan beheren en test een prototype met de bijbehorende stappen voor gegevensverwerking. Deze thesis toont de vereisten en ontwerp van dit systeem met het oog op schaalbaarheid. Een deel hiervan werd geïmplementeerd in een prototype dat de haalbaarheid van de datastroom in het veld test. Extra performantietesten tonen dat het systeem 64 observatoria met een meetsnelheid van 1 Hz kan ondersteunen, wat opgeschaald kan worden door het dataverwerkingssysteem over verschillende toestellen te verdelen.

Deze thesis toont de haalbaarheid van NV-gebaseerde monitoring van het aardmagnetisch veld, zodat deze technologie op grotere schaal kan worden ingezet.

1 Introduction

The Earth's magnetosphere is constantly being influenced by space weather, the group of phenomena that influences the space environment between the Sun and the Earth. The biggest contributor to these phenomena is the Sun, through its constantly emitted solar wind and more sporadic events such as solar flares and coronal mass ejections. The interaction between space weather and the magnetosphere can lead to energetic charged particles from space entering the Earth's atmosphere. Here, these particles can interact with Oxygen and Nitrogen, emitting light in the form of the Aurora Borealis & Aurora Australis. While this is a mostly innocent and appealing display, space weather can lead to more destructive outcomes as well. When enough of these charged particles enter the magnetosphere, they can cause a disturbance in the magnetosphere known as a geomagnetic storm. When intense enough, these geomagnetic storms can alter satellites' orbits, disrupt GPS signals and telecommunication, and damage power grids.

The impact of not anticipating and respecting these events has been shown repeatedly. They caused SpaceX to lose 38 Starlink satellites in 2022 and disabled the power grid of Quebec for 12 hours in 1989. Thankfully, the impact of these geomagnetic storms can be diminished when their occurrence is known in advance, which provides value to the ability to predict these events. To do so, there is a need for magnetic field tracking, which can be used to both predict upcoming geomagnetic storms and improve the models used to do so. This is done with both ground-based observatories and satellites in space, where satellites are more capable of measuring the magnetic field from multiple locations and ground-based observatories excel in measuring the changes in the magnetic field in the same location over time. These ground-based observatories can be found in larger magnetometry networks such as IMAGE, INTERMAGNET, and SuperMAG, which provide a more easily accessible format for users than these individual observatories would be able to on their own.

The main detectors used by both satellites and observatories are fluxgate magnetometers. While these are powerful detectors, they are reaching their limits with regards to size, weight and power and require regular (manual) recalibration. The former limits their use for space applications, whereas the latter makes them less fit for ground-based observatories in isolated locations. A promising technology that could overcome both of these hurdles are Nitrogen Vacancy (NV) quantum diamond magnetometers. While these detectors rely on physical principles that avoid this recalibration need and enable them to become smaller and more power-efficient than fluxgate magnetometers, their data output also requires extra processing, which raises a number of new challenges.

This thesis asks how this novel technology can be used in a global magnetic field tracking system. To answer this question, it proposes an architectural design for such a system based on a set of clearly defined requirements that consider the NV quantum diamond magnetometer's technology-specific challenges. A limited prototype of this system is then implemented and tested in real-world conditions, aiming to prove the feasibility of such a system's data processing and identifying potential practical issues that still have to be addressed. Lastly, it tests this prototype's performance limits when it comes to handling larger volumes of data, comparing multiple communication options for the processing system's data ingest service and identifying the current bottlenecks of the system when processing large volumes of data.

1.1 Positioning

This research was performed within the OSCAR team, an interdisciplinary student team at imo-imomec – the joint materials research institute of imec and UHasselt [1], [2]. The OSCAR team aims to develop novel quantum magnetic field sensors, using NV centers in diamond. The primary objective of the team is to implement this technology and bring it to space where it can be used for Earth observation and navigation purposes, but the technology can be used here on Earth as well. The OSCAR team provided the quantum diamond magnetometer used in this project. To support the prototype’s testing, a small mobile observatory using this magnetometer is developed as well. This was done by students from the Bachelor’s in Physics, which then took this observatory with them on their annual excursion to the polar region, as part of the physics+ course unit. This course focuses on offering students the opportunity to get hands-on experience with designing and implementing sensor technologies, in this case specifically for the monitoring of auroral activity.

This physical design consists of a waterproof box with the required battery packs and power convertors to power all components, the detector itself, a Raspberry Pi which served to read-out the detector, a dongle responsible for a 4G internet connection and a GPS-signal to augment the data, and a physical button to switch between two different, preprogrammed measurement modes. The software used by the Raspberry Pi started from an already developed program from the OSCAR team to read out the detector, which was then further modified as part of this project and will be discussed later.

1.2 Thesis Aims

While a number of Earth magnetic field measurement & space weather tracking systems already exist, the goal of this thesis is to explore the use of quantum diamond magnetometry for these purposes. To do so, 3 subtopics are specified:

1. The design of a system architecture for a back-end system for NV Diamond-based magnetometry, ensuring scalability, reliability and real-time data processing. This includes a communication API for data collection from geographically dispersed sensor locations and a storage mechanism for real-time magnetic field monitoring.
2. The development of data analysis algorithms for the processing of the collected NV Diamond magnetometer data, including the creation of a module to map the visibility of the Aurora Borealis and a user notification system API that provides accurate geographic predictions of potential Aurora visibility.
3. The comprehensive testing and validation of the entire system, including the back-end, API, data analysis, and user notification components, ensuring accuracy, performance, and reliability in real-world conditions.

These goals conclude in a central research question: “How can a central processing system be used to record, process and provide Earth magnetic field data for the detection of aurora borealis in a scalable and modular way”, and this this thesis hopes to pave the way for the practical use of quantum diamond-based magnetometry for large-scale Earth magnetic field observation.

2 State of the art

2.1 Space Weather

Space weather encompasses the phenomena that influence the space environment between the Sun and Earth. The main contributor to space weather is the Sun, as it constantly emits a stream of fully ionised particles flowing outwards into the solar system, which is known as the **solar wind**. It primarily consists of protons and electrons in the form of plasma, carrying the interplanetary magnetic field[3], [4], [5].

When the solar wind approaches Earth it will slow down when reaching the bow shock, which is located at around 14.5 Earth radii upstream from the center of the Earth. After this, it will encounter the Earth's magnetosphere, the region of space where the magnetic field is primarily determined by the Earth's own magnetic field. The Earth's magnetosphere is an irregular bubble that is crucial to life on Earth, as it protects the planet from having its atmosphere eroded away by the solar wind and shields us from cosmic radiation. It is generated by molten iron and copper in the Earth's outer core and is compressed on the Sun-facing side of the planet due to the constant pressure exerted by the solar wind. Conversely, it extends out significantly further in a magnetotail on the night-side (Fig. 1) [6]. The place where the solar wind meets the Earth's magnetic field is known as the magnetopause, where the pressure from the Earth and the Sun are in balance. The exact location of the magnetopause will vary based on the pressure exerted by the solar wind, but lies around 10 Earth radii upstream from the Earth's center [7].




Fig. 1: The Magnetosphere shaped by the solar wind, with its compressed shape on the Sun's side and its tail extending into space [8]

At the magnetopause, a phenomenon known as magnetic reconnection can take place, where the terrestrial and solar-wind magnetic fields merge. The efficiency of this phenomenon is strongly influenced by the orientation of the interplanetary magnetic field, and it is the primary way through which some of the mass, energy and momentum carried by the solar wind enters the magnetosphere. Otherwise, the magnetosphere shields Earth from space weather by deflecting most of the charged particles emitted by the Sun [8], [9].

When enough of these particles enter the magnetosphere, they can cause a geomagnetic storm. This is a short-term disturbance in the magnetosphere caused by space weather where energetic charged particles

from space precipitate into the upper atmosphere. While trapped in the Earth's atmosphere, these particles can interact with oxygen and nitrogen. This causes the Northern & Southern Lights, respectively known as the Aurora Borealis & Aurora Australis, but geomagnetic storms can cause technological disruptions as well [9]. The Aurora Borealis can mostly be seen between the latitudes of 55 degrees North and 80 degrees North [10].

Fortunately, the regular solar wind is typically insufficient to cause major technological disruptions. However, disruptions can be caused by more extreme events such as Coronal Mass Ejections (CME). These occur when a large volume of coronal mass of the Sun is violently ejected into interplanetary space in the form of plasma, which can reach higher speeds than the solar wind itself[11]. They are caused by solar storms, where sunspots created by abnormalities in the Sun's magnetic field collapse. The occurrence of such events varies based on the solar cycle, which reaches an activity peak once every 11 years[12]. One of the more famous examples of such an event and its impact on Earth was the Carrington Event in 1859, named after the astronomer who observed a number of collapsing sunspots, which was followed hours later by visible auroras as far south as Cuba and failing telegraph lines in Europe and North America. Next to these directly noticeable effects, the Carrington Event was so intense that it may even have damaged the ozone layer at the time. More recently, what was most likely a combination of CMEs and a solar flare caused a geomagnetic storm in March 1989 disabling satellites in orbit and causing a 12 hour long blackout in Quebec, Canada [12].

There are more recent, yet perhaps less extreme, examples of unusual space weather as well. At the end of October 2003, multiple solar flares caused one of the strongest geomagnetic storms in recent history. This geomagnetic storm lasted from October 29th until October 31st, causing auroras as far south as Texas and the Mediterranean. It also forced plane flights to divert away from the polar regions to avoid radiation, and impacted an estimated 59% of NASA's Earth and space science satellites in orbit at the time. The storm also introduced an unusual uncertainty in satellite location, which increases the risk of satellite collisions. Thankfully, none of these occurred, but the number of satellites in orbit has increased sharply since [13], [14], [15].

In July 2012, an even more severe solar storm occurred. Thankfully, this eruption that has been likened to the Carrington event missed Earth, although it would not have if the eruption happened a week earlier. It is estimated that, if it had hit, the resulting geomagnetic storm would have been twice as intense than the aforementioned storm in 1989 [16].

Even more recently, solar activity has repeatedly hit the news with notable storms in February 2022 and May 2024. The storm in 2022 caused the loss of 38 Starlink satellites by making the de-orbit as atmospheric drag changed [17], [18], while the 2024 geomagnetic storm was the first storm to reach the same classification as the one from. This led to the most southern aurora sightings on record in 5 centuries, some of them being visible in Belgium [19].

These examples demonstrate some of the present-day risks that space weather poses. The charged particles that enter the magnetosphere pose issues on their own, as they can damage and degrade satellites directly. The exchange can also heat up the upper atmosphere, which influences the atmospheric drag and may change the orbit of satellites and demands changes to avoid collisions. On top of that, the entered solar wind particles cause fluctuations in the Earth's magnetic field capable of causing further disruptions. On a lower scale, this involves degradation of telecommunication signal quality, such as GPS-signals and high frequency radio. The anomalous magnetic field can also generate induced currents. Depending on the size of the geomagnetic storm, this can cause damage in the electronics of satellites directly, but also lead to voltage instabilities in power grids, causing blackouts[6], [8], [9], [20].

Being able to predict and anticipate these issues is a practical advantage of monitoring the Earth's magnetic field, but there is also a scientific interest. The measuring of the Earth's magnetic field with the goal of monitoring space weather is one of these. For this purpose, fluctuations on time scales down to minutes can show the interactions between the magnetosphere and the solar wind. These perturbations

increase the value of magnetic field tracking near the polar ionosphere, as many current systems close through it. Next to this field of research, there is also value in tracking the Earth's magnetic field on longer time scales, which focuses on its internal source and primarily looks for variations over many years[4].

There are multiple ways to characterize geomagnetic storms. One way of doing so is the Disturbance Storm Time (DST), which measures the decrease in the magnetic field's horizontal component. The DST is expressed in nanotesla, with a DST of -50 nanotesla indicating the visibility of auroras near the polar circle while a DST of -600 nanotesla was measured during the aforementioned 1989 storm that damaged Quebec's power grid [16], [21], [22].

Another way to quantify a geomagnetic storm's intensity is the K-index, which measures how much the horizontal component of the magnetic field deviates from an expected variation over a 3-hour period. What deviation corresponds to what K-value varies based on the observatory's location and is determined by the International Service of Geomagnetic Indices (ISGI), but using a similar approach to calculating the K-index can still be used to identify the presence of auroras. K_p, the planetary K-index, is based on this K-index. The K_p index takes the arithmetic mean of the K-index from 13 observatories, 11 of which are located between 44° and 60°N, and 2 between 44° and 60° S. Because of this, the index is heavily weighted towards North America and Europe. This index expresses the intensity of geomagnetic intensity and geomagnetic storms, as shown in Table 1 [23], [24]. The United States National Oceanic and Atmospheric Administration (NOAA) uses this index to issue warnings for geomagnetic storms and their impact. This starts at "G1" for a K_p value of 5, where minor fluctuations in the power grid's voltage and auroras at higher latitudes can occur, to "G5" at a K_p value of 9. When this occurs, spacecrafts start to experience surface charging and issues with orienting themselves, auroras may be visible as far south as Sardinia at 40°N, satellite based navigation such as GPS may be degraded for multiple days, and power grids may experience widespread voltage control issues that can lead to large-scale blackouts and damaged transformers [25], [26].

Table 1: NOAA's Geomagnetic storms scale [25]

Scale	Physical measure (Kp)	Description	Effect
G 5	9	Extreme	Voltage control and protective system problems may occur in power systems. Power grid systems can encounter blackouts. Satellites face problems with communication, orientation and extensive surface charging. High frequency radio may be unusable for over a day. Auroras can be seen as far South as 40°N.
G 4	8, including 9-	Severe	Widespread voltage control problems in power systems. Satellite navigation may be degraded for hours. Auroras can be seen as far South as 45° N.
G 3	7	Strong	Power systems may require voltage corrections. High frequency radio communication may be intermittent. Satellite components may experience surface charging. Auroras can be seen as far South as 50°N.
G 2	6	Moderate	Voltage alarms may occur in high-latitude power systems, alongside transformer damage during longer storms. Auroras can be seen as far South as 55°N (New York). Spacecrafts may require corrective actions to orientation.
G 1	5	Minor	Power grids may experience weak fluctuations. Spacecrafts may be minimally impacted.

2.2 Space weather monitoring

Space weather can be monitored by measuring the magnetic field, but there are different ways in and locations from which this could be done. One of these ways is by placing a detector further upstream in the solar wind, closer to the Sun. A mission that plans to do so is SW-MAG, which will measure the magnetic field from Lagrange 1 (L1) after its planned launch in 2029. Lagrange points such as L1 are points where there is an equilibrium between the Earth and the Sun's gravitational forces, allowing this satellite to measure the solar wind in a stable location. By detecting space weather this far upstream of the solar wind, it provides a critical lead time of 10 minutes to an hour to anticipate the negative impact of a geomagnetic storm. It has been argued that for faster solar winds, this lead time would be insufficient, and that other Lagrange points would be more fit for this purpose. Regardless of which Lagrange point is used, it would suffer from the same shortcomings – while it does offer its lead time, being outside the magnetosphere would make it unable to determine what will happen precisely inside the Earth's magnetosphere [27], [28], [29].

A second way to measure the magnetic field fluctuations is by placing a detector in low Earth orbit (LEO). The first mission to do this with a vector magnetometer was NASA's Magsat satellite in 1979. The data gathered by the following LEO missions was crucial for the creation of better magnetic field

modeling and the discovery of tiny oceanic tide-generated magnetic fields, proving its value. The data collected in LEO is still used as the primary data in the creation of reference models such as the International Geomagnetic Reference Field (IGRF) [30]. Just like getting a mission to place a detector at a Lagrange point, getting a detector to LEO is expensive [31]. Inherent to being in LEO is also a level of spatio-temporal ambiguity [32]: as the magnetic field changes based on both location and time, a moving detector cannot identify changes caused solely by the time-component. There are currently multiple projects with satellites in LEO that measure the Earth’s magnetic field, such as Macau Science Satellite 1 (MSS-1) that launched in 2023 with 2 high-precision geomagnetic satellites [33] and the ESA Swarm mission with its 3 identical satellites in near-polar orbits aiming to measure the Earth’s magnetic field with unprecedented accuracy [34], [35].

A third option is the use of ground-based measurement stations. There are many of these magnetic observatories and their networks, but these will be discussed later. As these do not require a detector in space, they are significantly cheaper than both previous options. Another advantage they pose is continuous and uninterrupted monitoring over an extended period of time, up to decades [36]. Unlike space-located measurements, the geomagnetic field recorded at ground level changes significantly less, making the measurements more stable, and thus more sensitive to small changes. This shows in their use for identifying rapid core field fluctuations. They also are used to investigate magnetospheric and ionospheric magnetic field sources, which shows in the investigations of auroral substorms and geomagnetic indices such as the kP index and Auroral Electrojet (AE) index [32], [37], [38]. Deep electromagnetic studies also require a global coverage of ground-based stations over the course of multiple years [32].

Most modern observatories revolve around a triaxial fluxgate magnetometer, which is set up to function for several years without interference. That said, this type of magnetometer is not absolute and thus requires regular recalibration. This recalibration is called baseline correction, which is done with a Declination-Inclination fluxgate magnetometer (“DI-flux”) and a scalar magnetometer (such as a proton magnetometer) to correct any offset drifts in the triaxial fluxgate magnetometer [32].

Baseline correction requires human intervention and typically occurs on a weekly basis to maintain the triaxial fluxgate magnetometer’s accuracy and is the primary challenge posed for fully stand-alone observatories in inaccessible regions. Being able to automate this, or avoid the requirement for this calibration, would facilitate the creation of observatories in regions that lack permanent staff [32].

2.3 Current Space Weather Observatory Systems

Currently, a number of magnetic field tracking systems are already in existence. Systems such as IMAGE, INTERMAGNET, SuperMAG, NOAA’s Space Weather Prediction Center and AuroraReach all serve their own purpose.

2.3.1 IMAGE

The International Monitor for Auroral Geomagnetic Effects (IMAGE) [39] is a network that aims to study auroral electrojets and moving two-dimensional current systems using magnetometer stations spread between geographic latitudes from 51 to 82 degrees north. The network is maintained by institutes from Finland, Germany, Norway, Poland, Russia, Sweden, Denmark, and Iceland, although the membership of Russia’s institutes has been suspended in 2022 [39]. An overview of the included stations can be found on their website, showing 3 stations that are no longer in use, and 57 active stations (Fig. 2). The IMAGE-network as a whole has a 10 second time resolution, but most of its measurement stations also provide data with a 1 second time resolution. The network does not mention sub-second time resolution [39]. Its website also mentions that its final magnetometer data are freely available for scientific use, with most stations providing their own data under either a Creative Commons 4.0 or a Creative Commons for non-commercial use 4.0 license [39]. The network provides its data in three

formats. Firstly, there is the real-time data which provides graphs with 10-second averaged data, automatically updated once a minute. The second method is the preliminary data, showing data from the past 30 days that has not yet been checked for potential issues such as missing values or error spikes. Lastly, there is the option to request and download monthly data for larger and older timespans[39].

