Aalto University
School of Electrical Engineering
Master's Programme in Life Science Technologies

Heikki Sinisalo

Waveform control for transcranial magnetic stimulation

Master's Thesis

Espoo, October 22, 2018

Supervisor: Prof. Risto Ilmoniemi

Thesis advisor: Jaakko Nieminen, D.Sc. (Tech.)

AALTO UNIVERSITY
School of Electrical Engineering
Master's Programme in Life Science Technologies

ABSTRACT OF THE
MASTER´S THESIS

Author: Heikki Sinisalo

Title of the thesis: Waveform control for transcranial magnetic stimulation

Number of pages: 54 + 7

Date: October 22, 2018

Major or Minor: Biosensing and bioelectronics

Supervisor: Prof. Risto Ilmoniemi

Thesis advisor: Jaakko Nieminen, D.Sc. (Tech.)

Transcranial magnetic stimulation (TMS) is a non-invasive tool for stimulating the brain via induced electric fields generated by driving a strong current pulse through a stimulation coil. The TMS group at Aalto University has developed a multi-locus TMS (mTMS) device which exploits linear superposition of electric fields. This allows for the rotation or movement of the stimulated area to be controlled electrically by having a set of overlapping coils instead of physically moving a single coil. The device utilizes an H-bridge topology which can be used to carefully control the current through a coil by changing the path of current through the stimulator circuitry with insulated gate bipolar transistors (IGBTs), and the aim of this Thesis was to develop a method for controlling these currents in such a way that a given reference current pulse (i.e., waveform) could be approximated.

For better IGBT control, new controller cards had recently been designed, and one objective of this Thesis was to test them. Additionally, in preparation for more coils to be added to the system, a coil identification system utilizing digital temperature sensors and a microcontroller was prototyped. The bulk of this Thesis, however, consists of the algorithm that was developed for IGBT control. The idea is to calculate a timing sequence for the IGBTs in such a way that a waveform generated by a lower initial voltage reference pulse is effectively simulated by periodically driving current to the coil from a high-voltage source. The non-idealities present in the circuit, however, pose a problem for approximation accuracy, and this was compensated for by further developing the model by adding a back-prediction module that tries to predict a better input sequence for the system based on previous measurements.

The controller cards for the IGBTs were found to be satisfactory, and the prototyped coil identification system seems like a feasible solution even in the presence of strong magnetic fields. The waveform approximation was found to give rising-phase predictions with 0.3—7.7% relative difference in maximum amplitude compared to actual output for the sequences tested, depending on the chosen correction parameters. The falling-phase predictions varied significantly due to lack of parameter data. The tools developed in this Thesis give a good starting point for further development of waveform control in TMS.

Keywords: Transcranial magnetic stimulation, Waveform control, IGBT, H-bridge, mTMS

Publishing language: English

AALTO-YLIOPISTO
Sähkötekniikan korkeakoulu
Master's Programme in Life Science Technologies

DIPLOMITYÖN
TIIVISTELMÄ

Tekijä: Heikki Sinisalo

Työn nimi: Aaltomuodon hallinta transkraniaalisessa magneettistimulaatiossa

Sivumäärä: 54 + 7	Päivämäärä: 22.10.2018

Pääaine: Biosensing and bioelectronics

Valvoja: Prof. Risto Ilmoniemi

Ohjaaja: Jaakko Nieminen, TkT

Transkraniaalinen magneettistimulaatio (engl. transcranial magnetic stimulation, TMS) on ei-invasiivinen menetelmä aivojen stimulaatioon. Menetelmä perustuu indusoituihin sähkökenttiin, jotka luodaan ajamalla suuri virtapulssi stimulaatiokäämin läpi. Aalto-yliopiston TMS-ryhmä on kehittänyt uuden sukupolven mTMS-laitetta (engl. multi-locus TMS), jossa sähkökenttien superpositioon perustuen stimulaatioaluetta voidaan siirtää ja kääntää sähköisesti hyödyntäen useita keloja. Tyypillisesti stimulaatioalueen siirto tai kääntö toteutetaan fyysisesti kelaa liikuttamalla. Yliopistolla kehitetty laite perustuu sähköiseen H-siltakytkentään, jossa sähkövirran kulkureittiä voidaan hallita kytkemällä IGBT-transistoreita (engl. insulated gate bipolar transistor) päälle tai pois päältä. Tämän diplomityön tavoite oli kehittää menetelmä piirin virrankulun hallintaan siten, että haluttu referenssipulssi (eli aaltomuoto) voidaan mallintaa.

IGBT-transistoreiden parempaa hallittavuutta varten ryhmässä oli aiemmin kehitetty uudet ohjainkortit, joiden testaaminen oli yksi tämän diplomityön osa-alueista, ja lisäksi uusien kelojen lisäämistä silmälläpitäen järjestelmälle valmistettiin prototyyppi kelojen tunnistusta varten. Pääpaino työssä oli kuitenkin kehittää IGBT-transistoreiden hallintaan algoritmi, jolla kelan läpi kulkevaa virtaa voidaan tarkasti hallita. Perusidea tässä algoritmissa on, että suurella alkujännitteellä ajetaan stimulaatiokelaan virtaa vain hetkittäin, jolloin voidaan efektiivisesti simuloida tilannetta, jossa alemmalla alkujännitteellä ajetaan virtaa kelan läpi jatkuvasti. Laitteen epäideaalisuudet johtivat hyvin epätarkkaan approksimaatioon, mitä varten kehitettiin ennustusmoduuli, joka pyrkii aiempiin mittauksiin perustuen antamaan paremman ennusteen aaltomuodon käyttäytymisestä.

IGBT-ohjainkortit toimivat mittausten perusteella hyvin, ja kelojen tunnistusjärjestelmä vaikuttaa ainakin ensiarvioiden perusteella hyvältä, vahvoista magneettikentistä huolimatta. Approksimointialgoritmi testatuille aaltomuodoille antoi 0.3—7.7% suhteellisen eron maksimiamplitudien välille riippuen korjausparametreista. Työssä esitetyt työkalut antavat hyvän lähtökohdan aaltomuotojen hallinnan jatkokehitykseen TMS:ssä.

Avainsanat: Transkraniaalinen magneettistimulaatio, Aaltomuodon hallinta, IGBT, H-silta, mTMS	Kieli: Englanti

Preface

This Thesis was written while working with Aalto University's TMS group at the Department of Neuroscience and Biomedical Engineering (NBE).

I would like to thank my advisor, Jaakko Nieminen, for constructive feedback and discussions about the topic and this Thesis in general, and additional thanks goes to my supervisor, Risto Ilmoniemi, for giving me the opportunity of working with the TMS group. I would also like to extend my gratitude to the TMS group in general as well as friends and family.

Otaniemi, Espoo, October 22, 2018. Heikki Sinisalo

Contents

Glossary

FPGA Field-programmable gate array. 14, 17, 26, 28, 34, 36, 37, 38, 43, 49, 50

IGBT Insulated gate bipolar transistor. 9, 10, 11, 12, 13, 14, 16, 17, 18, 20, 21, 23, 26, 27, 35, 36, 37, 39, 40, 41, 48

Matlab MATrix LABoratory. 5, 26, 37, 48, 49, 50

MCU Microcontroller unit. 24, 25, 42, 49

mTMS Multi-locus transcranial magnetic stimulation. 9, 13, 15, 24, 25, 28, 41, 51

SPICE Simulation program with integrated circuit emphasis. 16, 28, 35, 36, 39, 47, 48, 50, 51

TMS Transcranial magnetic stimulation. 1, 2, 3, 4, 6, 7, 8, 9, 12, 15, 16, 23, 24, 25, 38, 41, 42, 48

1 Introduction

This work examines the theoretical aspects behind transcranial magnetic stimulation (TMS), a method of stimulating the brain with magnetic fields [1]. The focus is on the electronic implementation of one such stimulation device developed at Aalto University [2].

Based on the underlying behavior of the electronic circuitry, a model is proposed to offer the possibility of generating close approximations of given reference waveforms in real time. The significance of this is tied—but not limited—to the use of multiple coils simultaneously in order to rotate or move the area of stimulation in a predictable manner without physically moving the coil; in a typical setup the need to physically move the coil to adjust the point of stimulation is a limiting factor, especially in a confined environment.

A big portion of the work is devoted to dealing with non-idealities arising from the physical implementation of the device, but testing and characterization of the recently renewed controller cards is also a part of this Thesis. Additionally, a method for identifying between multiple coils is presented.

Section 2 presents the general background of TMS and describes the stimulation device developed at Aalto University. Section 3 presents the methods for characterizing the current TMS system, describes a method for coil identification and proposes a model for approximating reference waveforms. In addition, the fundamentals behind compensating for non-idealities are explained. Section 4 presents the results of the work in an orderly manner, and Section 5 provides viewpoints for what worked and what could be improved. Section 6 concludes this Thesis.

2 Background

This section describes the theoretical aspects behind TMS and the current state of our stimulation device, and proposes a model for controlling the waveform of the current driven through the stimulation coil.

2.1 Transcranial magnetic stimulation

TMS is a promising, non-invasive stimulation method that provides access to the brain. As the magnetic field penetrates through the skull to the brain almost uninterruptedly, it offers a unique way to interact with the brain. The neurons of the brain communicate with each other electrically, chemically and physically [3]. To understand better the underlying principles of brain activation via TMS, we will next take a look at the neuron and the way it sends and receives information via action potentials. [3, 4]

2.1.1 Neuron

Without going too much into detail of the anatomy of the cell, the typical neuron is a specialized, multi-branched cell with a large central portion called soma and outward spreading branches called neurites. These neurites can be divided into dendrites, which act as communication inputs for the neuron, and axons which act as outputs. Figure 1 shows a general depiction of the typical neuron. One should keep in mind that the depiction given here is very generalized—there are very different types and shapes of neurons, not all of which fit the description here, and a nice discussion of neuronal classification is available in [3, pp. 46–49].

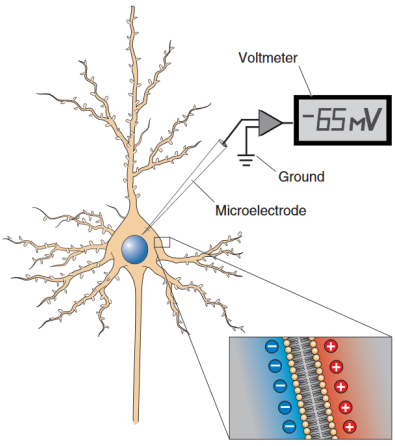

Figure 1: The typical neuron has a central body (soma) and branching neurites reaching outwards. Most of these neurites are small, treelike dendrites, which act as communication inputs, but one neurite is markedly larger and does not branch out. It is the axon. [3]

The neuronal membrane that separates the neuron's internal environment from the extracellular environment is filled with molecular pumps and channels. These pumps and channels act to establish a delicate potential difference across the neuronal membrane; the channels allow the ionic species that physically fit through them to flow against the direction of the concentration gradient whereas the pumps are using molecular energy (adenosine triphosphate, ATP) to drive ionic species in the direction of the concentration gradient. In addition the membrane is quite thin so the ions are affected by the ions on the other side of the membrane (i.e., a positive ion accumulation on one side would repel positive ions on the other side but attract negative ions), which sets up an ionic charge equilibrium across the membrane maintained by said pumps and channels. The potential difference, or the *resting membrane potential*, is approximately $-70\,\mathrm{mV}$. [3, 5]

The axonal membranes are also filled with another type of channels that are not always allowing ionic species to pass, so-called voltage-gated channels. These channels are in a non-conducting state at resting membrane potential, but as the neuronal membrane starts to depolarize (i.e., the voltage goes more positive), a specific sensing element alters the channel's morphology and switches the channel into a conducting state allowing ions to pass. This process typically starts at a region of the neuron called the *spike-initiation zone*, which could be located for instance at the axon hillock[1]. When the intraneuronal space starts depolarizing, at a certain threshold (approximately $-55\,\mathrm{mV}$, although this can depend on the neuron in question) the voltage-gated channels for sodium (Na^+) start opening and letting in more sodium ions which in turn start progressively turning the intra-axonal region more depolarized activating more voltage-gated channels further down the axon. This propagation of depolarization is what is known as the neuronal spike, or action potential. The depolarization does not, however, stay forever, as the voltage-gated sodium channels are not the only channels to activate during action potential initiation. With a small delay also voltage-gated potassium (K^+) channels are opened which let potassium ions move against the concentration gradient: out from the cell. This causes the neuronal membrane to be polarized again towards the resting potential. The action potential is shown in Fig. 2. [3, 5]

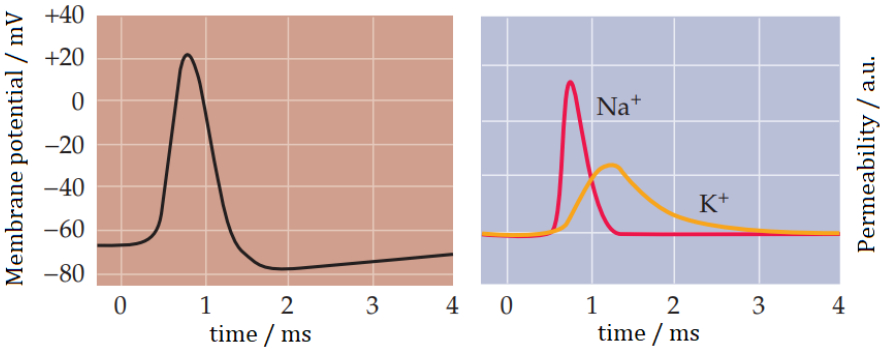

Figure 2: Action potential and membrane permeability for sodium and potassium. Image adapted from [5].

[1] The portion which connects the axon to the soma.

2.1.2 Stimulation

In TMS, the stimulation process is based on generating an electric field across an axon. This field will start to drive positive charge to the membrane, thereby depolarizing it. After the membrane has been sufficiently depolarized (typically around 15 mV increase from the resting potential [5]), the threshold for an action potential is reached and an action potential is fired. This electric field can be induced with a magnetic field produced by driving a sufficiently large time-dependent current $i(t)$ to a coil placed near the scalp. The mutual existence of the electric and magnetic fields is described by the Maxwell–Faraday equation:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{1}$$

where $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic field and t is time [6]. Related to these electric and magnetic fields, a force

$$\mathbf{F} = q\mathbf{E} + q\mathbf{v} \times \mathbf{B} \tag{2}$$

is acting on the ionic charge in the brain tissue where q is the magnitude of the ion's charge and $\mathbf{v}$ is its velocity. The magnetic force component can be considered insignificantly small[2]. This force[3] causes charge accumulation at neuronal membranes, disturbing the equilibrium by driving ionic species along the electric field. As a consequence, an action potential can be triggered. [4, 6]

To generate such an electric field in the cortex, a large capacitor (for example $\sim$100 µF in commercial devices) is typically used to drive the magnetic field generating current as a short-duration pulse [4]. A generic circuit diagram of a TMS device is depicted in Fig. 3, where R denotes the intrinsic resistance of the coil's wiring (and other possible sources of resistance, such as busbars), C is the capacitor's capacitance and L is the coil's inductance.

[2] If we consider the blood flowing through the aortic valve to have among the highest velocities found in the human body (approximately 100 cm/s [8]), with a field strength of 2 T the magnetic field related force component is at least two orders of magnitude weaker than that caused by an electric field of 100 V/m.

[3] Also known as Lorentz force.

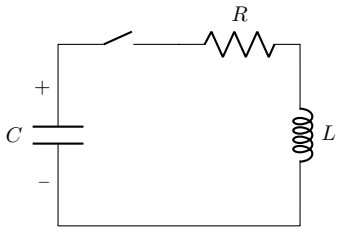

Figure 3: A basic TMS circuit consists of a resistance, an inductance and a capacitance connected in series. The resistance is due to the intrinsic resistance of the stimulation coil's wiring. The switch in a typical device can be implemented with a thyristor.

This circuit turns into an RLC resonator when the switch is closed, and this can be solved by applying Kirchhoff's voltage law, which states that the total voltage over a closed loop must be 0 [7]. As this circuit is an undriven[4] circuit, we can express the voltages over the individual components the following way:

$$V_{\mathrm{R}} = Ri(t) \tag{3}$$

$$V_{\mathrm{L}} = L\frac{di(t)}{dt} \tag{4}$$

$$V_{\mathrm{C}} = \frac{1}{C}\int i(t)dt \tag{5}$$

where $i(t)$ is the current as a function of time and V_{R}, V_{L} and V_{C} denote the voltages over the resistor, the inductor and the capacitor, respectively [7]. These equations can be combined to a more solvable form by noting that

$$i(t) = -\frac{dq(t)}{dt} \tag{6}$$

where q is charge. By combining Eqs. 3–6 we can state Kirchhoff's voltage law over the closed RLC loop:

$$\frac{R}{L}\frac{d}{dt}q(t) + \frac{1}{LC}q(t) + \frac{d^2}{dt^2}q(t) = 0 \tag{7}$$

The full analytical solution can be, for instance, computationally solved with software such as Matlab (The MathWorks, Inc., MA, USA) or Maple (Maplesoft, a division of Waterloo Maple Inc., Ontario, Canada) by setting the initial conditions to $i(0) = I_0$ and $q(0) = Q_0$. The resulting solution for $i(0) = 0$ has, to a high accuracy if the resistance is *small*[5], the form of [9, pp. 987–988]

[4] There is no active source driving the circuit.

