

Dynamical cancellation of eddy-current transients in ULF-MRI applications

Ilkka Rissanen

School of Science

Thesis submitted for examination for the degree of Master of
Science in Technology.

Espoo 25.11.2019

Supervisor

Prof. Risto Ilmoniemi

Advisor

Koos Zevenhoven



Aalto University
School of Science

Copyright © 2019 Ilkka Rissanen



Author Ilkka Rissanen

Title Dynamical cancellation of eddy-current transients in ULF-MRI applications

Degree programme Master's Programme in Engineering Physics

Major Engineering Physics

Code of major SCI3056

Supervisor Prof. Risto Ilmoniemi

Advisor Koos Zevenhoven

Date 25.11.2019

Number of pages 43

Language English

Abstract

Magnetic resonance imaging (MRI) is a medical imaging technique that is widely used for both clinical and research purposes. In recent years, an alternative approach to MRI known as ultra-low-field (ULF) MRI has attracted increasing interest, as performing measurements in the ultra-low-field regime has been demonstrated to have certain advantages over conventional MRI. These advantages include improved tissue contrast in certain applications, the capability to perform measurements on objects containing metal, potential lower cost, and capacity for combined use with other electromagnetically sensitive imaging techniques such as magnetoencephalography (MEG).

Because ULF-MRI signals are very delicate, measurements are highly sensitive to external magnetic noise. To alleviate this problem, ULF-MRI measurements are usually performed inside a magnetically shielded room (MSR). However, because MRI involves magnetic pulses, performing measurements inside an MSR typically means the ULF-MRI pulse sequence induces eddy currents in the MSR walls. These eddy currents can be strong enough to generate transient magnetic fields that prevent measurements until the transients have died down sufficiently.

While it is possible to simply wait for the transients to die down on their own, doing so can be impractical because the ULF-MRI signal also decays with time. To attain the best signal-to-noise ratio possible, transients should be eliminated from the system as quickly as possible. Transient cancellation techniques have been developed to accomplish this. In a novel approach known as dynamical cancellation, transients are nullified by adding a specifically designed cancellation pulse to the ULF-MRI pulse sequence. This pulse is designed to couple with eddy-current modes in the MSR walls in such a way that the ULF-MRI sequence as a whole induces minimal eddy currents.

In this thesis, I present an algorithm for finding efficient pulse waveforms for dynamical transient cancellation, along with measurements demonstrating the efficacy of dynamical cancellation in practice. Further improvements are also discussed, as it can be expected that, with various improvements to the implementation, dynamical cancellation can be used to even greater effect than demonstrated in this thesis.

Keywords ultra-low-field, ULF, magnetic resonance imaging, MRI, eddy current, transient, magnetically shielded room, MSR, dynamical shielding, DynaCAN, transient cancellation, SQUID



Tekijä Ilkka Rissanen

Työn nimi Pyörrevirtatransienttien dynaaminen kumoaminen ultramatalan kentän MRI:ssä

Koulutusohjelma Teknillisen fysiikan ja matematiikan tutkinto-ohjelma

Pääaine Teknillinen fysiikka

Pääaineen koodi SCI3056

Työn valvoja Prof. Risto Ilmoniemi

Työn ohjaaja Koos Zevenhoven

Päivämäärä 25.11.2019

Sivumäärä 43

Kieli Englanti

Tiivistelmä

Magneettikuvaus (MRI) on lääketieteellinen kuvantamismenetelmä jota käytetään laajalti sekä kliinisiin- että tutkimus-sovelluksiin. Viime vuosina vaihtoehtoinen lähestymistapa MRI:hin, nimeltään ultramatalan kentän MRI (ULF MRI) on ollut kasvavan kiinnostuksen kohteena, sillä MRI-mittausten tekemisellä ultramatalassa kentässä on havaittu olevan tiettyjä etuja tavalliseen MRI:hin verrattuna. Näihin hyötyihin kuuluvat mm. parempi kudiskontrasti tiettyissä sovelluksissa, kyky kuvata metallia sisältäviä kohteita, mahdollisesti halvemmat kustannukset sekä yhteensopivuus muiden sähkömagneettisesti herkkien kuvantamismenetelmien kuten magnetoencefalografian (MEG) kanssa.

ULF MRI -signaalit ovat hienovaraisia, ja siten herkkiä magneettisille häiriöille. Ulkoisen häiriön minimoimisen vuoksi ULF MRI -mittaukset toteutetaan usein magneettisuojahuoneen (MSR:n) sisällä. MSR:n käyttämisestä voi kuitenkin seurata oma ongelmansa. MRI-mittauksia varten on pulssitettava magneetikenttiä, ja tämän tekeminen MSR:n sisällä indusoi pyörrevirtoja huoneen seiniin. Nämä pyörrevirrat tuottavat magneetikenttätransientteja, jotka voivat olla niin voimakkaita, että ne estävät mittauksien tekemisen kunnes pyörrevirrat ovat riittävästi heikentyneet.

Vaikka pyörrevirtojen voi vain odottaa heikentyvän itsestään, odottaminen on usein epäkäytännöllistä, sillä ULF MRI -signaali myös heikkenee ajan kuluessa. Paras signaalikohinasuhde (SNR) saavutetaan, kun transientit poistetaan mahdollisimman nopeasti. Tätä varten on kehitetty transienttien kumoamismenetelmiä. Dynaaminen kumoaminen on uusi kumoamismenetelmä, jossa transientit kumotaan lisäämällä ULF MRI -pulssisekvenssiin erityinen kumoamispulssi. Tämä kumoamispulssi laaditaan kytkeytymään MSR:n seinien pyörrevirtamoodeihin tavalla, joka minimoi koko pulssisekvenssin yhdessä tuottamat pyörrevirrat.

Tässä opinnäytetyössä esitän algoritmin tehokkaiden kumoamispulssien löytämiseen dynaamista kumoamista varten, sekä mittauksia näyttämään dynaamisen kumoamisen toimivuutta käytännössä. Käyn myös läpi jatkokehitysmahdollisuuksia, sillä on odotettavissa, että kehittämällä käytettyjä menetelmiä edelleen on mahdollista saavuttaa vielä parempia tuloksia kuin mitä tässä työssä on esitetty.

Avainsanat ultramatala kenttä, ULF, magneettikuvaus, MRI, pyörrevirta, transientti, magneettisuojahuone, MSR, dynaaminen suojaus, DynaCAN, transienttien kumoaminen, SQUID

Preface

No thesis is an island, and so, many people deserve my thanks for helping me reach this point.

First, I would like to thank my supervisor, Professor Risto Ilmoniemi, and especially my advisor, Koos Zevenhoven, for their guidance, which has been extremely helpful in improving my writing skill. Through the many revisions of this work, I've learned many things about scientific writing that will surely help me with future writing endeavors. I would also like to thank the MEG–MRI research group as a whole for providing an excellent working environment for my thesis project and their patient assistance with my struggles with the measurement setup used in this project.

My time in this university has lasted over six years, and these years would've certainly been much less enjoyable were it not for the friends I've made along the way. I would like to extend my thanks to these fellow students who have shared this sometimes arduous journey with me. Whether it be banging our collective heads against a particularly tough physics exercise or spending the evening in a relaxing get-together, I truly appreciate all the times we've spent together during our time at Aalto University.

Finally, I would extend my gratitude towards my family for supporting me throughout all these years. Thanks to that support, both emotional and financial, I have been able to face the challenges of life with confidence and steadily forge ahead on my path in life.

Otaniemi, 25.11.2019

Ilkka J. Rissanen

Contents

Abstract	iii
Abstract (in Finnish)	iv
Preface	v
Contents	vi
Abbreviations	vii
1 Introduction	1
2 Background	3
2.1 Ultra-low-field magnetic resonance imaging	3
2.2 SQUID sensors	6
2.3 Magnetically shielded rooms	8
2.4 Eddy currents induced by MRI pulses	9
3 Theory	11
3.1 Dynamical Coupling for Additional dimeNsions – an overview	11
3.2 Eddy-current transients	14
3.3 DynaCAN based on inverse step responses	16
4 Methods	21
4.1 DynaCAN pulse waveform generation	21
4.1.1 Characterizing the eddy-current system	21
4.1.2 Finding a cancellation pulse waveform	22
4.1.3 Iterating residual transient cancellation	24
4.2 Error analysis	26
4.3 Measurement setup	29
5 Results	31
6 Discussion	37
7 Conclusions	40
References	42

Abbreviations

MRI	magnetic resonance imaging
NMR	nuclear magnetic resonance
ULF	ultra-low-field
MEG	magnetoencephalography
MSR	magnetically shielded room
DynaCAN	Dynamical Coupling for Additional dimeNsions
SQUID	superconducting quantum interference device
BFGS	Broyden–Fletcher–Goldfarb–Shanno (numerical optimization method)
RMS	root-mean-square

1 Introduction

Magnetic resonance imaging (MRI) is a biomedical imaging method that was brought to clinical use in its conventional form in the 1980s [1, 2]. The method is based on the phenomenon of nuclear magnetic resonance (NMR), which takes place when an external magnetic field imposes a net precessing motion upon the spins of the protons in the nuclei of the imaged object. This nuclear spin precession can be used to produce a measurable signal, from which the structure of the imaged object can be reconstructed. As MRI is safe, non-invasive, and capable of producing highly accurate 3D scans of human organs, it is the gold standard for structural biomedical imaging, widely used for both clinical and research purposes [1]. The various applications of MRI range from detecting tumors to measuring brain activity by studying the flow of blood in the brain.

Over the years, conventional MRI has developed towards using ever larger magnetic fields. In modern MRI devices, the strength of the main field \vec{B}_0 , which orients the proton spins, typically ranges from 1.5 T to 3 T [3]. A stronger main field has its advantages such as higher signal amplitude, which in turn results in clearer images. However, producing such strong magnetic fields requires expensive, heavy and power-consuming equipment. These inconveniences inherent in cutting-edge conventional MRI have encouraged the development of ultra-low-field MRI (ULF MRI) in recent years. In ULF MRI, the main field strength is many orders of magnitude less than that of conventional MRI, with typical B_0 strength ranging from 1 μ T to 100 μ T [3].

Convenience is not the only benefit expected from ULF MRI, as it has been shown that MRI in the ultra-low-field regime can outperform its conventional counterpart in some circumstances. Certain types of tissue contrast are higher in ULF MRI than in conventional MRI [4]. In addition to improved contrast in some applications, ULF MRI can image objects containing metal, whereas conventional MRI faces severe safety and artefact issues when metal is present in the imaged volume [3].

Perhaps one of the most interesting properties of ULF MRI, however, is the possibility of making it compatible with another biomedical imaging method of great clinical importance, magnetoencephalography (MEG) [5]. MEG is a neuroimaging method in which brain activity is imaged by measuring the faint magnetic fields produced by electrical activity in the brain [6]. Effective use of MEG requires knowledge of the structure of the subject's head, which is usually obtained from an MRI scan. As the same measurement setup can be used to measure both ULF MRI and MEG signals, a hybrid device could perform an ULF MRI scan of the subject's brain and an MEG measurement in one sitting. Achieving this would be a major step forward in MEG imaging, as current MEG setups use MRI scans made with a separate MRI imaging device, requiring the hospital to have access to two separate expensive measurement devices, and running the risk of introducing inaccuracies in the MEG estimate due to errors in mapping the MRI measurement geometry to the separate MEG coordinate system. In addition, using two separate devices makes the whole process more inconvenient and time-consuming for both patients and hospital staff.

Countering these upsides, the implementation of ULF MRI involves difficulties

that conventional MRI does not have to contend with [3]. Among these issues is the problem of the delicate ULF-MRI signal easily being drowned out by environmental magnetic noise. To alleviate this problem, ULF-MRI measurements are often performed inside a magnetically shielded room (MSR). While an MSR can eliminate disturbances caused by magnetic noise, using an MSR can cause a problem of its own. Magnetic pulses that are essential for MRI induce eddy currents in the MSR walls, which in turn generate transient magnetic fields that interfere with measurements. These magnetic-field transients can be many orders of magnitude larger than the measurement range of the detectors used in ULF MRI, making measurements impossible until the transients have died down.

Since the MRI signal also decays as time passes, simply waiting for the transients to diminish is at best suboptimal, at worst completely unfeasible. Because of this, effective transient cancellation methods are necessary for reaching the full potential of ULF MRI. While it is possible to design the ULF-MRI coils in a way that greatly diminishes the magnetic fields generated in the MSR walls, completely eliminating all the transients with coil design alone is often unattainable in practice [7].

An alternative approach, developed by Koos Zevenhoven in early 2010s [8], is to use Dynamical Coupling for Additional dimeNsions (DynaCAN), or dynamical cancellation, to nullify harmful transients [9]. In dynamical cancellation, eddy-current transients are cancelled by including an additional magnetic pulse to the ULF-MRI sequence that reduces the total eddy currents generated by the sequence as much as possible. While the individual pulses of the sequence would generate transient eddy currents in the MSR walls, the added cancellation pulse takes advantage of eddy-current dynamics and opposes the currents that would be induced by the rest of the sequence. As a result, the pulse sequence as a whole generates minimal transients.

The eddy currents induced by ULF-MRI pulses are a superposition of numerous overlapping spatial current patterns, or modes, that decay exponentially with different time constants [10]. Each of these eddy-current modes couples with different coils in different ways, making it difficult to find a DynaCAN pulse that generates the desired transient response when driven into the cancellation coil. However, it is possible to use the time behavior of the pulse to manipulate how it couples to different eddy-current modes. This approach of coupling to the dynamical eddy-current system with temporal pulse waveform features is the fundamental operating principle of dynamical cancellation with DynaCAN.

Since the cancellation efficacy of a DynaCAN pulse is mostly determined by the shape of the pulse waveform, DynaCAN allows for great flexibility in nullifying transients. With DynaCAN, the geometry of the cancellation coil can be chosen relatively freely, as the shape of the pulse waveform can be designed to generate the desired eddy-current response. This also means it is possible to use a single cancellation coil to cancel multiple different transients from different sources by applying multiple DynaCAN pulses into the coil, each designed to cancel a different transient. This thesis presents an algorithm for finding efficient cancellation pulses, along with measurements demonstrating the efficacy of DynaCAN in practice.

2 Background

2.1 Ultra-low-field magnetic resonance imaging

Most organic matter is rich in hydrogen atoms, the nuclei of which consist of only single protons. Protons, among other subatomic particles, exhibit a vector quantity known as the spin. The spin of a charged particle generates a magnetic field. A change in the orientations of proton spins causes a changing magnetic field. Magnetic resonance imaging (MRI) is a technique that uses this phenomenon to cause the hydrogen nuclei of an imaged object to generate a magnetic signal [1]. This MRI signal can then be measured in order to determine structural information about the imaged object.

In the presence of a magnetic field \vec{B} , the nuclear spins of an object tend to orient themselves parallel to the field. This phenomenon of the spins collectively orienting themselves in the same direction on average is referred to as spin polarization. If these polarized spins are then perturbed away from \vec{B} , they start precessing around \vec{B} similarly to how a slightly tilted spinning top behaves. This spin precession causes a changing magnetic field that can be detected from outside the imaged object. This signal, generated by the bulk precession of the entire ensemble of polarized protons, can be used for constructing high-resolution images of, for example, the human brain. This synchronized collective behavior of nuclei is known as nuclear magnetic resonance (NMR).