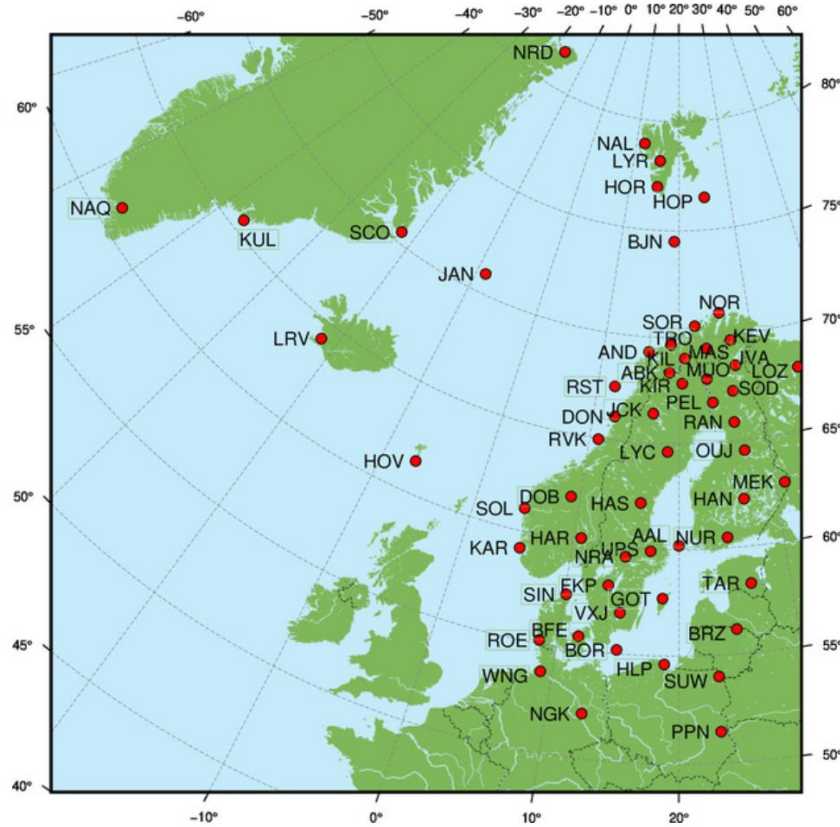


Fig. 2: IMAGE magnetometer network [39]

2.3.2 INTERMAGNET

INTERMAGNET plays a central role in the distribution of magnetic observatory data (Fig. 3). This is a global network that was established in the 1980s to create a standard in a period where observatories converted to digital measurements. It has since grown into a network with over 100 observatories spread over approximately 50 institutions, and aims to establish and maintain observatories in remote and developing regions. On top of this standard, INTERMAGNET also provides advice to institutes that aim to set up new observatories. Compared to other observatories, their standard tends to be higher and evolves to meet users' needs. It expects data to be transmitted at a near-realtime rate to its central database, after which the data is made available to the scientific community. The data is only provided under the condition that it is not used for commercial gain, although participating institutes do have the right to use their own data to cover costs. On top of the near-realtime data, participating institutions also commit to submitting definitive data annually. Data quality-wise, INTERMAGNET expects data with a sampling frequency of 1 Hz for directional fields and an accuracy of under 5 nanotesla, and a total magnetic field at 0.033 Hz with an accuracy of under 1 nanotesla. This data should be sent to INTERMAGNET within 72 hours of recording. While the values recorded in practice are typically variation measurements, INTERMAGNET recommends converting these to absolute values using the best available baseline. During this and other processing steps, the original 1Hz recorded data can be converted to one-minute means [32], [40], [41], [42].




Fig. 3: INTERMAGNET's full network, showing both observatories that are operational (red) and those that are currently closed (grey) [43]

2.3.3 SuperMAG

One of the contributors to INTERMAGNET is SuperMAG, a global network of nearly 600 ground based magnetometers (Fig. 4). Its observatories that provide absolute measurements are INTERMAGNET observatories, but most of their network provides relative measurements. As their focus is on the currents in the ionosphere and magnetosphere, it subtracts the slow-changing Earth magnetic field from its results. Their dataset is available in 2 formats, 1-second measurements and 1-minute measurements, as not all of its measurement stations are capable of providing 1-second data. SuperMAG aims to help experts, students and non-experts explore their global dataset and overcome certain obstacles related to magnetic field data handling such as correcting artefacts and errors, identifying and converting coordinate systems, and baseline removal. They do not directly operate any measurement stations, but focus on the processing of the data provided by its contributors instead [36], [44].




Fig. 4: SuperMAG's network [45]

2.3.4 NOAA's Space Weather Prediction Center

The United States' National Oceanic and Atmospheric Administration (NOAA) also has its own Space Weather Prediction System. Compared to IMAGE, it is more practically oriented than just magnetic field data collection, and is actively used for space weather forecasting and the issuing of specific warnings. This includes three specific scales for geomagnetic storms, solar radiation storms and radio blackouts respectively [25]. NOAA also provides a 27 day prediction with the expected planetary A index and largest Kp index expected for each day, a three-day forecast with more precise values, and a retrospective report with weekly recorded results [46], [47], [48]. To make these predictions, it uses near-realtime ground-based and space-based observations. In total, the Space Weather Prediction System uses 13 different models, including a 30-minute aurora forecast model[49], [50], [51], [52].

2.3.5 AuroraReach

The above-mentioned examples primarily focus on the scientific goals to work, but there are commercial services that track space weather as well. These materialise in applications that try to notify & inform users of aurora visibility. One of these is the Swedish application AuroraReach. As this is a commercial product, it is harder to identify its inner workings, but it seems to combine NOAA's models with a localised Kp-index, visibility and space weather to provide users with notifications from a certain probability of visible aurora onwards, as well as users' input of aurora pictures. It promises up to a 4 day forecast and partners with tourist destinations to help users find destinations to go view auroral activity[53], [54].

2.4 Magnetometry Techniques

To actually measure the magnetic field, there is a need for a detector. Multiple technologies can be used for this, each with their benefits and shortcomings.

One of these technologies are fluxgate magnetometers, as used by NASA's Tandem Reconnection and Cusp Electrodynamics Reconnaissance Satellites (TRACERS) mission that launched in the Summer of 2025. Fluxgate magnetometers are vector magnetometers that have a ferromagnetic core through which they detect changes in magnetic flux. They benefit from their reliability but their measurements are not

absolute, which means that the detector needs to be calibrated and re-calibrated on a regular basis. They can also be expensive compared to other magnetometers but can reach picotesla sensitivity and resolution. Their weight and power consumption can also be high compared to alternative technologies. Lastly, they can struggle with rapidly changing fields due to their low bandwidth and are sensitive to temperature changes[32], [55], [56], [57], [58].

Next to its fluxgate magnetometers, the TRACERS mission also makes use of a search coil magnetometer, which are vector magnetometers with a relatively low power consumption, small size and low weight. This makes them appealing for use in satellites. They use a high permeability core wrapped in wire, which causes a voltage to be induced when a varying magnetic field passes through the coil. This enables the device to measure fast-changing magnetic fields, but makes it equally incapable of measuring the absolute magnetic field as fluxgate magnetometers[59], [60].

Another magnetic field detection technology is the proton precession magnetometer. This is a scalar magnetometer that, unlike vector magnetometers, measures the magnitude of the magnetic field instead of a directional magnetic field. To do so, an artificial magnetic field is applied to protons (generally in water or a hydrocarbon) using a coil. When this magnetic field is removed, the precession of the protons changes based on the surrounding magnetic field, which creates a signal in the coil with an intensity proportional to the magnetic field. Modern systems using this technique manage to reach a sensitivity of 0.1 nanotesla but have a high power consumption due to its coil. They also fall short when it comes to the detection of rapidly changing fields [58], [61], [62], [63].

Optically Pumped Magnetometers (OPM) are also scalar magnetometers that have surpassed proton precession magnetometers. They are named after their laser that forces atoms into a certain state and operates at room temperature, showing a remarkable sensitivity. To reach this, they make use of the atomic spins of a vaporized alkali metal and how those react to a present magnetic field. This technology does have some limits, such as the potential for a heading error that occurs depending on the orientation of the detector relative to the measured magnetic field. It also requires proper electromagnetic shielding to avoid issues with electromagnetic radiation, which would otherwise introduce interference and disrupt the measurement of weak magnetic fields [64], [65].

Lastly, there are also magnetoresistive sensors, which rely on the way in which the resistivity of ferromagnetic materials changes depending on the angle between the applied magnetic field and the direction of the electric current running through them. Thanks to their robustness, low price and temperature stability, they are used in many ways including mobile phones, Earth magnetic field sensing in compass applications, and the read heads in magnetic hard disk drives. That said, they are also limited in how they are temperature dependent [66], [67].

2.5 Quantum diamond magnetometry

Fluxgate magnetometers are often used for space weather monitoring: ground-based stations rely on them for their main detectors, and they are often used in satellites as well. That said, they are far from perfect. Their sensitivity to stress and temperature, in combination with their recurring calibration needs, demands the support of other detectors and human intervention. Their power draw and unavoidable size due to their ferromagnetic core also leave room for improvement [32], [58], [68].

One technology that could improve on this, are quantum diamond magnetometers. This is a novel type of magnetic field detection that exploits the characteristics of a Nitrogen-vacancy center (NV center), consisting of Nitrogen and an adjacent lattice vacancy defect, in a diamond's crystal structure (Fig. 5). There are more techniques being researched that use this same concept of a magnetically sensitive atom embedded in a crystalline structure, but the best developed sensors using this technique at time of writing are those using NV diamonds. When such a diamond is illuminated by a green light with a wavelength of around 532 nm, it will photoluminesce and emit red light with a wavelength of 637nm. This photoluminescence comes from the excitation and subsequent relaxation of the pair of electrons in the NV- center [67], [70].




Fig. 5: Diamond structure with an NV center [68]

The intensity of the emitted light varies based on their spin state, as there is a different relaxation sequence that does not emit visible light that can occur from a $|mS=\pm 1\rangle$ excited state to a $|mS=0\rangle$ ground state (Fig. 6). This is the only transition in this system that is not state-preserving. This means that exciting a group of NV- centers in a ± 1 -state will reduce the intensity of the photoluminescence compared to a group of NV- centers in a 0-state, which can be measured optically, while moving more of the defect to a $|mS=0\rangle$ ground state [69].




Fig. 6: Energy diagram of the NV center, showing both possible relaxation paths [68]. It shows both the luminescent relaxation path that emits red light, as well as the alternative relaxation path from $|mS = \pm 1\rangle$ that does not emit any visible light.

The next step is to apply a microwave signal to this system. When the applied frequency resonates with the $|m_S = 0\rangle \leftrightarrow |m_S = \pm 1\rangle$ transition, it has higher probability to transfer some of the original $m_S=0$ -state to the $m_S=\pm 1$ -state with its diminished photoluminescence. Doing this for a range of frequencies around 2.87GHz and plotting the photoluminescence results creates an optically detected magnetic resonance (ODMR) spectrum, where luminescence dips indicate resonant frequencies of the present NV- centers. Without any external applied magnetic field, there will be a single extremum at 2.87 GHz. This can be split into two separate peaks for the $|+1\rangle$ and $|-1\rangle$ states by applying a static magnetic field, which is known as the Zeeman effect (Fig. 7).




Fig. 7: ODMR spectrum with and without Zeeman splitting [68]. The first graph shows the expected ODMR spectrum with no magnetic field present and a single peak in luminescence at 2.87 GHz. The second graph shows how these peaks split once a magnetic field is applied.

The distance between these peaks is dependent on the orientation of the NV- center and the strength of the used magnetic field. More precisely, it considers the strength of the magnetic field along the quantization axis of the NV- center as described by equation (1):

$$B_z = \Delta\nu / 2 \gamma_{\text{NV}} \quad (1)$$

with the magnetic field B_z , the distance between the two detected peaks $\Delta\nu$, and the constant $\gamma_{\text{NV}} = 28$ MHz/mT [68], [71], [72].

To further improve on this concept, an ensemble of NV- centers can be used instead of a single one. A single defect can only measure the magnetic field along the quantization axis of the NV-center, but as the NV- centers themselves have four possible orientations in the diamond's crystal structure, these four orientations can be used to measure the magnetic field alongside these four axes. This can then mathematically be transformed into a desired 3-dimensional frame[71].

These physical principles provide some very beneficial properties to this magnetometry technique. The earlier discussed formulas show that the calculation of a detected magnetic field does not need calibration in a traditional sense, and does not require the periodic recalibration demanded by fluxgate magnetometers. Instead of detector calibration, they do require characterization, which will be discussed later. Performance-wise, they have a wide dynamic range of 0.1 T and a theoretical sensitivity of down to around $10 \text{ fT}/\sqrt{\text{Hz}}$, although practical implementations have not reached this yet [68], [73].

While this technology is still in development, it does show promise for space applications as well. Where fluxgate magnetometers are reaching their limits with regards to sensitivity, size, weight, and power consumption, NV-based detectors show potential to outperform them on all of these, given further development. Additionally, these detectors have also been proven to not suffer from degradation from microgravity, temperature fluctuations and vibrations in space, making them even more fit for space applications [68], [68].

3 Materials & Methods

3.1 System design goals

As described earlier in chapter 2.5, NV- based quantum diamond magnetometers show potential as vector magnetometers and fit the role of measuring the Earth’s magnetic field. They overcome some issues posed by fluxgate magnetometers, like the need for recurring manual calibration and thus human intervention [32]. In space, where this is not possible, the information from multiple detectors is used to establish a baseline [32], but the specific approach and success of doing so varies between missions. This lack of recalibration makes NV- magnetometry interesting for observatories in remote locations. Despite their benefits, these magnetometers do pose certain challenges, which get exacerbated when trying to scale them up to a global-capable data collection system. The first step of designing such a system is to achieve a proper understanding of the requirements posed by both the input it aims to work with and the output it wishes to provide. The goal of this section is to create a brief overview of the top-level goals and limitations around which the system is built, which then gets elaborated upon and addressed in section 4 and 5.

3.1.1 The detector

Critical to any data collecting system is a data source, in this case the quantum diamond magnetometer. As this technology is rapidly evolving, the system should aim to work with currently existing detectors in a way that is coupled loosely enough to support future implementations of the technology as well. While it is near impossible to develop specific requirements based on technologies that do not exist yet, it is still possible to determine properties using the requirements inherent to the use of NV diamonds and the currently existing detectors (Fig. 8). This does mean that, to a certain extent, the requirements of the detector itself fall outside the scope of this thesis.



Fig. 8: OSCAR’s NV diamond detector

One of the key hurdles of fluxgate magnetometers that NV diamonds aim to overcome is the need for human intervention and calibration to create more remote-capable observatories. What this brings alongside it is a need for some form of communication from and to the magnetometer. This is to facilitate the collection of data with limited delay, but to allow remote reconfiguration of the system as well. This data collection speed is crucial to identify magnetic field changes ahead of geomagnetic storms and aurora’s formation, which is needed for the system to have any predictive capabilities. Remote reconfiguration of the device is something that should be a rare occurrence, but unusual space weather conditions could warrant changing to a more precise but resource-intensive measuring mode for a short period of time and would be valuable to ensure the feasibility of remote observatories as well. It is important to understand that there are two categories of observatories this system aims to support. For one, there are the fully self-sufficient detector modules (for example, the device discussed in 3.2).

These are observatories that can be mobile and should be treated as if they are. Practically, this means that the measured data from the detector needs to be appended with other data, such as identification, location, and orientation data, which the system then uses during its processing to reach usable results. The other category of observatories are those whose magnetometer has a set location and orientation. These omit orientation and location info, as they only send out the identifying and detector data. As location and orientation are relevant fields for processing

This magnetometer-generated data generated by both of these observatories' detectors (or "payload"-data) can be sent in two ways as well. On one hand, there is the option to send this payload data as the bytestream that was created by the detector itself; on the other hand the data can also be sent as individual, parsed fields. While the performance impact of this is tested in section 3.5 and the results of this test are discussed in section 5.4, the system should be able to support both, as certain observatory setups may lack the processing abilities for unpacking and parsing the bytestreams.

Beyond these points, other requirements related to the detector will be based on the performance of currently existing NV systems and currently existing magnetic field data collection systems such as INTERMAGNET, as one of the goals is to be compliant with or outperform these systems. This would demand a measuring speed of at least 1Hz. Specifically, INTERMAGNET demands directional magnetic field measurements at 1Hz with an accuracy of 5 nanotesla and total magnetic field measurements at 0.03Hz with an accuracy of 1 nanotesla for the total magnetic field [32].

3.1.2 System output

Next to the origin of the used data, the second critical source of requirements is the system's output. The goal is to provide both a scientific and a general audience-output – a scientific output consisting of both processed and unprocessed magnetic field data, and a general audience-output that provides information about potential auroras. While the former focuses on historical and long-term data and mostly needs raw data that is intended for re-processing elsewhere, aurora predictions and notifications introduce a limit on how long the system can take between the time of a measurement and it being fully processed.

3.1.3 Other general system goals

A core element of this system is how many detectors it should be able to support. While more data is always better, the number of detectors also raises practical constraints (such as costs). To determine the number of detectors this system aims to support, it is first assumed that this system considers any location viable for a detector to be placed, regardless of it being in the open sea or in an otherwise difficult to reach location. Then, a target detector density has to be determined. For this, the locations of the IMAGE magnetometer network's detectors are used to create multiple baselines to test around. How this is done is elaborated upon in section 4.1.

As this is a developing technology, the data processing steps continue to improve as well. For this reason, it is critical to design the system in a modular way that would enable easy replacement of current processing logic with more accurate or better performing versions in the future. Next to being modular, it is also critical for this system to be scalable. While the requirements should be designed to deal with certain load and performance goals, this should not necessarily be satisfied by single instances of the processing services or a first implementation of the system.

3.1.4 SuperMAG's data flow

A final piece of material to consider when designing the requirements and architecture of this system is how previous magnetometry data collection services function. For this, the SuperMAG data processing pipeline was analysed [36].

The process used by SuperMAG aims to address a number of issues caused by the lack of a global system before it. Some of the network's properties are critical to keep in mind for this thesis:

- it uses an identical time resolution between observatories;
- there is a shared baseline removal approach that filters out the Earth's magnetic field, to only retain information from the desired magnetospheric & ionospheric magnetic field sources;
- the system utilises a shared coordinate system.

SuperMAG utilises a unique approach to identifying the baseline for its stations. It does this to filter out the magnetic field from any sources outside the ionosphere and magnetosphere that it aims to record. To determine this baseline, it is split into 3 components: a daily baseline, a yearly trend and a residual offset. The daily baseline is used to remove fluctuations caused by the Solar quiet (Sq) current and the yearly trend aims to remove the slowly changing Earth magnetic field, which leads to a result that retains the data from geomagnetic storms and substorms [36].

For its coordinate system, SuperMAG expects all data to be in a North East Down (NED) frame. This poses some complexities as the N- and E-component are based on the current magnetic north, which can change over time. When looking at the resulting magnetic field data in these individual dimensions, it has been shown that the distribution of the data over multiple days is not symmetrical. For the N-component, this is shown as an increased number of negative perturbations, credited to the ring current, the auroral electrojet, and geomagnetic storms.