[5] Also called weak damping. The resistance can be considered small if $R^2/4L^2 \ll 1/LC$ which stems from the full analytical solution.

$$q(t) = Q_0 \exp\left(-\frac{R}{2L}t\right)\cos\omega_0 t \tag{8}$$

$$i(t) = \omega_0 Q_0 \exp\left(-\frac{R}{2L}t\right)\sin\omega_0 t \tag{9}$$

where

$$\omega_0 = \sqrt{\frac{1}{LC}} \tag{10}$$

Plotting the obtained solutions with respect to time we can see the time evolution of the pulse (the current) and the charge at the capacitor, which is shown in Fig. 4. As can be seen in the graph, the current oscillates back and forth in the RLC circuit as the energy content initially stored in the capacitor's electric field is released and stored in the inductor's magnetic field instead. When the capacitor's charge has been transferred from the positive plate to the negative, the current in the circuit reverses direction and the capacitor's original positive plate starts to recollect charge again, only this time some of the power has dissipated as heat because of the circuit resistance, and the maximum charge level obtained at the capacitor is lower than initially. The right side of the image shows the shape of the electric field which is proportional to the derivative of the current. [6, 9]

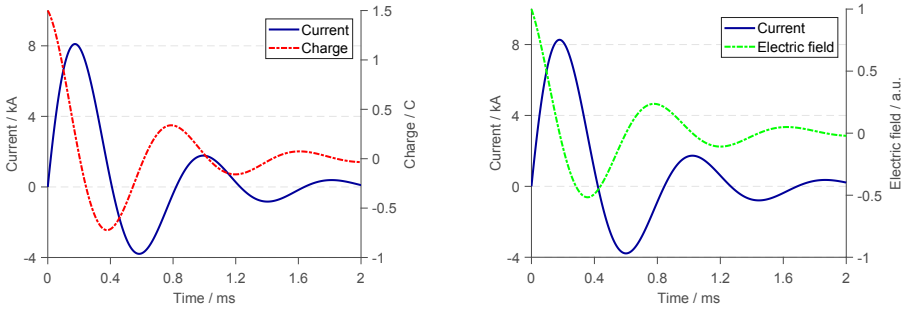

Figure 4: Undriven RLC circuit waveforms. Left: Oscillatory behavior of current (i) and capacitor charge (q) in the RLC circuit of Fig. 3 with the switch closed. Right: The behavior of the electric field (normalized) plotted against current. The parameters used were $R = 63\,\mathrm{m\Omega}$, $L = 17\,\mu\mathrm{H}$ and $C = 1\,\mathrm{mF}$.

Typical TMS devices are driven with either monophasic or biphasic pulses, although also polyphasic pulses are possible but not generally available. A monophasic pulse of current is essentially the first[6] positive phase of current seen in Fig. 4, and a biphasic pulse is the positive phase followed by a negative phase. As can be expected, the electric field induced by a monophasic waveform changes direction once and the field induced by a biphasic waveform changes

[6]Although the form of the pulse is not necessarily that of a sinusoid, it depends on the circuit topology used.

direction twice as the field is proportional to the time derivative of the current. [4]

A monophasic pulse can be produced by implementing the switch in the circuit of Fig. 3 with a thyristor. After the thyristor has been opened with a trigger signal, it will allow current to continue flowing even after the removal of the trigger until the current drops to 0, after which the thyristor closes. An idealized monophasic stimulation circuit is shown in Fig. 5. Placing a resistor and a diode in series parallel to the coil allows for the current to slowly decay in the coil. [4, 7]

A biphasic stimulation circuit can be obtained with a bypassing diode over the thyristor in reverse direction, also shown in Fig. 5. During the positive phase of the pulse the current passes through the thyristor normally, but when the current drops to 0 and the thyristor stops conducting current, the diode allows the current to travel in reverse direction, which also has the effect of partially recharging the capacitor. [4, 7]

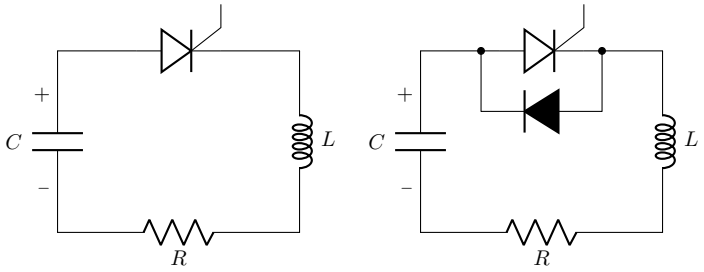

Figure 5: Idealized circuit diagrams for monophasic (left) and biphasic (right) stimulation. The bypass diode in the biphasic stimulation circuit allows the current to reverse direction while also partially recharging the capacitor.

As stated previously, the induced electric field is proportional to the rate of change in the coil current, or the second derivative of $q(t)$. This electric field $\mathbf{E}$ can, in turn, be used to calculate the depolarization of the neuronal membrane according to

$$\Delta V(t) \propto \int_{-\infty}^{t} dt' \cdot E(t') \cdot \frac{1}{\tau_{\mathrm{m}}} e^{(t'-t)/\tau_{\mathrm{m}}} \tag{11}$$

where τ_{m} is a membrane time constant related to the rate at which accumulated charge leaks from the membrane and $\Delta V(t)$ is the change in membrane potential [10]. τ_{m} has been estimated to be between 150 and 200 µs [10, 11] whereas the typical duration of a TMS pulse can be in the neighborhood of 60–120 µs regarding its rising period. The TMS pulse is thus short compared to the rate of leakage and as such the coalesced charge at the neuronal membranes leaks away relatively slowly compared to the rate of accumulation. One should also note that outside TMS experiments the neuronal time constant has been estimated to be from hundreds of microseconds to dozens of milliseconds, depending on the neuron [12]. [13]

Essentially, the changing magnetic flux of a TMS coil produces a current in the conducting brain tissue in a direction that produces a magnetic field opposing this change of flux, which is the direction opposite to the rate of change of the current in the coil. These currents induced in the tissue, however, are very miniscule but they have a secondary effect in that when the charge moves in such a way that it starts accumulating at neuronal membranes (so essentially perpendicular to the axon's axis), the membrane becomes depolarized. When it is sufficiently depolarized above certain threshold, an action potential is fired.

2.1.3 Coil types

Different coil types can be used in TMS to produce different spatial profiles of electric fields in the brain tissue. The most common ones used are the circular coil and the figure-of-eight coil. [4]

For the circular coil the generated electric field is circular and the stimulation area is very large and blurry[7] whereas the figure-of-eight coil has a so-called *hotspot* under the center point of the two adjacent windings where the electric field is maximal [14]. This is due to the current traversing in the same direction in both windings next to this position, which causes an electric field superposition of high magnitude. The figure-of-eight coil's well defined focal point allows for more precise stimulation than what is possible with a circular coil. From Fig. 6 one could visualize setting side by side two of the left-side electric field amplitude profiles (just like the rings above the right-side profile are side by side), which would cause the center portion to overlap causing a linear superposition of the amplitude. Utilizing different coil shapes allows one to customize the electric field profile. [4]

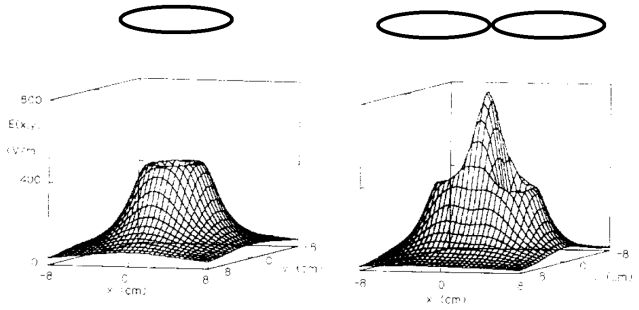

Figure 6: Induced electric field amplitudes under two coil shapes. Left: Circular coil has a broad electric field profile. Right: Figure-of-eight coil has a well-defined focal point. Images adapted and modified from [14].

[7] In other words, badly focused.

2.1.4 Multi-locus TMS

In a typical TMS setup, a major disadvantage is the need to physically move the coil when the stimulation focus point needs to be moved or rotated. This can be especially cumbersome in situations where there is limited space to physically interact with the device, such as inside the bore of a magnetic resonance imaging (MRI) scanner with the diameter of the bore being typically approximately 60 cm. It also has to be remembered that in MRI the head is not actually at the edge of the scanner bore, but deeper in the machine where the static field is maximally homogeneous making it hard to physically reach and interact with the coil. [15, 16]

Because the electric field induced in tissue is a superposition of all the fields generated by a set of coils, varying field patterns can be generated by having a coil set consisting of different shapes of coils and driving suitable currents through them. This kind of multi-locus TMS (mTMS) transducer was introduced in Ref. [17].

Figure 7 shows the electric field profile in a figure-of-eight coil's hotspot. When two such coils are placed on top of each other in a perpendicular orientation, the electric field can be changed by controlling the relative current in the coils [17].

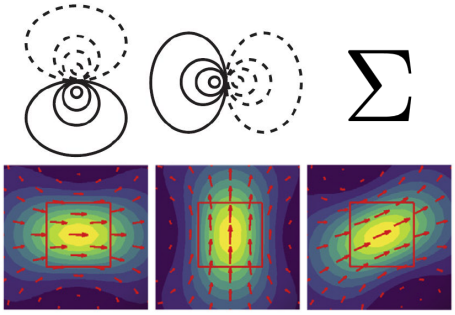

Figure 7: The electric field can be rotated by utilizing two perpendicularly arranged figure-of-eight coils on top of each other. Top row: The orientation of the coil. Bottom row: Resulting electric field profile in the hotspot. The rightmost field profile is a result of running almost equal currents through both coils. The higher the relative current through one coil compared to the other, the more the spatial profile of the total electric field matches that coil's individual electric field profile. Brighter color indicates higher electric field amplitude. Image adapted from [17].

2.2 mTMS device at Aalto University

The TMS research group in Aalto University has developed an mTMS device which can drive simultaneous current pulses through two coils in order to move or rotate the electric fields.

The power electronics setup for each coil of the stimulation device is an H-bridge

circuit, shown in Fig. 8. The insulated gate bipolar transistors (IGBTs, denoted by numbers 1–4) are controlled to alter the circuit topology, as current can only flow through them if they are in a conducting state. The freewheeling diodes (D1–D4) provide a path for the current to travel in the reverse direction past the transistors. The inductor in the middle represents the stimulation coil. The resistive component in the system is due to the intrinsic resistance of the coil wiring (and other possible sources of resistance, such as busbars). As the most significant contribution to this resistance, and common to all circuit topologies possible, comes from the coil, a single resistor in series with the coil is sufficient to represent the bulk of the resistance. [18, 19]

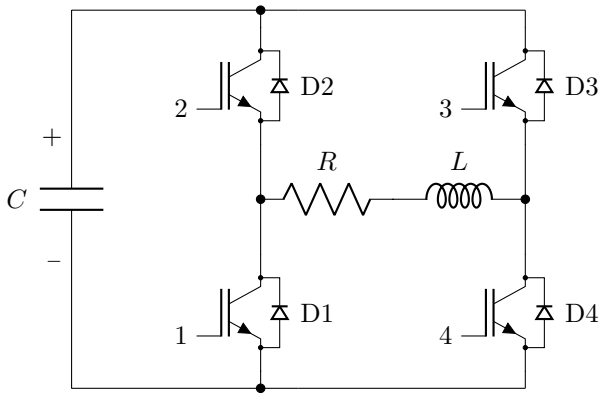

Figure 8: H-bridge circuit allows for the control of the pulse waveform by altering the state of the IGBTs. This physically changes the circuit topology between an RLC and an RL circuit. The component values in the circuit are approximately as follows: $C = 1.02\,\mathrm{mF}$, $L = 17.5\,\mu\mathrm{H}$, $R = 63\,\mathrm{m}\Omega$. The resistance and the inductance depend on the coil. Diodes D1–D4 are the freewheeling diodes of the IGBT modules (numbered 1–4), and allow the current to flow in the opposite direction over the transistor. [18, 19]

From Fig. 8 it can be seen that there are three possible modes for the current to flow from the perspective of the capacitor. When the capacitor is charged and IGBTs 2 and 4 are opened (i.e., their channels are open and thus the transistors are conducting), the current pathway goes from the capacitor through IGBT 2, through the coil and via IGBT 4 to the negative pole of the capacitor, producing a positive current slope henceforth known as the *rising phase*. This is an RLC circuit, which follows the equations presented earlier (e.g., Eq. 7), and the path of the current through the circuit can be seen in Fig. 9 (left). [19]

When IGBT 2 is closed[8] but IGBT 4 is still conducting, the current through the inductor must remain continuous (see Eq. 4) and it starts circulating in the loop coil – IGBT 4 – D1 since the freewheeling diode D1 of IGBT 1 allows the current to pass in this direction. The amplitude of the current decreases over time since power is dissipated as heat due to the intrinsic resistance of the wiring. The system is now a typical RL circuit with a well-known solution for the behavior of the current:

[8]Set to a non-conductive state.

$$i(t) = i(0)e^{-\frac{R}{L}t} \tag{12}$$

where $i(0)$ is the current after IGBT 2 has been switched off. The solution can be obtained from Eqs. 3 and 4 by applying Kirchhoff's voltage law over the closed loop[9]. [7] This mode of operation is henceforth known as a *holding period*, and the path of the current through the circuit during this mode is shown in Fig. 9 (middle). If a single long holding step follows the peak amplitude of the current waveform, that holding step is referred to as the *peak hold period* in this Thesis.

The last mode of operation is with all the IGBTs switched off. When there already is a current circulating in the system (such as during a holding period), switching all the transistors off will cause the current to flow via two of the diodes back to the capacitor, partially recharging it, as the current through an inductor cannot change instantaneously. This phase will henceforth be known as the *falling phase*. It should be noted that when the IGBTs are in a non-conducting state, the current cannot flow from the capacitor[10] since diodes only allow for unidirectional passage of current (unless they undergo catastrophic failure and turn into a short circuit [20]). The path the current takes during a falling phase is shown in Fig. 9 (right). [19]

Figure 9 shows the current pathways possible and Fig. 10 the waveforms of the current during the different phases. It should be noted that the current can also be circulated in the opposite direction through the coil if different IGBTs are manipulated for the process (IGBTs 1 and 3). Another point of note is the RL phase, during which the current can be circulated in either the upper or the lower loop depending on which transistor is switched off.

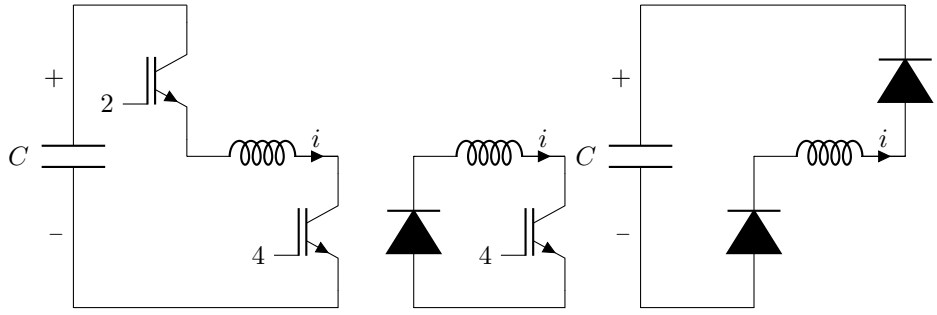

Figure 9: Circuit topology changes due to IGBT switching. Left: Rising-phase circuit (RLC). Middle: Holding-period circuit (RL). Right: Falling-phase circuit (RLC).

The shape of the pulse generated by the sequential switching of the three aforementioned circuit states once is referred to as a *trapezoidal pulse* from here on, and such a pulse is visualized in Fig. 10. The rising phase follows the exponentially decaying sinusoid of Eq. 9 and the peak hold period follows the RL

[9]It should be noted that for simplicity the effect of transistor and diode voltage drop is considered negligible.

[10]In other words, start a new cycle.

decay of Eq. 12. The falling-phase current can be calculated by reversing the capacitor's charge in the RLC equation (Eq. 7). In this work the word *phase* is used for the circuit states that cause large changes in the waveform (e.g., the RLC state causes a rapid rise or decay of current through the coil), whereas the word *period* indicates a state where the waveform stays relatively constant (e.g., the RL state causes a slow decay of current through the coil compared to the RLC state).

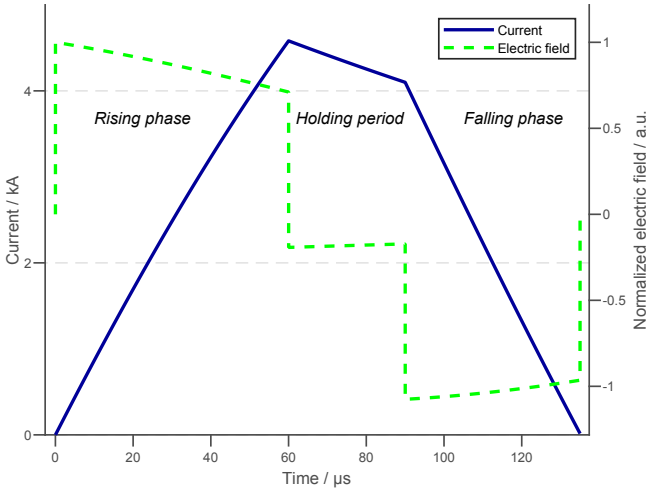

Figure 10: A trapezoidal pulse can be generated by driving the system of Fig. 8 in the three phases mentioned in a sequence: rising phase, holding period, falling phase.