NMR generates a periodic signal, the frequency of which is determined by the frequency of the spin precession. The frequency of proton spin precession in the presence of a main field \vec{B}_0 is known as the Larmor frequency

$$f_L = \gamma B_0 ,$$

where γ is the gyromagnetic ratio of the proton [1]. The signal responses of the detection coils used in conventional MRI are frequency-dependent, with higher frequencies yielding stronger signals. Since the Larmor frequency is proportional to main field strength B_0 , higher B_0 leads to higher signal amplitude in conventional MRI.

The increase in Larmor frequency is not the only benefit a high main field strength offers in conventional MRI. Another factor that increases signal amplitude is the total magnetization \vec{M} of the protons, which is a measure of how polarized their spins are. Because inhomogeneity in spin orientations causes interference that reduces the total signal, higher spin polarization improves signal strength. The stronger the magnetic field affecting the spin directions, the greater the magnetization \vec{M} will be. Therefore, a high B_0 leads to high signal not only through higher Larmor frequency, but also through higher bulk magnetization. These two factors lead to a B_0^2 scaling of signal strength in conventional MRI [2].

In ultra-low-field (ULF) MRI, in which B_0 is many orders of magnitude smaller than in conventional MRI, this B_0^2 scaling would greatly reduce the signal-to-noise ratio of measurements. To overcome this, ULF MRI has methods for circumventing both the problem of low frequency and that of low magnetization. The issue

of low Larmor frequency is rendered moot by performing measurements with Superconducting QUantum Interference Device (SQUID) sensors, since they have a frequency-independent response. SQUID detectors and their operating principles will be examined in Sec. 2.2.

Even though using frequency-independent SQUIDs resolves the problem of low Larmor frequency, the issue of small B_0 resulting in low magnetization M still remains. This problem can be alleviated by prepolarizing the spins by pulsing a polarizing field \vec{B}_p that is orders of magnitude stronger than \vec{B}_0 , typically 10 mT to 100 mT [3]. As the magnetization, and thus signal strength, now scales with B_p , the small B_0 is no longer an issue. However, as B_p is still orders of magnitude weaker than the main field of conventional MRI, signal strengths in ULF MRI will inevitably be weaker.

The polarizing field \vec{B}_p is produced by a coil separate from the coil used to generate the main field \vec{B}_0 . This polarizing field is generally made perpendicular to the main field, and only pulsed long enough to properly polarize the proton spins. Once the spins are suitably polarized, the polarizing field is ramped down and the spins are brought to a precessing motion around \vec{B}_0 . There are two methods for ramping down B_p [7]. In the non-adiabatic method, B_p is ramped down quickly, leaving the magnetization perpendicular to the main field, and thus causing precession around \vec{B}_0 . In the adiabatic method, the polarizing field is ramped down slowly. This causes \vec{M} to turn towards \vec{B}_0 , following the direction of the total magnetic field $\vec{B}_p + \vec{B}_0$. The ramp down is still sufficiently fast that, while the direction of \vec{M} changes, its magnitude remains roughly constant during the adiabatic ramp down. Once the magnetization has aligned with the main field, an excitation coil is pulsed to generate an excitation field \vec{B}_1 to tilt the magnetization away from \vec{B}_0 .

The non-adiabatic process has the potential advantages of allowing imaging to begin almost immediately after polarization and of making the excitation coil unnecessary, thereby simplifying the system. However, in the adiabatic process, the use of an excitation coil provides more control over the behavior of the magnetization. Furthermore, stabilizing the magnetization parallel to \vec{B}_0 before measurement improves the directional homogeneity of \vec{M} , as the main field can be designed to be much more homogeneous than the polarizing field.

The MRI signal results from the changing magnetic flux pattern generated by the precessing magnetization. Notably, the signal is generated by the bulk of protons within the imaged object; each of the precessing protons generates its own magnetic field, which means the measured field is a sum of the individual fields generated by individual proton spins. To distinguish between spins at different locations in order to determine the precise structure of the imaged object, gradient fields are used to encode spatial information into the MRI signal by having precession frequency depend on the location of the proton.

Three gradient fields, \vec{G}_x , \vec{G}_y and \vec{G}_z , are used to alter the Larmor frequency $f_L = \gamma B$. The magnitude of the total field becomes $B_{\text{tot}}(\vec{r}) = B_0 + \Delta B(\vec{r})$, where $\Delta B(\vec{r})$ is the contribution of the gradient fields, assuming (inaccurately) that there is no change in the field direction. If the main field is defined as $\vec{B}_0 = B_0 \hat{e}_z$, where \hat{e}_z is a unit vector along the z -axis, the gradient is defined as the vector $\vec{G}(t)$ with

components $G_i(t) = \frac{\partial B_z}{\partial i}(t)$, $i = x, y, z$. We can then define the spatial dependence of the change in B_z as $\Delta B(\vec{r}) = \vec{r} \cdot \vec{G}(t)$, resulting in

$$f_L(\vec{r}, t) = \gamma [B_0 + \vec{r} \cdot \vec{G}(t)] . \quad (1)$$

With f_L being the frequency of the precession of magnetization \vec{M} , and ignoring relaxation effects for now, the rotating component \vec{M}_{xy} can be expressed as a complex number

$$\widetilde{M}_{xy}(\vec{r}, t) = M_{xy} e^{-i \int_0^t f_L(\vec{r}, t') dt'} = M_{xy}(\vec{r}) e^{-i\gamma B_0 t} e^{-i\gamma \int_0^t \vec{r} \cdot \vec{G}(t') dt'} . \quad (2)$$

Defining $\omega_0 = \gamma B_0$ and $\vec{k} = \gamma \int_0^t \vec{G}(t') dt'$, the above equation can be simplified into

$$\widetilde{M}_{xy}(\vec{r}, t) = M_{xy}(\vec{r}) e^{-i\omega_0 t} e^{-i\vec{k} \cdot \vec{r}} . \quad (3)$$

One can construct a total complex signal Ψ , which is proportional to $\widetilde{M}_{xy}(\vec{r}, t)$ integrated over the volume of the imaged object. As the factor $e^{-i\omega_0 t}$ has no spatial dependence, it can be taken outside the integral, resulting in

$$\Psi \propto e^{-i\omega_0 t} \int_{\mathbb{R}^3} M_{xy}(\vec{r}) e^{-i\vec{k} \cdot \vec{r}} d^3\vec{r} , \quad (4)$$

where we assume an ideal detector with a uniform sensitivity profile. In signal processing, the constant frequency factor $e^{-i\omega_0 t}$ can be removed with a demodulation technique known as quadrature detection [1]. This demodulated signal Ψ_{QD} can be expressed as

$$\Psi_{\text{QD}} \propto \int_{\mathbb{R}^3} M_{xy}(\vec{r}) e^{-i\vec{k} \cdot \vec{r}} d^3\vec{r} , \quad (5)$$

which notably turns out to be the three-dimensional spatial Fourier transformation of $M_{xy}(\vec{r})$, where \vec{k} is a vector in spatial frequency space referred to as \vec{k} -space. By manipulating \vec{k} as measurements are made, MRI essentially directly measures the Fourier transform of the imaged object.

To get an image out of the signal, Ψ has to be measured at a large number of points in \vec{k} -space. As a rough reconstruction, after measuring over a grid in \vec{k} -space, a discrete inverse Fourier transform turns the spatial frequency data of \vec{k} -space into an image. The resolution of the image depends on the number of points in \vec{k} -space we have measurements of—the more points we have, the more complete a picture of the frequency-domain image we get, and thus the better the image resolution will be. To acquire \vec{k} -space data, the gradient $\vec{G}(t)$ is applied to alter the value of \vec{k} , which was defined as $\vec{k} = \gamma \int_0^t \vec{G}(t') dt'$. Essentially, the gradients allow us to travel in \vec{k} -space with velocity $\gamma \vec{G}$.

Traditional tissue contrast in MRI arises from proton-density contrast and the relaxation effects we have ignored until now. Because the total magnetization is proportional to the proton density, proton-density contrast can be seen directly from $\vec{M}_{xy}(\vec{r})$ acquired from the inverse Fourier transform of \vec{k} -space data. However, in soft

tissue imaging, proton-density contrast is much weaker than tissue contrast obtained from relaxation effects in magnetization [1].

Relaxation effects cause the magnetization component magnitudes M_{xy} and M_z to change as \vec{M} precesses around \vec{B}_0 . From the typical starting point of turning the magnetization $\vec{M} = M_0\hat{e}_z$ by 90° to the xy -plane so that the initial magnetization is $\vec{M}_0 = M_0\hat{e}_x$, there are two relaxation factors that affect \vec{M} . M_z relaxes with time constant T_1 from zero to equilibrium value M_{eq} determined by the applied magnetic field, and M_{xy} relaxes with time constant T_2 towards zero. As different tissues have different values of T_1 and T_2 , tissue contrast can be improved far beyond mere proton-density contrast by designing the MRI pulse sequences so that the generated signal greatly depends on T_1 or T_2 . Using contrast based on relaxation times is the standard method of achieving tissue contrast in most applications of biomedical MRI.

2.2 SQUID sensors

Superconducting QUantum Interference Devices (SQUIDS) are highly sensitive magnetic flux detectors, the operating principles of which are related to two phenomena of superconducting electronics: Josephson tunneling and flux quantization [11].

When a material capable of superconductivity is transitioned into a superconducting state by cooling it below its critical temperature T_c , its electrons form Cooper pairs that can carry a current with zero resistance. A Josephson junction consists of two superconductors separated by a thin insulating barrier. This barrier is thin enough that, in a superconducting state, Cooper pairs can tunnel through it and carry current across the junction. Applying a current through a Josephson junction does not generate voltage until a critical current is reached, after which a voltage appears in the circuit.

Flux quantization refers to the phenomenon of the total flux passing through a superconducting loop always being an integer multiple of the flux quantum $\Phi_0 = h/2e \approx 2.07 \cdot 10^{-15} \text{ Tm}^2$, where h is Planck's constant and e is the elementary charge [11]. The direct current (dc) SQUID is a superconducting loop with two Josephson junctions in parallel [12]. For a superconducting loop like this, an external magnetic field that is not an integer multiple of Φ_0 induces a phase difference in the wave function of the Cooper pairs across the Josephson junctions. Since the magnitude of the phase difference depends on the deviation of the external magnetic field from the nearest integer multiple of the flux quantum, the phase difference is periodic with period Φ_0 .

A constant bias current I_b , larger than the critical current, is applied across the SQUID loop, causing a voltage to appear due to the Josephson junctions. The magnitude of this voltage depends on the magnetic flux through the superconducting loop. This is because the phase difference induced in the Cooper pair wave function across the Josephson junctions alters the relation between voltage across and current through the loop. With the bias current remaining constant, this results in the voltage across the loop being a function of the external magnetic flux. The voltage oscillates with a period of Φ_0 in accordance with the aforementioned flux quantization.

Notably, if the voltage across the superconducting SQUID loop is measured,

the response is dependent on the magnitude of the magnetic flux, not its time derivative like in an induction coil. Because of this, a SQUID can be thought of as a flux-to-voltage transducer that permits direct measurement of magnetic flux by measuring the voltage across the loop. In measuring biomagnetic signals, however, it is conventional to operate the SQUID in a flux-locked loop, in which a feedback loop keeps the voltage constant, and thus also the flux [12]. What is then measured is the signal that goes into the feedback loop. This gives a direct measure of the flux without leading to issues that can arise from voltage response periodicity, such as loss of accuracy near the minima and maxima of the periodic response function.

As seen in Fig. 1, the superconducting loop with Josephson junctions and the feedback coil L_f with its related electronics are not the only building blocks of the basic dc SQUID setup. The small size of the SQUID loop makes the flux through it extremely small. To enhance the signal, a flux transformer is used to increase the effective area of the SQUID loop. A simple flux transformer consists of two connected coils, the pickup coil L_p and the input coil L_i . The larger pickup coil is used to capture the faint MRI signal, which is then replicated by the smaller input coil near the superconducting loop. This allows the SQUID to record the flux from a larger area than the SQUID loop itself encloses, improving signal strength. As the superconducting parts have to be cooled down below the critical temperature of the conductor, the SQUID setup has to be placed in a dewar containing a suitable coolant. The coolant used for SQUIDs is usually liquid helium [12], which boils at 4.2 K, well below the critical temperature of the common superconductor niobium, which has $T_c = 9.2$ K [13].

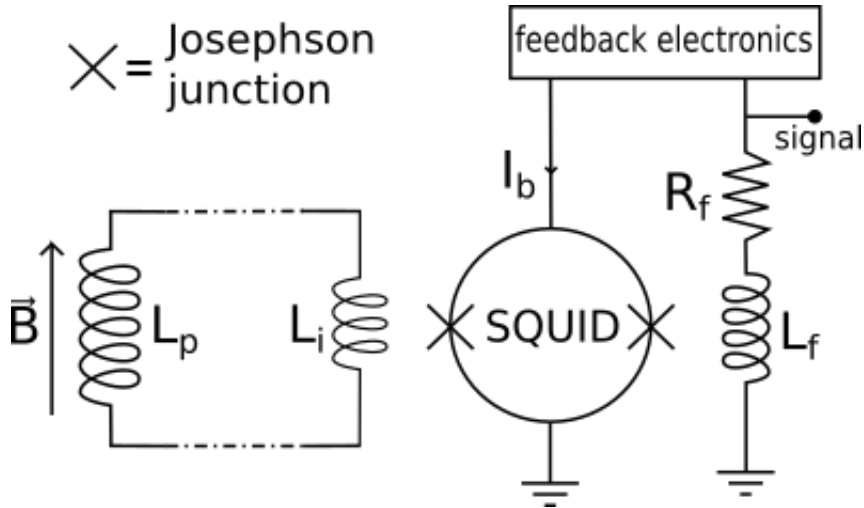


Figure 1: The schematic of a dc SQUID sensor. The external field \vec{B} is picked up by the flux transformer pickup loop L_p and the signal is replicated close to the SQUID loop by the input coil L_i . The feedback coil L_f keeps magnetic flux through the SQUID loop constant, with the input to the feedback coil being the SQUID signal.

2.3 Magnetically shielded rooms

A properly optimized SQUID sensor can detect tiny changes in magnetic fields, such as changes measured in femtoteslas or even attoteslas [12]. This extremely high resolution is necessary for properly measuring faint ULF-MRI signals. However, high resolution is of little use in the presence of external magnetic noise equal or greater in magnitude compared to the signal to be measured. The Earth's magnetic field alone can produce random noise of a magnitude measured in nanoteslas [14], and activity from nearby electronics and moving objects can result in even greater disturbances. Furthermore, the magnitude of the Earth's magnetic field on Earth's surface ranges from roughly 20 μT to 65 μT [15]. Since that magnitude is comparable to the typical strength of an ULF MRI main field, the Earth's magnetic field can severely distort ULF-MRI measurements. For these reasons, ULF MRI requires shielding from external magnetic interference for accurate measurements to be possible. A common approach to shielding is to perform measurements inside a magnetically shielded room (MSR).