As a whole, SuperMAG's data flow is founded on the requirements that it should be fully automated, flexible to new additions, not have any hardcoded data handling assumptions, and no need for external information. It expects data to be collected and resampled to a time resolution of 1 minute by a collaborator, after which it is rotated into a shared reference frame. Then, the baseline is determined and subtracted from the measurement. At each step, housekeeping data is plotted and validation checks are done to verify the quality of the data. SuperMAG's full data flow is depicted in Fig. 9.

Next to all the automated steps, SuperMAG also utilises time-extensive manual error correction & data validation. When all these steps are completed, the resulting files are moved to a different server for visualization and use on its website.

Looking at SuperMAG's data flow shows a number of critical elements for any magnetometry data gathering system, but also leaves some concerns unattended for this system as the two systems address different steps of the data processing pipeline – SuperMAG mostly serves as more of a 'second step' for data processing, whereas the system designed in this work aims to primarily handle hurdles that arise before that. This includes processing that is not present when using non-NV diamond magnetometers. Both systems share an initial data pool with potentially varying time resolutions and an internal coordinate system without any error filtering done, which this system will aim to unify. While the baseline calculation and removal system used by SuperMAG could also prove valuable for this use case, as the resulting data reflect space weather-based fluctuations more accurately, it should be done separately from the base magnetic field calculation as both results serve their own purposes.

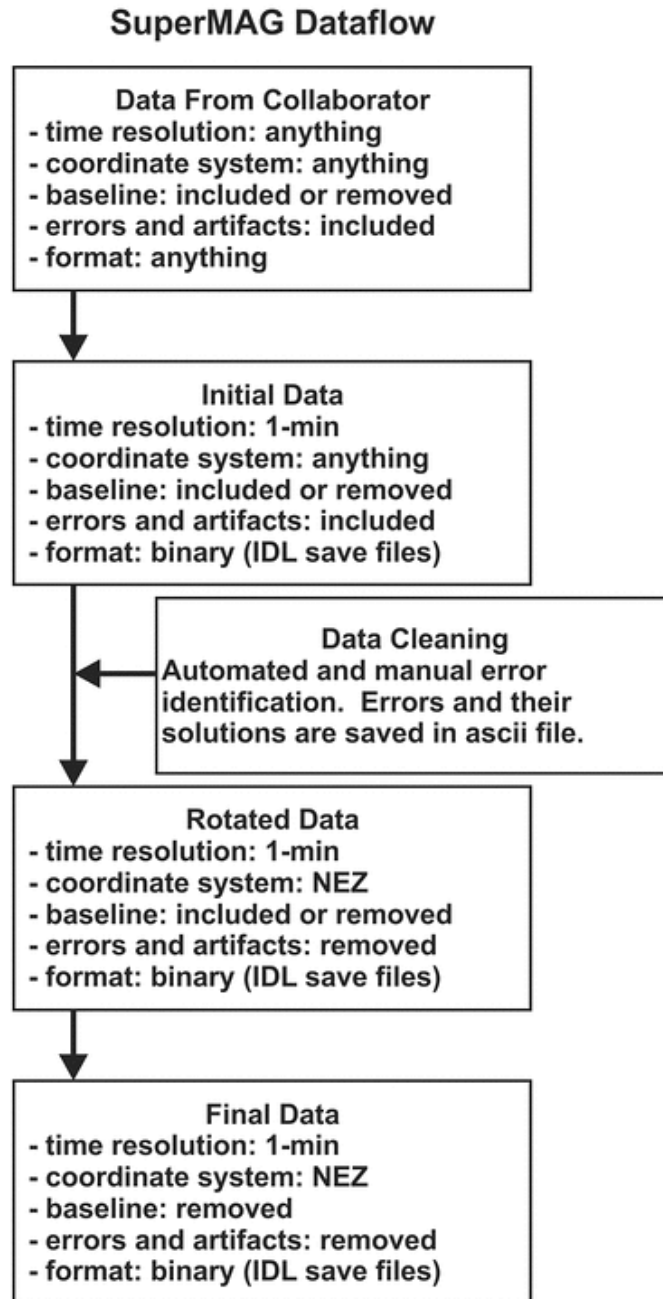


Fig. 9: SuperMAG’s data flow, showing the different steps data goes through between a collaborator’s submission and its final form. Between these, the data gets averaged, cleaned from errors and rotated to the correct reference frame [36].

3.2 First prototype design

A core element of this thesis is the design, implementation and testing of a prototype of the designed system. This prototype was then tested with a real-world setup. Due to practical concerns and limitations, this test run focuses on the functionality of the used communication methods, storage methods and data flow. For this reason, it intentionally only aims to address a subset of the full system’s requirements, specifically those directly related to the data flow and processing steps to reach a final total magnetic field value. The details of this implemented prototype will be elaborated upon later in section 4.2.

During March 2025, a group went on the aforementioned annual excursion to the polar region. As part of this excursion, this server architecture prototype was tested. A prototype mobile observatory was made in collaboration with members of the OSCAR-team. The measuring device at the core of this was an already developed quantum magnetometer, developed by the OSCAR-team. The physical design and construction of this observatory consists of a waterproof box with the required battery packs and power convertors to power all components, the detector itself, a Raspberry Pi which served as a read-out for the detector, a SIM7600-H 4G dongle responsible for a 4G internet connection and a GPS-signal to augment the data, and a physical button to switch between regular ODMR measurements and a broadsweep measurement mode, which measures the same sweep but with 8 times more datapoints. To read out the detector, the Raspberry Pi used a read-out program that started from the read-out program that was already developed by the OSCAR-team. This program was then modified to enable the button-control to switch measurement modes, data export to a remote server and performance optimizations. These will be further discussed in section 4.2 alongside the other configuration changes on the Raspberry Pi.

3.3 Data processing background

As described earlier, the output of a quantum diamond magnetometer is an ODMR-spectrum, a graph showing the intensity of red light returned when subjecting the diamond to a certain frequency microwave signal alongside illuminating it with a green laser. This is not the desired 3-dimensional magnetic field, so a number of processing steps are required to get there as shown in Fig. 10. In theory, the diamond's structure will contain NV- centers in each of the 4 possible orientations within the diamond's crystal lattice structure. The resonant microwave frequencies of the four pairs of observed intensity dips will then correspond to the magnetic field strength in the direction of each of those four directions. When looking at the ODMR-spectrum, a distance of 56 MHz between two peaks would correspond to 1 mT, which poses two problems. Firstly, it becomes impossible to identify which axis this shift would apply to, as it could be any of the four orientations. Secondly, all peaks being this close to each other makes data processing near impossible because the peaks become indistinguishable from one another. To resolve this, a bias magnet is added to the system. By adding a strong magnetic field such as a permanent NdFeB magnet to the system in the correct position, it will shift all 4 peak pairs to clearly distinguishable locations [68]. To do this, a magnetic field with different enough strengths for each of the four diamond lattice directions has to be applied to the diamond. Once this is done, all eight peaks are clearly detectable. However, one of the downsides of this approach is that this base field now has to be subtracted from any measurements.

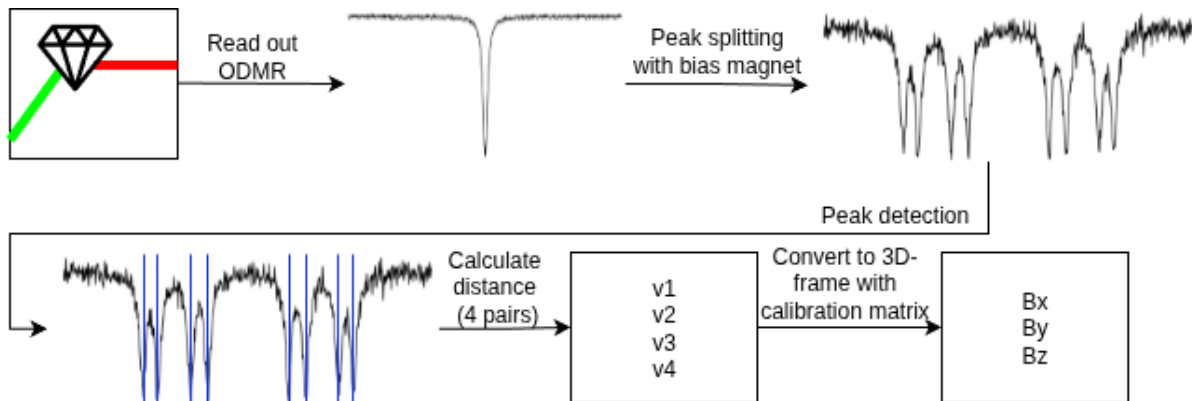


Fig. 10: High-level flow of information from the NV diamond to vector magnetic field. The original ODMR-spectrum barely splits so a bias magnet is added to split the peaks and identify the 4 peak pairs. The distance between those peak pairs can then be converted into the vector magnetic field.

With the difference in peak frequencies induced by the bias magnet known, the remaining distance between two corresponding peaks can be used with the earlier mentioned formula to calculate the magnetic field strength in that particular direction. After doing this to all four peak pairs, the resulting distances can be transformed to a regular 3D-frame. To do so, a calibration matrix is required to accurately map the four diamond lattice axes to the desired 3D frame, which is generally a chosen internal frame for the detector. From there on, the required manipulation steps are analogous to non-NV- based magnetometer data [74].

To get this calibration matrix, a known magnetic field is applied to the diamond and its effect on the peak locations is analysed, which is done with a Helmholtz coil. This is a device used to generate a uniform magnetic field through the use of two equal electromagnets positioned on the same axis. As the goal is to identify the influence of the magnetic field in three-dimensional space, three of these coil-pairs are combined to generate a 3D magnetic field vector. This is then used to generate a vector in one plane (for example, the XY-plane), which is then rotated to identify at which orientation a maximum and minimum in peak location shift can be found for each peak pair. At these orientations, this rotation is repeated for the plane consisting of this initial vector's extrema and the third axis to find the actual NV- center orientation.

As a project of this scale has to process a large volume of ODMR-sweeps, this is a process that requires proper and efficient automation. The commonly used approach to so is depicted in Fig. 11 and starts with filtering the ODMR readings with a savgol-filter, then taking the average of multiple sweeps to identify the relative extrema in it using the `sciPy` library's `argrelextrema`-function. When this finds eight valid peaks, each sweep will be individually analysed around these peak locations by iterating through its derivative. This starts at the location of the average's peak, around which a zone of interest is established between the closest local minimum and maximum of the measurement's derivative. Once these extrema are found, a third-degree polynomial fit is applied to the section, which then provides a more precise value for the exact location of the peak.

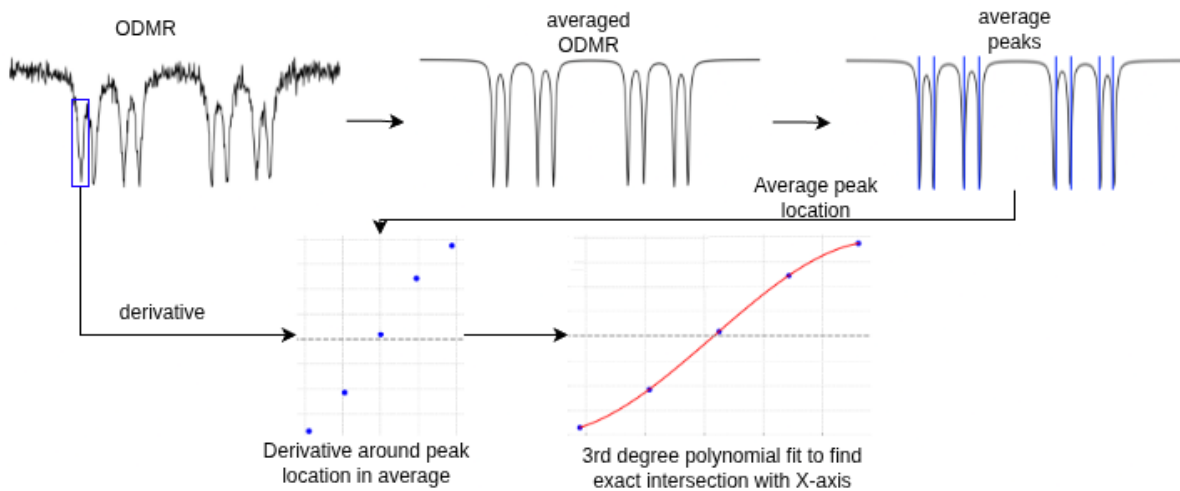


Fig. 11: The commonly used approach to go from ODMR data to exact peak locations. It starts with averaging and smoothing the data to find rough peak locations, which are then used on specific measurements to find a region around them and fit a third-degree polynomial on its derivative to find the true peak location.

As a part of the data processing, the data can be averaged or resampled to homogenise a dataset, increase the data quality, and reduce noise. This is not extensively used in this work as it does not deeply focus on that part of the data processing, but there are still some techniques that are worth mentioning. One example of downsampling is boxcar averaging, which is used by SuperMAG to downsample its data from a temporal resolution between 0.5s and 60s to a unified temporal resolution of 60s. This method groups adjacent data points and takes the average of both their dependent and independent variables

[75]. Another option is used by INTERMAGNET, which refines its data to become one-minute means centered on the minute. It achieves this by applying a gaussian filter to 19 samples of 5-second data around the minute [76].

3.4 Software tools

To develop the system designed in this thesis and its prototype, a wide range of technologies are used. Cron, Docker, Apache Web Server, InfluxDB, gRPC and the KNMI's HAPI Timeline Viewer are used directly as part of the system's development, while the CHAOS field model, Wireshark and H3Geo are used in a less direct way to support or validate things.

3.4.1 Cron

Cron is a unix utility to execute scheduled commands. This tool makes it possible to run programs periodically or upon system boot. Within the context of this project, it is primarily used to ensure certain processing steps run at the correct interval without keeping a dedicated separate process alive to do so. It is used in the observatory prototype as well to ensure all processes start correctly without further intervention beyond powering on the system [77].

3.4.2 Docker

Docker is a containerisation tool that aids developers by setting up isolated environments for programs they aim to run, making the developed program more portable and consistent as it becomes less reliant on the underlying operating system. Through this containerisation, it also brings system stability, as containers can be forced to restart upon a failure and the impact of such failures to the system as a whole is decreased. This combination of portability and system robustness makes the use of Docker in this project appealing by rendering the services more distributable and fault-tolerant respectively [78].

3.4.3 Apache Web Server

As this project utilises multiple web services and endpoints that share a host server, some form of routing or web server is required as well. For this, Apache Web Server is used. This is an open source HTTP server, and one of the three biggest web server technologies next to Nginx and Cloudflare Server [79]. Its capabilities when it comes to virtual hosts is what makes it fit this project, as there are multiple services that have to be externally available running on a single server and certain practical issues made the use of different ports unreliable [80].

3.4.4 InfluxDB

At the core of the system, there has to be a storage solution to save all recorded and processed data. For this, influxDB 2 is used. While influxBD 3 was released at time of writing, it was not yet publicly available at the start of the prototype's development. InfluxDB is a time series database (TSDB) often used for tasks that require high-speed read and write operations, as well as Internet of Things (IoT) sensor data. TSDB are often used for data that is timestamped such as application performance data and detector data, and specialise in compressed append-only time-centric data. As time stands at its core, they are optimised to handle queries looking for specific time ranges of data (such as the past minute, which will be relevant later), and handle the quickly accumulating time series data more efficiently than non-TSDB alternatives. This increase in efficiency is critical, as the collected data is expected to be high in volume and the starting ODMR-data itself is relatively large in comparison to regular magnetometry data. One of the ways in which InfluxDB reaches these performance improvements is by dividing its dataset into separate shards based on age. This is by default done once a week, but when handling data

volumes that are too large, this can still lead to performance hits and has to be tuned correctly based on use case. Creating too many shards can cause delays on queries as well, when data from multiple shards has to be fetched. Therefore, the expected query size should be considered when determining this value [81], [82]. These reasons made InfluxDB a fitting choice for the initial data ingestion of the system, and it was reused for the storage of the processed data, this time for the fast read speeds. The specifics of these decisions and use cases will be elaborated upon later.

3.4.5 gRPC

For the communication between the core server and the detectors, gRPC was used. This is a high performance Remote Procedure Call framework that uses protocol buffers to ensure that both the sender and receiver are using the exact same data structure. The shared .proto-file serves as a binding contract between sender and receiver that, while creating a highly coupled structure, ensures both sides are following the same data structure expectations and enables disconnected development of either side's internal code. Its multi-language support also facilitates disconnected development. However, using it can also bring some challenges with it, as protocol buffers bring extra complexity, the protocol relies on HTTP/2 and the highly coupled nature can complicate system debugging [83], [84]. This performance and the option for disconnected development between the detector side and server side are critical for this application.

3.4.6 KNMI's HAPI Timeline Viewer

Specifically during the real-world test, there is a need for a web interface that can show whether or not the system is functioning properly. As a starting point for this, the KNMI HAPI Timeline Viewer is used – an open source web application for visualising time series data from the Royal Netherlands Meteorological Institute. It uses the Heliophysics Application Programming Interface (HAPI) specification, and was slightly modified to replace all its current graphs and data sources with the relevant ones for the prototype test [85]. Alongside this, a small layer of middleware was developed to convert the database's output to the correct format for this visualiser.

3.4.7 Other non-development tools

Next to the tools used to develop this system directly, there are a number of supporting tools worth mentioning as well. One such tool that is briefly used is Wireshark. This open source network protocol analyser is briefly utilised to help estimate the networking load for the later performance tests [86].

In this project, there is also a need for tooling to validate magnetic field measurements. For this, the CHAOS-model is used. This is a model that describes the Earth's magnetic field for a specific location on a certain date, created by the Technical University of Denmark (DTU) [87].

A final tool that was used is h3geo, a geospatial indexing system developed by Uber. It divides the world up into similarly sized hexagons (and a low number of pentagons to make all shapes fit) at various resolution levels. For this project, it is important to be able to estimate the number of detectors a spread-out network would demand and what a certain density would look like. Where h3geo only supports specific resolutions, it is used to visualise and verify the other estimates made by calculating the average distance between hexagon centers and the number of hexagons present in a region of interest at its resolution of 2 and 3 [88], [89].

3.5 Performance Testing methodology

The final field of this thesis is the performance testing of the designed system. The aforementioned polar region expedition led to the development of a pipeline prototype that received the data of a single detector, processed this data, and served it through a webpage. Following that functional test, this section aims to test the performance of the system by starting at a low load with limited functionality, and increasing load until the system's performance degrades beyond a certain threshold. For these tests, this threshold was put at 95% data retention. The central tested device was the server provided by the OSCAR-team on which the prototype's ingestion, storage, processing, backup and visualiser service were deployed. This server has 64 Gb of memory and a cpu with 28 cores running at 4800 MHz. The visualisation of the data through the webserver was considered out of scope for these tests because it would be run on a different server in a production environment.