2.2.1 IGBTs

The switches in the H-bridge circuit are realized with IGBTs. The IGBT is a voltage-controlled power device capable of conducting high-amplitude currents when in the conducting state and withstanding high voltages without dielectric breakdown in the non-conducting state. [21]

From our perspective, the IGBT is especially advantageous due to its low switching losses and fast turn-on rate, which allows the current waveform to be precisely controlled in the microsecond range. The model of IGBT in the current setup is ABB 5SNA 1500E330305, which has characteristic turn-on and turn-off times of 0.1–1 µs when operating within specifications, depending on the collector current [22]. The maximum continuous current stated in the data sheet is 1.5 kA, but when the TMS device is operated in short bursts, much higher currents can be driven in the circuitry if the device is operating at temperatures under 150 °C [23].

2.2.2 Snubber circuits

The H-bridge configuration and sequential switching of the transistors can put the IGBTs under tremendous stress. When switching from one state to an-

other, the current through the coil must be continuous under all circumstances according to Eq. 4. This can cause a high-voltage transient to develop between the collector and the emitter of an IGBT which can cause the transistor to be operated beyond its ratings. [21]

Adding a snubber circuit across the IGBT allows it to take on a portion of the transistor's current when the IGBT is switched off, and additionally the snubber can help suppress the ringing of the collector-emitter voltage, V_{CE}, of the IGBT caused by parasitic inductances [23]. The unfortunate side effect of adding snubber circuits in parallel with the IGBTs is a slight loss of control over the waveform as the snubbers cause a rounding effect, which is especially noticeable during the first steps of repeated fast[11] switching applications. [21]

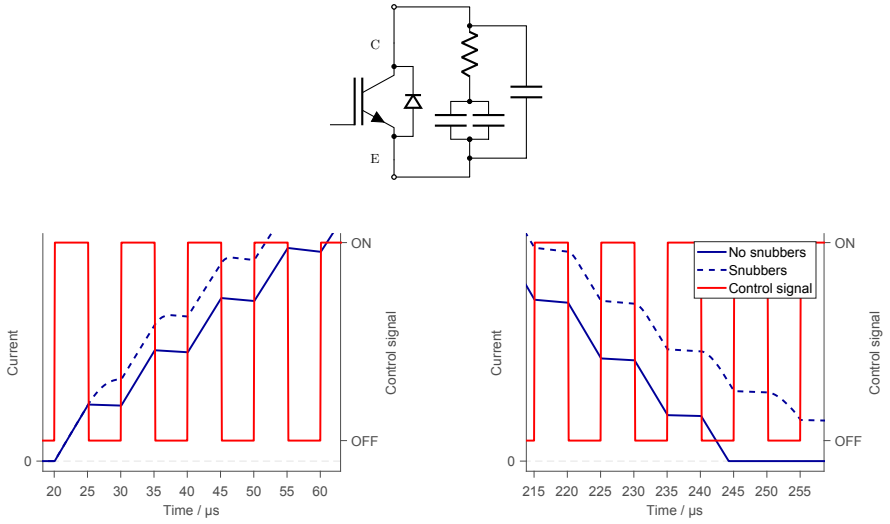

Figure 11: Top: Circuit diagram of the RC snubber design used in the device. C and E denote the IGBT's collector and emitter, respectively. The resistor has a value of $R = 1\,\Omega$ and the capacitors' values are $C = 0.33\,\mu\mathrm{F}$. Bottom: A simulation in LTSpice (Linear Technology Corp., Milpitas, CA, USA) of the coil current in the H-bridge (see Fig. 8) with and without the snubbers in place during fast switching, plotted against the control signal. Left side shows the rising phase where the transistor is switched on to drive more current into the coil. Right side shows the falling phase where the transistor is switched off to drive the current back into the capacitor.

Figure 11 shows the snubber design (RC type) that is in use in the mTMS device and the rounding effect it has on current pulses through the coil during fast switching mentioned earlier. With longer-duration pulses, the rounding effect of the first step is less and less significant (not visualized, but the rounding effect can be observed to diminish after the first few steps), and when the first step of a pulse series has a duration of 20 µs the rounding has vanished completely. This effect is estimated to be due to the snubber capacitors acting as a virtual short circuit until they are fully charged when the transistors are not conducting,

[11] ∼5 µs or less between switches.

and when the IGBTs are turned into a conductive state, the snubber capacitors release their charge content via the IGBT's channel. Similarly, during the falling phase the waveform in Fig. 11 without the snubbers can be seen to match exactly with the control signal, whereas the waveform with the snubbers has similar rounding effect that, in this case, causes the current to decay slower. [21]

2.2.3 Controller cards

The H-bridge IGBTs are controlled by dedicated controller cards that take an optical input from a field-programmable gate array (FPGA), a specialized microprocessor whose internal wiring, and thus logic, can be programmatically altered for application-specific purposes [24, p. 67]. The controller cards have gone through several iterations over the years, and the current version's circuit layout is shown in Fig. 12. The nodes labeled G and E are connected to the IGBT's gate and emitter, respectively.

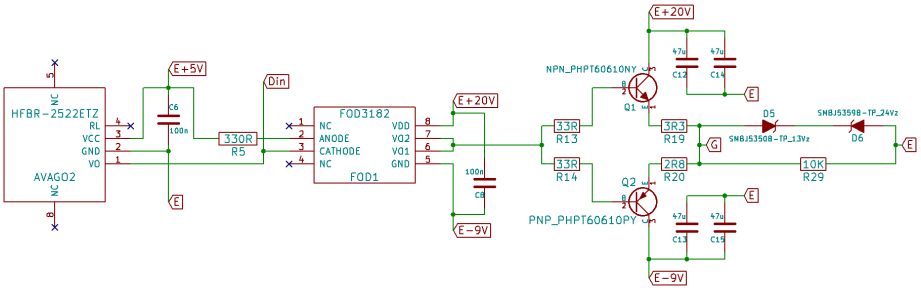

Figure 12: The current version of the controller cards. Missing from the diagram are the actual DC/DC converters and an experimental IGBT short circuit detection scheme that contains a high-voltage divider and a logic section (node D_{in}).

The controller card takes an optical input signal via an optical receiver,[12] which switches the internal transistor of the device on, providing a path for the current to ground from the E+5 node via an optical coupler[13], which in turn has an internal light emitting diode (LED). Without input signal, transistor Q2 is in a conducting state (PNP base low voltage) and Q1 in non-conductive state (NPN base low voltage), which sets the IGBT's V_{GE} (node G connected to E−9) negative, and the device thus in its non-conductive state. When the optocoupler's LED turns on, the transistors switch states which provides the IGBT with high V_{GE}, switching it on. It should be noted that the circuit diagram only depicts the driver side of the controller card. Additionally, the circuit contains three DC/DC (5, −9 and +20 volts) converters[14] for providing power to the various components, and node D_{in} is a branch for an experimental short-circuit detection scheme. The voltage clip seen at the far right limits the voltage between the gate and the emitter between 24 and −13 volts. The

[12] Avago HFBR2522ETZ. [25]
[13] FOD3182. [26]
[14] Murata MGJ2D152005SC, Murata MEJ2S1509SC and MEJ2D1505SC. [27, 28]

complete circuit diagram for the controller card can be seen in Appendix A. [21, 25, 26]

2.3 Effective voltage control

As explained previously (see Section 2.1.4), an mTMS transducer consists of multiple coils of different shapes that can be utilized to rotate or move the electric field hotspot. As the electric field is a superposition dependent on the relative current waveforms[15] driven through each coil, the initial capacitor voltages need to be adjustable if trapezoidal pulses as in Fig. 10 are used.

While recharging the capacitor is possible before each pulse[16], it is cumbersome and slow to discharge the capacitor through a resistor to generate lower-amplitude waveforms, and afterwards recharge the capacitor again if higher-amplitude waveforms are needed. Another possibility is to generate these lower amplitude waveforms by switching between circuit states, carefully controlling the current through the coil. This way we can produce low-current waveforms with some initial high capacitor voltage. As an additional benefit, a model based on the analytical solution for the current behavior can take into account variation of the initial capacitor charge that is partially depleted after each pulse. This approach of simulated voltages relies on the fact that the neuronal membrane's time constant is long compared to the characteristic times of the TMS waveform (see Section 2.1.2 and Eq. 11).

Controllable pulse parameters have been introduced previously. For instance, Peterchev et al. proposed this kind of device in Ref. [29], and a TMS device utilizing H-bridge topology for enhanced waveform controllability has been proposed in Ref. [18].

2.3.1 Electric field desynchronization

The desynchronization of electric fields is a possible problem in the system when the simulated waveforms are used, for instance, to rotate the electric field that is a superposition of the individual fields of two perpendicular coils. Whereas the reference electric field points only in a single direction during a given phase, the stepped approximation produces a stronger field (due to steeper current slope) for a short duration during the rising phases of the steps and a weak reverse field during the holding periods of the steps. Essentially we could end up with twice the field strength for half the time during a rising phase.

When we are running two of these stepped pulses to rotate the field superposition at a specific location, in a worst-case scenario we are running the exact opposite circuit state in each coil: when one coil is running a rising ramp the other is running a holding period and vice versa. This turns the induced electric field we are trying to generate into a non-superposition in the tissue, and instead we have alternating electric fields that, due to Lorentz force (see Eq. 2), cause the charge movement to change direction.

[15] Or rather, the slope of the current.
[16] Unless the inter-stimulus interval is very short, e.g., 0.5 ms.

Let us consider the rotation of the electric field in a corner like area of an axon (or a hypothetical cell), where the electric field can be oriented towards the junction wall, depicted in Fig. 13. In a system consisting of two perpendicular figure-of-eight coils the electric field in the focal point can be rotated, and a 45° angle would optimally drive the charge towards the junction, depolarizing the membrane. On the other hand, if we have alternating electric fields[17] driving charge to their respective perpendicular membranes, a corner region joining those walls would be a common region that would still collect charge accumulation from both electric fields, albeit to a lesser extent compared to a 45° constant field. As neither of the fields is pushing charge directly away from the parallel membrane, but rather moving the already accumulated charge at the membrane more towards the corner region, the desynchronization of the electric fields should not pose a problem.

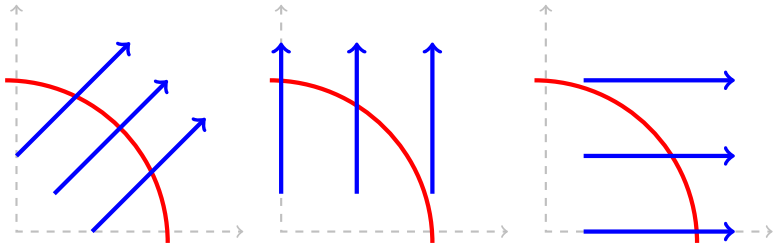

Figure 13: A membrane junction and associated electric fields. The red curve represents a neural cell membrane and the blue lines represent the electric field. Left: True superposition drives the charge optimally towards the junction. Middle and right: In a worst case scenario the electric fields come one at a time and there is no superposition. The corner region collects charge accumulation from both electric fields.

3 Methods

3.1 Snubber circuits

The behavior of current in the stimulator electronics (see Fig. 8) and the snubber circuits (see Fig. 11 (top)) was inspected with the SPICE (simulation program with integrated circuit emphasis) models described in Section 3.4.3. The control sequence used consisted of 5-µs-long rising phases followed by 5-µs-long holding periods. The stray inductances were removed from the model for clarity, as the resonant effects caused by them tend to mask the underlying phenomena. Observations show that the actual TMS device and the SPICE model generate highly similar waveforms when the same input sequence is used, partly justifying the simulation's trustworthiness.

[17]The miniscule reverse component caused by the RL decay of the other coil that slightly rotates the total electric field is not considered, as it is assumed to be insignificant.

3.2 Controller cards

The new controller-card design was not tested before, so the switching performance of the card under long trains of short pulses was tested with a bench-top setup. The IGBT was simulated by using a capacitive load of 100 nF between the gate and the emitter terminals, and an Arduino Zero (Arduino LLC, Somerville, MA, USA) was programmed to send long pulse trains every few seconds via an optical link to the controller card, utilizing an optical transmitter from the same manufacturer[18]. The V_{GE} level was monitored with an MSO-X 3024A oscilloscope (Agilent Technologies, Inc., Santa Clara, CA, USA). As Arduino's default library delay functions were not accurate enough, direct clock interrupts were utilized to generate more precise 5-µs delays. The pulse trains and the respective V_{GE} behavior was captured with the oscilloscope, where V_{GE} was measured from the added capacitor's legs. Electrical-to-optical conversion from Arduino to the controller card was built following the test-circuit schematics from [25, p. 5, Fig. 2] with slight modifications.

3.2.1 Short-circuit detection

The new controller cards have also been implemented with a specialized short-circuit detection scheme that provides optical feedback from the card itself back to the FPGA (or other analysis system). As there is a severe lack of documentation of this system, the qualitative function of the circuit was estimated from the circuit diagrams, scattered documentation and bench-top measurements.

The complete analysis circuit is shown in Fig. 14, and it seems to monitor the voltage level of the optocoupler input (FOD1 in Fig. 12, pin $\mathrm{D_{in}}$) and the voltage between the collector and the emitter of the IGBT and converts these to logical values, which are then compared to determine the status of the input (Is the controller card receiving an input signal?) and the status of the IGBT's channel (Is the voltage over the IGBT above some threshold?). The scheme utilized here is also known as *desaturation detection*, discussed in detail in [30]. Under normal operation, the voltage across the IGBT, when switched to its conductive state, is low and current is high. When a short circuit occurs at the load, the supply voltage will appear over the IGBT—the only impedance in the system. This voltage is detected with the V_{CE} divider and when it is sufficiently high, it is flagged by the comparator setup and compared with the optical input.

[18] Avago HFBR1522ETZ. [25]

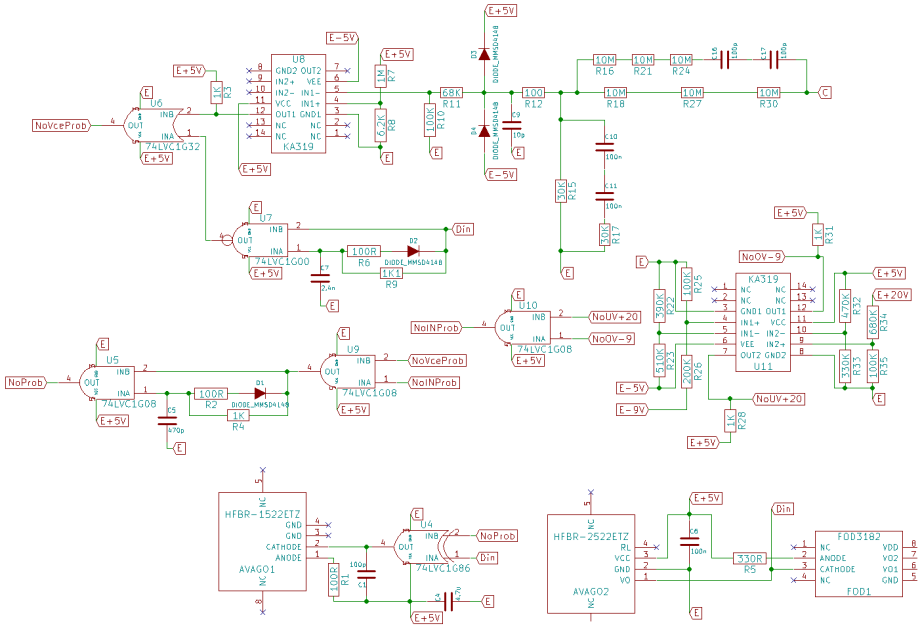

Figure 14: The short circuit detection scheme as a whole. The optical input side in the lower right corner is the input stage of the circuit represented in Fig. 12, and not part of the short-circuit detection. It is, however, included in the figure, as the D_{in} pin is located there.

Based on this information, a logical optical signal is sent out from the system to indicate whether everything is working normally or whether there is an error somewhere. The optical output is apparently[19] meant to be sent as the opposite phase of the input, i.e., no input signal should give an optical output signal if everything is functioning normally, and sending an optical input signal to the controller should result in no optical output signal. If an error occurs, the optical output should be in the same phase as the received input.

The voltage between the collector and the emitter of the IGBT is monitored with a high-voltage divider (HVD) connected to a comparator (U8, model KA319 [31]). The HVD consists of the typical voltage divider with each branch also having parallel capacitors to provide frequency compensation, as the system is proposed to be driven with high effective frequency where stray capacitances can affect the division's frequency response. A compensated divider is based on pole-zero cancellation which aims to make the transfer function independent of frequency by adjusting the transfer function of a resistive divider under certain conditions; Fig. 15 shows the typical resistive divider, a compensated divider and a more complex case utilized in the short circuit detection of this circuit.

[19]This was mentioned by the designer of the fault detector whom I interviewed.

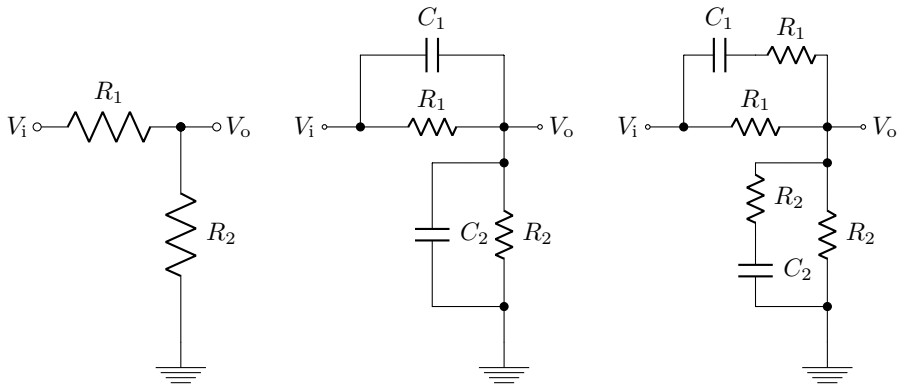

Figure 15: Voltage dividers. Left: Standard resistive voltage divider, can be problematic with high-frequency applications, as stray capacitance can cause a frequency dependent transfer function. Middle: Compensated divider, frequency independency when $R_1C_1 = R_2C_2$. Right: More complex design, series resistors are added to capacitor branches to prevent current shoot through.