An MSR can be designed to protect against different types of magnetic interference [7]. One type of magnetic interference is high-frequency magnetic interference, in which the magnetic field changes rapidly. As this type of interference is characterized by a changing magnetic field, it can be shielded against with eddy currents. Plates of metal with high electrical conductivity respond to the changing magnetic field by forming eddy currents that generate opposing magnetic fields as per Lenz's law. The higher the conductivity of the material, the stronger are the induced eddy currents, resulting in increased shielding capability.

Another type of magnetic interference that can be shielded against is low-frequency magnetic interference and (almost) static background fields. Since eddy currents are the result of a changing magnetic field, eddy-current shielding becomes less effective the lower the rate of change in the external magnetic field is. Shielding against low-frequency interference or static fields is instead accomplished by providing magnetic field lines an easier path around the shielded volume than through it. This is usually done by using layers of highly permeable ferromagnetic metals in the MSR walls. One common metal used for this purpose is an alloy of nickel and iron referred to as μ metal, the permeability of which can be on the order of $10^4\mu_0$ [16].

Depending on the application of the MSR, one form of shielding might be preferred over the other. An MSR intended for low-frequency measurements can be constructed almost entirely of μ metal, as reducing high-frequency interference is unimportant. For applications where static external fields or low-frequency interference are not a problem, the MSR could be made entirely of a highly conductive metal such as aluminium.

Regardless of what type of interference the MSR is tailored to shield against, there are certain factors to MSR design that have to be taken into account. For maximal shielding efficacy, the MSR plates should enclose the shielded area as completely as possible. This means that the MSR should consist of as few individual plates as possible, as seams between plates leave small gaps that reduce shielding efficacy. Especially in eddy-current shielding, it is conventional to weld seams to let current

flow as freely as possible from plate to plate and from wall to wall [17]. In addition to seams, the use of electrical equipment inside the MSR can necessitate holes and pass-through arrangements in the MSR walls for various cables, which has an effect on the shielding efficacy. Another problem factor is the necessity of having a door. As the seams connecting the door to the rest of the MSR cannot be welded, the existence of a door creates a weak point in the MSR that cannot be completely eliminated.

The physical shape of the MSR also has an effect on its shielding factor. Spherical shapes are sometimes considered more effective at reducing electromagnetic interference than cubic shapes [11, 18]. However, most MSRs are cube-shaped, as spherical rooms are inconvenient to build and have to be larger for there to be the same amount of operational space inside the room.

Tailoring the MSR to excel at shielding against one type of interference, even at the cost of inefficacy in shielding against other types, is often useful. Unfortunately ULF MRI, especially hybrid MEG-MRI, requires shielding against both high- and low-frequency interference. Static fields affect the magnitude, direction and homogeneity of the main field B_0 , when all of those parameters should be carefully fixed. High-frequency noise makes accurately measuring the delicate signals impossible. As such, the design of an MSR for ULF MRI often incorporates layers of both μ metal and aluminium for cancellation against both types of interference. This is especially true for ULF MRI combined with MEG, as MEG measurements involve frequencies almost down to dc. Applying multiple layers of different materials for MSR construction is not particularly unusual, as even MSRs designed to protect against one type of interference are often made of layers of appropriate metals. It has been demonstrated that adding layers improves the shielding factor more than increasing layer thickness does [19].

While eddy-current shielding is necessary to prevent high-frequency external noise from interfering with the sensitive ULF-MRI measurements, having conductive plates near the ULF-MRI coils during measurement can be a problem of its own. ULF-MRI pulses such as the polarizing pulse induce eddy-current transients in the MSR walls. While these transients may not always significantly interfere with the ULF-MRI pulses themselves, they can easily prevent SQUID measurements for a period of time after pulsing the ULF-MRI coils. Alleviating this problem by restricting eddy currents is often not viable, since shielding against external high-frequency interference is based on induction of eddy currents—preventing eddy currents from forming in the walls of the MSR would compromise its shielding efficacy.

2.4 Eddy currents induced by MRI pulses

Changing magnetic fields induce eddy currents in conductive surfaces such as MSR walls. At any point in the conductive surface, the current can flow in any direction tangent to the surface. The only restriction for eddy currents is that the current patterns must consist of closed current loops; eddy currents cannot cause charge to accumulate in one point in the conductor [10]. Because there is an endless variety of current distributions that fulfill that condition, there is an infinite number of possible

eddy-current patterns an MRI pulse sequence can induce in the MSR walls.

The eddy-current patterns induced by ULF-MRI pulse sequences can differ significantly in complexity. Some patterns might be as simple as current travelling in a straight path across the walls of the MSR, encircling the room. Other patterns might involve complicated current distributions spanning the floor and ceiling of the MSR in addition to the walls. The current in each of these characteristic spatial patterns, or modes, decays exponentially with its own time constant [10]. These time constants can vary greatly between different modes, ranging from less than a millisecond to several seconds.

Several factors affect the time constants of eddy-current modes. Notably, the geometry of an eddy-current mode plays a significant part in determining its time constant. Simple, straightforward current modes consisting of large current loops have relatively long time constants, whereas more intricate modes tend to decay first [7]. Another significant factor that affects time constants is the thickness of the conductive layers of the MSR, with thicker walls resulting in longer time constants [10]. The conductivity of the walls also affects eddy currents. The less resistance an eddy current encounters, the longer it will take for it to decay.

Through these factors, the design of the MSR has a significant effect on the eddy currents induced in its walls. In addition to the straightforward effects of the choice of material and the thickness of the walls, the design of seams between conductive plates affects the behavior of eddy currents in the MSR walls. If the seams between the conductive plates are not highly conductive, eddy currents passing between plates are weakened due to the extra resistance at the seams. Since eddy currents spanning multiple plates tend to cover more surface area and thus have longer time constants, resistive seams affect eddy-current modes with long time constants the most, whereas modes with short time constants might be unaffected. How small areas the eddy currents are confined to depends on the design of the plates the walls consist of. If each wall is a single plate, there can be enough space for large current loops with long time constants to form even if the seams between walls are highly resistive. If walls consist of a number of weakly connected smaller plates, currents can be confined in a way that does not leave room for modes with long time constants.

The variety of possible eddy-current modes makes the behavior of eddy currents in conductive surfaces complicated. What makes it even more complex is that magnetic pulses do not simply excite a single, complicated eddy-current mode with a singular time constant. Instead, the pulse can couple to any number of modes, causing the resulting eddy currents to be a superposition of multiple modes with different time constants [9]. Because the current in each of these individual eddy-current modes decays at its own rate, the time behavior of the transients as a whole cannot be characterized with a single time constant, and the spatial pattern changes with time.

3 Theory

Transient magnetic fields generated by eddy currents induced by MRI pulse sequences present a significant hurdle for making accurate ULF-MRI measurements. These exponentially decaying transients can be many orders of magnitude larger than the signal produced by ULF MRI. Due to the limited field range of SQUID sensors, large magnetic-field transients make measurements impossible until the fields have decayed sufficiently, at which point the MRI signal has significantly diminished as well. In this chapter, I provide an overview of the transient cancellation technique known as dynamical cancellation, or Dynamical Coupling for Additional dimeNsions (DynaCAN), along with an exploration of the physics of eddy-current transients and how they relate to DynaCAN.

3.1 Dynamical Coupling for Additional dimeNsions – an overview

As established in Sec. 2.4, the transient eddy currents induced in MSR walls by ULF-MRI pulse sequences appear as a combination of multiple eddy-current modes that decay at different rates. While the presence of multiple separate modes with different time constants makes the transient behavior more complicated overall, it can also be made use of in dynamical transient cancellation. This method of nullifying transients was first developed by Koos Zevenhoven in 2011 [8], and first implemented by Zevenhoven *et al.* in 2015 [9].

Each pulse in an ULF-MRI sequence induces a unique transient response that is a linear combination of various eddy-current modes [10]. How a pulse couples to different modes depends on the design of the coil it is driven to, the design of the MSR, and the shape of the pulse waveform [7]. The objective of dynamical transient cancellation is to nullify the transient eddy currents induced by ULF-MRI pulses by adding a new pulse into the ULF-MRI sequence [9]. This additional pulse waveform is designed to couple to the eddy-current modes in a way that reduces the transients induced by the ULF-MRI pulse sequence as a whole. By adding a cancellation pulse that on its own would generate eddy-current transients that are the opposite of those induced by the rest of the sequence, the eddy currents induced by the pulse sequence as a whole are minimized.

Because the coils used for MRI pulses are usually not designed to couple strongly with the eddy-current modes of the MSR, the cancellation pulse waveform is driven into a separate cancellation coil. This can also reduce the magnetic disturbance caused by the pulse in the imaging volume. Because the design of the coil plays a part in determining how a pulse couples with different modes, driving the same waveform into two different coils will not induce the same eddy currents. This means that simply using a cancellation pulse identical but opposite to another pulse in the sequence, multiplied with some suitable scaling coefficient, is unlikely to produce the exact opposite transient response required to properly nullify eddy currents. Instead, the cancellation pulse uses a specifically designed waveform in order to couple with the eddy-current modes of the MSR in a way that cancels the transients generated

by the rest of the sequence.

The coupling of a magnetic pulse waveform to eddy-current modes with different time constants depends on the features of the waveform at different time scales [9]. For example, low-frequency features of the magnetic pulse couple stronger to eddy-current modes with long time constants than high-frequency features, which in turn couple stronger to eddy-current modes with short time constants. By constructing the cancellation waveform out of various features with different time scales, it is possible to design a pulse that induces eddy currents with desired magnitudes to particular eddy-current modes. This way, harmful transients induced by ULF-MRI pulses can be eliminated by finding a waveform that couples with the relevant eddy-current modes in a way that induces the opposite current to them. This approach of using pulse features with different time scales to manipulate eddy-current modes is known as DynaCAN. Figure 2 shows an example of a DynaCAN transient cancellation pulse waveform designed to zero the 12 eddy-current modes with the longest time constants induced by its accompanying polarizing pulse [8].

What gives transient cancellation with DynaCAN an advantage over approaches to transient cancellation based on coil or MSR design is its flexibility. Tackling eddy currents by designing the MSR and the coil setup in such a way that minimal transients are produced, while effective, does not lend itself well to tweaking the transient response should the need arise. Should changes be made to the coil system such as changing the location of a coil, or should a new measurement scheme have different requirements for transients, the affected coils have to be redesigned and reconstructed in order to retain optimal cancellation efficacy. With DynaCAN, there is no need to make changes to the cancellation coil, as the transient response can be

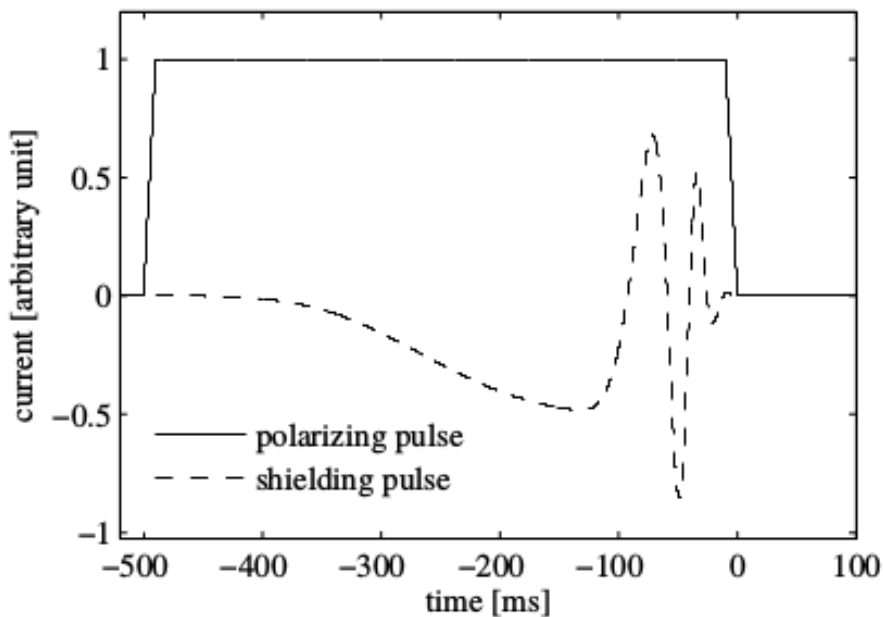


Figure 2: An example pulse for cancelling transients generated by a polarizing pulse. Figure reproduced from Ref. [8].

optimized to meet new requirements simply by changing the waveform of the pulse driven in the coil accordingly.

While there is leeway in cancellation coil design thanks to the flexibility of the pulse waveform, some coil designs work better for DynaCAN than others. The stronger the field generated by the coil is at the MSR walls, the stronger its coupling with eddy-current modes will be. Because of this, having the DynaCAN coil be close to the walls can improve the efficacy of dynamical cancellation while also reducing interference caused by the DynaCAN pulse in the center of the room. Additionally, the shape of the magnetic field generated by the coil determines how easy it is to couple to particular modes. The closer the shape of the magnetic field produced by the cancellation coil is to the magnetic field produced by another coil, the easier it is to generate transient responses similar to what that coil could produce. Since the shape of the magnetic field is determined by coil design alone, it is useful to design the coil to generate a similar magnetic field as the coils that cause the most harmful transients during the ULF-MRI pulse sequence. As the polarizing pulse is orders of magnitude stronger than other magnetic pulses used in ULF MRI, this usually means that designing the DynaCAN coil to be concentric with the polarizing coil is preferable for efficient transient nullification. An illustration of this type of coil setup can be found in Fig. 3.

The flexibility of DynaCAN can also enable a greater cancellation efficacy than what can be accomplished with coil design alone. While it is theoretically possible to construct the coils and the MSR so that eddy-current transients are practically nonexistent, space limitations and other practical considerations impose constraints on MSR and coil designs that restrict the cancellation efficacy that can be achieved with such methods in practice. Of course, DynaCAN is not without its practical limitations either—for example, the amplifier used for driving the pulse into the DynaCAN coil limits the maximum amplitude and slope of the waveform. Still, in suitable conditions, it can be expected that dynamical cancellation can be made

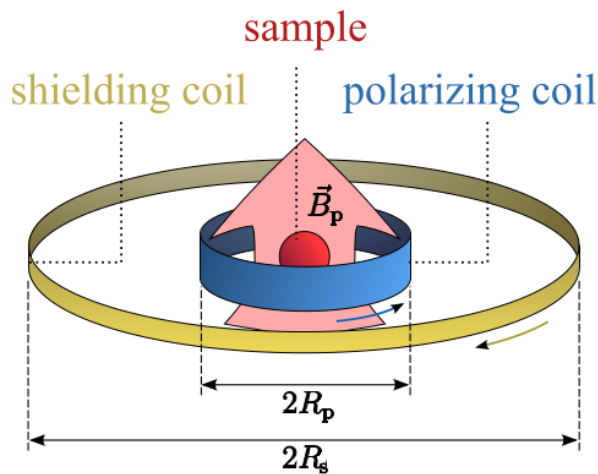


Figure 3: A schematic depiction of an example coil setup showing a cancellation coil coplanar with the polarizing coil. Figure reproduced from Ref. [8].

more effective at cancelling transients than methods based on coil design alone.