To execute these tests, mock modules for the detector are required to simulate the desired data volume. This could be done by developing a Docker container that mimics the behavior of a detector and recreates its data sending rate. However, the limitation of this approach arises in how stress testing would require a large volume of these containers. Because each container brings a significant amount of overhead, this could quickly start hurting system performance and thus the validity of the tests. Therefore, a single program that sends data at a set (but variable across different tests) rate is used instead. This corresponds to an equivalent number of detectors at a certain sending rate. The used rates are 1Hz and 3Hz, which are INTERMAGNET's standard and a rounding of the observed data sending rate during the polar region excursion respectively. To ensure any performance limit is not caused by the data senders, tests that display a limit that could be caused by them are re-ran using multiple sending devices that combine to the same sending rate.

The initial test needs an estimate of what the system's limit would be to properly understand what constraints should be tested around. For this, the data transmission capabilities of a single sender are analysed, for which the transmission load of a single packet is captured using Wireshark. This shows the individual messages sent and received, as well as their size, for a total exchange of 4919 bytes for a single measurement. To reach the limits of a hypothetical 50 Mbps connection, this would require 133 packets per second. Before the first test, this is recalculated based on the observed network speed at the time and the test targets are adjusted accordingly.

3.5.1 Asynchronous gRPC call performance test

The first test aims to assess the performance of gRPC itself. As shown by Fig. 12, the data senders, data ingest system and raw data-database are running for this test, with an increasing data sending rate to identify when this fails. This test compares two sending modes, based on the sync and async options provided by gRPC. The first option sends a synchronous call and actively checks the value returned by the server. When this errors, it will actively try to re-transmit the packet. The second approach uses the async method to send data, and omits any active checking on the response provided by the server. Because gRPC is based on HTTP/2 which by nature will verify connection status and result, the expected difference is minimal, but the synchronous sending could lead to overloading the server less gracefully and causing issues on the detector's side as well.

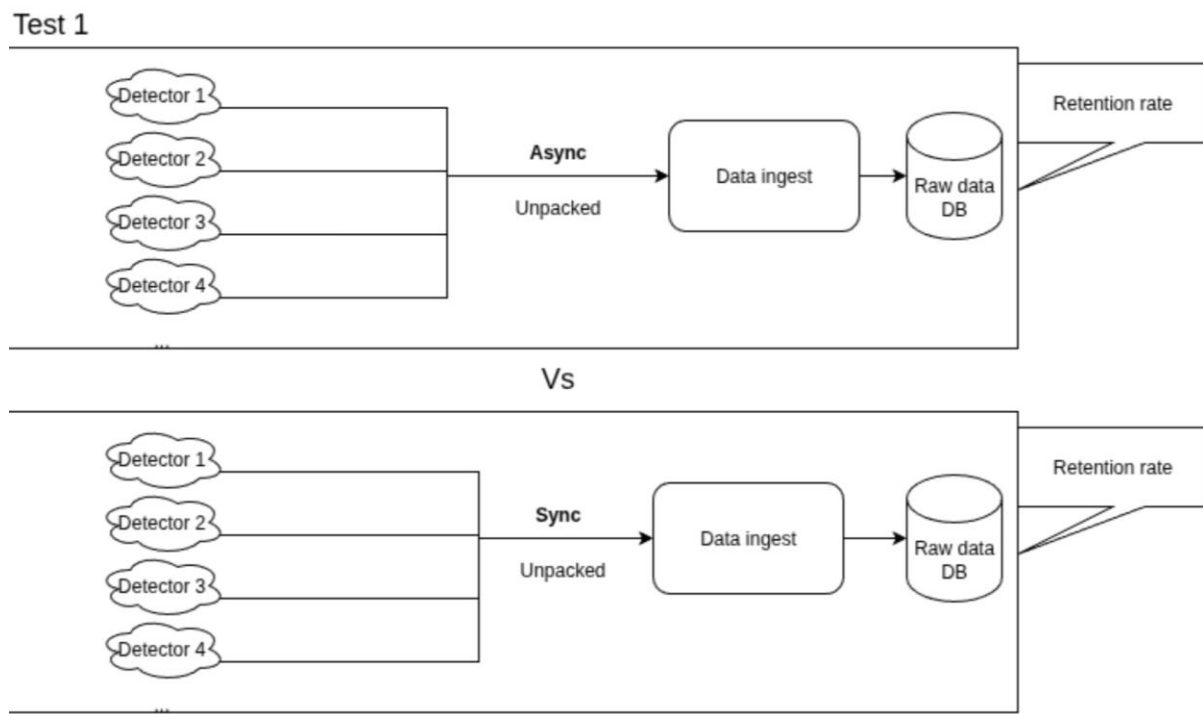


Fig. 12: Setup for the first test, which compares the use of gRPC's synchronous and asynchronous calls.

3.5.2 Unpacked data performance test

The second test will start from the results of the first test, but test the nature of the transmitted payload instead. The approach for test 1 is to unpack all data into human-readable values before sending them through the gRPC-method, while this data originates from a byte stream provided by the detector. In this test, the performance between the best results from test 1 are compared with similar configurations that transmits the unprocessed data, as seen in Fig. 13. The hypothesis is that the transport load will decrease as the "packed" data is stored in a more efficient manner, and the system will become more performant as the load provided by unpacking server-side is smaller than this loss. This could, however, be wrong when the bottleneck during test 1 is the processing of data already, and not the networking capabilities of the server.

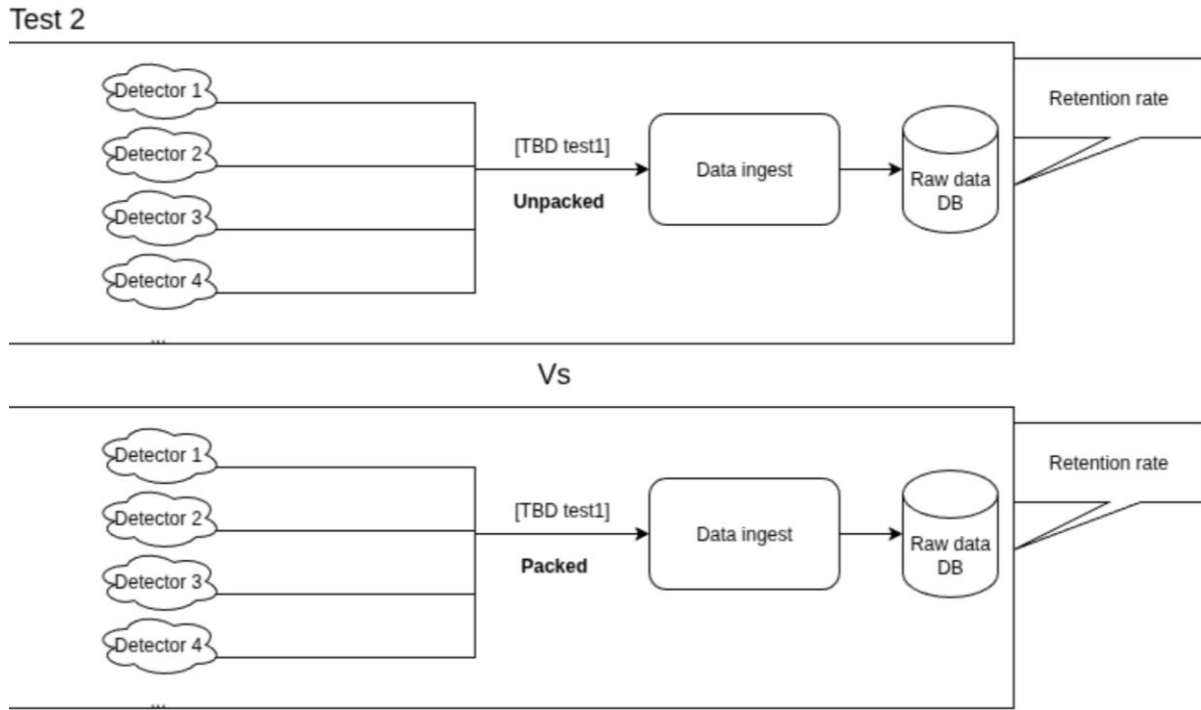


Fig. 13: Setup for the second test, which compares the performance difference between sending the detector's byte stream or the processed data through gRPC.

3.5.3 Full pipeline performance test

A major load factor of this system is the further processing of the received data, which the third test aims to assess. As shown by Fig. 14, the processing of the raw data is added to the system for this test. To validate how much data was fully processed, results provided by the processed data database are analysed to see how many measurements get processed each minute. Next to this, the processing time and raw data database's retention rate are recorded as well. For communication style and configuration, the more performant options during test 1 and 2 are reused.

This test intentionally leaves out multiple factors. For one, human load is left out because of its unpredictability. Secondly, the back-up system is not activated for this test, as it runs so periodically that it would need its own, separate test with an already properly populated database to assess well. Lastly, the recombination steps of the notification service and map generation are excluded from this test. This is because up until the point of the ODMR processing step, the large volume of transported data strongly incentivises trying to run the connected services on the same hardware to avoid networking-related performance limits. These steps could also be easily distributed over multiple servers, as they all concern the data of a specific set of detectors. The peak and diagnostics databases are included for ease of validation of the test, but the recombination steps would require the data from all used detectors, which strongly encourages moving these services to a different device as soon as any load concerns arise. In practice, moving the magnetic field database to a different server can also be considered to avoid an unexpected user load interfering with the system's data collection, which will be discussed in the following sections.

Test 3

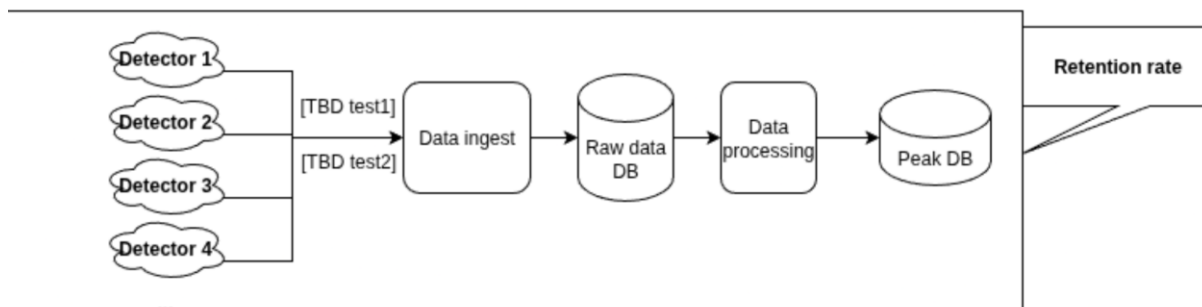


Fig. 14: Setup used for the full pipeline performance testing. Here, data is sent out by (mock) detectors, received by the ingest node, stored and then further processed.

4 Experimental setup

4.1 System design approach

As described in section 3, there are a wide range of factors to consider when designing the system and its requirements. This starts with a set of system-wide requirements revolving around the system's general functionality, after which the specific subsystems will be discussed.

4.1.1 System-wide design approach

Above anything else, the system should be designed in a way that is scalable to fit the needs of the data network as a whole. To reach this goal, the server architecture is designed to be horizontally scalable for all steps except the final one, which actively recombines data from different detectors and thus needs to be more centralised. This means that, if a single data ingest system were to fail at handling a desired number of detectors, it would be possible to add a separate server that is responsible for the ingestion of data from a subsection of the detector pool. To combine this with the possibility to run the system in a distributed manner, the server-side code was split up in five chunks (Fig. 15):

- Firstly there is a raw data ingestion and storage system. These are not separated, as doing so would incur excessive network traffic. This chunk is responsible for taking in the raw data, and storing it for the desired amount of time.
- Secondly, there is the raw data's backup and data cleaning service, which could be split off if performance reasons demand it. This chunk is responsible for creating periodic backups of the raw data database and the removal of outdated data. As is usual for backups, their output should be stored on a different device.
- The third chunk is run periodically and responsible for the raw data processing. This can be offloaded to a different device, as this is likely one of the most load-intensive elements present in the entire system and could be replaced with even more intensive, AI-based systems later.
- Fourthly, there is the diagnostics & processed data database, which can be accessed through a web interface. This should be split off to avoid any data collection or processing issues as a consequence of users requesting too much data and overloading the database.
- The last chunk is the data recombination system, another periodically running process responsible for any steps that require multiple detectors' data to succeed (e.g. interpolation).




Fig. 15: A simplified overview of the high-level system architecture, with the backup & cleanup system running regularly at a slower pace and the raw data processing system running regularly at a faster pace. The processing can run in a distributed manner except for the data recombination, which requires data from all detectors in the system.

Next, the expected detector load has to be calculated. This is done as introduced in section 3.1 by using the IMAGE magnetometer network's estimates for the average, minimum and maximum distance between two neighbouring detectors. This results in an estimated minimum, average and maximum distance of 52.88km, 220.25km and 700.62km. The downside of this approach is that it only considers a magnetometer's closest neighbour and fails to properly account for isolated clusters of observatories: hypothetically, 4 detectors could have a very low average nearest neighbour distance when positioned as two pairs of adjacent detectors with a large distance between them. For this reason, these numbers are validated by calculating the area covered by IMAGE to determine the average area each detector would cover if they were evenly spaced out. Simply looking at the minimum and maximum longitudes and latitudes to assert a region would provide poor results, as the most western-located observatory (Narsarsuaq at 45° West) would create a large zone that IMAGE does not aim to cover (Fig. 2). That is why a polygon is generated instead. This approach is limited as well, as it ignores the region covered outside the most distant magnetometers and empty, uncovered zones within it. This results in detector-circles with a radius of 182.78km, or a nearest neighbour-distance of 365.57km – significantly higher than the estimated average and minimum distance.

From these distances, three scenarios are determined for this system: as bare minimum, IMAGE's most remote observatory has to be considered, which gets rounded to 700km as the nearest neighbour distance. To estimate the upper limit of the system's needs, a distance of 200km is chosen, as it falls between IMAGE's minimum distance of 52km and the average nearest neighbour distance of 220km. This is intentionally on the lower end to compensate for locations that are too inhospitable to support any kind of observatory. Lastly, as a middle ground between these two, a distance of 350km is chosen. This final distance is not used within the system requirements, but will be used to contextualise load capabilities.

Using these distances, the detectors needed to fill the aurora-relevant region between 55 degrees north and 80 degrees north is calculated. In any situation where an even distribution of detectors is assumed, the given area is filled with equally sized circles. These circles have a radius that is half the nearest neighbour-distance, and have a 0.9069 packing efficiency factor applied to them [90]. This factor corresponds to the efficiency of hexagonal sphere packing and, while not perfect, should be sufficient for the estimates this project aims for. Following this reasoning and the aforementioned numbers results in a fillable region of $42.2 \cdot 10^6 \text{ km}^2$, which results in a detector count of 1483, 485, and 122 for a nearest neighbour distance of 200km, 350km, and 700km respectively. What such a spread could look like is shown in (Fig. 16) for 1483 detectors in this region and (Fig. 17) for 122 detectors in this region.




Fig. 16: An example of what a detector spread containing 1483 detectors would look like if the detectors are evenly spread out.




Fig. 17: An example of what a detector spread containing 122 detectors would look like if the detectors are evenly spread out.

As this approach has gone through multiple assumptions and estimates, these results are compared with the number of tiles h3geo would place in the same region. Where the earlier used circles could use the diameter to indicate the distance between stations, this approach instead relies on an estimate of the distance between a hexagon's center and the middle of its edges. For the relevant h3geo resolutions this

results in 92 detectors with a spacing of 826.5km at resolution 1, 544 detectors with a spacing of 315.0km at resolution 2, and 3545 detectors with a spacing of 118.93km at resolution 3.

4.1.2 Detector & observatory requirements approach

As discussed earlier, this system revolves around an observatory that generates data. Using a network connection, this data then has to be transmitted to the central system within a limited timeframe to ensure “real-time” capabilities of the system similar to those of IMAGE as described in section 2.3. The first observatory-related hard requirement this poses is the maximum delay permissible between recording and transmitting, which is 30 seconds. It also poses a soft requirement set at 10 seconds for a more responsive system. The data should be sent in either a packed or unpacked form, as described in section 3.1.1.

This same network connection is also the method through which the detector should be configurable, which is the second observatory-related requirement. The specific parameters to configure are left unspecified, as these may change depending on the detector and future developments in measuring methods.

The final requirements directly imposed on the detector are its measuring speed and accuracy. The aforementioned requirements imposed by INTERMAGNET would demand a measuring speed of 1Hz from the detector. That said, the total magnetic field values after processing can be at a significantly lower time resolution than the measurements themselves. This means processing techniques such as those discussed in section 3.3 and 5.3.2 (averaging and filtering) can be used to improve on the single-measurement accuracy. While this 1Hz is a minimum data collection speed and should be sufficient, it should be possible to increase this measuring speed to improve the data quality during unusual conditions.

In summary, the observatories are expected to be able to reach a high enough measuring speed, transmit this data fast enough after it is recorded, and receive reconfiguration commands remotely.

4.1.3 Data ingest requirements approach

Once data is produced, it needs to be sent to the system. For this, a data ingest system is needed, responsible for receiving all data produced by the detectors and storing it. Aligned with section 3.1.1, this system must be capable of receiving packed and unpacked payload data, and process both location-augmented and payload-only packets. If a packet does not contain location data, they should be added in this step, as all data after this step should be uniform. This step also records both the observatory’s measurement time and the system ingest time, with the latter serving as a back-up to avoid complete data loss in the case of observatory misconfiguration issues. The data ingest system also has a performance requirement, as it should be able to handle the data of all the earlier-mentioned detectors at their expected measuring speeds. This data, after properly being ingested, will then be saved in the raw data database.

4.1.4 Data storage requirements approach

The storage system has multiple responsibilities, each providing similar but slightly different requirements. There are four categories of stored data: recent unprocessed data, long-term data, processed data and diagnostics data.

The recent unprocessed data has to be easily accessible by the data processing system and should therefore be stored in a quick-reading database. This database should be able to store two weeks worth of data for all detectors at an average measuring rate of 5Hz. Data that exceeds two weeks in age should be removed from this database.

Every day, this database should export a backup of the past 24 hours, which is also saved to the long-term data storage system. This data should be stored forever, but this can be done in a more compressed

and less efficiently accessible manner. This data should also be accessible through an API that allows a direct download of the compressed files.

The processed data database once again requires quick read- and write-capabilities, and should be capable of storing 15 years worth of data at 5Hz. This data should also be backed up daily to the long-term data storage system. Any data older than two years in the processed data database can be resampled to a time resolution of 0.03Hz to reduce storage requirements as well. Next to the directly processed data, this database is also responsible for the storage of recombined data, for which the same standards apply.

Lastly, there is the diagnostics database, which is responsible for monitoring how many measurements are received and properly processed. It also keeps track of any unexpected issues that occur during any processing steps. The former should have a retention period of two weeks, whereas the latter should be retained indefinitely.

In general, no explicit storage demands are used in this section, as these can vary based on the compression capabilities of the used database and compression algorithms.