As a quick literature survey did not reveal the frequency behavior of the employed HVD on the board, the frequency dependent transfer function is derived here[20]. From the figure above, the leftmost case is the familiar standard resistive voltage divider with transfer function of

$$H = \frac{R_2}{R_2 + R_1} \tag{13}$$

and adding parallel capacitors to it we get

$$H(j\omega) = V_\mathrm{out}/V_\mathrm{in} = \frac{Z_2}{Z_2 + Z_1} \tag{14}$$

with Z_1 and Z_2 denoting the complex impedance of each branch and the capacitor impedance is

$$Z_C = \frac{1}{j\omega C} \tag{15}$$

with ω denoting the frequency in radians, C the capacitor's capacitance and j is $\sqrt{-1}$. A branch impedance is thus

$$Z = R /\!\!/ Z_C = \frac{R}{j\omega RC + 1} \tag{16}$$

for each branch with their respective values and the transfer function after a bit of rearranging terms (Eq. 14)

[20] A proof to myself more than anything, as obviously this has been done before countless times. The simple compensated divider calculations can probably be found from most literature, such as Ref. [7].

$$H(j\omega) = \frac{R_2}{R_2 + R_1 \cdot \frac{j\omega R_2 C_2 + 1}{j\omega R_1 C_1 + 1}} \tag{17}$$

which reveals that with $R_1C_1 = R_2C_2$ the transfer function is frequency independent. For the divider in the rightmost case (and our actual circuit) of Fig. 15 we have branch impedance of

$$Z = R /\!\!/ \left(\frac{2}{j\omega C} + R\right) = \frac{R \cdot \left(\frac{2}{j\omega C} + R\right)}{R + \frac{2}{j\omega C} + R} \tag{18}$$

where the 2 in the numerator comes from our specific circuit having two series capacitors in each branch[21], see Fig. 16. Substituting R and C with the values from each branch we get the complex transfer function

$$H(jw) = \frac{R_2 \cdot \left(\frac{2}{j\omega C_2} + R_2\right)}{R_2 \cdot \left(\frac{2}{j\omega C_2} + R_2\right) + \frac{\left(R_2 + \frac{2}{j\omega C_2} + R_2\right) \cdot R_1 \cdot \left(\frac{2}{j\omega C_1} + R_1\right)}{R_1 + \frac{2}{j\omega C_1} + R_1}} \tag{19}$$

where we can make an inspired guess loosely based on our previous finding: $R_1C_1 = R_2C_2 = A$. Substituting this and rearranging the terms a bit, we arrive at

$$H(jw) = \frac{(j\omega A + 1) \cdot (j\omega A + 2) \cdot R_2}{(j\omega A + 1) \cdot j\omega A R_2 + R_2(2 + j^2\omega^2 A R_1 C_2 + 2j\omega R_1 \cdot (C_1 + C_2)) + j\omega A R_1 + 2R_1} \tag{20}$$

$$= \frac{(j^2\omega^2 A^2 + 2j\omega A + j\omega A + 2) \cdot R_2}{j^2\omega^2 A^2 R_2 + R_2(2 + j^2\omega^2 A R_1 C_2 + 2j\omega R_1 C_1 + 2j\omega R_1 C_2) + j\omega A R_1 + 2R_1} \tag{21}$$

$$= \frac{R_2 \cdot (j^2\omega^2 A^2 + 3j\omega A + 2)}{j^2\omega^2 A^2 \cdot (R_2 + R_1) + 3j\omega A \cdot (R_1 + R_2) + 2 \cdot (R_1 + R_2)} \tag{22}$$

$$= \frac{R_2 \cdot \cancel{(j^2\omega^2 A^2 + 3j\omega A + 2)}}{(R_1 + R_2) \cdot \cancel{(j^2\omega^2 A^2 + 3j\omega A + 2)}} = \frac{R_2}{R_2 + R_1} \quad \text{when} \quad R_1C_1 = R_2C_2 \tag{23}$$

which shows the divider employed in the V_{CE} detection scheme to be frequency-independent. The HVD converts the voltage level from 1.5-kV range to 1.5-V range, and the signal is further passed through a low-pass filter and a voltage clamp as a protective measure, through another divider (purely resistive) to a comparator which determines if the IGBT voltage is low or high, converting the result into a logical value. The HVD section of the circuit is shown in Fig. 16. [32]

[21] Calculated this way for the sake of the actual circuit having this exact configuration.

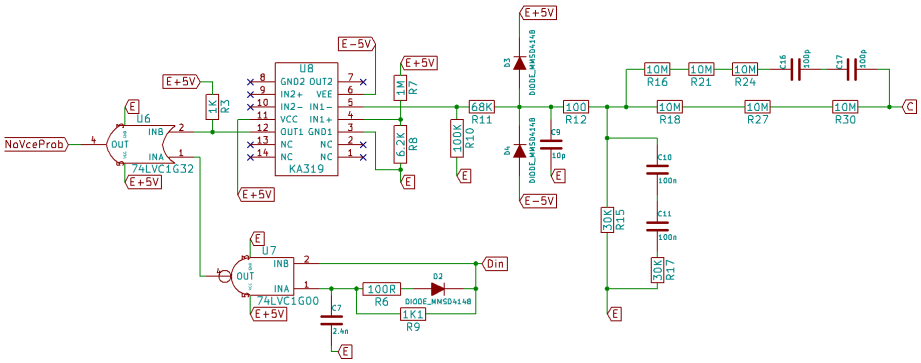

Figure 16: The high-voltage divider combined with the secondary divider converts the kilovolt voltage range to the one-volt range, allowing the level to be monitored by a simple comparator. The capacitors in the divider provide frequency compensation and the following low-pass filter directs transients to the ground. A voltage clamp is utilized for limiting the voltage to values between +5 and −5V, and a resistive secondary divider converts the voltage to the 100-mV range.

The comparator output of the high-voltage divider stage is compared to the status of D_{in} (which is HIGH for no optical input and LOW for positive input) in a logical circuit (OR port U6 and NAND port U7 [33, 34]) which determines whether the status of V_{CE} is correct[22] relative to the input signal. The timing circuit (D2–R9–R6–C7) seems to delay the NAND signal switch, thus starting the monitoring of V_{CE} slightly after the optical signal, possibly to account for the switching time of the IGBTs; this delay was also mentioned in the documentation but not discussed further. The delay is possibly a means to avoid an initial voltage spike of the HVD which could otherwise cause a false flag at the following comparator.

Interestingly, this timing circuit was indicated to have a delaying effect of ∼3.1 µs. This value can be calculated if we consider the time constant of the delay circuit to consist of the 1.1-kΩ resistor and the 2.4-nF capacitor, and the fact that the NAND gate's latching according to the data sheet occurs at $0.7V_{CC}$, where V_{CC} denotes the logical HIGH voltage assigned for the port [34]. From these values we would arrive at

$$0.7V_{CC} = V_{CC}(1 - \exp(-t/RC)) \Longrightarrow t = -\tau \ln 0.3 \simeq 3.178 \text{ µs} \qquad (24)$$

for the NAND latching delay where $\tau = RC$. However, there is another resistor in series with the 1.1-kΩ one in a preceding stage: a 330-Ω resistor is located between the D_{in} pin and the E+5 source, and this would increase the RC time constant of the delay circuit from 2.64 to 3.43 µs, and the total delay of the sub circuit to 4.13 µs.

The +20- and −9-volt inputs are also tested using a logic circuit shown in Fig. 17. The +20-V line is tested for undervoltage and the −9-V line for overvoltage using

[22] Conductive state during input and non-conductive otherwise.

a dual comparator (KA319, [31]), as they are used for the IGBT V_{GE} control. According to the comparators' input dividers, the overvoltage was set to trigger at −6.5 V and the undervoltage at 16.1 V, but a slight oversight can be seen at the negative reference, as the inputs should be inverted in this case. This was fortunately easy to fix by carefully cutting the card's surface traces and using short jumper wires to cross the inputs. The result is shown in the upper left corner of Fig. 17. [35]

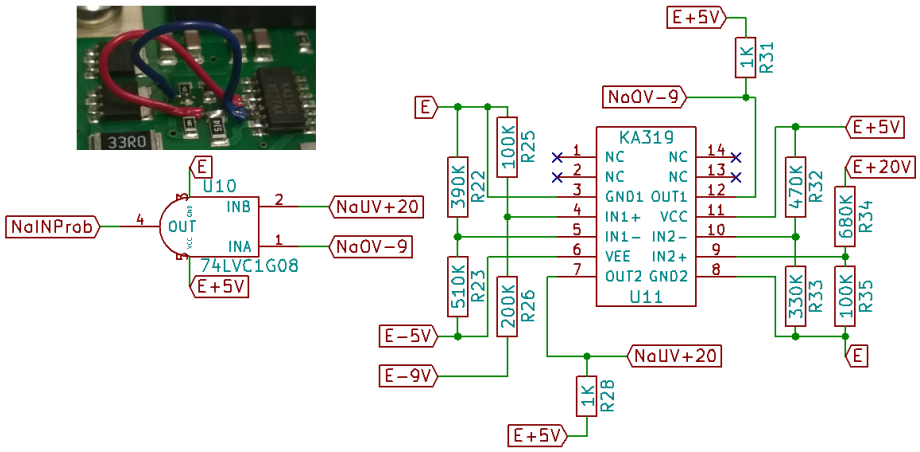

Figure 17: Under- and over-voltage detection for the +20- and −9-V inputs. The monitoring is done with a KA319 dual comparator and set to trigger at 16.1 and −6.5 volts, respectively (the diagram shows the original, uncorrected circuit). Top left shows the trace cut and cross soldering of the inputs at the negative monitoring stage. [31]

Both the input monitor and the V_{CE} status monitor are further fed to a delay circuit designed to keep the error signal active for at least 0.5 µs (τ = 0.47 µs, AND latch at low level $0.35V_{CC}$ maximum according to [36]) according to the existing documentation. In the short-circuit detection's output stage, an exclusive or (XOR) port [37] takes the D_{in} signal and compares it with the error signal and sets the output optical status accordingly. The delay circuit and the output stage are shown in Fig. 18.

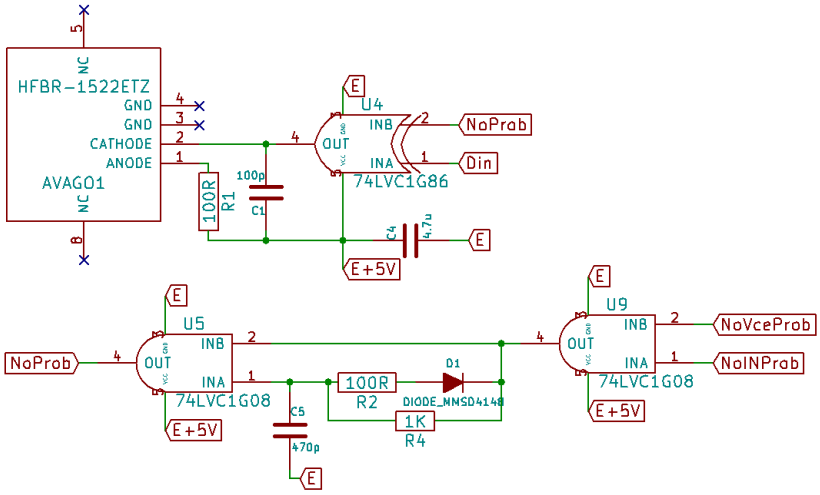

Figure 18: Error signal delay circuit (bottom) and the subsequent output stage (top) of the short-circuit detection module. When both inputs to the XOR port are either LOW or HIGH, the output is connected to ground which provides a path for the current to ground via the transmitter, lighting up the LED.

The logical complexity of the circuit along with the almost non-existant documentation make the short-circuit detection scheme quite hard to follow, and it also feels that the detection might not necessarily work as it is supposed to. For instance, if the optical input is OFF, the D_{in} signal is HIGH, and if at the same time the V_{CE} voltage was over the threshold that flags it high (i.e., the subsequent OR input set to low), the problem signal would be triggered and optical output set to OFF (the same phase as the input, indicating an error). But optical input OFF and high V_{CE} combination should be a valid IGBT state. On the other hand, if the optical input signal is OFF and at the same time V_{CE} is low meaning that the IGBT is in a conducting state, the resulting optical output signal is ON, indicating that everything is working as intended. But if there is no input signal the IGBT should not be in a conducting state. This function was also confirmed in a bench-top measurement where the collector port was left floating (essentially connected to the emitter via the HVD's resistors) and optical input was sent in a pulsed manner to the controller card which always returned the optical signal in the opposite phase even though the collector–emitter voltage was low and constant. The high-voltage behavior of the controller card was not tested in a bench-top measurement due to difficulty and limited amount of controller cards available (they were needed in the actual system). It should be noted that, as the documentation for the logic behind the system is quite incomplete, an erroneous assumption at some point during the circuit analysis is a distinct possibility.

The analysis system was tested on one controller card in the actual TMS system, on which the negative comparation side of the overvoltage detector was fixed with a setup depicted in Fig. 19. An Avago receiver was set to receive the optical signal from the controller card and the response was monitored with an oscilloscope along with the current waveform measured from the coil. In the test

setup, the receiver stage was powered with a 9-V battery and the current path was set to go through an LED to give a rough indication whether the battery had a sufficient voltage level. When no optical input is received, the measurement point (MP in Fig. 19) voltage level is high (1.95 V and upon receiving an optical signal, the voltage level drops to approximately 225 mV.

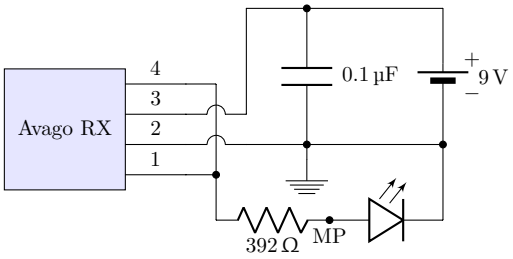

Figure 19: The test setup for the reception of the optical signal from the controller card's short-circuit detection. The receiver pins are labeled as in the data sheet, and a 9-V battery was used to power the receiver stage. Without optical input, the LED stays on, but with optical input, the current path is redirected to ground, turning the LED off. MP indicates the measurement point from which the response was measured with an oscilloscope. The scheme was adapted from [25].

3.3 Coil ID and temperature

To avoid user error, a method to identify digitally between coils in an mTMS device would be highly beneficial. A digital ID could be assigned to each coil, and this kind of recognition would allow the system to know which coil is in which channel.

DS18B20 (Maxim Integrated, Inc., San Jose, CA, USA) is a digital thermometer, which can operate on a single communication pin via 1-Wire interface by utilizing serial communication. The sensor looks like the typical transistor in a TO-92 package[23], and it has onboard memory in the form of read-only memory (ROM), which stores a unique 64-bit serial code for every device. This serial code could act as a unique identifier for each TMS coil. As many as 256 1-Wire devices can be placed in parallel in the same bus, so the polling circuit topology can be kept simple. An additional benefit is the possibility of getting a rough estimate for the coil temperature, which could be monitored to ensure operation within safety limits. [38]

The temperature sensor was tested by prototyping a test setup with an Arduino-like microcontroller unit (MCU) which was programmed to poll the DS18B20 every two seconds for the temperature [24]. As reading the temperature levels from a computer's console via a serial link between the MCU and a computer is tedious to say the least, a 128×64 resolution organic light emitting diode (OLED) display was attached to the setup via inter-integrated circuit (I^2C)

[23] Although other packaging options are also available, see [38].

serial bus, which allows important information to be readily available during the testing. The prototype setup is shown in Fig. 20.

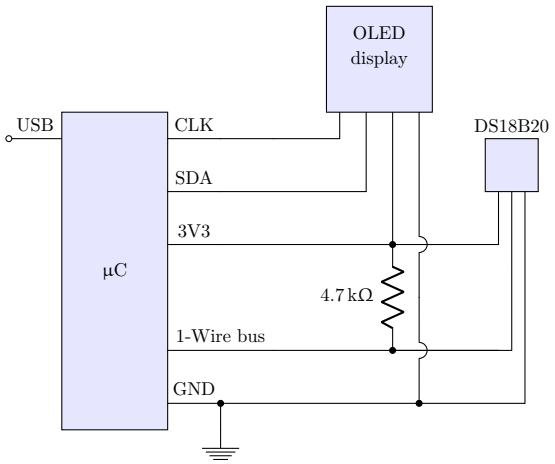

Figure 20: Prototyped test setup for the DS18B20. The microcontroller gets its power from a USB power source and provides a 3.3-V line for the peripherals. Two bus lines labeled SDA (serial data) and 1-Wire bus are used to serially communicate with the temperature sensor and the display. CLK denotes the reference clock passed on to the display. [35, ch. 9][38]

1-Wire devices can be searched from the bus and polled after they report their unique ID (i.e., their address [38]), but the program was configured to poll a hard-coded address to also test that the address does not change due to the strong magnetic fields present near the sensor. In the test setup, the sensor was placed approximately 0.5 cm under the middle point of a figure-of-eight coil, i.e., in the coil's hotspot. The connection to the MCU was done via a twisted triple-wire cable to minimize magnetic flux coupling and the test performed by giving TMS pulses at maximum intensity.