3.2 Eddy-current transients

In this section, we examine the theory of a system of eddy currents in more detail to see that the eddy currents remaining after magnetic pulses indeed are a combination of various exponentially decaying modes with their individual time constants.

The physics of eddy currents in MSR walls has been studied in detail by Zevenhoven *et al.* in Refs. [10] and [7]. Approximating the layers of aluminium and μ metal in the MSR walls as ohmic conductors, eddy-current transient generation can be thought of as a time-invariant linear process. Expressing eddy currents as a linear combination of n basis functions that each have their own amplitude and mutual inductance with the MRI pulses, transient behavior is modeled as an n -dimensional system with input $\mathbf{g}(t)$ and state $\mathbf{j}(t)$. Both $\mathbf{g}(t)$ and $\mathbf{j}(t)$ are vectors of size n where each element corresponds to a particular eddy-current basis function. The state vector $\mathbf{j}(t)$ contains the current amplitudes of each basis function, and the input vector $\mathbf{g}(t)$ displays the electromagnetic induction of new current into the various basis functions. A simple model for a time-invariant linear system with n degrees of freedom is

$$\frac{d}{dt}\mathbf{j}(t) = \mathbf{A}\mathbf{j}(t) + \mathbf{g}(t) , \quad (6)$$

where \mathbf{A} is an $n \times n$ constant coefficient matrix that describes the properties of the system. Because there is an infinity of possible eddy-current patterns, a perfectly accurate model would require n to be infinitely large. However, in practice, the number of eddy-current modes that produce a notable transient is limited. Simply taking into account the ten to twenty most significant modes can be enough to model the transient behavior of the system with reasonable accuracy. Therefore, n can be limited to a computationally manageable number.

Equation (6) has a known solution

$$\mathbf{j}(t) = e^{t\mathbf{A}}\mathbf{j}(t_0) + \int_{t_0}^t e^{(t-s)\mathbf{A}}\mathbf{g}(s) ds , \quad (7)$$

where the first term describes homogeneous behavior with initial condition $\mathbf{j}(t_0)$ and the second term shows the effect of outside input. In the case of eddy-current transients in ULF-MRI applications, this outside input is the induction of new eddy currents caused by the ULF-MRI pulse sequence as a whole.

If \mathbf{A} is diagonalized as $\mathbf{A} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^{-1}$, where $\mathbf{\Lambda}$ is a diagonal matrix consisting of the eigenvalues λ_k of \mathbf{A} and matrix \mathbf{S} contains the eigenvectors \mathbf{s}_k corresponding to λ_k , Eq. (7) can be expressed as

$$\mathbf{j}(t) = \mathbf{S}e^{t\mathbf{\Lambda}}\mathbf{S}^{-1}\mathbf{j}(t_0) + \mathbf{S} \int_{t_0}^t e^{(t-s)\mathbf{\Lambda}}\mathbf{S}^{-1}\mathbf{g}(s) ds . \quad (8)$$

The above expression describes transient behavior in terms of the vector quantities $\mathbf{j}(t)$ and $\mathbf{g}(t)$, which describe, respectively, the current magnitude and the induction of new current at any given time separately for every eddy-current basis function.

However, if in practical measurements the measured quantity is a scalar magnetic field component and the input is a combination of scalar pulses driven into coils, it is more useful to look at these scalar qualities as opposed to the vectors $\mathbf{j}(t)$ and $\mathbf{g}(t)$.

The magnitude of a scalar magnetic field component generated by eddy-current transients at a particular point of interest can be expressed as

$$B(t) = \mathbf{q}^* \mathbf{j}(t) , \quad (9)$$

where $*$ denotes the conjugate transpose and vector $\mathbf{q}(t)$ describes the coupling between the n eddy-current patterns and a magnetic-field component $B(t)$ at a point of interest.

Since the input $\mathbf{g}(t)$ of the system arises from a set of coils each with their own pulse $I(t)$ and coupling vector \mathbf{p} , the total input can be expressed as a sum over all the coils in the system,

$$\mathbf{g}(t) = \sum_i \frac{dI_i}{dt} \mathbf{p}_i , \quad (10)$$

where vectors \mathbf{p}_i describe the coupling between the scalar pulses $I_i(t)$ in particular coils and the n eddy-current patterns of the system. Because electromagnetic induction is caused by change in magnetic field, the vector input $\mathbf{g}(t)$ is expressed in terms of the time derivatives of the pulses $I_i(t)$.

By applying the definitions in Eqs. (9) and (10) into Eq. (8), we get the scalar expression

$$\begin{aligned} B(t) &= \mathbf{q}^* \mathbf{S} \left(e^{t\Lambda} \mathbf{S}^{-1} \mathbf{j}(t_0) + \sum_i \int_{t_0}^t e^{(t-s)\Lambda} \mathbf{S}^{-1} \mathbf{p}_i \frac{dI_i}{dt}(s) ds \right) \\ &= \sum_i \tilde{\mathbf{q}}^* \left(e^{t\Lambda} \tilde{\mathbf{j}}(t_0) + \int_{t_0}^t e^{(t-s)\Lambda} \tilde{\mathbf{p}}_i \frac{dI_i}{dt}(s) ds \right) , \end{aligned} \quad (11)$$

where \mathbf{S} has been absorbed into redefinitions of the coupling vectors and the system state as $\tilde{\mathbf{p}}_i = \mathbf{S}^{-1} \mathbf{p}_i$, $\tilde{\mathbf{q}}(t) = \mathbf{S}^* \mathbf{q}(t)$ and $\tilde{\mathbf{j}}(t) = \mathbf{S}^{-1} \mathbf{j}(t)$.

We can examine the transient response of a single coil by setting the current in other coils to zero. The same essential physics of eddy-current transients will be found regardless of whether the system has one or more coils. In a single-coil system, $\mathbf{g}(t) = \frac{dI}{dt} \mathbf{p}$, where $I(t)$ is the pulse driven into the coil and \mathbf{p} the eddy-current coupling vector corresponding to the coil. This simplifies the above expression of $B(t)$ into

$$B(t) = \tilde{\mathbf{q}}^* \left(e^{t\Lambda} \tilde{\mathbf{j}}(t_0) + \int_{t_0}^t e^{(t-s)\Lambda} \tilde{\mathbf{p}} \frac{dI}{dt}(s) ds \right) . \quad (12)$$

Writing Eq. (12) as a sum over the n eddy-current modes k results in

$$\begin{aligned} B(t) &= \sum_k \tilde{q}_k^* \left(e^{\lambda_k t} \tilde{j}_k(t_0) + \int_{t_0}^t e^{\lambda_k(t-s)} \tilde{p}_k \frac{dI}{dt}(s) ds \right) \\ &= \sum_k \tilde{q}_k^* e^{-\frac{t}{\tau_k}} \left(\tilde{j}_k(t_0) + \int_{t_0}^t e^{\frac{s}{\tau_k}} \tilde{p}_k \frac{dI}{dt}(s) ds \right) , \end{aligned} \quad (13)$$

where $\tau_k = -1/\lambda_k$. This expression for the generated magnetic-field component $B(t)$ shows that the transient is composed of exponential components with time constants τ_k . The scalar $\tilde{j}_k(t_0)$ is the state of transient mode k at the starting time t_0 , whereas the integral $\int_{t_0}^t e^{s/\tau_k} \tilde{p}_k \frac{dI}{dt}(s) ds$ expresses the induction of new current into mode k caused by the changing magnetic field resulting from dI/dt . For a pulse $I(t)$ that ends at t_1 , and choosing a t_0 before any eddy currents have been induced in the walls yet, Eq. (13) can be expressed as

$$B(t) = \sum_k \tilde{q}_k^* e^{-\frac{t}{\tau_k}} \int_{t_0}^{t_1} e^{\frac{s}{\tau_k}} \tilde{p}_k \frac{dI}{dt}(s) ds \quad \text{when } t \geq t_1, \quad (14)$$

since $\tilde{j}_k(t_0) = 0$ for all k and $\frac{dI}{dt}(t) = 0$ for $t \geq t_1$. As the integral is no longer a function of t , it is clear that without external magnetic pulses inducing new currents into the system, the magnetic-field transient is a linear combination of exponentially decaying components. The magnitudes of these transients are determined by the pulse generated by $I(t)$ and its coupling with the various eddy-current modes.

All coils in the ULF MRI setup produce transients in accordance with Eq. (13). What differs between coils is the coupling \tilde{p} between the pulse $I(t)$ and the eddy-current modes $\tilde{j}(t)$, resulting in the same pulse driven to different coils producing a different transient response. The aim of removing transients with DynaCAN is to find a cancellation waveform $I_c(t)$ that, when driven into a separate DynaCAN coil, excites $\tilde{j}(t)$ in a way that opposes the transients produced by the other pulses in an MRI sequence, despite the different coupling \tilde{p} of the DynaCAN coil. In ideal cancellation, when ULF-MRI measurements are being performed, $\tilde{j}(t)$ has been reduced to zero by the cancellation pulse.

When directly measuring eddy currents $\tilde{j}(t)$ in the MSR walls or accurately characterizing the coupling \tilde{p} between coil pulses and eddy-current modes is difficult, it is simpler to find a suitable DynaCAN waveform by focusing on the easily measurable magnetic-field transient rather than the eddy currents themselves. By finding a waveform that produces a magnetic-field transient that is the opposite of the transient produced by ULF-MRI pulses, the total transient field in the system becomes zero when the cancellation pulse is added to the ULF-MRI sequence. Because the magnetic-field transient is generated by the multiple eddy-current modes in the MSR walls, the field being reduced to zero means that the eddy currents generating the magnetic field have been nullified.

3.3 DynaCAN based on inverse step responses

As established in Sec. 3.2, eddy currents in the MSR walls are a superposition of various modes of exponentially decaying modes with overlapping spatial patterns. To find effective DynaCAN waveforms for cancelling these transients, we characterize the behavior of the system of eddy currents. If the transient response of the cancellation pulse can be calculated from the applied pulse waveform, various computational methods can then be used to find DynaCAN pulses that nullify the harmful transients generated by ULF-MRI pulses.

The generated magnetic-field transient of any coil can be expressed as a function of the pulse driven into the coil by characterizing the transient response with what is called the inverse step response (ISR) of the coil [9]. Since the inverse step response of a coil is relatively easy to determine with simple measurements, this allows us easy access to a mathematical model for magnetic-field transients generated by pulses in a particular coil. In this section, we go through the derivation of the inverse step response expression for the magnetic-field transient, as established by Zevenhoven *et al.* in Ref. [9]. Furthermore, we use the expression to demonstrate a fundamental property of eddy-current transients that makes DynaCAN possible: that eddy-current modes with different time constants respond differently to pulse features at different time scales.

The various possible eddy-current patterns in the MSR walls are coupled with the MRI coils and other eddy-current patterns via mutual inductances [10]. Describing the mutual inductances between different eddy-current modes with matrix \mathbf{M} and the mutual inductances between modes and a particular coil with vector \mathbf{m} , the dynamics of the system can be described as

$$\mathbf{M} \frac{d\mathbf{j}}{dt} + \mathbf{R}\mathbf{j} + \mathbf{m} \frac{dI}{dt} = 0 , \quad (15)$$

where the vector $\mathbf{j}(t)$ describes the amplitude of each eddy-current mode, the diagonal matrix \mathbf{R} contains the respective resistances of the current modes and $I(t)$ is the pulse driven to the coil of interest [9]. This expression can be used to describe the eddy-current dynamics of any coil in the system, with each coil having its own unique \mathbf{M} and \mathbf{m} .

By applying a Fourier transform to the above equation, the frequency-domain solution for $\mathbf{j}(t)$ becomes

$$\hat{\mathbf{j}}(\omega) = -(i\omega\mathbf{M} + \mathbf{R})^{-1}\mathbf{m} \times i\omega\hat{I}(\omega) , \quad (16)$$

where $\hat{\cdot}$ denotes the Fourier transform. Next, we show that $(i\omega\mathbf{M} + \mathbf{R})^{-1}\mathbf{m}$ is the frequency-domain eddy-current response induced by a particular type of pulse, the inverse step pulse. By using the inverse step pulse $I(t) = 1 - \theta(t)$ where $\theta(t)$ is the Heaviside step function, we get $dI/dt = -\delta(t)$, with $\delta(t)$ being the Dirac delta function. Inserting this pulse into Eq. (15) and applying a Fourier transform, and noting that $\hat{\delta} = 1$, we get

$$\hat{\mathbf{j}}_{\text{ISR}}(\omega) = (i\omega\mathbf{M} + \mathbf{R})^{-1}\mathbf{m} , \quad (17)$$

where $\hat{\mathbf{j}}_{\text{ISR}}(\omega)$ is the frequency-domain expression for how the currents in the eddy-current modes respond to the inverse step pulse. By inserting Eq. (17) into Eq. (16) and applying an inverse Fourier transform, we get

$$\mathbf{j}(t) = - \left(\mathbf{j}_{\text{ISR}} \otimes \frac{dI}{dt} \right) (t) , \quad (18)$$

where the inverse Fourier transform has turned the product $\hat{\mathbf{j}}_{\text{ISR}}(\omega) \times i\omega\hat{I}$ into a convolution.

As established in the previous section, the total magnetic-field transient at any point in space is a linear combination of the individual fields generated by the individual modes. Therefore, a component of the magnetic-field transient at a particular point in space can be calculated as $B(t) = \mathbf{q}^* \mathbf{j}(t)$, where \mathbf{q} is a vector describing the magnetic coupling of the individual modes to a magnetic field component at particular point in space, and $*$ denotes the conjugate transpose. Multiplying Eq. (18) with \mathbf{q}^* , the equation becomes

$$B(t) = - \left(B_{\text{ISR}} \circledast \frac{dI}{dt} \right) (t) , \quad (19)$$

where the inverse step response $B_{\text{ISR}}(t)$ is the magnetic-field transient induced by an inverse step pulse driven in a coil of interest, with every coil having its own $B_{\text{ISR}}(t)$. This result shows that, if the ISR of a coil is known, the magnetic-field transient produced by any arbitrary pulse $I(t)$ driven in the corresponding coil can be calculated for any point in time. This is convenient, as the ISR of a coil can be determined with relatively simple measurements, whereas properties such as mutual inductances can be very difficult to measure. Because of this, Eq. (19) is much easier to use for calculating transient responses than, for example, Eq. (14).