4.1.5 Data processing requirements approach

There are multiple ways to process the NV- magnetometers' data, some of which require methods such as averaging nearby data to reduce noise. For this reason, fully processing all data right when it arrives is challenging. Instead, this subsystem shall run periodically to convert the collected ODMR data into magnetic field data, apply the desired filters, transform it into the correct reference frames, and potentially remove a baseline before storing its results in the processed data database. These results are both the absolute magnetic field measured and the magnetic field measured after applying baseline correction as described in 3.1.5.

To maintain the system's goal of providing near real-time data, there is an expectation that this processing happens often enough and fast enough.

Next to these processing steps, there are also the post-processing or recombination steps. These use multiple processed measurements, potentially spread out over multiple detectors, and recombine them in some way. This subsystem is responsible for a map with interpolated magnetic field values and the detection of abnormal deviations in the recorded magnetic field.

The regular processing and recombination steps shall both be implemented as stand-alone modules, as it is critical that these can be replaced with as little downtime as possible to improve or test new processing algorithms. These steps are also expected to provide output to the diagnostics database for a quick overview of the system's general wellbeing. Next to missing data or failing processing steps, unusual magnetic field fluctuations above a certain threshold should be sent here as well. This ability to re-process data and test new processing algorithms is also one of the reasons why it is critical to keep the raw ODMR stored for this long after its initial processing.

4.1.6 Data output requirements approach

The expected outputs of this system can be classified into three categories of output: the raw data output, the administrator interface, and the web interface.

Firstly, there is the aforementioned raw data output. These are accessible as compressed files that can be downloaded directly from an API and are primarily directed towards research purposes.

Secondly, there is an administrator interface. This is solely accessible to privileged users, and shows the data collected by the diagnostics database. Next to the diagnostics data, this interface also can be used to send commands to detectors to reconfigure them.

Lastly, there is a publicly available web interface with multiple data views. This contains a graph showing the magnetic field by time-graph for individual stations and the interpolated map. This web interface exists next to an app that shows the same visualizations, but also receives push notifications

when a deviation above a certain threshold is recorded for a selected location. This threshold is based on the fluctuation in magnetic field compared to a baseline, and may be recalibrated based on the observatory's latitude.

4.2 First prototype implementation

4.2.1 Collection server prototype

The primary goal of this prototype is to prove the viability of this system. As such, what it aims to show is a process of remotely collected ODMR-data that gets processed at near real-time speeds and displayed through a web interface. For this, the following services were created, each in their own docker container (Fig. 18):

- The data ingest service, a python script that collects data from the detectors over gRPC and passes the data to the database. While it was written to also support simple command transmission and “packed” data, these options were not used during this field test.
- The influxDB container stores the collected data. For this prototype, the same database is used for the raw data and processed data, although both get an influxBD bucket for their own data.
- The data processing container, a python script that converts ODMR data into directional magnetic field data and uses GPS-based data to correctly transform it to a universal reference frame, namely ENU (East/North/Up). To do this, it uses the “course”-value provided by the GPS-sensor, which demands that the detector is always oriented correctly relative to its movement direction. When no course value is present, it is assumed that the detector is stationary and oriented in a way that avoids the need for this rotating step, as agreed upon with those responsible for handling the observatory prototype.
- A processed data output API, which provides the processed magnetic field data used by the web visualiser. This is added to allow for more data flow control and to avoid publicly exposing the database itself.
- The web visualiser itself, using the data API to get its data. This is a modified & simplified version of the aforementioned KNMI HAPI Timeline Viewer trying to visualise the X, Y, Z and total magnetic field values over a certain period.

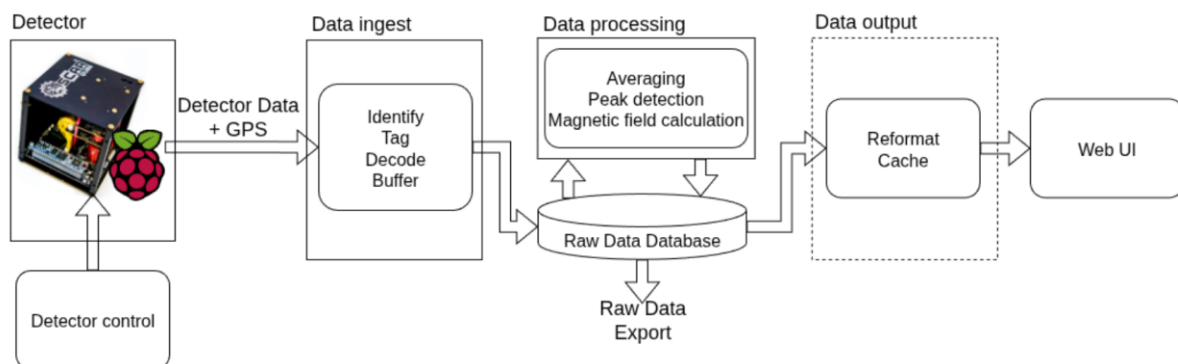


Fig. 18: High-level architecture of the prototype used to prove the viability of the system. It uses a single database for all data from detector to visualiser.

Due to firewall issues, all of these services outside the database itself were rerouted through port 8080 using Apache's virtual hosts. This becomes more complicated for the ingest service, as gRPC requires HTTP/2. The ingest service also demands modifications on the client's side, as the standard gRPC python library does not support directly adding a subpath to the destination URL. To make this work, it has to be added to each method call directly. To streamline this process, a `prefixPathInterceptor` intercepts the used channel and modifies all method calls that are sent.

As the data processing container does not run continuously, it needs a system to repeatedly start it. For this, a cronjob is used, running this container once every minute and looking at the past 90 seconds of data. This is to ensure that no accidental data loss occurs in this phase. As this does introduce extra load, influxDB will de-duplicate this data upon receiving it.

4.2.2 Mobile observatory prototype

Next to the server that takes, processes, and shows data, this prototype also needs a portable observatory to provide its data. The physical setup and regular detector read-out code for this were provided, but the software changes to make this mobile observatory work do fall under this thesis. At the core of the earlier described box is a raspberry pi, who receives data through an ethernet connection from the detector and receives GPS and cellular signal through USB. Both of these connect through IP interfaces, but which receives priority can vary. The first detected device will be configured as "default" with highest priority, but this could lead to a situation where the detector is seen as the default gateway to the internet. To avoid the potential lack of cellular connection this could cause, a script runs 100 seconds after boot to remove the default interface of the detector if this were to be the one with priority. 20 seconds later, the detector readout code starts. This ensures that the data from the detector comes in correctly while the Raspberry Pi also remains capable of properly transmitting its data to the server.

This detector readout code also required multiple changes. The original system uses a PyQt interface, but this was removed from this version as it could cause performance issues under specific configurations. The next major change to the readout code was the addition of gRPC. The code already had an export option that saved its collected data to a file, but this is now also being sent to the server described in section 4.2.1. To avoid potential performance issues at high measuring rates, these gRPC calls are done asynchronously.

Next, a button with a small LED indicator was added to the raspberry pi. This LED functions as a status indicator to show that the observatory was still running without relying on the full pipeline. When pressed, the button will change the measuring mode from regular ODMR sweep to broadsweep as mentioned in section 3.2. Changing modes will also initialise different blinking behaviour in the LED. This measurement mode switching bases itself on two locally stored configuration files, but in a later iteration this should be doable remotely instead.

4.2.3 Data quality verification

With this prototype up and running, the next step is to look at the output of the processing pipeline to check the performance of the system. The first step here is to simply look at whether or not the pipeline runs. Once this is confirmed, the nature of the output data is looked at. This showed some issues which will be discussed in section 5.2, but these were mostly related to the detector and observatory setup itself. Afterwards, a specific period with a moving detector was analysed. This data was then visualised to identify a correlation between the measured magnetic field using NV-, the measured magnetic field using the reference magnetometer in the device, and the expected magnetic field values from IGRF. From there on, the data is further processed and analysed to identify issues in the data processing itself.

4.3 Testing implementation

The tests are executed as described in section 3.5, and reuse parts of the server-side code from the prototype. The primary modifications are that the processing node is changed to properly support multi-detector input, and the database is reconfigured to use sharding more efficiently as this was misconfigured during the prototype's test due to an oversight. Instead of the default shard size for a database with no automatic retention limit, the raw data bucket is now configured to use a shard size of 5 minutes. For the first two tests, only the ingest node and raw data database are turned on. For the third test, the data processing service is added as well.

The test is considered a success for a certain load if the data retention rate over the entire test manages to remain above 95% for a test of 10 minutes. The data senders are set up to provide data for 15 minutes, but approximately the first 2.5 and last 2.5 minutes are omitted to avoid potential desync between senders leading to an unfairly reduced retention rate. This becomes a bigger concern once coördination between multiple sending devices is introduced. After the test is finished, the data stored in the database gets analysed for retention percentage per minute.

The data sender itself is a python script that reads an output file from the prototype's test to get a realistically shaped payload. This payload is then transmitted at a variable rate to the data ingestion service. The code itself is split into two files as shown in Fig. 19, with one responsible for the timing and read-in options, and the other responsible for any gRPC-specific elements. This makes it so this test suite could be easily re-implemented using a different technology by rewriting the second file. All it needs to have for this are methods responsible for setup, payload processing and the data sending itself. This last call is made in a multithreaded manner, using a thread pool of 16 workers. This is to mimic the code flow from the prototype device and to reach a higher load before requiring multiple sender devices.

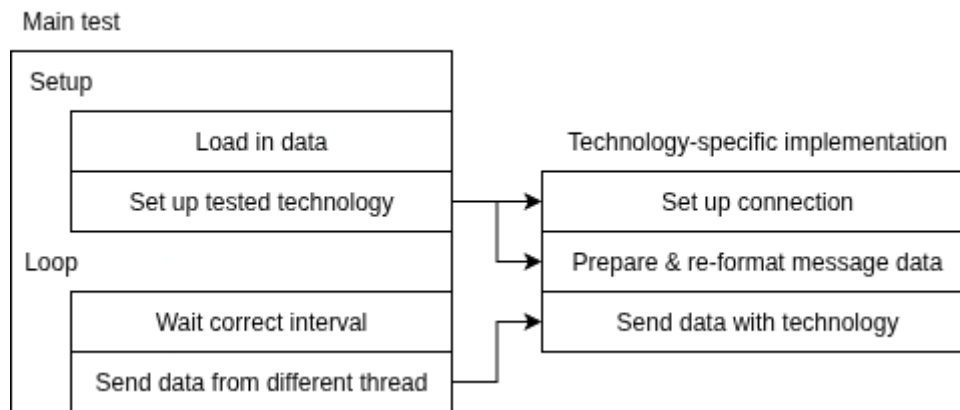


Fig. 19: High-level structure of the code used for the test, with the technology-specific elements split off to a separate file to both ensure equal conditions for all tests and enable easier testing with different technologies in the future.

For the first test, there is an expected failure for a single sending device at 133 Hz due to networking limitations. For this reason, the test starts with the synchronous communication at a rate of 8 Hz and doubles until the system starts to degrade. Once this is observed, load volumes around the failing value will be performed again from two different devices simultaneously to ensure that the observed limit is on the receiving end, and not an issue with the data senders. Afterwards, the asynchronous communication is tested around this point as well, and the volume is increased or decreased based on the observed performance.

The second test repeats this approach, but reuses the results of the superior version in test 1 and observes the retention rate of the packed version at the same sending rate as what the previous test concluded. Then, this is decreased or increased to find an accurate limit as well.

The third test starts at a data sending rate of 64Hz after which the data rate increases or decreases at a similar rate once more. The distributed attempts are only considered for this step when the data sender limits from test 2 are reached. Instead of having just a direct retention rate to test for like this previous tests, this tests 3 factors. The first is the retention rate in the raw data database just like test 1 and 2. The second is this same retention rate, but in the processed data database to validate how much data manages to actually get processed. The third factor to observe is processing speed. For this, all received data is normalised, and the delay between arriving in the raw data database and arriving in the processed data database is analysed.

5 Results & Discussion

5.1 System design & requirements

As discussed in section 4.1, this system poses a set of specific requirements. The specific list of these can be found in Appendix A, although a number of these demand some further justification. The requirements are split up into 6 groups: system-wide requirements, detector requirements, data ingest requirements, data storage requirements, data processing requirements, and data output requirements.

5.1.1 System-wide requirements

This section concerns more high-level requirements of the system. This starts with the demand to support at least 122 detectors and aim to support 1483, as discussed in section 4.1.1. The system also aims to fulfil INTERMAGNET's requirements, which demand a 1 nanotesla accuracy on total magnetic field measurements at a speed of 0.03 Hz.

Next to these, the system also demands a certain speed for the full processing of data. As such, there is a hard requirement to have a measurement fully processed in at most 10 minutes, and a soft requirement to be able to do so in 5 minutes or less.

5.1.2 Observatory requirements

The observatories and their detectors within this system also have their own requirements. One of these is the aim to comply with the requirements posed by INTERMAGNET – the aforementioned 1 Hz time resolution with a 5 nanotesla accuracy for the directional magnetic field measurements.

5.1.3 Data ingest requirements

The data ingest system is responsible for receiving and processing all raw data, and is expected to do so with a retention rate of at least 95%. The number of measurements per second it should be able to satisfy is calculated by assuming either 122 or 1483 detectors as calculated earlier, and assuming these provide data at an average speed of 5Hz. This leads to either 610 or 7415 measurements per second. We thus specify 610 measurements per second as a hard requirement and 7145 as a soft requirement.

An important matter to note on these requirements is the distinction between a data ingest node and the data ingest system. Whereas a data ingest node refers to a single instance which may handle one or more detectors as described in section 4.1.1, The data ingest system refers to all present data ingest nodes in the system when fully deployed.

5.1.4 Data storage requirements

The data storage requirements reuse the same load numbers as the data ingest requirements. This is combined with the aim of retaining all unprocessed data for 2 weeks, which leads to 737.856.000 and 8.969.184.000 stored measurements respectively. Starting from these same detector counts, there is also the aim to store the processed data for 15 years at a time resolution of .5Hz, leading to 28.2 billion and 342 billion stored measurements respectively.

5.1.5 Data processing requirements

There is an expectation that this system can provide near real-time data, which shows itself in the 5- and 10-minute requirement to fully process data. To support this, there is an expectation that the data processing system is run at least once a minute. Depending on hardware performance and further testing

of the system's implementation, it is possible to run this more often as well, but this should be managed to balance between system speed and processing capabilities. The result of these processing steps shall then be saved to either the magnetic field database or diagnostics database, the former keeping all correctly processed data, and the latter storing any observed problems during processing.

There is also an expectation that the data processing system can be replaced with different processing logic without impeding the system's processing speed requirements, namely the aforementioned 5 and 10 minutes between measurement and processing.

5.2 System architecture

All this information leads into the proposed system architecture shown in Fig. 20. It starts from a remotely configurable detector which generates data in the form of ODMR-spectra, as well as housekeeping data and potentially GPS data. This data is then ingested and processed into a homogenous form by the data ingest system, which forwards this data to the raw data database. This database receives regular backups, provides an API to download this unprocessed data, and passes it on to the data processing system. This is a system that runs periodically to take the ODMR-data and transform it into the directional and total magnetic field values the system aims to generate. Diagnostics results of this process and the system in general are sent to the diagnostics database, and the magnetic field data itself is forwarded to the processed data database. The diagnostics database provides its data to the administrator interface, allowing a maintainer to properly overlook the system's wellbeing. Meanwhile, the magnetic field database provides its data both directly to the web UI, as to the data recombination service. This final service recombines data from multiple detectors to identify abnormal fluctuations and interpolate magnetic field values. Abnormal fluctuations are flagged and trigger a notification system, while the interpolated maps are served to the web UI as well.

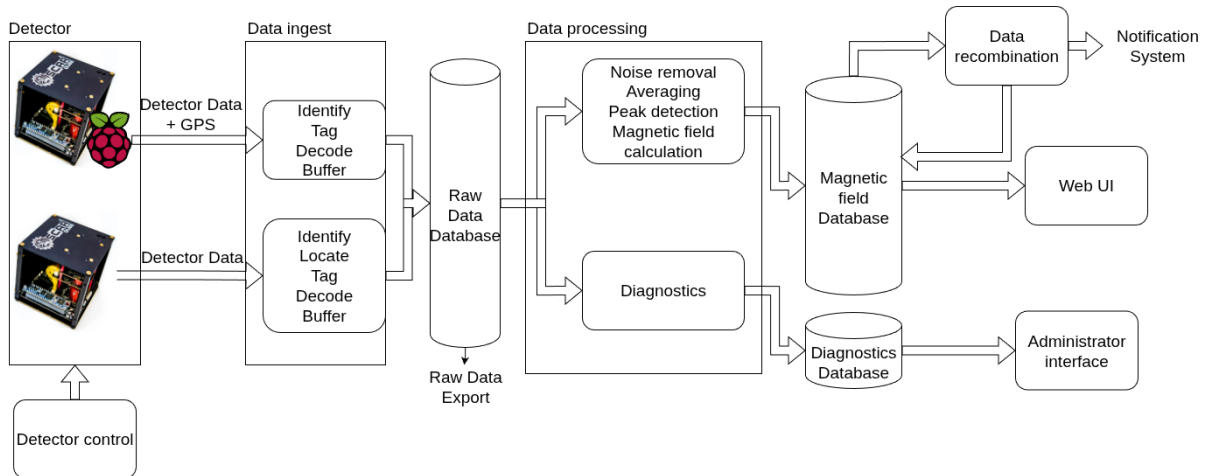


Fig. 20: The full system architecture, showing individual processing & storage steps

5.3 Prototype & data verification results

5.3.1 Prototype performance

The primary goal of this prototype was to validate the viability of this system. To test this, it used two measurement modes which could be selected using a button on the mobile observatory. The full protocol buffer used by gRPC can be found in Appendix B, although this version supports more communication than was actually used during this test.

Overall, this test was successful in showing the viability of this processing pipeline. The data gathered was properly received, processed and displayed in the visualiser (Fig. 21) with a processing speed conforming to the requirements. Because the observatory was battery-powered and due to recurring connection loss, there are some lapses in the data, but all data that was acquired, was properly transmitted and processed. That being said, this test also showed a number of shortcomings within both the observatory and the system.




Fig. 21: Directional and total magnetic field data recorded by the NV magnetometer during the polar circle excursion as visualised by the web visualiser.

At first, the observatory failed to properly send any data to the server. After analysing this issue, the most likely cause of this failure was a misconfiguration of the Access Point Name (APN) for the SIM-card used by the dongle, which made it unable to properly connect to the local network and transmit any data. As the APN is the central interface between the internet and a mobile device, that would explain the failing connection [91]. This issue got resolved by replacing the SIM-card used by the dongle.

Once this issue was resolved, the observatory started sending data to the server. This data included both regular ODMR measurements and broadsweep measurements. The current implementation of the processing node was only designed to properly process non-broadsweep measurements. Therefore broadsweep measurements will not be further discussed here. The goal for these measurements was solely that the ingest system and database would be able to receive them properly.