3.4 Effective voltage control

Next, let us examine the methodology of producing the proposed *simulated* voltage levels with the mTMS device.

3.4.1 Ideal model

For now the aim is to approximate a trapezoidal[24] pulse waveform of a given initial condition (voltage and current) and given durations for the rising phase, the peak hold period and the falling phase. When the initial voltage of our system is suitably higher than that of the reference, the slope of the RLC phase is markedly steeper and the current amplitude rises very quickly past that of the reference. At some point the current has surpassed that of the reference by

[24] See Section 2.2.

some margin, and to keep the approximation feasible, we need to switch to the RL phase where the current slowly decays over time until we are again below the reference current. At this point the RLC circuit can be activated again. The two subsequent circuit states consisting of an RLC state and followed by an RL state is henceforth referred to as a *step* in this thesis. Repeating these steps we can effectively replicate the slope of the reference waveform with different initial conditions, and the method is here called *stepped approximation*. This approach can be used for both the rising and the falling phases, whereas the middle holding period can be replicated as is. The stepping procedure is essentially performed backwards for the falling phase, as the RLC circuit state redirects the current to the capacitor, thereby decreasing the waveform amplitude. Figure 21 shows a stepped approximation for a trapezoidal waveform.

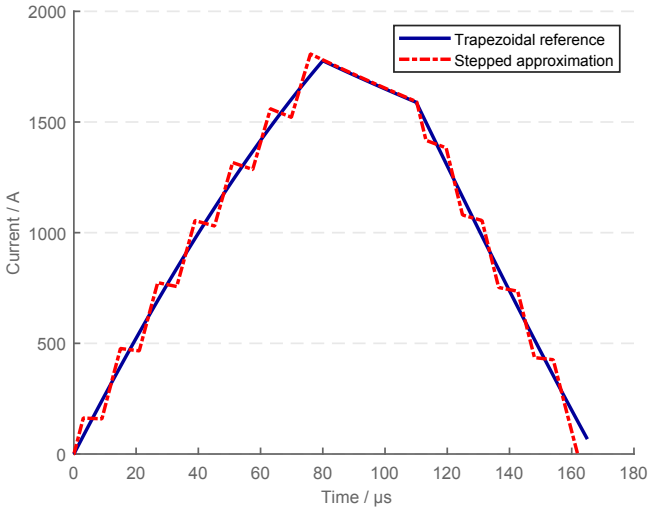

Figure 21: Stepped approximation of a trapezoidal current waveform. By switching between the RLC and the RL states of the circuit, we can manipulate the waveform and approximate other shapes than the basic exponentially decaying sinusoid of an RLC circuit.

The switching of the IGBTs (i.e., the circuit states) in our device is done by communicating a *timing vector* to the FPGA. This timing vector is essentially a vector of pairs that tells the microcontroller which IGBTs are switched into their conductive states and for how long.

The Labview (National Instruments, Corp., Austin, TX, USA) environment used to control the system can be integrated with Matlab, and this offers the possibility of developing programs for various functionalities such as customized, on-the-fly calculated timing vectors for IGBT switching. To study the waveform behavior during an RLC circuit state, Eq. 7 was analytically solved with Matlab. The solution was also verified with Maple and turned into functions that could be used in calculating precise waveform approximations. The current and charge behavior during the different circuit states was also verified with LTSpice simulations, as the analytical functions are quite messy in Matlab format. The behavior of current is well known for the RL circuit state, as mentioned before;

see Eq. 12.

The trapezoidal reference's behavior of current is calculated as a single, combined vector of all the three phases, i.e., all the values of current during the pulse's different phases are appended to a single vector. For each individual step of the waveform approximation, an RLC and an RL vector is calculated according to the analytical solutions discussed in the previous paragraph. These vectors are iterated and referenced to the trapezoid's vector value at the same time instant until a sufficient pair of time and value of current is found. During an RL hold period, the capacitor's charge level is maintained, as it is disconnected from the circuit. When a suitable value is found, the phase (i.e., the circuit state) is switched, a new vector of values calculated based on what the new initial conditions are (initial charge and current have changed which calls for the full analytical solution of the differential equation (Eq. 7)), and the process repeated. Simply put, the RLC phase takes us from below the reference to above the reference. Then, the circuit state is switched, and the RL phase takes us from above the reference to below it.

One limitation for the algorithm is the minimum on- and off-time for the IGBTs, which is estimated to be approximately 2 µs, so each circuit state (i.e., half step) must be at least this long. The algorithm developed also has a minimum stray distance (in amperes) from the reference to prevent biasing[25] the approximation towards one side of the reference.

The best approximation for any waveform would be the waveform itself, and in the case of the peak hold period this is possible. It is, however, unlikely that we end up at the exact corner point with the approximation where the trapezoidal waveform changes from the RLC phase to the RL peak hold period. This has been solved by back projecting the reference's RL decay current from the cornerpoint and by looking for a collision with either the decay current or the back projection. The RL decay's current behavior follows Eq. 12, and as the capacitor charge is not involved, the back-projected current is simply calculated by

$$i_{\mathrm{b}} = i(0)e^{-\frac{-T}{\tau}} = i(0)e^{T/\tau} \tag{25}$$

where $-T$ is the time back from the corner point, $i(0)$ is the current at the corner point and τ is the time constant of the circuit. In this algorithm, the last step calculated is periodically compared against the RLC/RL corner point of the reference. Simple logic is used to determine whether it is feasible to start a new RLC rising step which will collide with the back-projected current, or if we should backtrack the previous RL period and instead continue the RLC step before that. In either case, we end up with an RLC rise which will at some point surpass either the back-projected current or the actual RL-decay current. The collision point's distance from the actual corner point is saved and added to the reference's peak hold duration and by this addition we will end up with a peak-hold period of similar[26] length—a very feasible approximation. A solution

[25] This biasing, or asymmetric weighing, is the result of the minimum on- and off-times for the IGBTs, magnitude of the selected minimum stray distance and the difference in the slopes of the approximation and the approximatee.

[26] But not exactly the same.

might not be found this way if the last step is below the minimum step length, so the final rising step duration is inspected as well. If this duration is too short, the approximation's starting point is delayed for a short time period and the procedure is repeated. A default delay value of 0.5 µs was set, as it seemed to produce good results, but this is user configurable.

After the peak hold period, the falling phase approximation is done in a similar fashion to the rising phase approximation, and if a solution is not found, the preceding peak hold period is extended by some value, defaulting to 0.5 µs similarly to the rising phase.

3.4.2 Advanced model

The performed experiments with the mTMS device indicate that when the delay between circuit-state switches is short (i.e., less than 15 µs), the resulting waveform is quite different compared to the analytical solutions for the behavior of current. This is caused by stray inductances, which were inspected with SPICE simulations (see Section 3.4.3), and the snubber circuits.

The stray inductances are largely unknown for the moment, and numerous configurations could result in the behavior experienced, so it seems like a good solution would be a function that acts as a link between experimental data and the ideal model. The function would take the ideal model's timing vector as input and modify it in such a way that when the modified vector is sent to the FPGA, the coil current would follow our approximated ideal waveform. Because both the ideal waveform and the real waveform follow largely the same fundamental rules, they can be closely matched by manipulating the switching times of our circuit. The problem becomes an inverse problem where we have a measured waveform with some parameters, and we want to match it to our expected ideal waveform by denoting some Δt_i for each corner point[27]. This value indicates how much backward or forward the corner needs to be moved for the waveforms to match.

There is a complication, however, that in the pure lab data the waveforms are quite distorted, and cannot be reproduced with the ideal RL/RLC equations. Taking a closer look at the microscale variations of the actual pulse with high-resolution oscilloscope sweeps and an investigation of the SPICE model (see Section 3.4.3) reveal that various combinations of stray inductances can cause these modifications. The non-idealities observed can be seen in Fig. 22. The rounding effect observed on the first step, however, is not due to stray inductances, it is caused by the snubber circuits. Similar non-idealities can be observed on the falling phase, which seems to have a severe rounding effect and a relative increase in hold-step length as the current level approaches 0. It should also be noted that in experimental data the current level can not go below 0 as the direction of current is limited by the diodes. The non-idealities are shown in Fig. 22.

[27] Circuit state switch between RLC and RL.

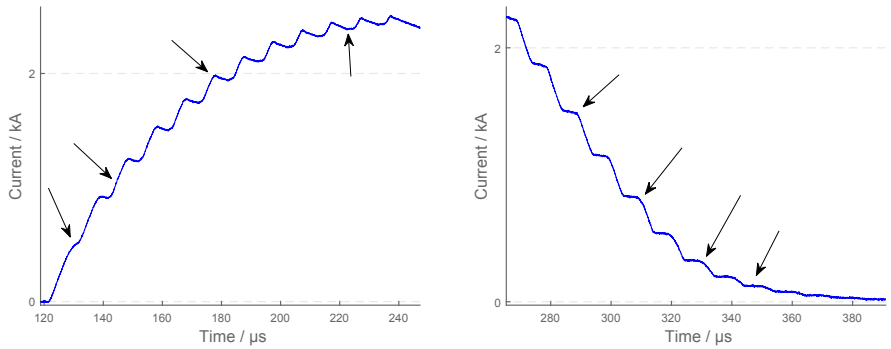

Figure 22: High-resolution measurement of a stepped waveform. Left: Rising phase. Non-idealities include snubber rounding (first step), RLC rise distortion (second arrow from left), current spiking during transistor switch (third arrow) and additional, evolving angles (fourth arrow). Right: Falling phase. Non-idealities during this phase include heavy, progressive rounding due to snubber circuits (first three arrows from left) and minor distortions (last arrow, poorly visible at this scale).

To parametrize the measurement, we need to find a closely matching waveform that can be generated with our equations; the distortions are replaced with modified step durations in such a way that the modified waveform stays true to the one measured. This was done by creating a semi-automated script that analyzes the measured signal, finds the "reproducible" corner points and reconstructs the measured waveform with the new timing values found from these corner points. This reconstructed signal is the waveform that is used to obtain the values for the parameter data; for our purposes the reconstructed waveform *is* the measured waveform to a close-enough degree.

For the rising phase, the idea is to first pass the *inverted*, high-pass filtered signal through a peak finder to find all the troughs. The peak finder simply tries to find the highest data point surrounded by lower points on both sides in a deterministic way. The troughs indicate a general neighborhood where an RL to RLC switch has occured, but not necessarily the exact location, see Fig. 23. The black triangles represent the loci found by the inverted peak finder, and, as can be seen, instead of the expected sharp corner after it, non-idealities cause the waveform to have two wider corners before the expected rising phase. Also a current overshoot which masks the actual circuit switching point can be observed after each rising phase; this is also caused by the circuit non-idealities.

As the durations of the steps are very small compared to the RLC/RL time constant, the behavior of the current is almost linear, and we can use this to our advantage. From the found point we can move a fraction of the holding period back, select a small segment (also a fraction of the holding period) of the data points and create a linear fit. Next, we move a fraction of a rising step forward from the trough point and similarly select a fraction of those data points, calculating another linear fit. From these linear fit functions' crossing point we have found a reproducible RL to RLC switching point, and the procedure can be seen in Fig. 23. It should be noted that the inverted peak finder actually

does (generally) find exactly the real switching point, but if this was used for the correction model, the approximations would be grossly oversized due to the model's inability to produce the non-ideality effects that limit the current growth.

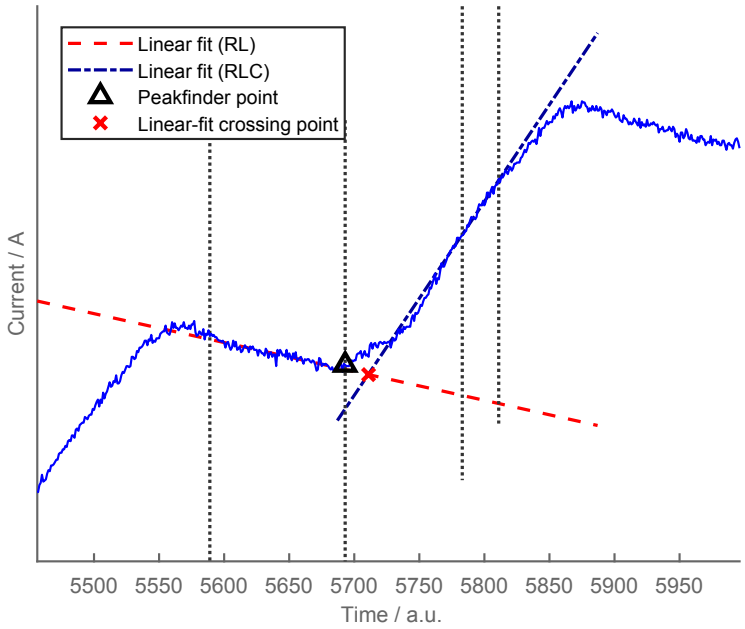

Figure 23: Selecting sections of the signal that do not include distortions allows us to find a linear fit crossing point which represents a hypothetical switching point from the RL circuit to the RLC circuit.

From this crossing point we can use the previously discussed current back-projection (Eq. 25) to calculate the RL decay backwards, and from a *previous* linear fit crossing point we can calculate an RLC rise forward. The crossing point of these two lines is our reproduced RLC to RL switching point of the measured signal. Of course we still need an initial starting point for this process, but that is simply the starting point of the waveform, found by an edge detector. This part of the process is depicted in Fig. 24. The last rising step's holding period does not have a corner point, so the analysis is performed on the first $n-1$ steps where n is the number of steps.

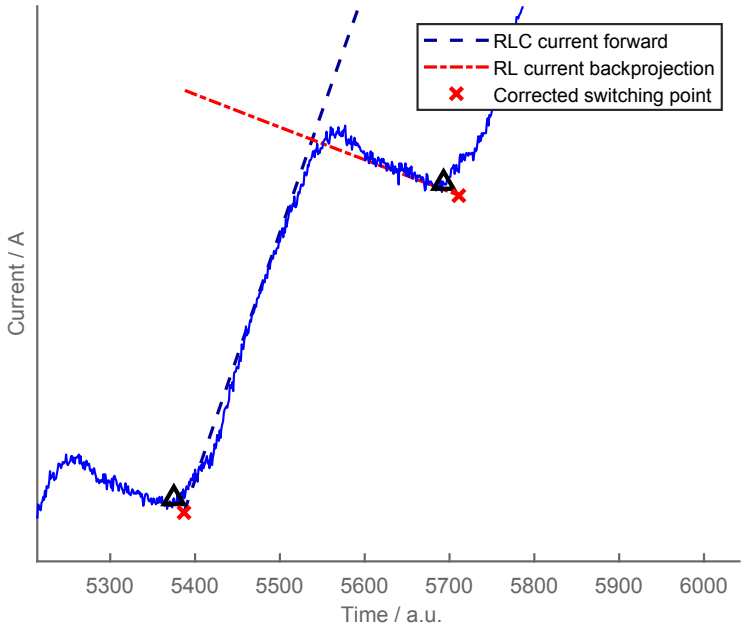

Figure 24: Between the linear fit crossing point (see Fig. 23) and a previously found crossing point is another switching point. This can be found by calculating the crossing point of an RL back-projection and an RLC rising phase.

Figure 25 shows the originally measured signal, the reconstructed signal based on the process described previously and the waveform from the ideal model.

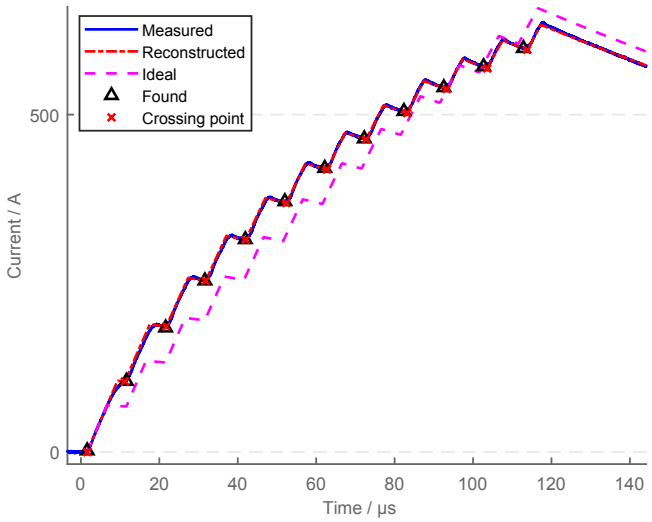

Figure 25: A reproduction of the measured signal's rising phase compared against the measured signal and the ideal model output. The triangles and the crosses represent the peak finder points and the linear fit crossing points, respectively.

With short steps the first step's rounding effect in the waveform (see Section 2.2.2) prevents the subsequent RL to RLC switching point from being found with a peak finder, but observations show that the total duration of the step and the subsequent holding period is almost the same as the values in an ideal system. This way we can set an initial guess manually at this location and process the step like the other steps.