Equation (19) can then be used to great effect in dynamical cancellation. In a system that has a known $B_{\text{ISR}}(t)$ for the cancellation coil, a cancellation pulse $I_c(t)$ can be found for a harmful transient by using Eq. (19) to find a waveform that produces the opposite transient. Defining the harmful transient as $B_{\text{tr}}(t)$ and the transient generated by the DynaCAN pulse $I_c(t)$ as $B_c(t)$, a suitable waveform can be found by minimizing the total transient

$$|B_{\text{tot}}(t)| = |B_{\text{tr}}(t) + B_c(t)| = \left| B_{\text{tr}}(t) - \left(B_{\text{ISR}} \circledast \frac{dI_c}{dt} \right) (t) \right| \quad (20)$$

over a period of time relevant to measurements, so that $|B_{\text{tot}}(t)|$ is as small as possible for all t in the chosen time interval. This can be accomplished in practice by, for example, minimizing the root-mean-square of $|B_{\text{tot}}(t)|$ in a time interval that starts from the end of the pulse sequence and lasts until the end of measurements.

If both $B_{\text{tr}}(t)$ and $B_{\text{ISR}}(t)$ are known, various mathematical methods can be applied to use Eq. (20) in order to find a cancellation pulse $I_c(t)$ that minimizes $|B_{\text{tot}}(t)|$. However, minimizing $|B_{\text{tot}}(t)|$ without constraining the shape of the pulse waveform can result in a waveform that is unusable in practice. For example, the maximum amplitude of the waveform can be larger than what the electronics used to drive the pulse into the cancellation coil can produce. Because of this, what is minimized in DynaCAN waveform generation is a penalty function that includes penalty based on pulse shape in addition to penalty based on $|B_{\text{tot}}(t)|$. Detailed discussion on practical implementation of this approach to DynaCAN pulse optimization can be found in Sec. 4.

In ideal cancellation, $|B_{\text{tot}}(t)|$ is reduced to zero in the relevant time interval, meaning $B_c(t)$ is the exact opposite of the harmful transient. This in turn means the eddy currents induced by the cancellation pulse are the exact opposite of those

induced by the rest of the ULF-MRI sequence, resulting in zero eddy currents induced by the pulse sequence as a whole.

While it is not readily apparent from Eq. (19), this model can be used to demonstrate a key property of transient behavior that makes DynaCAN possible: that eddy-current modes with different time constants respond differently to pulse features with different time scales. Fig. 4 demonstrates how sine pulses with different frequencies couple with different eddy-current modes by showing pulses and their transient responses calculated with Eq. (19) using an example ISR. It can be seen from the calculated transient responses that sine pulses with low frequency couple stronger to eddy-current modes with large time constants, whereas high-frequency pulses couple primarily to eddy-current modes with small time constants.

Because the convolution in Eq. (19) is a linear operation, a pulse that is a linear combination of different sine waves has a transient response that is a linear combination of the responses of the individual sine components. The individual components can, through their features at different time scales, cancel different modes of the harmful transients. Therefore, a cancellation pulse that incorporates multiple sine components can nullify a wide range of eddy-current modes, which would be impossible to accomplish with a pulse operating at only one time scale. This makes it possible to cancel highly complex transients by carefully manipulating the different eddy-current modes using components with varying frequencies in the cancellation pulse. This is the principle which forms the basis of DynaCAN, though a DynaCAN waveform need not necessarily be a linear combination of specific component waveforms, let alone sine components in particular. Any kind of pulse waveform with features at different time scales can be used as a DynaCAN waveform, regardless of how it is constructed.

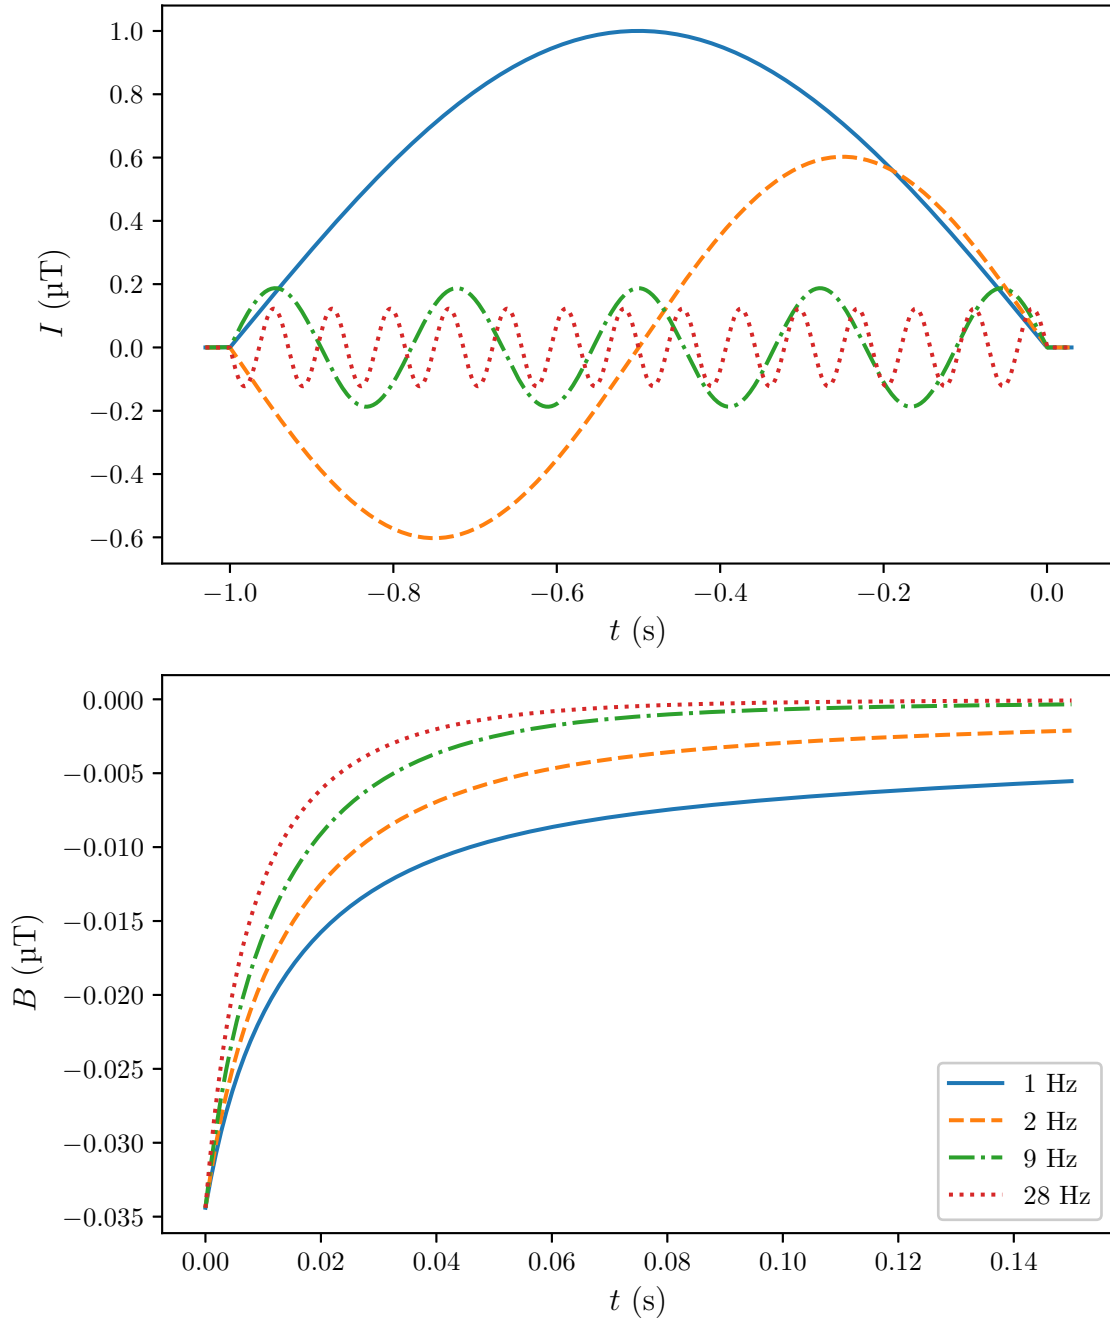


Figure 4: Sine pulses of different frequencies and their corresponding calculated magnetic-field transients. The pulse amplitudes have been scaled to produce a transient of the same magnitude at $t = 0$. As the frequency of the pulse increases, eddy-current modes with short time constants become increasingly significant compared to modes with long time constants.

4 Methods

Implementing DynaCAN for cancelling a given eddy-current transient in an MSR requires an algorithm for finding a suitable cancellation pulse to null the transient. In this section, I present an algorithm for finding effective DynaCAN waveforms, along with an analysis on how sources of error result in residual transients, and a description of the setup used for the measurements displayed in Sec. 5.

The algorithm for generating DynaCAN pulse waveforms for cancelling a known transient is a three-step process. First, the ISR of the cancellation coil is determined. Second, the ISR is used with Eq. (20) in order to find a first-iteration DynaCAN pulse to nullify the known transient. Third, the remaining residual transient after cancellation is measured, and a new pulse component is generated to nullify the residual transient. The third step is then iteratively repeated until suitable cancellation efficacy has been reached, or the efficacy can no longer be improved.

4.1 DynaCAN pulse waveform generation

4.1.1 Characterizing the eddy-current system

The inverse step response is used to calculate the transients generated by pulses driven in the cancellation coil by using Eq. (19). Accurately determining the ISR is necessary to ensure that the pulse waveforms found by the algorithm are effective in cancelling real transients. Differences between the ISR used for computational cancellation pulse optimization and the real ISR of the system result in the candidate DynaCAN pulses having a different transient response than predicted from Eq. (19). This means that DynaCAN pulse waveforms that are optimized to perfectly cancel harmful transients according to the mathematical model will leave residual transients when applied in real measurements. The more accurately the ISR has been determined, the smaller the residual fields will be. While residual fields can be reduced by iterating the pulse optimization algorithm as detailed later in Sec. 4.1.3, smaller residuals mean fewer iterations are necessary to achieve high cancellation efficacy. Furthermore, inaccuracies in the ISR can limit the maximal cancellation efficacy that can be reached by iterating the pulse optimization algorithm.

Since accurately determining the ISR is important for finding effective DynaCAN waveforms, it would be convenient if the ISR could be measured directly. In theory, this can be done by simply pulsing the cancellation coil with an inverse step pulse and measuring the induced magnetic-field transient. Unfortunately, this is not feasible if the practical electronics are unable to produce pulses that would be close enough to an inverse step pulse for our purposes.

When measuring the ISR itself is not an option, a less direct method must be used. The transient response from the DynaCAN pulse, $B_c(t) = -(B_{\text{ISR}} \otimes dI_c/dt)(t)$, is determined by the cancellation pulse $I_c(t)$ and the ISR of the cancellation coil, $B_{\text{ISR}}(t)$. Therefore, if we measure the response of a known pulse, the only unknown left is the ISR. This means we can numerically estimate the ISR by minimizing the difference between the measured transient $B_{\text{meas}}(t)$ and the transient calculated

using a candidate ISR. In this work, the difference to be minimized was expressed as the root-mean-square difference between the measured and calculated transients over a time interval $[T_1, T_2]$,

$$\sqrt{\frac{1}{T_2 - T_1} \int_{T_1}^{T_2} \left| B_{\text{meas}}(t) + f_{\mathbf{x}}(t) \otimes \frac{dI}{dt}(t) \right|^2 dt}, \quad (21)$$

where the time interval contains as much reliable transient measurement data as possible, and $f_{\mathbf{x}}(t)$ is the candidate ISR expressed as a linear combination of basis functions characterized by a set of parameters \mathbf{x} . As each basis function can be thought of as the magnetic-field transient generated by one exponentially decaying eddy-current mode with its own time constant, it makes sense to define our function basis as

$$f_{\mathbf{x}}(t) = f_{\mathbf{a},\mathbf{b}}(t) = \begin{cases} 0 & \text{when } t < 0 \\ \sum_i a_i e^{-b_i t} & \text{when } t \geq 0, \end{cases} \quad (22)$$

where a_i defines how strongly mode i appears, and b_i is its time constant. The number of modes should be chosen such that the important eddy-current modes are included. Figure 5 shows an example ISR that is a linear combination of exponential functions with different time constants, along with the individual exponential components it consists of.

The numerical method used in this work to optimize parameters \mathbf{a} and \mathbf{b} to minimize the expression in Eq. (21) was the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [20] as implemented in the Python package `scipy.optimize`. The RMS error was calculated over a time interval that began as soon after the pulse ramp down as possible and extended until the end of the measurement. Due to acquired measurement data being unreliable for roughly 5 ms after pulse ramp down, T_1 was set to 5 ms after ramp down.

4.1.2 Finding a cancellation pulse waveform

Once the cancellation coil ISR is known, the next step is to measure the harmful transient $B_{\text{tr}}(t)$ and apply Eq. (20) to find an effective DynaCAN waveform $I_c(t)$ to cancel $B_{\text{tr}}(t)$. Since the ISR of the DynaCAN coil can be used to calculate the transient responses of candidate cancellation pulses, a DynaCAN waveform for a known transient can be optimized entirely in software.

In a linear system, the total transient generated by a pulse sequence is the sum of the transients individually generated by the pulses. In this work, instead of using a full ULF-MRI pulse sequence, only the polarizing pulse was studied, as using a single pulse made measurements easier to perform. Since the polarizing pulse is much stronger than other ULF-MRI pulses, the transients generated by it are correspondingly larger, making the polarizing pulse a natural first choice for transient cancellation.

Assuming linearity, the polarizing pulse induces a harmful magnetic-field transient $B_{\text{tr}}(t)$ and the cancellation waveform induces its own transient $B_c(t)$, resulting in total




Figure 5: An example ISR and the four exponential components with different time constants it consists of.

field $B_{\text{tot}}(t) = B_{\text{tr}}(t) + B_c(t)$. The task is to find a DynaCAN waveform that opposes the harmful transient and brings the total field to zero. In perfect cancellation, this means $B_c(t) = -B_{\text{tr}}(t)$, where the waveform couples with the eddy-current modes in the walls in a way that perfectly nullifies the transients induced by the polarizing pulse.

Finding a cancellation pulse for a given transient is straightforward in principle. By minimizing a penalty function that depends on total remaining transient calculated with Eq. (20), we can find a waveform that minimizes the total generated transient. To ensure that the maximum amplitude of the pulse stays within reasonable limits, the penalty function can also depend on the maximum amplitude of the cancellation waveform. In this work, the penalty function $f_p\{I_c(t)\}$ was defined as a weighted sum of the root-mean-square total transient and the square of the cancellation pulse maximum amplitude

$$f_p\{I_c(t)\} = c_B \sqrt{\frac{1}{t_1 - t_0} \int_{t_0}^{t_1} |B_{\text{tot}}(t)|^2 dt} + c_I \left(\max_t I_c(t) \right)^2, \quad (23)$$

where $|B_{\text{tot}}(t)|$ can be found in Eq. (20), and c_B , c_I are weighing coefficients for the two types of penalty. The time interval from t_0 to t_1 is where transients would affect ULF-MRI measurements. In this work, t_0 was 15 ms and t_1 200 ms, where the

polarizing pulse ramp down ended at $t = 5$ ms. While it would have been preferable to have $t_0 = 5$ ms, the measurement data acquired was unreliable for roughly 10 ms after polarizing pulse ramp down, making a small delay necessary.