At first, the regular measurement data was shaped as the expected ODMR-spectra, but during the first night of data collection, this changed to spectra that lacked any peaks. This was resolved by restarting the detector, but shows an unexpected point of failure in the system. The most likely cause was an artifact coming from the physical setup, but it shows a point of failure that a later implementation should be able to handle: the identification of malformed data and the ability to either send out an alert for this, or to send a reset signal to the detector directly. This should be added to the next iteration.

5.3.2 Data processing verification

Next to these measurement-related issues, some hurdles within the data processing were identified as well. The first step in doing so was comparing the NV-based measurements with a reference magnetometer in the device and a model for a specific period while the detector was moving. This comparison is shown in Fig. 22, Fig. 23 and Fig. 24.



Fig. 22: The magnetic field values measured by the NV-based magnetometer before filtering over the course of a day during the polar circle excursion. An arbitrary unit was used, as the purpose of this figure is to show similarities in trends, and the bias magnet makes absolute measurements misleading.



Fig. 23: The magnetic field values measured by the reference magnetometer before filtering over the course of a day during the polar circle excursion. An arbitrary unit was used, as the purpose of this figure is to show similarities in trends, and the bias magnet makes absolute measurements misleading.



Fig. 24: The magnetic field values measured predicted by the CHAOS model for the same points as those displayed in Fig. 22 and Fig. 23 in nanotesla.

This shows a correlation between the reference magnetometer and the CHAOS-model, but a lesser connection between either of those and of the NV-based measurements, indicating an issue in the processing of NV-data.

To identify the source of this issue, the data and processing pipeline were further analysed. A first issue that showed itself here is that the bias magnetic field was not properly being removed from the data. While this should not influence the observed changes in magnetic field, it does impede the ability to calculate the total magnetic field. By measuring the magnetic field with the detector in different orientations, the influence of the bias magnet is identified, and subsequently removed from the measurements. As this still did not resolve the issue, some more aggressive processing and filtering steps were explored. One potential cause for problems in the data processing pipeline is an erroneous result during the peak detection algorithm. The ODMR-spectra show 8 peaks which move based on changes in the magnetic field around the detector, and because the observed magnetic field is primarily the Earth's magnetic field, there are limits to how much fluctuation can be expected. This can be used to filter out incorrectly filtered peaks by imposing an upper (orange) and lower (green) limit as shown in Fig. 25. Looking at individual ODMR-spectra for the cause of this, the failures in peak detection most often seem to originate from a peak erroneously getting detected in the middle of the spectrum. This causes peaks 5 to 8 to shift one peak location forward, explaining the numerous low outliers in Fig. 25. The resulting data of this process is shown in Fig. 26.




Fig. 25: All 8 peak locations from data collected during the polar region excursion, and the boundaries used for each to identify and filter out outliers.

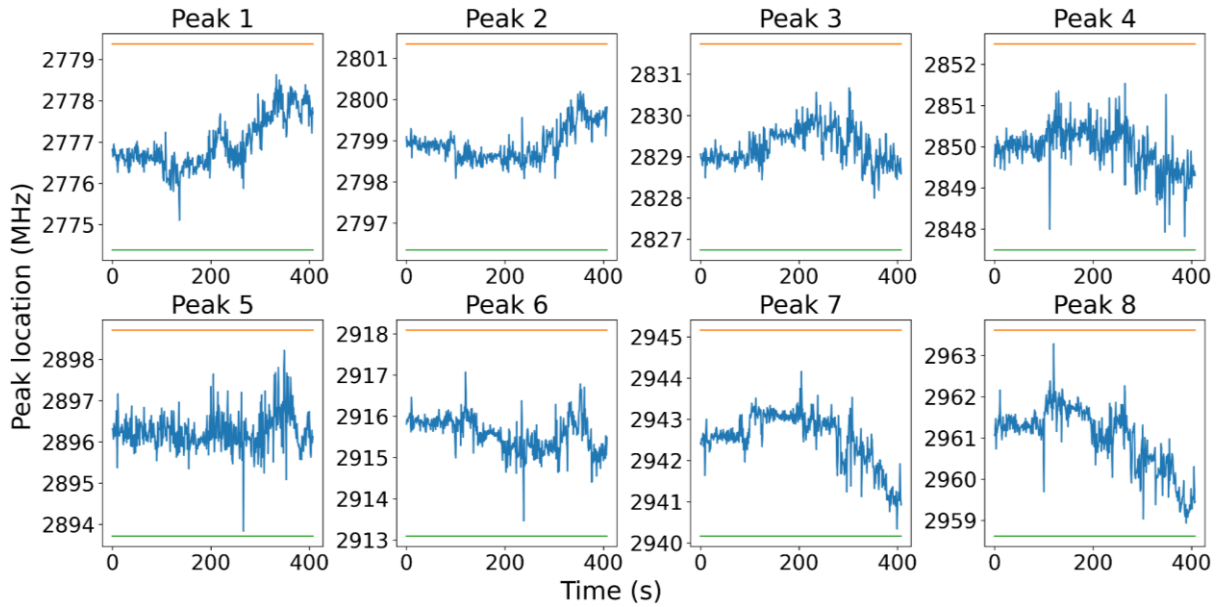


Fig. 26: All 8 peak locations from data collected during the polar region excursion after filtering out outliers.

After correcting for these measurements with failing peak detection, the directional magnetic field is plotted in Fig. 27. This shows that the X-axis, while noisier than the reference data, does show a proper correlation between reference and NV-based data. The Y-axis also clearly shows data, although this seems to be closer to an inversion of the data recorded by the reference magnetometer. The Z-axis, however, shows the most likely culprit of the issues so far, as the NV-based values seem to only show noise with no meaningful data or trend to speak of. This suggests that there is an issue with the used characterization matrix, that should transform the peak locations into correct magnetic field data, but seems to fail in this.

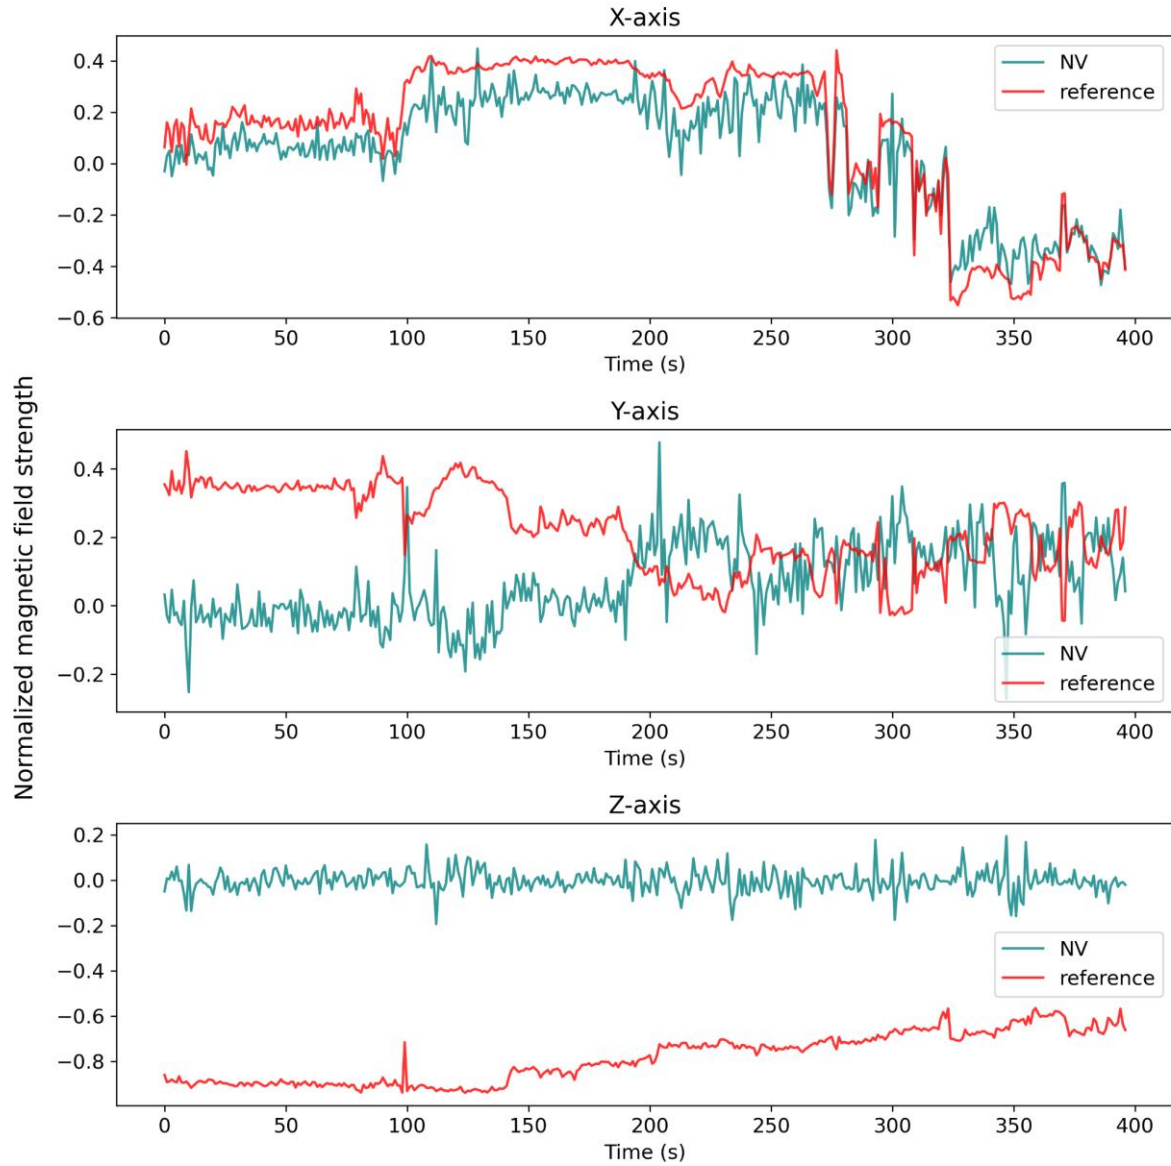


Fig. 27: Directional magnetic field measured during the polar region excursion as measured by the NV-detector and the reference magnetometer after filtering outliers during peak detection.

To validate this, a test with more controlled data is performed. This time, the detector was placed in the Helmholtz coil. The coil was then used to apply a magnetic field alongside each of the axes. The data of this test is recorded and both the peak locations and measured magnetic field are plotted. The peak locations, shown in Fig. 28, show a clear response for each peak pair on the applied magnetic field.

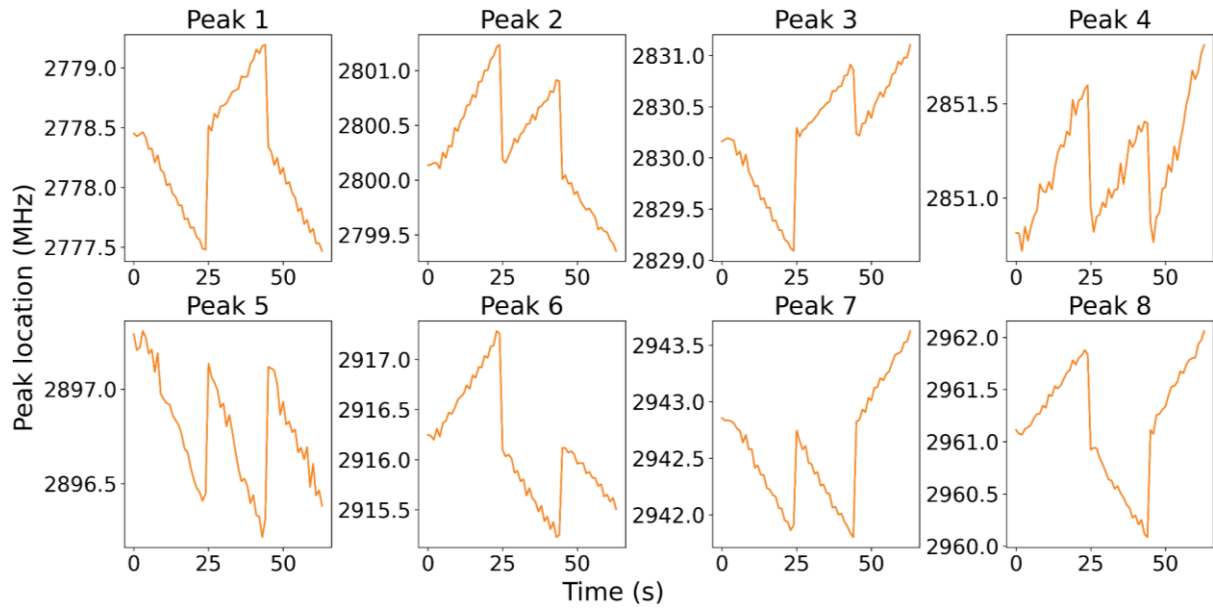


Fig. 28: All 8 peak locations from the Helmholtz coil-generated data, showing the response to the artificially applied external magnetic field along all 3 axes.

The directional magnetic field results, on the other hand, fall in line with the earlier observations. As shown in Fig. 29, the X-axis and Y axis show a correlation and inverse correlation respectively, but the Z-axis shows no response to the applied magnetic field. The fact that this does show in the peak locations shown by Fig. 28 proves the presence of an incorrect calibration matrix.

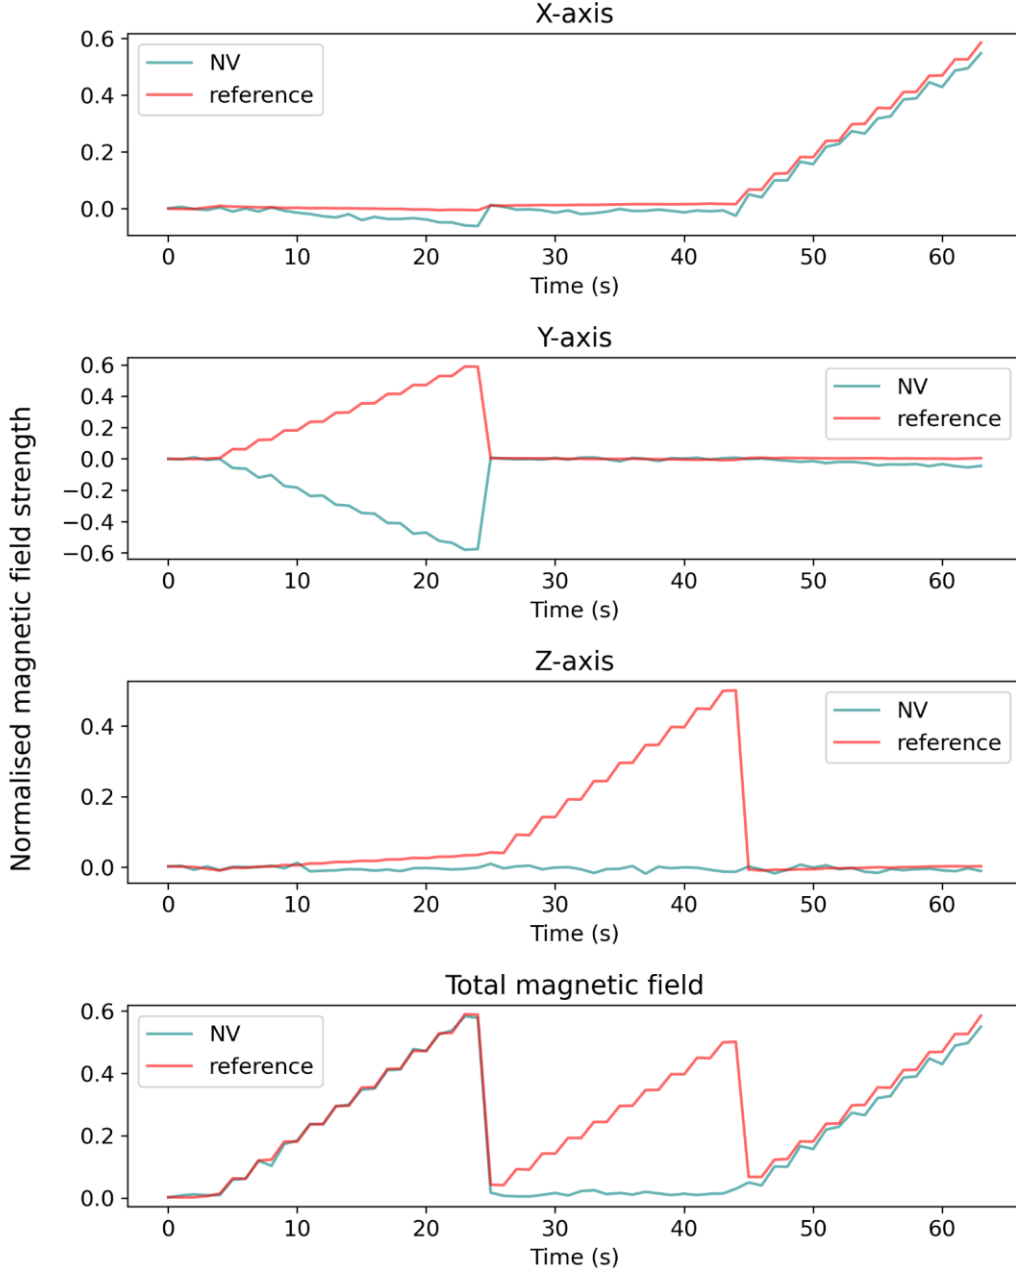


Fig. 29: Comparison of the directional and total magnetic field between the NV-based magnetic field measurements and the reference magnetometer's measurement under the same conditions generated by the Helmholtz coil.

5.4 Performance test results

5.4.1 Communication limit test result

The first test aimed to identify the limits of both gRPC's synchronous and asynchronous calls. The results of this test are shown in Fig. 30. This test shows a clear advantage to using the synchronous call for this use case, with a load limit between 185 Hz and 256Hz. The 71% retention rate at 256Hz suggests that the limit is around 181 Hz.