The next step is to link the reproduced waveform (i.e., the measured waveform) with the ideal model output. This can be done by parametrizing the reproduction and representing it as a system of equations. The link to the ideal model is a vector of Δt_is that represent how much each step must be lengthened or shortened for the reproduced waveform to match the ideal one. In this case, the matrix representation is

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{26}$$

where $\mathbf{A}$ is the parameter matrix for the waveform, $\mathbf{x}$ is a vector of unknowns and $\mathbf{b}$ is the vector of Δt_is. The system of equations was selected as

$$\begin{bmatrix} I_1 & n_1 & V_{\text{in},1} & \tau_{\text{R/H},1} & t_{\text{RLC,tot},1} & t_{\text{RL,tot},1} & 1 \\ I_2 & n_2 & V_{\text{in},2} & \tau_{\text{R/H},2} & t_{\text{RLC,tot},2} & t_{\text{RL,tot},2} & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ I_k & n_k & V_{\text{in},k} & \tau_{\text{R/H},k} & t_{\text{RLC,tot},k} & t_{\text{RL,tot},k} & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_7 \end{bmatrix} = \begin{bmatrix} \Delta t_1 \\ \Delta t_2 \\ \vdots \\ \Delta t_k \end{bmatrix}$$

where I is the measured current, n is the current step, V_{in} is the initial voltage, $\tau_{\text{R/H}}$ is the duration of the step, t_{RLC} is the cumulative duration the RLC circuit has been on before the step, t_{RL} is the cumulative duration the RL circuit has been on before the step and the last column is a constant offset term. In the actual implementation, the matrix was further divided into two submatrices, one of which contained the parametrization of the rising phases of each step and the other matrix containing the parameters for the holding periods of the steps.

For the falling phase a similar parameter matrix (also divided to separate submatrices for the falling phases of the steps and the related holding periods) was constructed, and a representation of the waveform reproducible with the ideal model's equations was calculated. Here, the rounding of the signal due to the snubber circuits is quite severe and causes heavy distortion resembling an exponential decay from macroscale viewpoint, shown in Fig. 25. These rounding effects cannot be reproduced with our equations, so we are more interested in doing a somewhat close approximation with a deterministic error, i.e., estimate either too far or too short. Peak finder could not be utilized as is, but multiplying the signal with a normalized time exponential

$$i_{\text{ref}}(t) = i(t)e^{3t/t_{\max}} \tag{27}$$

gives us a good sharpening effect to find peaks from after the signal is high-pass filtered. This gives us the locations on the actual signal as the corner points do not move on the time axis. The exponential also amplifies the noise levels, but

as we know the number of steps for this analysis beforehand, we can find the last actual switching point and cut out the rest of the signal from the analysis. The end of the last holding period cannot be found as it is part of the noise, so the matrix parameters are gathered for the first $n-1$ steps where n is the number of steps. The falling-phase current and the modified peak-finder reference are shown in Fig. 26.

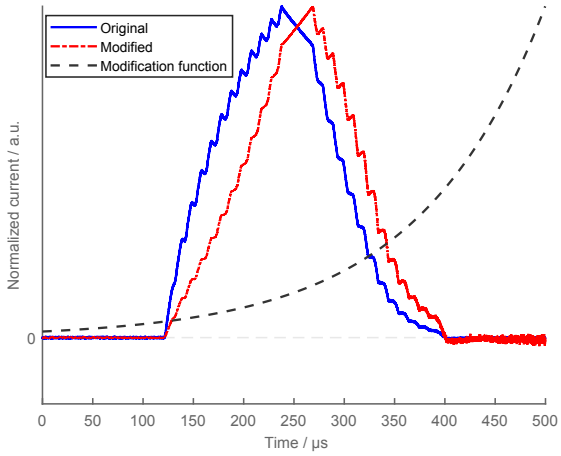

Figure 26: Modification process. Measured current is shown in blue and the modifying function in black. The resulting reference waveform passed to the peak finder is graphed in red.

Similarly to the rising-phase analysis, the measured falling-phase waveform is reconstructed with the ideal model's equations, aiming for as close an approximation as possible with the system. The heavy rounding effect present in the later steps poses a serious problem for the analysis, but the main idea is to translate the macroscale behavior of the system, i.e., the long trailing tail, to our timing vector modifications. In the end we are not actually interested in reproducing such a trailing tail, we are interested in bringing the current to 0, closely following the falling phase of a trapezoidal waveform. As this is the case, the corner points found are noted as the RL to RLC switches, and from each corner point an RL decay is back-projected and from each previous corner point found an RLC current forward is calculated. From these graphs a crossing point is calculated which indicates an RLC to RL crossing point at a certain time point with some current value. Due to the rounding effect the estimate grows very inaccurate as the pulse series goes on, but the macroscale behavior stays true to the measured signal: RL holds get longer whereas RLC steps get significantly shorter. The peak finder again automatically finds the first point needed to perform the analysis for the subsequent points as the first corner point's current and time are the same for the experimental signal and the approximation. Figure 27 shows the approximation found for the measured signal, and this is used to obtain the parameters needed for the falling-phase matrix following the form of Eq. 26.

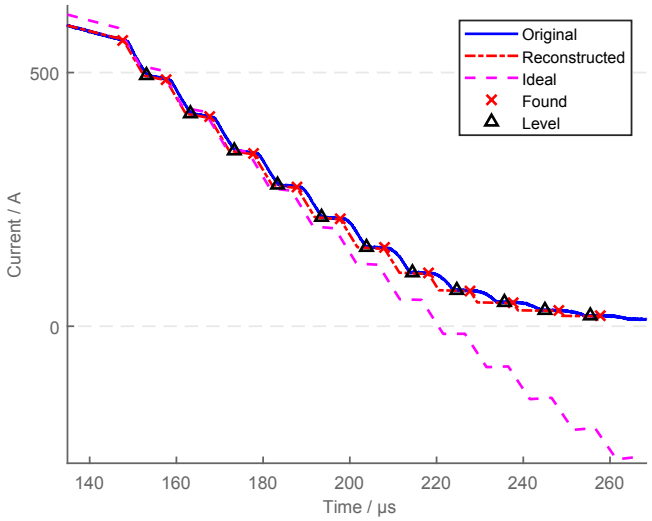

Figure 27: A reproduction (red) of the measured signal compared to the measured signal (blue) and the ideal model output (magenta). As can be seen, the approximation is very inaccurate as the pulse series continues, but the exponential decay-ish macroscale behavior stays similar. The red crosses indicate the points found by the peak finder and the black triangles indicate the current level after a previous falling step.

The found corner points are again referenced to the corner points obtained from the ideal model using the same timing vector, and thus another set of Δt_is is obtained based on how much and in which direction the ideal model's corner points need to be moved for the waveforms to match. The peak hold period was also parametrized into its own matrix, increasing the total number of parameter matrices to five.

Now that we have the parameter matrices ($\mathbf{A}$) and the changes in step durations ($\mathbf{b}$) for both rising and falling phases (and the peak hold period), we can solve the unknown variables in Eq. 26:

$$\mathbf{Ax} = \mathbf{b} \Longleftrightarrow \mathbf{x} = \mathbf{A}^{-1}\mathbf{b} \ , \tag{28}$$

where vector $\mathbf{x}$ will now modify our experimental timing vector in such a way that the original input vector is reconstructed. When we obtain a timing vector from the ideal model, this represents the final waveform we want, and to obtain this waveform we need to calculate what ideal waveform distorts to this shape in the system with Eq. 26, which can be done now with the known $\mathbf{x}$. Figure 28 (left) shows an ideal waveform and the measured signal that is obtained when the same timing vector is sent to the FPGA. Figure 28 (right) shows the reconstruction of the measured waveform with the process discussed in this section and a back-prediction of what the original input was, based on the solution vector obtained from the parametrization of the reconstructed waveform.

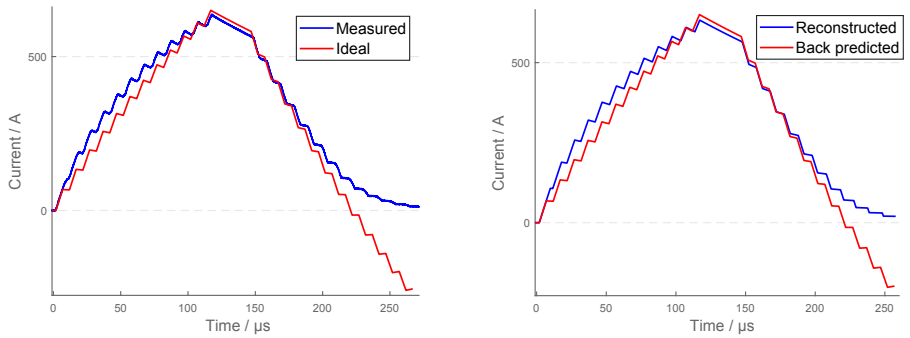

Figure 28: Back-prediction process. Left: Ideal waveform distorts in the actual system due to non-idealities. Right: Reconstruction of the measured waveform and subsequent parametrization allows for back-predicting the original input.

It should be noted that the approximated step lengths of the reconstructions are erroneous, but as we cannot replicate the rounding effects or other distortions with our model, error is going to be introduced one way or the other. By keeping the system (and the error introduced) deterministic, i.e., always calculate the approximation in the same way, we should still get a decent model for reproducing the measured waveform.

For the correction system a set of measurements consisting of 168 different waveforms were made, each measurement resulting in five parameter matrices (2 for the rising phase, 2 for the falling phase, 1 for the peak hold period) as described before.

The waveforms were generated by systematically going through varying step sizes (rising/falling phase of each step 4–20 µs, holding period 4–10 µs) with varying number of steps (2–14). The step size was kept constant during each waveform, and the rising phase of the waveform utilized the same step size as the falling phase. Each waveform was generated with initial capacitor voltages of 250, 500, 750, 1000, 1250 and 1400 V. The peak hold period was kept at 30 µs.

The waveforms were captured with a CWT6 Rogowski current waveform transducer (Power Electronic Measurements, Ltd., San Carlos, CA, USA) connected to an oscilloscope (Agilent MSO-X 3024A). The original timing vector was also saved and the captured waveform characterized with the semi-automatic analysis system described before. The resulting parameter matrices give general solution vectors for the expected behavior when solved for unknowns, and the individual waveform analysis data are also saved.

3.4.3 SPICE model

A series of SPICE models were also created with LTSpice XVII to understand better the behavior of the system. The main advantage of SPICE is the relative ease of modeling even complex circuits and, more importantly, making the implementation of large changes in circuit topology fast and easy.

For the model the active components were modeled with ideal components—the IGBTs were modeled with voltage-controlled switches with close-to-ideal properties. The on-state resistance was set to a few mΩ, off-state resistance to several hundred MΩ, and the switching time to 0.1 µs. The freewheeling diodes were modeled with a custom diode model loosely based on the diode specifications. The exact values are not that important, as we are more interested in the general behavior of the system and whether the actual system functions as expected. Snubber circuits were added as in Fig. 11 over the IGBTs, and a stray inductance component was added between each major node. The schematic of the circuit is shown in Fig. 29, where the snubbers have been simplified for clarity.

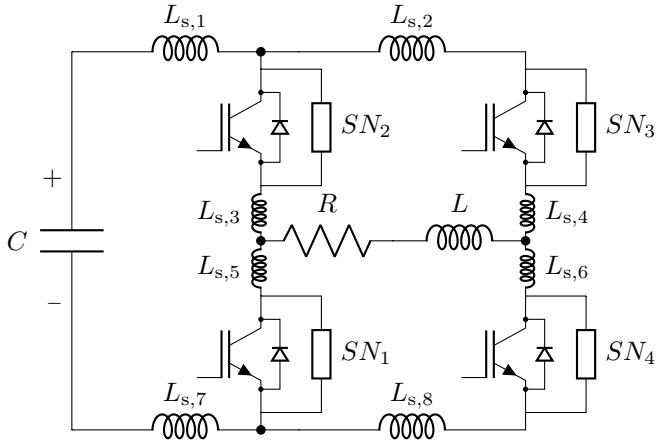

Figure 29: The SPICE model schematic (not actual program visuals) with all the relevant components for replicating experimental results. Inductors $L_{s,i}$ denote the various stray inductances in the circuit and impedances SN_j denote the snubber circuits. The FPGA is not included in this image; it is, however, part of the simulation and realized with pulsed voltage sources.

Even though the device has been designed for minimized inductance, there still seems to be inductive effects on the waveform when the current is passing through the coil. In the actual system the busbars connecting the transistors together and to the main capacitor are relatively long, and there are two approximately 20-cm-long cables connecting the coil to the coil cable. A quick measurement with an oscilloscope, a waveform generator and a series resistor showed that up to 0.2-µH stray inductance could be caused by the longer busbars. The value is only indicative at best, as the system was not (and could not) be operated as normal during the measurements, but it gave a starting point for evaluating the effects of stray inductances in the system.

The FPGA was modeled as a system of ideal voltage sources generating a pulsed square-wave sequence with a short risetime (set to 0.1 µs, although this is obviously changeable) and each IGBT gate receiving its own control signal as in the actual system.

Even though circuit analysis is fast with LTSpice, a non-linear time-domain

analysis[28] needs to be performed on the system with reactive components. This is done by numerically solving ordinary differential equations in SPICE, and it tends to be quite time consuming even with newer processors compared to our needs. Even with simpler circuits the calculation can take easily over 100 ms for mere 5 µs of circuit simulation. If longer processing times were feasible for our needs, a direct SPICE simulation integration with the system could provide a highly detailed model for the system's behavior—but for this detailed knowledge of the various stray inductances is necessary (i.e., their values and locations). The stray inductance values used for general simulations were iterated to

$$\begin{aligned} L_{\mathrm{s},1} &= 0.2\,\mu\mathrm{H}, \quad L_{\mathrm{s},2} = 0.05\,\mu\mathrm{H}, \\ L_{\mathrm{s},3} &= 0.07\,\mu\mathrm{H}, \quad L_{\mathrm{s},4} = 0.07\,\mu\mathrm{H}, \\ L_{\mathrm{s},5} &= 0.2\,\mu\mathrm{H}, \quad L_{\mathrm{s},6} = 0.2\,\mu\mathrm{H}, \\ L_{\mathrm{s},7} &= 0.1\,\mu\mathrm{H}, \quad L_{\mathrm{s},8} = 0.05\,\mu\mathrm{H} \end{aligned}$$

by comparing the behavior of the simulated waveforms to observed phenomena.

3.4.4 Advanced model—toolset

To test the advanced model, a set of tools were developed for Matlab. The tools include the actual Matlab program of the ideal model, a back-prediction submodule, a matrix manipulator module, an FPGA writeout tool and an oscilloscope interface utilizing the virtual instrument software architecture (VISA). Various plotting functionalities were also added for easier monitoring of the waveform during the approximation and modification processes.

The ideal model is the discussed Matlab model described in Section 3.4.1, and it takes a reference trapezoidal pulse as input, where the reference is described with a timing vector consisting of three (3) values: the durations of the rising phase, the peak hold period and the falling phase. In addition, the initial capacitor charge used to generate this pulse shape is given. To approximate this waveform, a certain set of parameters[29] are also needed, and these are the initial capacitor charge of the approximation, minimum pulse duration, minimum stray distance from reference and initial wait time. Other parameters are also possible in the model but were not touched during the lab session.

The initial capacitor charge for the approximation should be within reasonable limits compared to the trapezoidal reference's initial charge, as a huge relative difference will cause the approximation to overshoot the reference by a large margin and produce a suboptimal approximation. Similar suboptimality happens with too close charge values, as the difference in the waveforms' slopes is too similar to produce a feasible estimate. Minimum pulse duration was originally meant to characterize the IGBTs' minimum switching time, but it can be changed to alter the precision of the approximation. Minimum stray distance

[28] In other words, a transient analysis.
[29] Or constraints.

is the distance in amperes the stepped approximation has to at least be greater than (rising step) or less than (holding period) the reference at the same point in time before the circuit's state can be switched.

The back-prediction module tries to predict what the input to the system needs to be in order to obtain the waveform calculated with the ideal model. The module takes the approximated waveform's timing vector as input and processes it with the solution vector obtained from the parametrized measurements done before (see Section 3.4.2).

The matrix manipulator module's function is to generate combined matrices for the solution vector from the individual data matrices generated from previous measurements. These are the solution vectors that are used by the back-prediction module. All the measurements done previously are described by a few parameters which are saved in a file system individually, and this allows the oscilloscope data to be recalled and the ideal waveform counterpart to be reproduced and a parameter matrix to be generated. The individual parameter matrices can be combined in order to solve for a solution vector, and the matrices can also be augmented later after their generation by injecting more matrices from the measurements into the mix.

This matrix manipulator also has an experimental weighing system that injects the same parameter matrix over and over into the combined matrix and checks the back-prediction with this new matrix against the actual input until certain criterion is met. In this case, the criterion was set to be the peak waveform amplitude, which has to be inside certain thresholds (e.g., 1.5% maximum deviation from actual) for the function to move on to the next parameter matrix. When the weighing loop is run a few times, a general solution should (in theory) emerge in which the solution vector has converged to the *best fit for all input* if such a solution is possible. However, it needs to be remembered that as we inject the same data over and over into the parameter matrix, we are not adding any *new* data, and as such the solution vector might behave unpredictably when trying to predict anything else than waveforms similar to the learning data.

The FPGA writeout tool is mostly a convenience tool which takes the timing vector and a few other parameters and converts them to a format understandable by the Labview program that controls the FPGA. The oscilloscope interface is another convenience tool that utilizes the VISA interface and allows the oscilloscope to be used with the same computer that is controlling the TMS system.

With these tools it is possible to make an approximation of a trapezoidal waveform, generate a parameter matrix from the previous measurements or augment an existing matrix, calculate the solution vector from said matrix, write the control sequence of this back-predicted approximation for the Labview program, run the sequence, and capture the waveform with the oscilloscope. When the waveform is captured, it is plotted in a graph along with the original approximation and the back-predicted input.

In the experimental tests a 80/30/55 µs trapezoid was selected as the reference with an initial capacitor charge set to 0.25 C. The approximation settings were set to a charge level of 0.5 C with the step size chosen as 6/8/10 µs as these values

produce a decent approximation for the trapezoid, and the earlier measurements have a good ideal representation with similar stepsizes. The minimum drift was set to 10 A.