In order to minimize the penalty function with respect to $I_c(t)$, we parametrize $I_c(t)$ as a linear combination of a suitable set of basis functions. Since the cancellation pulse must be continuous and zero at boundaries, a sine basis is a natural choice, although it is by no means the only viable option. A DynaCAN waveform expressed as a linear combination of n sine functions bounded between t_{c0} and t_{c1} can be written as

$$I_a(t) = \begin{cases} 0 & \text{when } t < t_{c0} \text{ or } t > t_{c1} \\ \sum_n a_n \sin\left(n\pi \frac{t - t_{c0}}{t_{c1} - t_{c0}}\right) & \text{when } t_{c0} \leq t \leq t_{c1} . \end{cases} \quad (24)$$

With this parametrization, one can find an \mathbf{a} that minimizes Eq. (23) in order to find a cancellation pulse. However, finding a suitable waveform with numerical optimization is not without its difficulties. Minimizing the total transient with an iterative optimization method such as BFGS can result in the optimization ending in a local minimum with suboptimal cancellation efficacy even if better waveforms exist. For example, the more sine basis functions were included in waveform optimization, the more likely it was for the optimization to end prematurely in a local minimum. In order to find efficient DynaCAN waveforms without restricting the algorithm to using only a small number of basis functions, the waveform optimization was performed in steps with increasing numbers of sine basis functions. First, pulse optimization was performed with only a small amount of sine bases. After an initial round using only a few basis functions, successive rounds of optimization were performed using the results of the previous iterations as their starting points, with each round using more basis functions than the previous one. This process was repeated until a specified number of basis functions was reached. An example of a DynaCAN waveform produced with BFGS minimization of Eq. (23), along with the sine components it consists of, is shown in Fig. 6.

4.1.3 Iterating residual transient cancellation

Even with great care taken when performing the transient measurements and the DynaCAN pulse optimization, the waveform obtained in the previous step of the algorithm may not reach a sufficient cancellation efficacy such as 0.999. Even if the calculated efficacy of the found DynaCAN pulse was as high as 0.999, various sources of error can cause the actual transient cancellation efficacy to be lower. Errors in the ISR estimate cause the calculated transient response of the DynaCAN coil to differ from its real transient response. Errors in the measurement of the transients induced by the polarizing pulse directly result in imperfect cancellation. Furthermore, the assumption that eddy currents in the MSR are a linear system might not be entirely accurate.

The cancellation efficacy can be improved by iterating the pulse waveform optimization on the residual transient left after initial cancellation. By measuring the




Figure 6: An example cancellation pulse and the individual sine components it consists of. For clarity, only the first five sine components are included, one of which is close to zero.

remaining field and treating it as a new transient to be cancelled, we can find a new component to the DynaCAN waveform that cancels the residual field. Assuming the eddy currents are a linear system, we can simply take the sum of the original DynaCAN pulse and the residual cancellation pulse to create a new waveform that aims to cancel also the leftover transient. If efficacy is still not sufficient, this process of measuring the remaining transient and adding a new pulse component to cancel it can be repeated until the desired efficacy is reached or until the sources of error prevent further increase of cancellation efficacy. An example waveform after one iteration of residual nullification can be found in Fig. 7.

Adding multiple residual cancellation pulses to the original DynaCAN pulse runs the risk of each iteration increasing the pulse amplitude, eventually resulting in amplitudes higher than permissible. To prevent this from happening, the residual removal optimization algorithm does not impose amplitude penalties based on the amplitude of the residual nullification pulse by itself. Instead, in the penalty function shown in Eq. (23), the total cancellation pulse was used as $I_c(t)$ for calculating amplitude penalty

However, it is not always necessary to completely prevent residual cancellation components from increasing the total pulse amplitude. As long as care is taken to

ensure the pulse amplitude does not increase beyond acceptable limits, permitting some increase in signal amplitude can result in more efficient transient removal. The pulses shown in Fig. 7 are an example of this. The best waveform found by the first iteration of DynaCAN pulse optimization had a maximum amplitude of roughly 4 A, whereas the electronics driving the pulse into the cancellation coil were capable of generating amplitudes up to 10 A. Because of this, the residual cancellation component was allowed to increase the total pulse amplitude, as the new maximum amplitude of 7.5 A was still well within the limits of what the pulse-generating electronics could produce.

4.2 Error analysis

Measurement errors at various stages of the DynaCAN pulse optimization process can cause transient cancellation to work less efficiently than the calculations would indicate. Errors in the measurements of $B_c(t)$ for determining the ISR results in an inaccurate $B_{\text{ISR}}(t)$, which will cause error in all attempts to generate DynaCAN pulses with said ISR. Measurement errors also occur when measuring the harmful

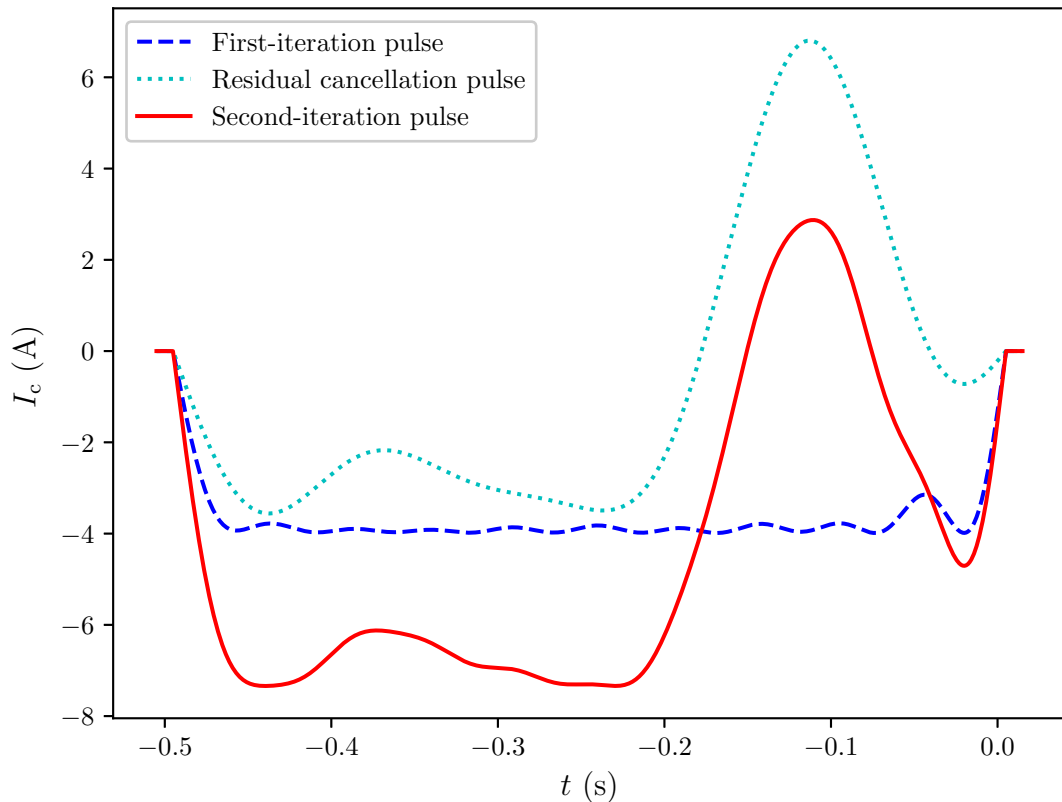


Figure 7: An example second-iteration DynaCAN pulse, consisting of a first-iteration cancellation pulse and a single residual cancellation pulse. Because the amplitude of the first-iteration pulse was low, the addition of a residual cancellation pulse was allowed to significantly increase the maximum amplitude of the pulse as a whole.

transient to be cancelled, $B_{\text{tr}}(t)$. All of these errors result in residual transients.

The effect of measurement errors in ISR estimation can be found by examining Eq. (19), which describes the transient generated by a pulse given the pulse waveform and ISR of the pulsed coil. In ISR estimation, the transient generated by a known pulse is measured, and Eq. (19) is used to determine the ISR of the coil. In case of an erroneous measurement $B_{\text{meas}}(t) = B_{\text{real}}(t) + \delta B(t)$, where $B_{\text{real}}(t)$ is the real transient and $\delta B(t)$ is an arbitrary measurement error, Eq. (19) takes the form

$$B_{\text{real}}(t) + \delta B(t) = - \left(B_{\text{ISR}}^{\text{est}} \otimes \frac{dI}{dt} \right) (t) , \quad (25)$$

where $B_{\text{ISR}}^{\text{est}}(t)$ is the estimated ISR warped by the measurement error.

Expressing the ISR estimate as $B_{\text{ISR}}^{\text{est}}(t) = B_{\text{ISR}}^{\text{real}}(t) + \delta B_{\text{ISR}}(t)$, the equation above can be written as

$$B_{\text{real}}(t) + \delta B(t) = - \left(B_{\text{ISR}}^{\text{real}} \otimes \frac{dI}{dt} \right) (t) - \left(\delta B_{\text{ISR}} \otimes \frac{dI}{dt} \right) (t) , \quad (26)$$

where $B_{\text{real}}(t) = - \left(B_{\text{ISR}}^{\text{real}} \otimes dI/dt \right) (t)$ and $\delta B(t) = - \left(\delta B_{\text{ISR}} \otimes dI/dt \right) (t)$. This means that a measurement error $\delta B(t)$ in ISR estimation results in a corresponding error $\delta B_{\text{ISR}}(t)$ in the ISR estimate, where the shape of $\delta B_{\text{ISR}}(t)$ depends on the measurement error $\delta B(t)$.

DynaCAN waveforms for cancelling a harmful transient $B_{\text{tr}}(t)$ are found by solving for a cancellation pulse waveform $I_c(t)$ that produces the opposite transient so that $B_c(t) + B_{\text{tr}}(t) = 0$. With an erroneous ISR estimate, using Eq. (19) to find a DynaCAN pulse $I_c(t)$ that results in zero total transient gives us

$$\begin{aligned} B_{\text{tot}}^{\text{est}}(t) &= B_c^{\text{est}}(t) + B_{\text{tr}}(t) \\ &= - \left(B_{\text{ISR}}^{\text{real}} \otimes \frac{dI_c}{dt} \right) (t) - \left(\delta B_{\text{ISR}} \otimes \frac{dI_c}{dt} \right) (t) + B_{\text{tr}}(t) = 0 , \end{aligned} \quad (27)$$

where $B_{\text{tot}}^{\text{est}}(t)$ is the estimated total transient and $B_c^{\text{est}}(t)$ is the estimated cancellation coil transient response calculated with Eq. (19). The above expression can be written as

$$B_{\text{tot}}^{\text{real}}(t) = B_c^{\text{real}}(t) + B_{\text{tr}}(t) = \left(\delta B_{\text{ISR}} \otimes \frac{dI_c}{dt} \right) (t) , \quad (28)$$

where $B_{\text{tot}}^{\text{real}}(t)$ is the real total transient and $B_c^{\text{real}}(t)$ is the real transient produced by the candidate cancellation pulse $I_c(t)$. This means that instead of the pulse optimization finding a DynaCAN waveform $I_c(t)$ that results in $B_c(t) + B_{\text{tr}}(t) = 0$, the error in the ISR causes the optimization algorithm to find a waveform that leaves a residual field $\left(\delta B_{\text{ISR}} \otimes dI_c/dt \right) (t)$.

The analysis above examined the propagation of an arbitrary measurement error $\delta B(t)$ in ISR estimation, and the result was a general expression for the residual transient caused by the measurement error. However, this general expression is difficult to use to calculate residual transients given a particular measurement error. More specific results can be found by examining the effects of systematic error,

where the measurement error is a function of the real magnetic-field transient. The two most common types of systematic error are constant error and relative error. In constant error, the measurement is shifted by a constant δB_{const} , resulting in $B_{\text{meas}}(t) = B_{\text{real}}(t) + \delta B_{\text{const}}$. Measurements with relative error will result in $B_{\text{meas}}(t) = (1 + \delta b)B_{\text{real}}(t)$, where δb characterizes a shift in the measured magnetic field that is proportional to the magnitude of the real magnetic field.

Since the magnetic-field transients consist of exponentially decaying components, waiting long enough should cause the fields generated by any pulse sequence to die down. This means that with suitably long measurements, what remains after the measured magnetic field has stopped changing can be thought of as constant error. By shifting the measurement results so that the transients die down to zero, constant error can be dealt with before the measurements results are used for estimating ISRs or optimizing cancellation pulses. However, one must be cautious to make sure the transients have truly died down—some transient components can have very large time constants and thus decay very slowly. Mistaking transient components for measurement error results in constant error appearing in the data.

Unlike with constant error, there is no simple method for dealing with relative error by post-processing the measurement data. Because of this, relative error in measurements used for ISR estimation will propagate throughout the pulse optimization process, causing residual transients to appear. With erroneous measurement data $B_{\text{meas}}(t) = (1 + \delta b)B_{\text{real}}(t)$, we can find the effect this type of error has on the ISR estimate by multiplying Eq. (19) with $(1 + \delta b)$, resulting in

$$(1 + \delta b)B_{\text{real}}(t) = -(1 + \delta b) \left(B_{\text{ISR}}^{\text{real}} \otimes \frac{dI}{dt} \right) (t) = - \left((1 + \delta b)B_{\text{ISR}}^{\text{real}} \otimes \frac{dI}{dt} \right) (t) , \quad (29)$$

which can be expressed as

$$B_{\text{meas}}(t) = - \left(B_{\text{ISR}}^{\text{est}} \otimes \frac{dI}{dt} \right) (t) , \quad (30)$$

where the ISR estimate is $B_{\text{ISR}}^{\text{est}}(t) = (1 + \delta b)B_{\text{ISR}}^{\text{real}}(t)$. This means that a consistent relative error δb causes the same relative error to appear in the ISR estimate.

When this incorrect ISR estimate is applied to find a cancellation pulse with Eq. (19), the equation appears as

$$B_{\text{tot}}^{\text{est}}(t) = B_{\text{c}}^{\text{est}} + B_{\text{tr}}(t) = - \left((1 + \delta b)B_{\text{ISR}}^{\text{real}} \otimes \frac{dI_{\text{c}}}{dt} \right) (t) + B_{\text{tr}}(t) = 0 . \quad (31)$$

Since $- \left((1 + \delta b)B_{\text{ISR}}^{\text{real}} \otimes dI_{\text{c}}/dt \right) (t) = -(1 + \delta b) \left(B_{\text{ISR}}^{\text{real}} \otimes dI_{\text{c}}/dt \right) (t) = (1 + \delta b)B_{\text{c}}^{\text{real}}(t)$, the above expression results in

$$B_{\text{tot}}^{\text{real}}(t) = B_{\text{c}}^{\text{real}}(t) + B_{\text{tr}}(t) = -\delta b B_{\text{c}}^{\text{real}}(t) . \quad (32)$$

This expression can be further refined by substituting $B_{\text{c}}^{\text{real}}(t) = B_{\text{tot}}^{\text{real}}(t) - B_{\text{tr}}(t)$ to $B_{\text{tot}}^{\text{real}}(t) = -\delta b B_{\text{c}}^{\text{real}}(t)$, resulting in $B_{\text{tot}}^{\text{real}}(t) = -\delta b (B_{\text{tot}}^{\text{real}}(t) - B_{\text{tr}}(t))$, which in turn can be simplified into

$$B_{\text{tot}}^{\text{real}}(t) = \frac{\delta b}{1 + \delta b} B_{\text{tr}}(t) . \quad (33)$$

This means relative error in ISR estimation causes residual transients directly proportional to the relative error δb and the harmful transient $B_{\text{tr}}(t)$. With small δb , the residual can be approximated as $\delta b B_{\text{tr}}(t)$.