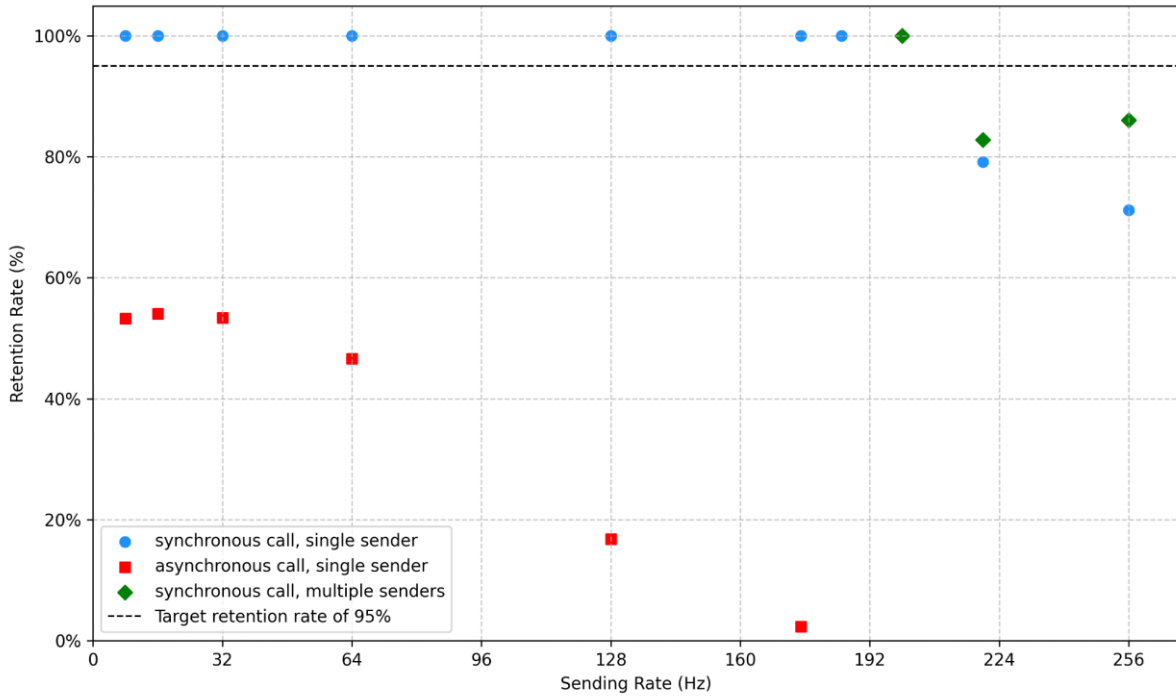


Fig. 30: Data retention rate (%) of the data ingest system when using synchronous calls from a single device, asynchronous calls from a single device and synchronous calls from multiple devices. The multi-device tests only start from the point where data loss started to show for the single-device tests to ensure the results are representative of the receiver's limits. This test has a target retention rate of 95%.

As it is plausible that this synchronous limit is a consequence of using a single sender, this test was repeated with two senders at 128Hz each. This showed a significantly improved retention rate of 86% which, while still falling short of the target 95%, is likely a more accurate representation of what the ingest node can process properly. It is worth noting that the single sender originally reached a similar speed and retention rate during its first minutes, but this decreased over the course of the test. The highest retention rate measured during this test was 13365 measurements in a minute, which would equate to approximately 223Hz. For this reason, the two sender-setup was used with a total sending rate of 200 Hz and 220Hz as well. At 200 Hz this showed a retention rate of above 99%, but a significantly lower retention rate of 82% showed for the test at 220 Hz. It is likely that this data point is malformed due to networking issues, but this will be elaborated upon in section 5.4.4.

5.4.2 Data structure test result

The second test compares the unpacked and packed data, with its results shown in Fig. 31. This test was performed using synchronous calls, given their superior performance in test 1. The results of this test show that the unpacked data is more performant. The packed data drops below the target retention rate of 95% around 150Hz, whereas the unpacked data manages to remain above it up until at least 185Hz. This indicates a performance limit on the side of the server caused by the unpacking of the data there. While this could be caused or exacerbated by the configuration of the data ingest system, this was not further explored in this work.




Fig. 31: Data retention rate (%) of the data ingest module when using either an unpacked or packed detector payload for various sending rates. This has a target retention rate of 95%.

5.4.3 Full pipeline test result

The third test set focuses on the performance of the full pipeline, including the raw data processing. The timing at which the raw data and the processed data enter the database is tracked and normalised to the number of expected measurements for one minute for the respective test. The processed measurement count has a small offset for which is compensated as well, caused by an artifact related to how this data is recorded. A third value is plotted as well, namely the number of minutes the processing is behind. This value should increase at the start of the test, after which it flattens out once the first processing tasks start to finalise. After a preliminary test, the processing node was changed to run once every 30 seconds instead of once every minute to more efficiently utilise the present system resources.

The results of these tests are shown in Fig. 32, Fig. 33 and Fig. 34, which show the results for three sending frequencies: 64, 80 and 100Hz. These graphs show how many measurements have been received, have already been processed and how many minutes worth of measurements are currently between them (i.e. how many minutes the system is “behind”). The first minutes of these tests consist of the data build-up and should thus be ignored. Similarly, the test devices stop sending data 15 minutes after starting, making it so the “minutes behind”-metric decreases from that point going forward.

While this test observed minor data loss when it comes to measurements received (under 2%), this does not exceed the target retention rate and may be related to other testing limitations further discussed in section 5.4.4. For the three tested sending speeds, only the 64Hz measurement manages to limit its delay to under 5 minutes, at 4.5 minutes. The 80Hz test manages to stabilize as well, but does so around a 5.5 minute delay. The 100Hz test’s delay seems to continuously increase until no new data is received, at which point it has over a 7.5 minute delay.




Fig. 32: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 64Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.




Fig. 33: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 80Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.




Fig. 34: The results of the full pipeline processing test showing the relative number of measurements received, processed and their difference for each minute interval given a data sending rate of 100Hz. The grey regions on both sides of the graph indicate the startup and finishing of the test and are less important.

5.4.4 Final performance testing discussion

There were some noticeable trends and behaviours during these tests that are worth mentioning. These should not be taken as fully verified results, but were noted as a starting point for future work.

Firstly, the gRPC communication between the data sending devices and the server under test was limited. After running a certain number of tests, any amount of communication would trigger an HTTP2/GOAWAY, prompting the senders to reduce the amount of data they send and rendering the test useless. This behaviour was device-specific at first, suggesting it may be related to a firewall or other networking-related issue. This is also the reason for the limited volume of test data.

The processing also showed a flawed expectation. Originally, it was hypothesized that the increased load from the processing would impede the data retention of the data ingest system. As the data processing step is a single-threaded process, this never happens as there are more cores than resource-intensive processes on the tested hardware. This makes the bottleneck of the full processing pipeline the time it needs to run, not any ingestion-related issues. This also implies that it would be best to keep the processing step on the same device as the raw data database, as the bottleneck is not processing power itself but a networking- or ingest-related limit.

As described earlier, only a middle section of the test data was actively considered to avoid misleading data during the test's start and finish, as these would provide different data loads. In some of the tests, this proved to be relevant, as the system was capable of retaining slightly higher loads during the startup phase.

These performance tests show that the data ingest system manages to convincingly remain above the target 95% data retention rate up until at least a measuring rate of 200 Hz, where the processing manages to remain under the target 5 minute delay at 64Hz. The system's requirements aim for 610 and 7415 measurements per second as bare minimum and target load, which would require 10 and 116 instances of the ingest system respectively using the current implementation at a measuring speed of 1Hz for each

detector. The limit on the ingest system is higher, at around 200Hz, which would reduce this to 4 or 38 instances respectively if the processing system can be improved enough to meet this ingest performance bottleneck, although all of these numbers could be improved upon by making the system itself more efficient as well.

6 Conclusion & future work

This thesis consisted of three core goals: the design of a system architecture for a back-end NV Diamond-based Magnetometry

This thesis consisted of three core goals: creating the requirements and architectural design of a global magnetometry system using NV- magnetometers, implementing a prototype of this system to validate the feasibility of this system in a real-world environment, and a set of performance tests on this system to identify its load limits.

The system's requirements were set up to cover the challenges inherent to a global magnetometer data collection system, with additional attention paid to the specific challenges posed by the utilised NV magnetometer. To do so, factors such as the expected system load, the resulting data's target audience and existing data collection systems were taken into account. With these in mind, an architecture design is proposed that focuses on modularity. This is to ensure the ease of replacing processing elements with as little impact on the system as possible as these improve over time.

With these in mind, a prototype was implemented to demonstrate the feasibility of this data processing pipeline, satisfying a part of the proposed architecture and requirements. This prototype was successful in proving the feasibility, but also exposed a number of challenges that a future implementation should aim to resolve. These challenges concern the detector and its physical setup, as well as the data processing. The observed abnormalities during the data processing step gave rise to further investigation, showing multiple ways in which the data quality could be improved. This analysis eventually concluded that there was an issue with the characterization matrix responsible for transforming the observed peak locations into a vectorial magnetic field value. The explored improvements in data processing, such as actively verifying that the data is valid ODMR data and identifying failures in the peak detection algorithm, are used on a part of the dataset but not implemented in the prototype yet.

This prototype was then slightly reconfigured to perform performance tests, comparing multiple communication solutions. This showed that the system's ingest node does not suffer from any performance degradation until at least 200 Hz, and that the processing system as a whole manages to process data in under 5 minutes until between 64 Hz and 80 Hz. As this was only a prototype, there is likely room to finetune the processing system to improve its performance by more efficiently using the available hardware resources. The potential improvements to the ingest system are likely less impactful and would have to be more invasive. Techniques similar to Nagle's algorithm to batch data before it is transmitted could reduce networking overhead and improve performance, but this would have to be tested. Similarly, gRPC could be compared to other communication protocols that may be more efficient for this specific use case, or other implementations of gRPC that have less overhead such as streaming. As this thesis proves the feasibility of an NV-based magnetometry network, it presents three clear angles to further explore. The first one focuses on improving this system's prototype and data processing element as described earlier. The second would look at the used mobile observatory and aim to make it more remote-capable. As the current system is limited by its battery life of under 10 hours, it would not be able to function as an observatory without human intervention. How well it would handle the strain provided by natural weather conditions remains an open question as well. Lastly, a long-term test to prove the functionality of the detector could be performed as well. While this thesis proves the feasibility of the data collection and processing systems, it does not demonstrate the ability of an NV-based magnetometer to function for a prolonged period of time without human intervention, and this should be tested as well.

References

- [1] UHasselt, “OSCAR-QUBE- UHasselt.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.uhasselt.be/en/aparte-sites-uhasselt-en/oscar-qube>
- [2] UHasselt, “imo-imomec,” UHasselt. Accessed: Aug. 15, 2025. [Online]. Available: <https://www.uhasselt.be/en/instituten-en/imo-imomec>
- [3] A. A. Vidotto, “The evolution of the solar wind,” *Living Rev. Sol. Phys.*, vol. 18, no. 1, p. 3, Apr. 2021, doi: 10.1007/s41116-021-00029-w.
- [4] S. E. Milan *et al.*, “Overview of Solar Wind–Magnetosphere–Ionosphere–Atmosphere Coupling and the Generation of Magnetospheric Currents,” *Space Sci. Rev.*, vol. 206, no. 1, pp. 547–573, Mar. 2017, doi: 10.1007/s11214-017-0333-0.
- [5] C. T. Russell, “The solar wind interaction with the Earth’s magnetosphere: a tutorial,” *IEEE Trans. Plasma Sci.*, vol. 28, no. 6, pp. 1818–1830, Dec. 2000, doi: 10.1109/27.902211.
- [6] NASA, “Magnetospheres - NASA Science.” Accessed: July 20, 2025. [Online]. Available: <https://science.nasa.gov/heliophysics/focus-areas/magnetosphere-ionosphere/>
- [7] D. G. Sibeck and K. R. Murphy, “Large-Scale Structure and Dynamics of the Magnetosphere,” in *Magnetospheres in the Solar System*, American Geophysical Union (AGU), 2021, pp. 15–36. doi: 10.1002/9781119815624.ch2.
- [8] Space Weather Prediction Center, “Earth’s magnetosphere | NOAA / NWS Space Weather Prediction Center.” Accessed: Dec. 11, 2024. [Online]. Available: <https://www.swpc.noaa.gov/phenomena/earths-magnetosphere>
- [9] EBSCO, “Solar wind interactions | EBSCO Research Starters.” Accessed: July 20, 2025. [Online]. Available: <https://www.ebsco.com/research-starters/physics/solar-wind-interactions>
- [10] N. US Department of Commerce, “Aurora Borealis (Northern Lights).” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.weather.gov/fsd/aurora>
- [11] N. Lugaz, W. B. M. Iv, and T. I. Gombosi, “The Evolution of Coronal Mass Ejection Density Structures,” *Astrophys. J.*, vol. 627, no. 2, p. 1019, July 2005, doi: 10.1086/430465.
- [12] EBSCO, “Geomagnetic storm | EBSCO Research Starters.” Accessed: July 20, 2025. [Online]. Available: <https://www.ebsco.com/research-starters/earth-and-atmospheric-sciences/geomagnetic-storm>
- [13] D. T. Phillips, “The Day Earth Lost Half Its Satellites (Halloween Storms 2003),” Spaceweather.com. Accessed: Aug. 15, 2025. [Online]. Available: <https://spaceweatherarchive.com/2021/10/28/the-day-earth-lost-half-its-satellites-halloween-storms-2003/>
- [14] National Centers for Environmental Information (NCEI), “Remembering the Great Halloween Solar Storms,” 2017. Accessed: Aug. 15, 2025. [Online]. Available: <https://www.ncei.noaa.gov/news/great-halloween-solar-storm-2003>
- [15] E. Gresser, “PPI’s Trade Fact of the Week: The number of working satellites in space has doubled since the Biden administration began,” 2023. Accessed: Aug. 15, 2025. [Online]. Available: <https://www.progressivepolicy.org/ppis-trade-fact-of-the-week-the-number-of-working-satellites-in-space-has-doubled-since-the-biden-administration-began/>

- [16] NASA, “Near Miss: The Solar Superstorm of July 2012 - NASA Science.” Accessed: Aug. 15, 2025. [Online]. Available: https://science.nasa.gov/science-research/planetary-science/23jul_superstorm/
- [17] D. D. Billett *et al.*, “The 2022 Starlink Geomagnetic Storms: Global Thermospheric Response to a High-Latitude Ionospheric Driver,” *Space Weather*, vol. 22, no. 2, p. e2023SW003748, 2024, doi: 10.1029/2023SW003748.
- [18] A. C. and S. Merkin, “NASA Scientific Visualization Studio | Geomagnetic Storm Causes Satellite Loss for FullDome,” NASA Scientific Visualization Studio. Accessed: Aug. 15, 2025. [Online]. Available: <https://svs.gsfc.nasa.gov/5214>
- [19] NASA, “How NASA Tracked the Most Intense Solar Storm in Decades.” Accessed: Aug. 15, 2025. [Online]. Available: <https://science.nasa.gov/science-research/heliophysics/how-nasa-tracked-the-most-intense-solar-storm-in-decades/>
- [20] NASA, “NASA Space Weather.” Accessed: July 20, 2025. [Online]. Available: <https://science.nasa.gov/heliophysics/focus-areas/space-weather/>
- [21] Space Weather Prediction Center, “Geomagnetic Storms.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.swpc.noaa.gov/phenomena/geomagnetic-storms>
- [22] Space Weather Live “Magnetometers,” SpaceWeatherLive.com. Accessed: Aug. 15, 2025. [Online]. Available: <https://spaceweatherlive.com/en/auroral-activity/magnetometers.html>
- [23] J. Matzka, O. Bronkalla, K. Tornow, K. Elger, and C. Stolle, “Geomagnetic Kp index.” GFZ Data Services, 2021. doi: 10.5880/KP.0001.
- [24] International Service of Geomagnetic Indices, “What are K Indices?.” Accessed: Aug. 15, 2025. [Online]. Available: https://isgi.unistra.fr/what_are_kindices.php
- [25] Space Weather Prediction Center, “NOAA Space Weather Scales.” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/noaa-scales-explanation>
- [26] Space Weather Prediction Center, “Planetary K-index.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.swpc.noaa.gov/products/planetary-k-index>
- [27] Southwest Research Institute, “New NOAA Magnetometers.” Accessed: July 21, 2025. [Online]. Available: <https://www.swri.org/newsroom/technology-today/new-noaa-magnetometers>
- [28] NASA, “SWFO-L1.” Accessed: July 21, 2025. [Online]. Available: <https://science.nasa.gov/mission/swfo-l1/>
- [29] A. Posner, M. Hesse, and O. C. St Cyr, “The main pillar: Assessment of space weather observational asset performance supporting nowcasting, forecasting, and research to operations,” *Space Weather*, vol. 12, no. 4, pp. 257–276, Apr. 2014, doi: 10.1002/2013SW001007.
- [30] E. Thébault *et al.*, “International Geomagnetic Reference Field: the 12th generation,” *Earth Planets Space*, vol. 67, no. 1, p. 79, May 2015, doi: 10.1186/s40623-015-0228-9.
- [31] SpaceX, “SpaceX.” Accessed: Aug. 15, 2025. [Online]. Available: <http://www.spacex.com>
- [32] A. Chulliat, J. Matzka, A. Masson, and S. E. Milan, “Key Ground-Based and Space-Based Assets to Disentangle Magnetic Field Sources in the Earth’s Environment,” *Space Sci. Rev.*, vol. 206, no. 1, pp. 123–156, Mar. 2017, doi: 10.1007/s11214-016-0291-y.
- [33] Q. Yan and H. B. Yao, “Recent geomagnetic storms observed by Macau Science Satellite-1,” *Earth Planet. Phys.*, vol. 8, no. 4, pp. 565–569, 2024, doi: 10.26464/epp2024047.