A value that *at some level* describes the goodness of our back-prediction is the relative difference between the calculated approximation's and the measured waveform's peak amplitude,

$$\text{Relative difference} = \frac{A_{\text{meas}} - A_{\text{approx}}}{A_{\text{approx}}} \cdot 100\% \tag{29}$$

where A_{meas} denotes the peak amplitude of the measured signal and A_{approx} is the peak amplitude of the waveform approximation we want to achieve. Of course the timeliness of these has to be considered as well—even a 0% deviation is not very good if the waveforms are dozens of microseconds apart.

4 Results

4.1 Snubber circuits

The results of the investigation of the snubber circuits performed with the SPICE model are visualized in Fig. 30. The simulated current through the snubber's resistor is plotted against the total coil current and the respective IGBT's control signal. This shows how the snubbing capacitors are discharged during the conductive state of the IGBT and recharged during the non-conductive states, effectively making the rising steps a bit longer than they should be. The "rounding" behavior (compare the blue line in Fig. 30 to the red control signal) is especially noticeable during the first step, which could be caused by the lower current (initially 0 A) through the coil and the fact that this current has to be continuous. The positive magenta exponential-decay curve shows the discharge current through the snubber circuit's resistor into the IGBT channel (compare against the red control signal), and the negative portion shows the snubber capacitors' recharge period. During the first step the recharge never finishes (the waveform only has the first half compared to the following waveforms), thus prolonging the actual rise period compared to the subsequent holding periods where the capacitors seem to recharge fully (see the 2nd and 3rd holding period). At higher current levels, the exponential growth preceding the snubber's charging state proceeds very rapidly, and the $1 - \exp(-t/\tau)$ waveform evolves almost directly to the characteristic $\exp(-t/\tau)$ shape of an RC circuit.

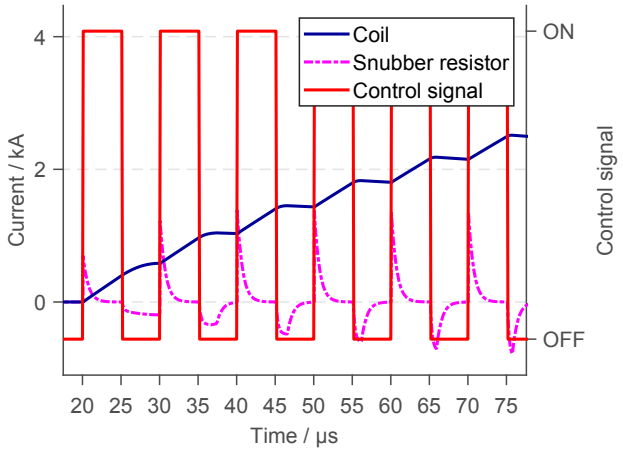

Figure 30: Simulated behavior of the current through the resistor in a snubber circuit of an IGBT that receives control signals. The charge of the snubber capacitors is released via the conductive path provided by the IGBT when it receives a control signal, and the snubber capacitors recharge when the transistor is in its non-conductive state.

4.2 Controller cards

Figure 31 shows the behavior of the gate–emitter voltage during a long pulse train with the pulse sequence shown at the bottom and the measured V_{GE} voltage shown at the top.

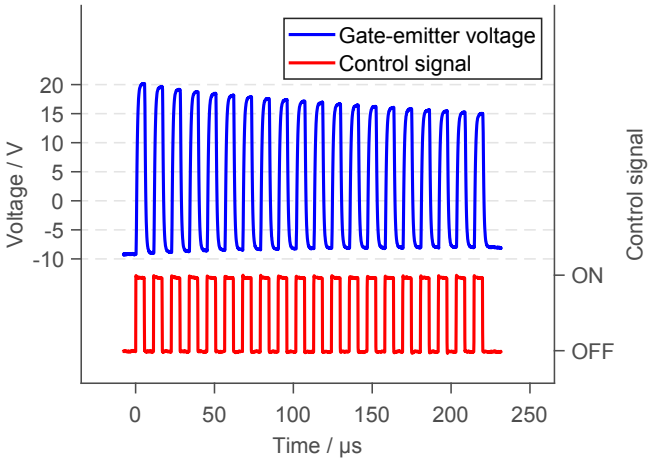

Figure 31: The voltage between the gate and the emitter terminals relative to the pulse train consisting of 20 separate pulses with a duration of approximately 5 µs.

As can be seen, the voltage level V_{GE} reaches during the on-period is dropping gradually when the pulse train continues. The maximal frequency of operation

for the controller cards was set in the initial requirements for the device to approximately 100 kHz with a maximum of 20 pulses, and during this time the voltage level drops to approximately 14–15 V, which should be sufficient to keep the IGBT switching normal. According to the IGBT datasheets, V_{GE} should be between 6 and 7 V minimum depending on the specimen [22, p. 2]. The effective frequency is approximately 100 kHz which exceeds the +20-V DC/DC converter's performance ([27, p. 2]) and could be (along with high current drain) thus the main culprit of causing voltage-level degeneration towards the end of the pulse train. Similar behavior can be seen on the negative voltage level during the pulse train where the negative level is set with another DC/DC converter.

4.2.1 Short-circuit detection

Laboratory tests with the mTMS device with the new controller cards in place indicate that at lower current levels the optical feedback signal from the short-circuit detector indeed changes states depending on the input signal and the voltage over the IGBT's collector and emitter terminal. This is shown in Fig. 32. The circuit seems to send the feedback signal as originally designed, i.e., the opposite phase compared to the input. However, upon close inspection the delay circuit does not seem to delay the signal switch at all, which begs the question whether the system is giving the correct feedback signal for the right reasons. The bench-top measurements do not correlate with these findings either, but it has to be remembered that the situation in the actual TMS device is a bit different.

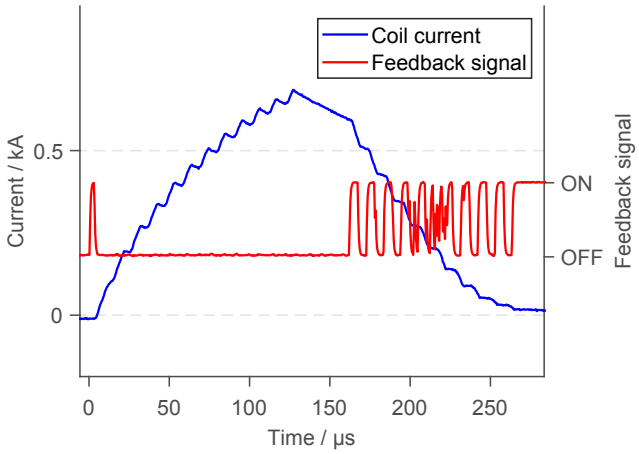

Figure 32: Feedback from the controller card. The status of the optical output seems to fluctuate heavily after the circuit states have been switched a few times.

Another problem is that when higher currents with longer sequences are driven through the IGBT, the output signal becomes quite scrambled, which can also be seen in Fig. 32. A possible reason for this is the +20-V DC/DC converter's output voltage which is known to drop below 16 V during long pulse trains (see Fig. 31). This would cause the over- and under-voltage comparation stage to

flag an error signal. As such the short-circuit detector does not seem very useful, as the device is intended to be run at up to 1.5-kV initial capacitor voltage and long pulse trains might be required.

4.3 Coil ID and temperature

Initial testing indicates that repeated TMS pulses at maximal intensity (∼1.4 kV initial capacitor voltage) do not alter the ID or functionality of the sensor. However, during and a brief moment after a pulse the temperature cannot be polled or the device does not return a value if it was polled just before the pulse. This is estimated to be due to the strong magnetic field which induces an electromotive force into the system's data line, thus corrupting the communication between the sensor's internal systems and also between the sensor and the MCU. This is especially noticeable if the sensor is set up to function at a 12-bit resolution, at which it takes 750 ms for the sensor to do the digital conversion and communicate it forward. Dropping the resolution to 9 bits offers a significant speedup in the conversion time where it only takes 94 ms and thus some responses can be obtained[30] from the device between moderately fast pulses (e.g., 2–3 s between pulses is plausible). The temperature values seemed to follow the experiment, although the actual values were not verified. The heat transfer from the coil to the outer surface of the coil former can be noticed with a small delay as expected, and the values seemed to be what would be expected of the former surroundings. This indicates that the sensor was still operating normally even after numerous pulses.

Figure 33 shows the prototype of the DS18B20 testing platform. Power is supplied by a universal serial bus (USB) power source to the MCU which in turn provides power for the attached devices (display and the sensor).

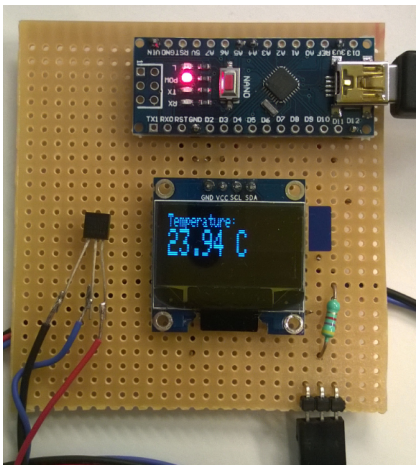

Figure 33: The prototype system polls the temperature sensor every 2 s and displays the temperature. Several error codes are displayed instead if the operation fails, courtesy of Maxim's sensor libraries.

[30] Technically values could be obtained every 100 ms, but for the moment the hard-coded polling rate was set to 2 s between requests.

4.4 Effective voltage control

4.4.1 Ideal model

It was observed that the model produces waveforms quite similar to experimental data when the sequences consist of long-duration pulses, and essentially identical results can be obtained with simple trapezoidal pulses by altering the circuit parameters a bit and accounting for the circuit switching times. With shorter pulse durations the difference to the ideal model grows drastically.

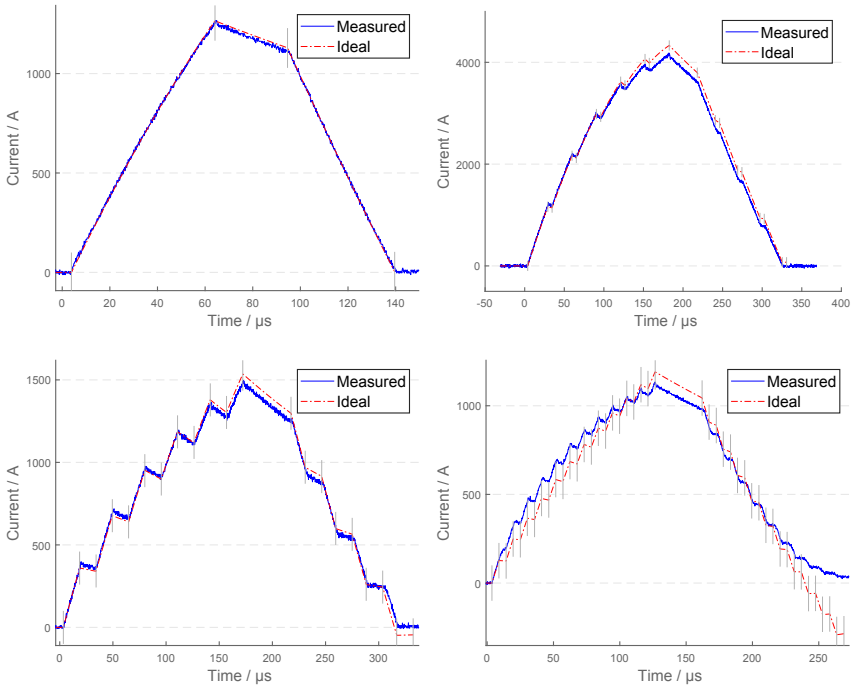

Figure 34: Waveform comparison between the ideal model (red lines) and experimental results (blue lines) obtained with identical timing parameters. Top left: Trapezoidal waveform. Top right: Long steps (25-µs rise, 15-µs hold). Bottom left: Shorter steps (15-µs rise and hold). Bottom right: Short steps (5-µs rise and hold). The gray lines indicate the switching times of the ideal waveform.

Figure 34 shows the behavior of current waveforms measured experimentally (blue line) compared to the analytical solutions used in the ideal stepped approximation model. The trapezoidal waveform is nearly identical and the long-pulse-duration waveform follows somewhat closely. The short-duration pulses with many steps, however, are quite different from the ideal model's results. Due to the current FPGA implementation, all circuit states are 50 ns longer than requested, and this is in line with the microcontroller's cycle time of 25 ns. The vertical gray lines indicate ideal switches in the stepped approximation model, and they correspond well with the measured waveform. While the results obtained so far look *decent* for the longer pulse duration, the shorter-pulse-duration

sequences deviate too much to provide a reliable approximation. This behavior is consistent regardless of the initial voltage used.

4.4.2 Advanced model

The first testing run was done with a solution vector calculated from general matrices that consist of all the individual parameter matrices from the previous 168 measurements. The minimum stepsize was set to 6 µs. Figure 35 shows the approximation (wanted waveform) in black, the back-predicted input in dashed red line (suggested input), and the actual measured result from the oscilloscope in blue (measured). Left side of the figure shows the reference waveform and its stepped approximation for illustrative purposes.

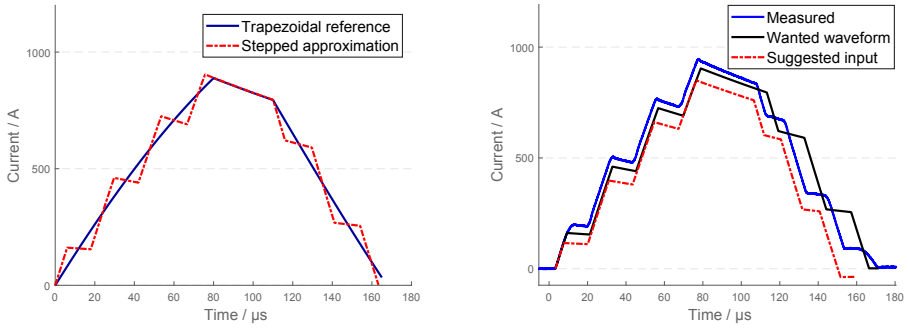

Figure 35: Stepped approximation and input prediction for achieving it. Left: Reference waveform and its stepped approximation. Right: General matrix prediction for 6-µs-minimum-stepsize waveform.

As can be seen in Fig. 35, the estimation is definitely in the right direction with a relative difference of 4.6 % (i.e., the measured response at its peak is 4.6 % higher than expected), but the individual phases of the steps are either a bit too short (rising phases) or too long (holding periods) to produce an accurate end result. There are at least two main reasons for this behavior. First, the parameter matrices are very general, and they have information from a wide range of waveforms where, e.g., the stepsize ranges from 4 to 20 µs. Second, where the original parameter matrices have a uniform stepsize (e.g., 5/5 µs) throughout the waveform, the approximated pulse shape now has a minimum stepsize on the first step, but because of the overshoot this causes, all the subsequent steps have markedly longer durations. Essentially, the original matrices lack the information on how this varying step size might affect the step behavior. The peak hold period on the graph seems to be a bit too long, but is actually reasonably predicted with the information we have (all the original measurements for the parameter matrices were done with a peak hold period of 30 µs), and only seems off due to the misprediction of the rising-phase steps. The falling phase is badly predicted due to lack of information in the original parameter matrices as well, but the main idea in this phase and at this stage of system development is to drop the current level to 0 (at least for now) instead of leaving large currents circulating in the system.

Increasing the minimum step size of the model seems to have a positive effect on the accuracy of the back-prediction which is expected, as it was observed before that with growing step size the relative error between the ideal model and measured response gets smaller. Results for the back-prediction of 8- and 10-µs-step-size approximations are shown in Fig. 36.

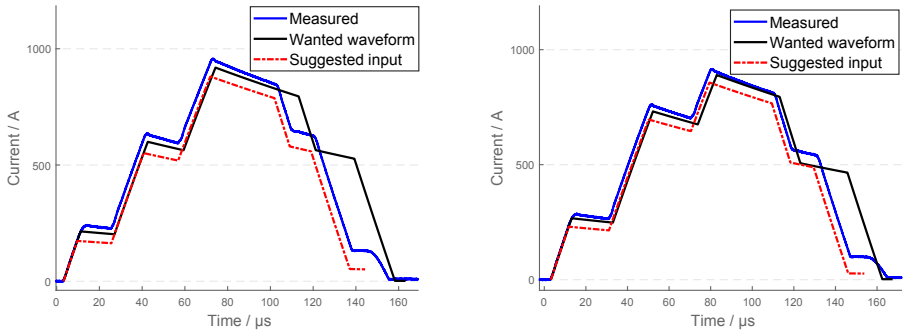

Figure 36: Input predicions for two approximations. Left: Minimum step size set to 8 µs produces an approximation with a relative difference of 4.2%. Right: Increasing the minimum step size to 10 µs produces an even more accurate waveform replicate with a relative difference of 2.9%.

As can be seen, the relative difference gets smaller as the step size gets closer to the 10-µs region. The prediction gets progressively more easy after the snubber rounding effect is not so prevalent, which is one reason for the better prediction accuracy, but another reason could also be that we have more surrounding matrix data for the 8- and especially the 10-µs step sizes. For the 8-µs step size the relative difference is 4.2 % and for the 10-µs step size the difference has dropped to 2.9 %. The left side graph of Fig. 36 also shows a significantly different duration for the peak hold period, which is caused by the initial parameter matrices only containing data for a 30-µs-duration peak hold. The falling phase is also quite different from the approximation due to insufficient parameter data. From the graphs it can be seen that simple reductions or additions to the step durations are sufficient to reproduce more or less an exact replica of the approximated waveform.