Errors in the measurement of the harmful transient $B_{\text{tr}}(t)$ convert into residual transients in a rather straightforward way. If the measured harmful transient is $B_{\text{meas}}(t) = B_{\text{tr}}(t) + \delta B_{\text{tr}}(t)$, where $\delta B_{\text{tr}}(t)$ is a measurement error of any shape, with perfect DynaCAN waveform optimization the cancellation pulse will produce the opposite transient, meaning $B_{\text{c}}(t) = -B_{\text{meas}}(t) = -B_{\text{tr}}(t) - \delta B_{\text{tr}}(t)$. This results in the total field becoming $B_{\text{tr}}(t) + B_{\text{c}}(t) = B_{\text{tr}}(t) - B_{\text{tr}}(t) - \delta B_{\text{tr}}(t) = -\delta B_{\text{tr}}(t)$, which means that a measurement error $\delta B_{\text{tr}}(t)$ will result in a corresponding residual transient $-\delta B_{\text{tr}}(t)$.

4.3 Measurement setup

The measurements were performed in the Aalto University MEG–MRI laboratory using parts of the ULF-MRI setup designed a hybrid MEG–MRI device. Measurements were performed with a two-coil setup; one coil for generating a harmful transient, and one coil to use as the cancellation coil. The polarizing coil would have been a natural choice for the transient generating coil, but since it was not accessible at the time the measurements were performed, a substitute polarizing coil was used instead. This substitute coil was a small coil placed where the head of a patient would be located in an MEG–MRI scan. Since there was no coil designed specifically for DynaCAN in the coil setup, a shielding coil designed to be operated in series with the polarizing coil was chosen as the cancellation coil. This shielding coil was selected because it was roughly concentric with the substitute polarizing coil—the center of the substitute polarizing coil was approximately 4 cm below the center of the shielding coil. The average radius of the substitute polarizing coil was 7 cm, and that of the shielding coil was 30 cm. These coils are shown in Fig. 8.

The MSR the measurements were performed in suits the needs of both MEG and ULF MRI. Its walls were made of layers of both aluminium and μ metal for shielding against high- and low-frequency interference. The walls of the MSR consist of two layers, both made of slabs of aluminium and μ metal. The thickness of aluminium in the layers was 8 mm, and the thickness of μ metal was 1 mm. The layers were constructed so that the μ metal formed the inner surface of each layer [7].

The detector used for the measurements was a three-axis fluxgate magnetic field sensor (MAG03MC1000, Bartington Instruments Ltd., England). The detector was placed in the center of the substitute polarizing coil. That was taken as the location where transient cancellation matters most. The measured magnetic field component was the component perpendicular to the plane the two coils were located in, as this was the strongest component of the magnetic-field transients.

The pulses driven into the coils were defined and generated by Python-based MRI sequencing software and analog-to-digital conversion hardware. The pulses driven into the cancellation coil were fed with an AE Techron 2105 gradient amplifier (AE Techron Inc., USA), whereas the pulses driven into the polarizing coil were fed with a Kepco BOP72-6ML bipolar amplifier (Kepco Inc., USA).




Figure 8: The coils used for measurements in this work. The small coil in the middle is the substitute polarizing coil applied to generate a harmful transient. Nearly concentric with the polarizing coil is the shielding coil used as the cancellation coil.

5 Results

The ISR used for DynaCAN pulse optimization was determined from three sets of measurements, each set using a different pulse driven in the cancellation coil. These sets of measurements were performed by driving a trapezoid pulse into the DynaCAN coil and averaging over 200 measurements with each pulse. The first pulse had amplitude 4 A and ramp time 2 ms, the second had amplitude 6 A and ramp time 4 ms, and the third had amplitude 8 A and ramp time 6 ms. All three pulses were 500 ms long, with each measurement lasting 1 second after ramp down. After averaging the measurements, zero-level correction was performed by taking the mean of the last 2000 points of each measurement time series, corresponding to roughly 0.2 seconds of measurement data, and subtracting these values from the measurements.

Figure 9 shows the ISR estimate that was found by minimizing Eq. (21), using the three pulses described above. The individual exponential components the ISR estimate consists of are also shown. Figure 10 shows results from the measurements with the three ISR optimization pulses together with their corresponding calculated transients determined with the found ISR estimate. The calculated transients closely match the measured transients for all three pulses.

Transient cancellation measurements were performed with a trapezoid polarizing




Figure 9: The ISR estimate of the cancellation coil and the four individual exponential components with different time constants it consists of.




Figure 10: The measured transients generated by the three trapezoid pulses used for ISR optimization, along with their corresponding calculated transients. The transients calculated with the ISR estimate accurately reproduce the measured transients.

pulse of amplitude 10 A and ramp time 10 ms. The pulse lasted 510 ms, ending at $t = 5$ ms. The transient generated by this pulse was cancelled by optimizing a DynaCAN waveform to minimize the generated field from $t = 15$ ms onwards. While it would have been preferable to minimize the transient immediately after the polarizing pulse ended, the measurements were not reliable immediately after ramp down. Three iterations of DynaCAN waveform optimization were performed. The cancellation pulses lasted 500 ms, ending when the polarizing pulse finished ramping down. The polarizing pulse and the three iterations of DynaCAN waveforms can be found in Fig. 11.

The cancellation efficacy of a DynaCAN pulse can be quantified in many ways. Two approaches were used in this work. The first was to define it as relative maximum cancellation $1 - B_{\max}^{\text{cancel}} / B_{\max}^{\text{orig}}$, where B_{\max}^{cancel} is the maximum absolute value of the total magnetic-field transient after ramp down with the DynaCAN pulse applied, and B_{\max}^{orig} is the maximum absolute magnitude of the field after ramp down without DynaCAN. Using this metric, a cancellation efficacy of 0.95 would mean the magnitude of the cancelled transient at its highest is only 5% of the maximum magnitude of the uncanceled transient. The other approach to characterizing efficacy was to look at

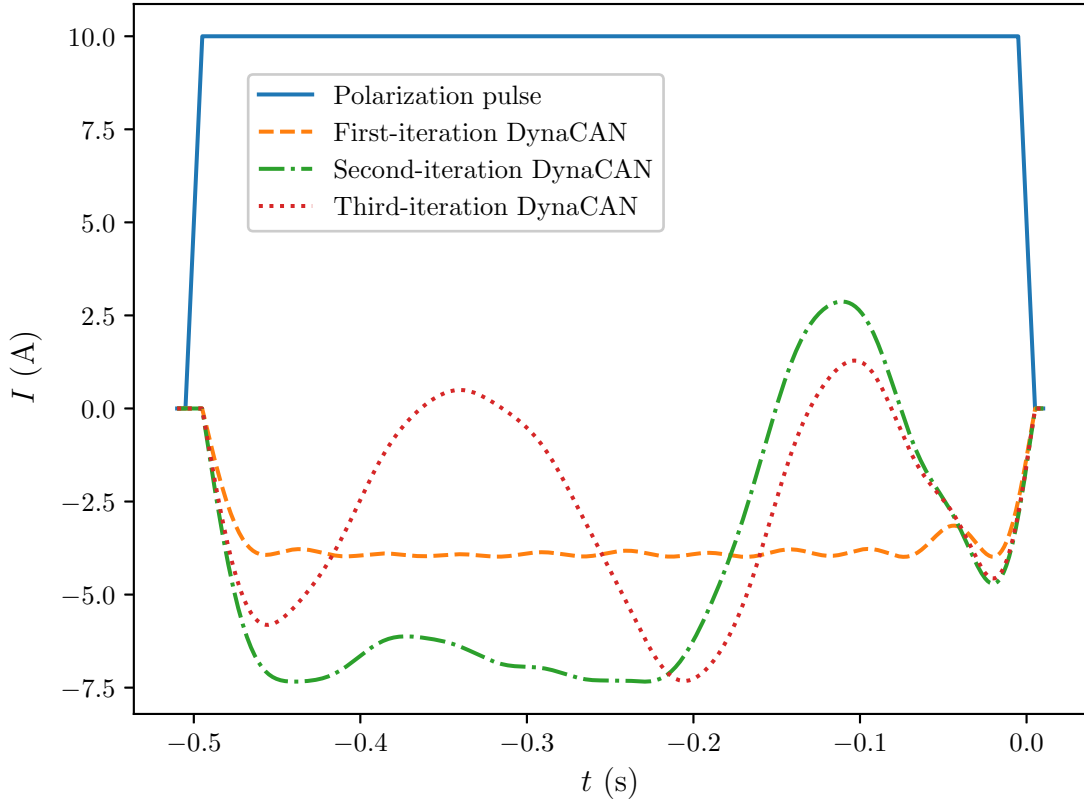


Figure 11: The polarizing pulse and the three iterations of DynaCAN waveforms used as cancellation pulses. The first-iteration cancellation pulse is similar in shape to the polarizing pulse, whereas the following iterations are more complex.

how long it took for the absolute value of the field to decay below a chosen threshold B_{thres} . This means finding t_d such that $|B_{\text{tot}}(t)| < B_{\text{thres}}$ for $t \geq t_d$. In this work, the threshold was chosen as 1% of the maximum measured absolute value of the uncanceled transient after ramp down, $B_{\text{max}}^{\text{orig}}/100$.

Because noise greatly distorts these metrics of cancellation efficacy, noise was removed by fitting multiexponential functions to the measurement data. While noise results in these exponential functions being imperfect representations of the real transients, their maximum values are expected to be closer to the real maxima than the maximum values of the raw measurement data. Estimating the point in time when the transient has decayed below a specific value also becomes easier when using multiexponential fits in place of raw measurement data.

The transient generated by the polarizing pulse alone and the transient generated by the polarization and first-iteration cancellation pulses together are shown in Fig. 12. The first-iteration DynaCAN waveform already significantly reduced the measured field magnitude. Its cancellation efficacy was 0.977, meaning the remaining transient was at worst 2.3% of the maximum value of the uncanceled transient.

The uncanceled transient has its maximum at the beginning of the data set at $t = 15$ ms. This transient becomes consistently less than 1% of its maximum value




Figure 12: The transients generated by the polarizing pulse without cancellation and with the first iteration DynaCAN waveform. The cancellation pulse significantly decreases the total field, but a noticeable residual transient still remains.

from $t_d = 429$ ms onwards. With the first-iteration DynaCAN pulse, the measured field becomes consistently smaller than 1% of the maximum uncanceled transient at $t_d = 127$ ms, roughly 300 ms earlier than without cancellation.

The second-iteration cancellation pulse was more effective than the first-iteration pulse. Figure 13 compares the remaining transients when the first- and second-iteration cancellation waveforms are applied. The overall shape of the transient response is very similar between the two cancellation pulses, but the second-iteration waveform successfully cancels a significant amount of the residual field from approximately $t = 20$ ms to $t = 200$ ms. However, a short-lived transient spike remains at the very beginning of the data set, where the remaining field is at its highest magnitude. The cancellation efficacy of the second-iteration pulse was 0.988, with the transient dropping below 1% of the maximum of the original transient almost immediately, at roughly $t_d = 15.3$ ms. If the spike at the beginning of the data set is ignored by skipping the first 2 ms of data, the cancellation efficacy becomes 0.994. This means that, from $t = 17$ ms onwards, the remaining transient is, at worst, 0.6% of the magnitude of the uncanceled field at $t = 17$ ms.

The transient responses of the polarizing pulse together with the second- and third-iteration DynaCAN waveforms are shown in Fig. 14. This figure also displays

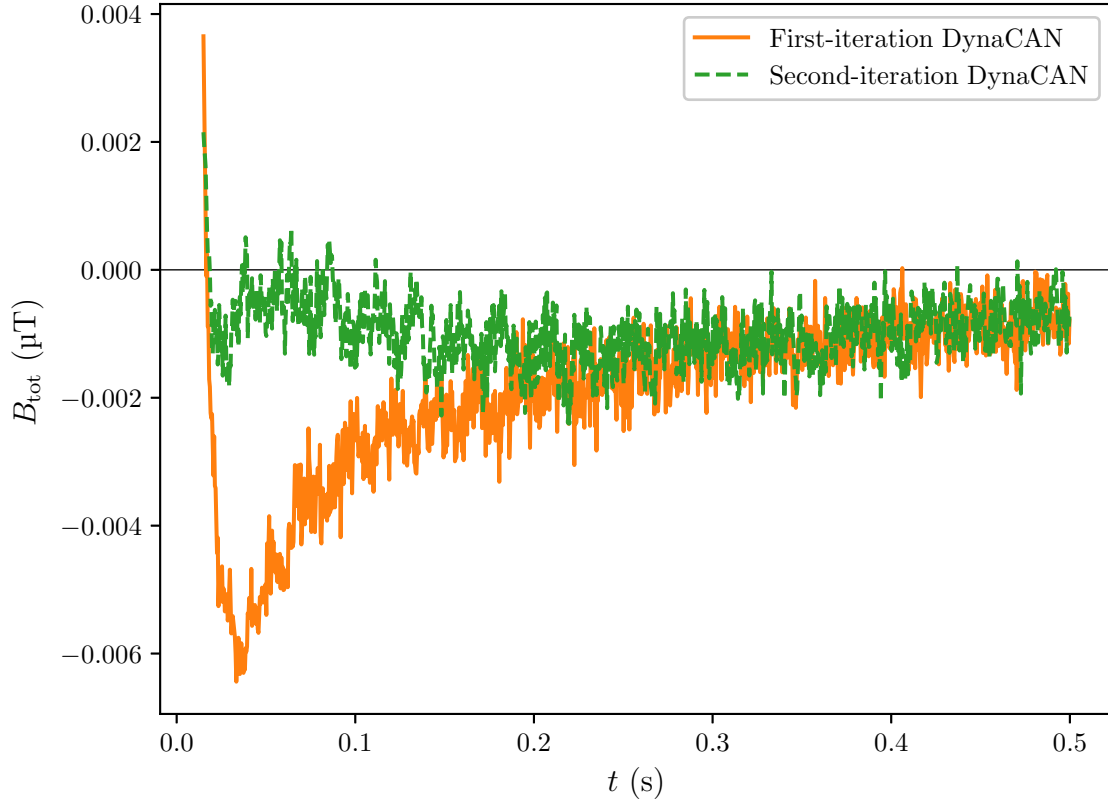


Figure 13: The transients generated by the polarizing pulse together with the first and second-iteration cancellation pulses. The second-iteration waveform reduces the early residual transient so that its magnitude remains small throughout the measurement interval.

the multiexponential fits to the measurement data that are used in efficacy calculation. According to calculations, the third-iteration pulse should have resulted in smaller residual transients than the second-iteration waveform from roughly $t = 20$ ms onwards. However, when the third-iteration cancellation pulse was applied and the transient response measured, the resulting residual field increased in magnitude compared to the results with second-iteration waveform until approximately $t = 135$ ms. From $t = 135$ ms onwards, the cancellation efficacy of the third-iteration pulse was superior to that of the second-iteration pulse.