- [34] S. Macmillan and N. Olsen, “Observatory data and the Swarm mission,” *Earth Planets Space*, vol. 65, no. 11, pp. 1355–1362, Nov. 2013, doi: 10.5047/eps.2013.07.011.
- [35] E. Friis-Christensen, H. Lühr, D. Knudsen, and R. Haagmans, “Swarm – An Earth Observation Mission investigating Geospace,” *Adv. Space Res.*, vol. 41, no. 1, pp. 210–216, Jan. 2008, doi: 10.1016/j.asr.2006.10.008.
- [36] J. W. Gjerloev, “The SuperMAG data processing technique,” *J. Geophys. Res. Space Phys.*, vol. 117, no. A9, 2012, doi: 10.1029/2012JA017683.
- [37] J. W. Gjerloev and R. A. Hoffman, “The large-scale current system during auroral substorms,” *J. Geophys. Res. Space Phys.*, vol. 119, no. 6, pp. 4591–4606, 2014, doi: 10.1002/2013JA019176.
- [38] M. Menvielle, T. Iyemori, A. Marchaudon, and M. Nosé, “Geomagnetic Indices,” in *Geomagnetic Observations and Models*, M. Manda and M. Korte, Eds., Dordrecht: Springer Netherlands, 2011, pp. 183–228. doi: 10.1007/978-90-481-9858-0_8.
- [39] International Monitor for Auroral Geomagnetic Effects, “IMAGE Magnetometer network.” Accessed: July 21, 2025. [Online]. Available: <https://space.fmi.fi/image/www/index.php>
- [40] INTERMAGNET, “1.1. Introduction — Technical Reference Manual 5.2.0 documentation.” Accessed: July 21, 2025. [Online]. Available: <https://tech-man.intermagnet.org/stable/chapters/intermagnet/intro.html>
- [41] INTERMAGNET, “2.1. Description — Technical Reference Manual 5.2.0 documentation.” Accessed: July 21, 2025. [Online]. Available: <https://tech-man.intermagnet.org/stable/chapters/oneminuteimos/description.html>
- [42] INTERMAGNET, “International Real-time Magnetic Observatory Network.” Accessed: Aug. 15, 2025. [Online]. Available: <https://intermagnet.org/>
- [43] INTERMAGNET, “INTERMAGNET Metadata.” Accessed: Aug. 15, 2025. [Online]. Available: <https://intermagnet.org/metadata/#/map>
- [44] SuperMAG, “SuperMAG: About SuperMAG.” Accessed: Aug. 15, 2025. [Online]. Available: <https://supermag.jhuapl.edu/info/>
- [45] SuperMAG, “SuperMAG.” Accessed: Aug. 15, 2025. [Online]. Available: <https://supermag.jhuapl.edu/>
- [46] Space Weather Prediction Center, “3-Day Forecast.” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/products/3-day-forecast>
- [47] Space Weather Prediction Center, “Subscription Services.” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/content/subscription-services>
- [48] Space Weather Prediction Center, “27-Day Outlook of 10.7 cm Radio Flux and Geomagnetic Indices.” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/products/27-day-outlook-107-cm-radio-flux-and-geomagnetic-indices>
- [49] Space Weather Prediction Center, “Space Weather FAQ (Frequently Asked Questions)” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/content/space-weather-faq-frequently-asked-questions>
- [50] N. US Department of Commerce, “Space weather forecasting and impacts: what you need to know.” Accessed: July 24, 2025. [Online]. Available: <https://www.weather.gov/news/040623-space-scientist>
- [51] Space Weather Prediction Center, “Models.” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/models>

- [52] Space Weather Prediction Center, “Products and Data” Accessed: July 24, 2025. [Online]. Available: <https://www.swpc.noaa.gov/products-and-data>
- [53] AuroraResearch, “Aurora sighting in Qikiqtarjuaq, Canada.” Accessed: July 23, 2025. [Online]. Available: <https://aurorareach.com/>
- [54] AuroraResearch, “AuroraReach.” Accessed: July 23, 2025. [Online]. Available: <https://aurorareach.com/>
- [55] The University of Iowa, “Instruments | TRACERS - The University of Iowa.” Accessed: July 21, 2025. [Online]. Available: <https://tracers.physics.uiowa.edu/instruments>
- [56] PNI Sensor, “Spacecraft Magnetometers and their Applications.” Accessed: July 21, 2025. [Online]. Available: <https://www.pnisor.com/spacecraft-magnetometers-types-discoveries-and-applications/>
- [57] PNI Sensor, “How Fluxgate Magnetometers Work: Applications and Limitations.” . Accessed: July 21, 2025. [Online]. Available: <https://www.pnisor.com/fluxgate-magnetometer/>
- [58] AV3 AEROVISUAL “Types of magnetometers, uses and characteristics.” Accessed: July 21, 2025. [Online]. Available: <https://av3aerovisual.com/en/types-of-magnetometers-uses-and-characteristics/>
- [59] G.B. Hospodarsky, “Spaced Based Search Coil Magnetometers,” in *Journal of Geophysical Research: Space Physics*, vol. 121, no. 12. , November 2016, doi: 10.1002/2016JA022565.
- [60] A. Roux *et al.*, “The Search Coil Magnetometer for THEMIS,” *Space Science Review.*, vol. 141, no. 1, pp. 265–275, Dec. 2008, doi: 10.1007/s11214-008-9455-8.
- [61] W. D. S. Ruhunusiri and M. K. Jayananda, “Construction of a Proton Magnetometer,” in *Proceedings of the Technical Sessions*, 2008.
- [62] V. Sapunov, A. Denisov, O. Denisova, and D. V. Saveliev, "Proton and Overhauser magnetometers metrology," *Contributions to Geophysics and Geodesy*, vol. 31, pp. 119–124, 2001.
- [63] GEM Systems, Inc., “All applications of proton precession magnetometers,” Ontario, Canada: GEM Systems, Inc., 2005. [Online]. Available: <https://www.gemsys.ca/pdf/Q2-A10-All-Applications-Proton-Precession.pdf>
- [64] B. Zhao, L. Li, Y. Zhang, J. Tang, Y. Liu, and Y. Zhai, “Optically Pumped Magnetometers Recent Advances and Applications in Biomagnetism: A Review,” *IEEE Sens. J.*, vol. 23, no. 17, pp. 18949–18962, Sept. 2023, doi: 10.1109/JSEN.2023.3297109.
- [65] G. Oelsner, V. Schultze, R. IJsselsteijn, F. Wittkämper, and R. Stolz, “Sources of heading errors in optically pumped magnetometers operated in the Earth’s magnetic field,” *Phys. Rev. A*, vol. 99, no. 1, p. 013420, Jan. 2019, doi: 10.1103/PhysRevA.99.013420.
- [66] L. Jogschies *et al.*, “Recent Developments of Magnetoresistive Sensors for Industrial Applications,” *Sensors*, vol. 15, no. 11, pp. 28665–28689, Nov. 2015, doi: 10.3390/s151128665.
- [67] M. Vopalensky and A. Platil, “Temperature Drift of Offset and Sensitivity in Full-Bridge Magnetoresistive Sensors,” *IEEE Trans. Magn.*, vol. 49, no. 1, pp. 136–139, Jan. 2013, doi: 10.1109/TMAG.2012.2220535.
- [68] Y. Beerden, R. Vandeboosch, A. Ermakova, M. Nesladek, and J. Hruby, “Long-Term Stability Assessment of Quantum Diamond Magnetometers in Low Earth Orbit,” in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, July 2024, pp. 465–468. doi: 10.1109/IGARSS53475.2024.10641354.

- [69] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, “The nitrogen-vacancy colour centre in diamond,” *Phys. Rep.*, vol. 528, no. 1, pp. 1–45, July 2013, doi: 10.1016/j.physrep.2013.02.001.
- [70] Qnami AG, “Technical note: Fundamentals of magnetic field measurement with NV centers in diamond,” 2020. Accessed: July 23, 2025. [Online]. Available: <https://qnami.ch/wp-content/uploads/2020/12/2020-12-07-Qnami-TN1-The-NV-center-1.pdf>
- [71] L. Rondin, J.-P. Tetienne, T. Hingant, J.-F. Roch, P. Maletinsky, and V. Jacques, “Magnetometry with nitrogen-vacancy defects in diamond,” *Rep. Prog. Phys.*, vol. 77, no. 5, p. 056503, May 2014, doi: 10.1088/0034-4885/77/5/056503.
- [72] J. S. Bennett *et al.*, “Precision Magnetometers for Aerospace Applications: A Review,” *Sensors*, vol. 21, no. 16, Art. no. 16, Jan. 2021, doi: 10.3390/s21165568.
- [73] J. Zhang *et al.*, “Diamond Nitrogen-Vacancy Center Magnetometry: Advances and Challenges,” Oct. 20, 2020, *arXiv*: arXiv:2010.10231. doi: 10.48550/arXiv.2010.10231.
- [74] E. V. Levine *et al.*, “Principles and techniques of the quantum diamond microscope,” *Nanophotonics*, vol. 8, no. 11, pp. 1945–1973, Nov. 2019, doi: 10.1515/nanoph-2019-0209.
- [75] “Boxcar Averaging,” Chemistry LibreTexts. Accessed: Aug. 15, 2025. [Online]. Available: [https://chem.libretexts.org/Bookshelves/Analytical_Chemistry/Supplemental_Modules_\(Analytical_Chemistry\)/Analytical_Sciences_Digital_Library/Courseware/Introduction_to_Signals_and_Noise/04_Signal-to-Noise_Enhancement/03_Digital_Filtering/03_Boxcar_Averaging](https://chem.libretexts.org/Bookshelves/Analytical_Chemistry/Supplemental_Modules_(Analytical_Chemistry)/Analytical_Sciences_Digital_Library/Courseware/Introduction_to_Signals_and_Noise/04_Signal-to-Noise_Enhancement/03_Digital_Filtering/03_Boxcar_Averaging)
- [76] INTERMAGNET, “2.4. Data Sampling and Filtering — Technical Reference Manual 5.2.0 documentation.” Accessed: Aug. 15, 2025. [Online]. Available: <https://tech-man.intermagnet.org/stable/chapters/oneminuteimos/samplingfiltering.html#data-sampling-and-filtering>
- [77] Linux Manual Page, “cron(8).” Accessed: Aug. 15, 2025. [Online]. Available: <https://man7.org/linux/man-pages/man8/cron.8.html>
- [78] Docker, “Accelerated Container Application Development.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.docker.com/>
- [79] Apache, “About the Apache HTTP Server Project - The Apache HTTP Server Project.” Accessed: Aug. 15, 2025. [Online]. Available: https://httpd.apache.org/ABOUT_APACHE.html
- [80] Apache, “Apache Virtual Host documentation - Apache HTTP Server Version 2.4.” Accessed: Aug. 15, 2025. [Online]. Available: <https://httpd.apache.org/docs/2.4/vhosts/>
- [81] T. Sandholm, “InfluxDB: The Good, the Bad, and the Ugly,” Medium. Accessed: Aug. 15, 2025. [Online]. Available: <https://medium.com/@thomas.sandholm/influxdb-the-good-the-bad-and-the-ugly-b409f8c370a3>
- [82] InfluxData, “Introducing InfluxDB 3.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.influxdata.com/products/influxdb-overview/>
- [83] F. Vantroys, “The Pros and Cons of Using gRPC in Modern Software Architecture,” Medium. Accessed: Aug. 15, 2025. [Online]. Available: https://medium.com/@frederik_62300/the-pros-and-cons-of-using-grpc-in-modern-software-architecture-93cca0a8c8fd
- [84] gRPC, “gRPC.” Accessed: Aug. 15, 2025. [Online]. Available: <https://grpc.io/>
- [85] Gitlab, “Royal Netherlands Meteorological Institute Open Source / Space Weather / HAPI timeline viewer.” Accessed: Aug. 15, 2025. [Online]. Available: <https://gitlab.com/KNMI-OSS/spaceweather/knmi-hapi-timeline-viewer>

- [86] Wireshark, “Wireshark • undefined.”-. Accessed: Aug. 15, 2025. [Online]. Available: <https://www.wireshark.org/>
- [87] Division of Geomagnetism and Geospace (DTU Space), “Welcome to the CHAOS-8 field model website — CHAOS website 1.0 documentation.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.spacecenter.dk/files/magnetic-models/CHAOS-8/>
- [88] H3, “Introduction.” Accessed: Aug. 15, 2025. [Online]. Available: <https://h3geo.org/docs/>
- [89] H3, “Tables of Cell Statistics Across Resolutions.” Accessed: Aug. 15, 2025. [Online]. Available: <https://h3geo.org/docs/core-library/restable/>
- [90] DataGenetics, “Optimal Packing.” Accessed: Aug. 15, 2025. [Online]. Available: <http://datagenetics.com/blog/june32014/index.html>
- [91] A1 Digital, “What is APN (Access Point Name)?.” Accessed: Aug. 15, 2025. [Online]. Available: <https://www.a1.digital/knowledge-hub/what-is-apn-access-point-name/>

Appendices

Appendix A: Requirements list

System wide requirements

Requirement ID	Description
AC_S_1	The system shall be capable of collecting data from 122 detectors
AC_S_2	The system should be capable of collecting data from 1483 detectors
AC_S_3	The system shall provide an export function for the unprocessed collected data
AC_S_4	The system shall provide an export function for the fully processed data
AC_S_5	The system shall provide a web visualiser showing the measured magnetic field data
AC_S_6	The system shall have a web interface showing information for administrative & maintenance-related purposes
AC_S_7	The system shall support a notification system, sending out push notifications once a fluctuation threshold is passed.
AC_S_8	The system shall provide a map visualiser, indicating the likelihood of visible auroras
AC_S_9	The system shall store raw data for 2 weeks for use within the system itself for use by processing steps
AC_S_10	The system shall store fully processed data for 15 years in a directly accessible manner, accessible through the web visualiser
AC_S_11	The system shall have a data back-up solution that stores all recorded data.
AC_S_12	The system shall have a processing pipeline to convert measured ODMR-data into magnetic field data
AC_S_13	The system shall conform with the requirements to be included in the INTERMAGNET-system.
AC_S_15	The system shall be capable of fully processing data from time of measurement to a visualised magnetic field in the web visualiser in 10 minutes or less
AC_S_16	The system should be capable of fully processing data from time of measurement to a visualised magnetic field in the web visualiser in 5 minutes or less
AC_S_17	The system shall store the per-hour averaged magnetic field data of each detector permanently.

Detector requirements

Requirement ID	Description
AC_D_1	The detector shall measure data at a rate of at least 1Hz, increasing based on measurement settings and regimen.
AC_D_2	The detector shall be configurable remotely through an internet connection
AC_D_3	The detector shall produce packets of data containing either <ul style="list-style-type: none"> - a timestamp, ODMR measurement data, measurement settings data, and an identifier of the measuring detector or <ul style="list-style-type: none"> - a timestamp, ODMR measurement data, measurement settings data, location data and orientation data
AC_D_4	The ODMR measurement data & measurement settings data shall be either <ul style="list-style-type: none"> - in its raw, undecoded format or <ul style="list-style-type: none"> - decoded into its individual fields
AC_D_5	The detector shall send its data to the data ingest system
AC_D_6	Any data measured by the detector shall be sent to the data ingest system in 30 seconds or less.
AC_D_7	Any data measured by the detector should be sent to the data ingest system in 10 seconds or less

Data ingest requirements

Requirement ID	Description
AC_I_1	The data ingest node shall take in packets of data containing either <ul style="list-style-type: none"> - a timestamp, ODMR measurement data, measurement settings data, and an identifier of the measuring detector or <ul style="list-style-type: none"> - a timestamp, ODMR measurement data, measurement settings data, location data and orientation data
AC_I_2	The data ingest node shall be able to add correct GPS-data to a packet of data that contains a detector identifier
AC_I_3	The data ingest node shall record both the timestamp included in the message, and the timestamp at which the data was received
AC_I_4	The data ingest node shall process any data it receives into the same format and store it in the raw data database
AC_I_5	The data ingest system shall be capable of processing 610 measurements per second with a data retention rate of at least 95%
AC_I_6	The data ingest system should be capable of processing 7415 measurements per second with a data retention rate of at least 95%

Data storage requirements

Requirement ID	Description
AC_DB_1	The data storage system shall have separate storage solutions for the following elements: <ul style="list-style-type: none"> - Unprocessed data - Processed magnetic field data - Diagnostics data
AC_DB_2	The unprocessed data storage system shall be capable of storing 737856000 measurements worth of data
AC_DB_3	The unprocessed data storage system should be capable of storing 8969184000 measurements worth of data
AC_DB_4	The unprocessed data storage system shall create a backup of its recently added data every 24 hours for long-term storage
AC_DB_5	The unprocessed data storage system shall remove any data older than 2 weeks
AC_DB_6	The unprocessed data storage system shall provide the option to download the data backup through an API
AC_DB_7	The processed magnetic field data storage system shall be capable of storing the processed data of 28.2 billion measurements.
AC_DB_8	The processed magnetic field data storage system should be capable of storing the processed data of 342 billion measurements.
AC_DB_9	The system's backup shall be stored on a different system to ensure redundancy.

Data processing requirements

Requirement ID	Description
AC_P_1	The Data Processing system shall run at least once every minute
AC_P_2	The Data Processing system shall process the measured ODMR-sweeps into directional magnetic field values and, when applicable, convert these to a universal frame
AC_P_3	The Data Processing system shall use techniques such as averaging, smoothing and noise removal to improve the quality of the output data
AC_P_4	The Data Processing system shall save the directional and total magnetic field values to the Magnetic Field Database alongside the measurement's timestamp, location and the detector's name
AC_P_5	The Data Processing system shall be implemented in a way that allows for a replacement of the used processing algorithm with little enough downtime to not exceed other processing time-related requirements
AC_P_6	The Data Processing system shall monitor abnormal behaviour such as missing data packets and failures in data processing, and store these in a dedicated Diagnostics Database.
AC_P_7	The Data Processing System shall combine the data of different detectors to generate a map with interpolated magnetic field values
AC_P_8	The Data Processing System shall identify and record any fluctuation above a threshold and flag it as an unusual deviation
AC_P_9	The Data Processing System shall track the K-index of each detector in time

Data output requirements

Requirement ID	Description
AC_O_1	The system shall provide an API through which the back-up of the raw data can be downloaded
AC_O_2	The system shall provide diagnostics interface through which a privileged user can get information on recorded abnormal system behaviour and current measurement settings
AC_O_3	The system shall provide a public web interface through which a user can see a map of recorded magnetic field fluctuations
AC_O_4	The system shall send notifications to users within the region where auroral activity would currently be visible or upcoming
AC_O_5	The system shall provide a public web interface through which a user can see the directional and total magnetic field measurements for a selected detector

Appendix B: Protocol buffer used for communication

```
□ syntax = "proto3";

import "google/protobuf/timestamp.proto";

package auroracatcher;

service auroraQube{
    rpc sendDataFromQubeToServer (addData) returns (confirmationResponse);
    rpc sendCommandFromServerToQube (sendCommand) returns (confirmationResponse);
    rpc      sendDataFromQubeToServerWithRawData(addRawData)      returns
(confirmationResponse);
}
message addRawData{
    google.protobuf.Timestamp timestamp = 1;
    bytes raw_data = 2;
    float gps_long = 3;
    float gps_lat = 4;
    string deviceName = 5;
    float gps_date= 6;
    float gps_time= 7;
    float gps_alt= 8;
    float gps_speed= 9;
    float gps_course = 10;
}
message addData{
    google.protobuf.Timestamp timestamp = 1;
    float MagnetoX = 2; //referencevalue, not ODMR-based
    float MagnetoY = 3;
    float MagnetoZ = 4;
    float accelX = 5;
    float accelY = 6;
    float accelZ = 7;
    float gyroX = 8;
    float gyroY = 9;
    float gyroZ = 10;
    float laserTemp = 11;
    float boardTemp = 12;
    repeated float ODMR = 13;
    float gps_long = 14;
    float gps_lat = 15;
    string deviceName = 16;
    int32 qube_operationMode      = 17;
    int32 qube_measurementFlag    = 18;
    int32 qube_lastPingTick      = 19;
    int32 qube_hasConnection      = 20;
    int32 qube_averageOptical     = 21;
    int32 qube_averageFPGA       = 22;
    int32 qube_tickCount          = 23;
    int32 qube_blockIndex         = 24;
    int32 qube_ADCSettings        = 25;
```

```

string qube_accelGyroSettings = 26;
string qube_magnetoSettings   = 27;
string qube_devicesInit       = 28;
int32  windowSpan             = 29;
float  windowStep              = 30;
float  windowBroadStep        = 31;
int32  APPID = 32;
bool   peak_config = 33;
bool   pulsed_config0 = 34;
bool   pulsed_config1 = 35;
int32  laser_status = 36;
int32  laser_error = 37;
int32  laser_potValue = 38;
int32  mw_minFreq = 39;
int32  mw_maxFreq = 40;
int32  mw_attenuation = 41;
int32  mw_steps = 42;
int32  mw_broadSteps = 43;
string mw_settings = 44;
int32  mw_freqSkips = 45;
int32  mw_broadSweepStep = 46;
float  gps_date= 47;
float  gps_time= 48;
float  gps_alt= 49;
float  gps_speed= 50;
float  gps_course = 51;
}

message sendCommand{
    int32 nextMode = 1; //call to modify OperationMode
}
message confirmationResponse{
    bool validRequest = 1;
}

```