The cancellation efficacy of the third-iteration pulse was 0.979, with the transient response dropping below 1% of the maximum uncanceled transient at approximately $t_d = 50$ ms. By those metrics, the third-iteration waveform is less effective than the second-iteration waveform. However, the goal of transient cancellation for ULF MRI purposes is to reduce transients enough to bring the magnetic field at the SQUIDS within the field range of the detectors. Whether the second- or third-iteration pulse better accomplishes this goal depends on how the field ranges of the SQUID sensors are situated. If, for example, the absolute value of the magnetic-field transient must be brought below 1 nT for measurements to be possible, the third-iteration pulse




Figure 14: The transients generated by the polarizing pulse together with the second and third-iteration cancellation pulses. The multiexponential fits to the measurement data that are used in efficacy calculation are also shown. The early residual transient was larger with the third-iteration pulse, but the field from $t \approx 135$ ms onwards is smaller than that of the second-iteration waveform.

accomplishes this earlier, at roughly $t = 200$ ms, whereas it takes until roughly $t = 400$ ms for the second-iteration pulse. If the SQUIDs can perform measurements with absolute transient magnitude 2 nT, then the second-iteration waveform is superior as it reduces the field magnitude below that threshold almost immediately, at approximately $t = 16$ ms, while the third-iteration waveform reaches low enough transient magnitude at roughly $t = 70$ ms.

6 Discussion

The measurement results demonstrate that dynamical cancellation with DynaCAN can significantly reduce transients induced in MSR walls by coil pulses inside the room. Notably, the second-iteration DynaCAN waveform reduced the magnetic-field transient below one hundredth of the original test transient starting at approximately 12 ms after the end of the polarizing pulse, whereas without cancellation it would have taken roughly 415 ms for the transient to decay to that extent.

However, the measurement results also deviated from the calculated transients acquired in pulse waveform optimization to varying degrees, especially at short time scales. Expected residual transients are compared with multiexponential fits to measured residual fields using the second- and third-iteration cancellation pulses in Fig. 15. The calculated transient response of the second-iteration DynaCAN pulse is fairly close to the multiexponential fit to the corresponding measurement data, with minor differences that have little effect on the cancellation efficacy of the pulse. However, the function fitted to the residual field measured when using the third-iteration DynaCAN waveform differs significantly from the expected transient,




Figure 15: The transients generated by the polarizing pulse together with the second and third-iteration cancellation pulses, along with the calculated transients found by the pulse optimization algorithm. For clarity, the measured transients are expressed as multiexponential functions fitted to the measurement data.

resulting in considerably less effective removal of transients than predicted.

These discrepancies between the calculated and measured transient responses could be the result of flaws in the ISR estimate. The ISR was estimated using three sets of measurements averaged over 200 measurements. Individual measurements lasted for one second after ramp down, which may have been too little time for eddy-current modes with the longest time constants to die down, resulting in inaccuracies in the ISR estimate, especially in the modes with long time constants. Furthermore, the acquired measurement data was unreliable for roughly 5 ms after pulse ramp down. These 5 ms of lost data could have resulted in eddy-current modes with very short time constants being mischaracterized or completely ignored in ISR estimation. Constant-frequency electronic noise visible in Fig. 14 can also be expected to distort the ISR estimate, especially modes with short time constants. As shown in Fig. 15, the differences between measured and calculated transients were at their largest in short time scales, suggesting that eddy-current modes with short time constants were indeed improperly represented in the ISR estimate. The above issues could be alleviated by performing longer measurements and by improving the measurement setup in order to reduce the delay between pulse ramp down and reliable data acquisition. Eliminating noise further by averaging over more than 200 measurements could also improve the accuracy of the ISR estimate.

It is also possible that refining the penalty function that is minimized to find an ISR estimate could result in better ISR estimates. The penalty function used in this work was a sum of RMS errors of transient responses calculated with the ISR estimate compared to measured transients. Other approaches to defining the penalty function may potentially yield better results. Furthermore, applying more varied pulses for the measurements used for ISR estimation could improve the estimate. In this work, only simple trapezoid pulses with fairly similar transient responses were used. It is possible that these trapezoid pulses failed to sufficiently excite some eddy-current modes that the DynaCAN waveforms did excite, resulting in those modes being ignored in the ISR estimate. Optimizing the ISR estimate to accurately reproduce a wide variety of different transient responses, generated by pulses that differ greatly from one another, would likely make ISR estimation more reliable.

In addition to the above measures, the accuracy of the ISR estimate could be further improved by performing measurements at multiple points inside the MSR. While the ISR of the cancellation coil will change depending on location, the ISR will always be a linear combination of the same eddy-current modes with the same time constants. This means that the different ISRs should still be linear combinations of the same set of exponential basis functions. Finding a set of time constants that can be used to construct ISR estimates in multiple locations makes it more likely that modes corresponding to those time constants are truly present in the real ISR of the coil.

Errors in the ISR estimate is not the only explanation for why the calculated transients could differ from the measured fields. The transient response of the room might not be perfectly linear. The electronics that drive the DynaCAN waveform into the cancellation coil might be unable to perfectly replicate the intended pulse, for example if the waveform has slopes steeper than what the electronics can generate.

While there is little that can be done about nonlinearity in the system, using a cancellation coil designed to couple strongly to the eddy currents in the MSR walls could help with pulse generation. A specifically designed cancellation coil allows for weaker pulses to achieve the same effects as waveforms with high amplitudes and slopes driven into coils not designed for dynamical cancellation. Using pulses with smaller amplitudes and slopes places less strain on electronics, making it less likely that the amplifier fails to generate the exact specified pulse. Using a coil designed for DynaCAN could also be expected to improve cancellation efficacy beyond what was shown in the measurements presented in this thesis. These measurements were performed using a shielding coil that was not designed for dynamical cancellation, which may have significantly limited how strongly the cancellation pulses coupled with various eddy-current modes.

The measurement setup is not the only part of the transient cancellation process that has room for improvement. The implementation of the cancellation pulse waveform optimization algorithm could also be refined further. While the basic principle of finding a DynaCAN waveform by minimizing the penalty function found in Eq. (23) is straightforward, making it work efficiently in practice can be challenging. For example, waveform optimization with BFGS using sine bases was likely to end in a local minimum that resulted in an ineffective cancellation pulse when more than 10 basis functions were used in optimization.

The implementation in this work decreased the likelihood of the optimization ending prematurely in local minima by performing multiple rounds of optimization with increasing numbers of sine bases in order to find an optimal pulse waveform to minimize Eq. (23). However, BFGS optimization of a sine basis might not be the most effective approach for finding efficient DynaCAN waveforms in the first place. A numerical optimization algorithm other than BFGS could be better suited for the task. Using a function basis other than a sine basis could make pulse generation more flexible, potentially improving cancellation efficacy.

Furthermore, the penalty function used for pulse optimization was rather simple, consisting only of a weighted sum of RMS remaining transient and maximum amplitude of the DynaCAN waveform squared. Refinements such as increasing the significance of early transients in penalty calculation could result in improved cancellation efficacy. It could also be worthwhile to explore alternatives to RMS transient penalty, such as penalty based on maximum transient amplitude. While the cancellation pulses presented in this work were effective, there is still considerable room for exploring alternative implementations of the DynaCAN waveform optimization algorithm.

7 Conclusions

Eddy-current transients induced by ULF-MRI pulses in magnetically shielded rooms limit imaging accuracy by preventing measurements until the transients have died down sufficiently. In order to get as accurate ULF-MRI images as possible, it is necessary to minimize the magnetic interference caused by these eddy currents. Dynamical cancellation with DynaCAN is a novel transient nullification technique that uses an individual DynaCAN pulse driven into a cancellation coil to minimize the eddy-current transients induced by the ULF-MRI pulse sequence. The pulse waveform makes use of features at different time scales to couple with the dynamical eddy-current system in a way that counteracts transients generated by the rest of the pulse sequence. Transients can be nullified with dynamical cancellation alone or by using dynamical cancellation together with other cancellation methods.

In this work, I presented an algorithm for finding efficient DynaCAN waveforms. First, the magnetic field inverse step response of the cancellation coil is determined. The inverse step response can then be used to calculate the transient responses of candidate DynaCAN pulses. This allows a measured harmful magnetic-field transient to be cancelled by using numerical optimization to find a DynaCAN pulse waveform that generates the opposite field. Once a suitable candidate cancellation pulse has been found, the pulse is included in the ULF-MRI sequence and the resulting field measured. If significant residual transients still remain, the cancellation pulse waveform optimization can be iterated on by finding a waveform that cancels the remaining transient. This residual nullification pulse can then be summed with the original waveform to create a new pulse that attempts to cancel both the original and the residual transient. This process of updating the DynaCAN pulse with residual cancellation waveforms can be repeated until more effective pulses can no longer be found.

Measurements were performed to demonstrate in practice the efficacy of dynamical cancellation, using DynaCAN waveforms generated with the pulse optimization algorithm. The transient to be cancelled was generated with a 500 ms long 10 A trapezoid pulse with a 10 ms ramp time driven into a polarizing coil. Three iterations of cancellation pulse optimization were performed, using waveforms that were 500 ms long linear combinations of sine waves that ended at the same time as the polarizing pulse. The resulting transients with each iteration of the algorithm were measured, measurements starting 10 ms after polarizing pulse ramp down due to technical limitations.

The first-iteration cancellation pulse reduced the transients in the measured interval down to 2.3% or less of the maximum measured magnitude of the transient without cancellation. The pulse reduced the field to under 1% of the maximum uncanceled field from 122 ms after ramp down onwards, whereas without cancellation it took 424 ms for the field to decay to 1% of its maximum magnitude. The second-iteration cancellation pulse was more efficient, with the transient reduced to below 1.2% of the maximum uncanceled field in the measurement interval. Notably, after 12 ms from ramp down, the remaining field was consistently below 0.6% of the maximum uncanceled transient. The third-iteration cancellation waveform was less

effective than the second-iteration waveform until roughly 130 ms after ramp down, after which it was more effective than the second-iteration pulse.

The measurements demonstrated that DynaCAN can be used for effective transient cancellation. The algorithm for finding DynaCAN waveforms was capable of finding cancellation pulses that greatly diminished the transients generated by the test pulse. However, the measurements also showed that there is still room for improvement in the implementation of dynamical cancellation. Both the measurement setup and the implementation of the pulse waveform optimization algorithm could be improved in numerous ways. With better cancellation coil design, better measurement equipment, and refinements to the implementation of the measurements and the waveform optimization algorithm, it can be expected that the efficacy of dynamical cancellation with DynaCAN could be increased well beyond what was demonstrated in the measurements presented in this work.

References

- [1] Liang, Z.-P., Lauterbur, P. C., *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*, IEEE Press Series in Biomedical Engineering, IEEE Press, Piscataway, New Jersey, USA (2000)
- [2] Brown, R. W., Cheng, Y.-C. N., Haacke, E. M., Thompson, M. R., Venkatesan, R., *Magnetic Resonance Imaging: Physical Principles and Sequence Design* (2nd ed.), John Wiley & Sons, Inc., Hoboken, New Jersey, USA (2014)
- [3] Kraus, R. Jr, Espy, M., Magnelind, P., Volegov, P, *Ultra-low field nuclear magnetic resonance: A new MRI regime*, Oxford University Press, Oxford, UK (2014)
- [4] Lee, S-K., Mößle, M., Myers, W., Kelso, N., Trabesinger, A. H., Pines, A., Clarke, J., “SQUID-detected MRI at 132 μ T with T1-weighted contrast established at 10 μ T-300 mT”, *Magnetic Resonance in Medicine*, **53**, 9-14 (2005)
- [5] Sternickel, K., Braginski, A. I., “Biomagnetism using SQUIDS: status and perspectives”, *Superconductor Science and Technology*, **19**, S160-71 (2006)
- [6] *Compilation from various authors*, edited by Hansen, P. C., Kringelbach, M., Salmelin, R., *MEG: An Introduction to Methods*, Oxford University Press, London, UK (2010)
- [7] Zevenhoven, K. C. J., *Solving transient problems in ultra-low-field MRI*, Master’s thesis, Aalto University, Helsinki, Finland (2011)
- [8] Zevenhoven, K. C. J., *Four-dimensional dynamic shielding: A novel method for reducing eddy-current transients in ultra-low-field MRI*, Special assignment, Aalto University, Helsinki, Finland (2011)
- [9] Zevenhoven, K. C. J., Dong, h., Ilmoniemi, R. J., Clarke, J., “Dynamical cancellation of pulse-induced transients in a metallic shielded room for ultra-low-field magnetic resonance imaging”, *Applied Physics Letters*, **106**, 034101 (2015)
- [10] Zevenhoven, K. C. J., Busch, S., Hatridge, M., Öisjöen, F., Ilmoniemi, R. J., Clarke, J., “Conductive shield for ultra-low-field magnetic resonance imaging: Theory and measurements of eddy currents”, *Journal of Applied Physics*, **115**, 103902 (2014)
- [11] Clarke, J., Braginski, A. I., *The SQUID Handbook*, Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany (2004)
- [12] Clarke, J., Hatridge, M, Mößle, “SQUID-Detected Magnetic Resonance Imaging in Microtesla Fields”, *Annual Review of Biomedical Engineering*, 9:389-413 (2007)

- [13] Peiniger, M., Piel, H., “A Superconducting Nb₃Sn Coated Multicell Accelerating Cavity”, IEEE Transactions on Nuclear Science, **32** (5), 3610–3612 (1985)
- [14] Khomutov, S. Y., Mandrikova, O. V., Budilova, E. A., Arora, K., Manjula, L., “Noise in raw data from magnetic observatories”, Geoscientific Instrumentation Methods and Data Systems, **6**, 329-343 (2017)
- [15] Thébault, E. *et al.*, “Intergenerational Geomagnetic Reference Field: the 12th generation”, Earth, Planets and Space, **67**:79 (2015)
- [16] Ruder, W., “New magnetic materials”, Proceedings of the IRE, **30**, 435-440 (1942)
- [17] Zimmerman, J. E., “SQUID instruments and shielding for low-level magnetic measurements”, Journal of Applied Physics, **48**, 702-710 (1977)
- [18] Cohen, D., “Large-volume conventional magnetic shields”, Revue de Physique Appliquée, **5**, 53-58 (1970)
- [19] Baum, E., Bork, J., “Systematic design of magnetic shields”, Journal of Magnetism and Magnetic Materials, **101**, 69-74 (1991)
- [20] Fletcher, R, *Practical Methods of Optimization* (2nd ed.), John Wiley & Sons, New York (1987)