As the first step's rise time is highly different from the subsequent steps, the next step was to tell the approximation model to wait a bit before starting the approximation, causing the rising step duration to be more uniform between steps. The wait time necessary to produce a feasible uniformity depends on various parameters, so there is no single best answer for a good value for this. Values for the wait time used in testing were 2.5 µs for the 6-µs steps and 4.5 µs for the longer steps. Figure 37 shows the waveform approximations, predicted input and the measured results from the tests.

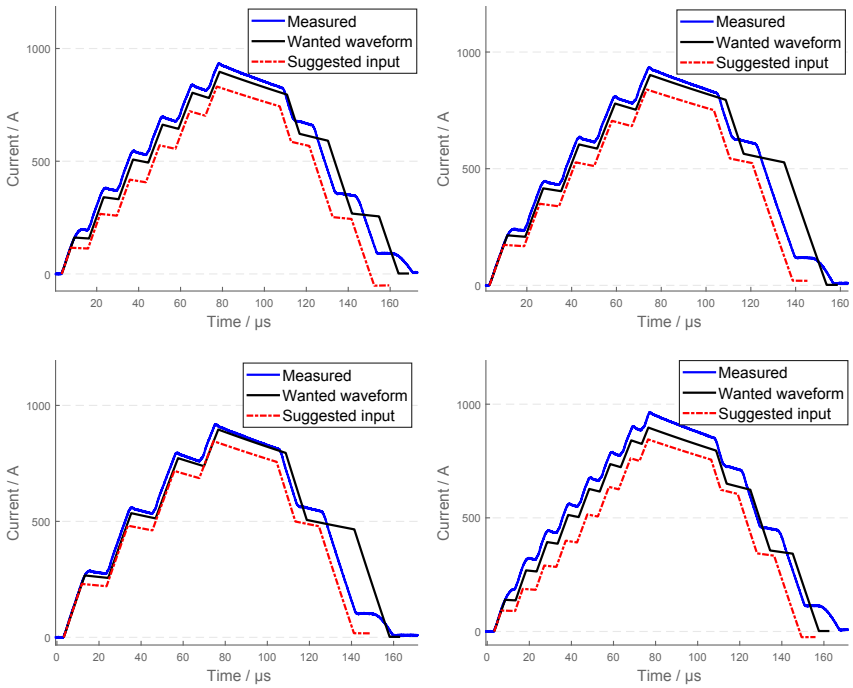

Figure 37: Different-step-size predictions with the solution vector from the general matrix. Top left: 6 µs duration of steps' rising phase, 4.16 % relative difference. Top right: 8 µs, 3.7 %. Bottom left: 10 µs, 2.7 %. Bottom right: 5 µs, 7.7 %.

As can be seen, the results follow a similar pattern to what was measured before without the initial delay before starting the approximation, albeit the relative difference is some fraction of a percent smaller (i.e., better in this case). The 6-µs step size approximation has a relative difference of 4.16 % compared to the 8-µs approximation's corresponding 3.7 %. The 10-µs step size approximation's relative difference is only 2.7 %. The high nonlinearity of the circuit results in the 5-µs step size approximation to be quite poorly predicted, and in this case the relative difference is little less than 7.7 %. The observed variance between runs was approximately 0.2 %.

The next step was to select a sufficient subset of matrices for more localized prediction instead of using a general matrix that might have irrelevant data for our needs. A 6/6 µs stepped approximation (i.e., 6-µs rising phases followed by 6-µs holding periods) was calculated with the ideal model, and this time the model was told to start the approximation with a holding period (duration of 4.5 µs) which causes the approximation to have more consistent step durations with minimal variation (instead of the short first step and longer subsequent steps as before). The minimum drift was also increased to 20 A.

For the 6/6 µs stepping the subset of parameter matrices were chosen from the surrounding sets: a combination of 5/5 µs and 7/5 µs parameter matrices. The initial results from this can be seen in Fig. 38 (left), and as can be seen the

prediction is quite bad. This is caused by a lack of information about the effect of variations in the hold step periods, and the matrix was augmented with this missing information by injecting parameter matrices from 5/4 µs and 5/8 µs measurements. The resulting waveforms from the augmented matrix are shown on the right side of Fig. 38.

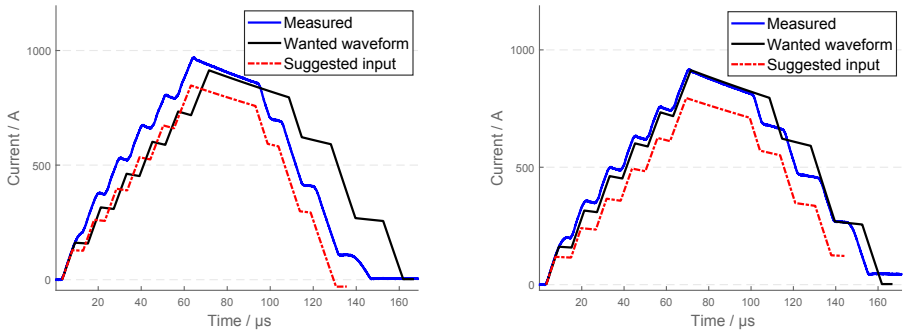

Figure 38: Left: Back-prediction done utilizing surrounding parameter matrices for the rising step only produces a significantly erroneus waveform due to lacking holding-step information. Right: Augmenting this missing information produces a very feasible waveform estimate for the rising phase.

The overall form of the augmented surround matrix approach gives a *very* feasible relative difference of 0.3 %, although the peak hold is once again quite a bit off due to the nature of the peak-hold parameter matrix as explained before. Even though the relative difference is very good for the back-prediction, the individual steps are not replicated perfectly, and as such even with this approximation we cannot be certain of the microscale evolution of the waveform. From a macroscale perspective the waveform estimation that could be attained with this approach, however, is very good. A crucial missing piece of information in the augmented matrix approach is the simultaneously altered rise/hold period matrices (e.g. 4/4 µs and 7/7 µs, etc.).

4.4.3 SPICE model

Various combinations of stray inductances can produce the experimentally observed waveform behavior, and two clarifying results are shown in Fig. 39. On the left side, a collection of waveforms is shown where three of the stray inductances are cycled with a list of values and, as can be seen, the macroscale behavior stays very similar concerning the shape of the waveform, whereas the peak amplitude can vary drastically. The first rising step's rounding behavior is virtually identical in all the graphs. On the right side, a closeup of some rising steps is shown, and, as can be seen, the microscale behavior is quite different at the corner points. By adjusting the SPICE model's stray inductance parameters and comparing the resulting microscale phenomena to experimentally measured waveforms the model can be tuned to more closely match the actual device.

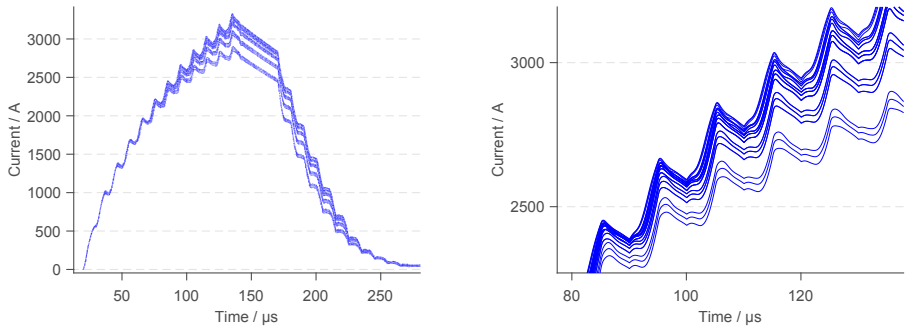

Figure 39: Simulated waveform behavior with stray inductance. Left: Multiple runs of the simulation with varying stray inductance parameters. Right: Close-up of several steps reveals significant differences at corner points. The stray inductances modified in the graph were $L_{s,1}$, $L_{s,6}$ and $L_{s,7}$ (see Fig. 29), and each of the inductance components got all the values from a list of [0.1, 0.2, 0.3] µH (27 runs total).

Quite interestingly, in the actual TMS device the waveform was observed to behave very differently at the microscale level depending on which path the current takes through the circuitry. Deducing from the SPICE model results it seems that the stray inductance is highly asymmetric depending on the current path—ideally (in this case) we would only want to use the current path that has the least distortions in the waveform. Without these distortions assumedly only tweaking the resistance, the capacitance and the inductance of the Matlab/SPICE model could be enough to simulate a waveform with high enough accuracy.

5 Discussion

The new controller cards seem to function well considering their IGBT switching control, and no issues so far have come up concerning that. The short-circuit-detection scheme, on the other hand, is not working as expected. Even though in the actual device the feedback signal is working for a few switches, the delay circuit seems to be ignored completely. The data scrambling that happens after several switches is also quite problematic, as this kind of system needs to be robust and reliable. The bill of materials is not necessarily very high compared to the total system costs (the individual components cost essentially some cents per component when bought in bulk), but a simpler solution would be recommended. Also the possibility of detecting the exact cause of error if over-/under-voltage detection is to be implemented would be useful, but as always, the solution is a compromise between cost, effort and usability/manufacturability.

According to preliminary results, DS18B20 seems like a valid choice for the purposes of providing a unique ID and simple to use temperature sensing capability for the coil system. The resolution can be set according to needs to support either better accuracy in temperature or faster response, although at a 9-bit

resolution the temperature can only be detected at 0.5 °C intervals. It should be emphasized, however, that if such a sensor were to be for instance embedded inside the coil former, the reliability should be carefully considered, as a TMS coil is a source of both mechanical, almost impact like vibration and elevated operation temperature; a dangerous combination from reliability point of view [39]. As automotive industry has quite strict requirements for components (also including solder compositions), a glance at the materials and/or techniques used there could offer some insight. Inside an MRI scanner's bore, the situation is even worse, as the forces the coil formers are subjected to are highly magnified[31] [16]. More tests need to be performed with the sensor under repeated pulses to ensure proper function.

The cost efficiency of this kind of system is high as the temperature sensors are priced at around €4 and a sufficient MCU can be obtained for €2–20, depending on the MCU and the vendor. In the complete system the ID polling MCU could be implemented as a separate unit that flags the connected coils and possible errors which the FPGA can subsequently poll. A generic existence call for 1-Wire devices can also be sent to the bus which would allow the registration of new sensors (i.e., coils) to be highly automated for the end product making the system modular and user friendly. [4, 38]

At the moment, the advanced model utilizing a general solution for the timing vector correction can approximate the output waveform with ∼5 % relative difference when inspected from the point of maximum amplitude of the current. While from this perspective the accuracy might be decent, closer inspection reveals that the individual steps are not approximated very precisely and thus the output waveform remains unknown if not measured.

If, instead of a general solution, a more localized matrix is used which contains both changing rising step durations and changing holding period durations, a very accurate (up to fractions of a percent) output waveform can be produced from a macroscale perspective. The individual steps, while still not entirely accurately predicted, are superior compared to the general-matrix solution.

The generation of matrices and solving a solution vector for them does not take too much time (maybe around 10 ms or less if files are preloaded into memory), so one possibility could be on-the-fly formation and solving of localized parameter matrices from a so-called supermatrix which would contain a good collection of submatrices with varying parameters. Of course, the problem with this approach is the sheer number of required measurements and their characterization, which at this point is slow, as human inspection is required for confirmation. Another enhancement to predicting proper input could be the separate handling of the first step compared to the following steps, as the snubber rounding has the most significant effect during it.

The approximator's run time is also mostly within a time window of 100 ms, with the run time ranging from 8 to 130 ms when the simulation is run at a 0.01-µs resolution depending on the number of iterations that have to be performed for a solution to be found. This run time would allow real-time calculation of timing vectors for short inter-stimulus interval experiments. Reducing the resolution by

[31] Depending of course on the coil orientation considering the static field direction and the cross sectional area of the coil loops. A worst case situation should be emphasized.

a decade (i.e., to 0.1 µs) speeds up the process significantly while also reducing the accuracy of the approximation, although this might not be significant with short pulse sequences as the cumulative error does not have time to add up. As low as 2.5-ms run times were recorded with the aforementioned lower resolution. Occasionally the Matlab would hickup and cause a very significant delay (500 ms to 3000 ms) in the run time and it is not exactly clear what causes this. Likely cause could be some additional software running in the background.

The current algorithm to generate waveform approximations uses a configurable stray distance from the reference, but a more refined algorithm could instead omit the use of such stray distance term, and instead use some kind of iterative method to find an optimal waveform approximation. One possibility could be a methodology of dividing the reference's rising phase into two parts by driving an RLC rise through the middle (let us call this point the *strike-through* point for now), thus approximating the phase with a single step (that starts with a long delay). Next, the two parts could be further divided into another two parts (now four total), and both of these fitted with an RLC rising phase driven through them, connected together with an RL period that is fitted through the previous approximation's RLC strike-through point. This approximates the reference with two smaller steps, and the validity of their length can be checked to see if we should continue the iteration. This process would then be continued until we meet some criteria, such as going under the minimum step length. In the end we should get a very nice, consistent approximation of our reference.

The oscilloscope interface module is usable as is, and it opens up new opportunities for automated analysis of the waveforms, as the scope channel digitization can be called inside a loop and the FPGA can be told to run a sequence that is automatically generated. This could be utilized to subsequently modify and run sequences until the output waveform matches the approximation, although some kind of robust algorithm is necessary to detect the matches and non-matches. The current algorithm developed and used in the waveform characterization is not robust enough and needs human inspection to confirm results in addition to it sometimes needing manual injection of significant locations.

The SPICE model developed offers good insight into the nature of the non-idealities (i.e., stray inductances), and this model (or some other circuit simulation model, such as one built in Matlab's Simulink) could also be utilized to generate more accurate waveforms. If the stray inductances and other parameters of the components were carefully measured and adjusted in the SPICE model, the output should essentially be very close to the actual output—this can already be seen in the simulations run in the model as all the non-idealities are present there but to varying degrees. The main problem with this appoach is the iterative nature of the circuit simulators: for a human the simulation is fast (i.e., less than a second or two) but for our needs the approximations need to be done even faster, ideally less than 100 ms. According to Ref. [40] SPICE would be a better choice than SimuLink if maximal speed is essential, although it should be noted that the paper is from 2011 and Matlab has evolved since then which might have narrowed the gap.

6 Conclusion

A model for approximating given reference waveforms in an mTMS device was proposed and tested in this work, and, related to the topic, a set of new controller cards were tested. A prototype for coil identification was also developed.

The controller cards work as expected and no issues have been experienced so far. The experimental short-circuit detection scheme, on the other hand, is a bit questionable. While the feedback signal in actual use seems to work as designed, at least longer pulse sequences and higher currents pose a problem with the scrambling of the signal. Another problem is the ambiguous nature of the feedback signal, as only the presence of an error is indicated but not the exact nature of the error. It should, however, be noted that some kind of an error detection scheme would be mandatory in a finalized product—any potentially hazardous error needs to be identified well in time and the system must act to counteract the hazard.

The temperature- and coil-identification system prototyped seems promising, and at least software-wise the implementation should be quite straightforward. More conclusive testing, however, needs to be performed before the system can be validated.

The effective voltage control, i.e., stepped approximation, is a bit lacking. As it turned out, the non-ideality compensation can only go so far if we want to have a general solution for any input. More localized solution vectors, on the other hand, result in better predictions at least for the rising-phase approximation. The peak-hold-period and the falling-phase predictions vary by potentially quite a large margin, but this seems to be due to lack of parameter data for the matrices.

As the oscilloscope interface and automatic matrix processing were found to be functional, one future endeavor could be to develop the system further by automating the whole process of generating sequences and capturing the resultant waveforms. A further step towards less human interaction during the waveform parametrization could be to replace the inverted peak finder with already known data, i.e., the ideal RL to RLC switching point. This would be justified by the fact that, as was found out with the SPICE model, even now the peak finder typically finds exactly this point, and the found point is only used as a general neighborhood to look for a linear fit. This kind of methodology would offer a solution for the problem of miniscule deviations, e.g., snubber rounding and stray-inductance effects, in the waveform shape confusing the peak finder.

The stepped-approximation model proposed provides a good starting point for further developments in mTMS, and with an automated waveform characterization interface and proper evaluation of the waveform highly accurate predictions sound more than plausible.

References

[1] Barker, A.T., Jalinous, R., Freeston, I.L., *Non-invasive magnetic stimulation of human motor cortex*, The Lancet, 325:1106–1107, 1985.

[2] Koponen, L.M., Nieminen, J.O., Mutanen, T.P., Stenroos, M., Ilmoniemi, R.J., *Coil optimisation for transcranial magnetic stimulation in realistic head geometry*, Brain Stimulation, 10:795–805, 2017.

[3] Bear, M.F., Connors, B.W., Paradiso, M.A., *Neuroscience: Exploring the Brain*, 4th ed., Wolters Kluwer, 2015.

[4] Bungert, A., *TMS combined with fMRI*, Doctoral thesis, University of Nottingham, 2010.

[5] Purves, D., Augustine, G.J., Fitzpatrick, D., Hall, W.C., LaMantia, A., McNamara, J.O., Williams, S.M., *Neuroscience*, 3rd ed., Sinauer Associates, 2004.

[6] Sihvola, A., Lindell, I., *Sähkömagneettinen kenttäteoria: 2. dynaamiset kentät*, 5th ed., Otatieto, 2003.

[7] Sedra, A.S., Smith, K.C., *Microelectronic Circuits*, 5th ed., Oxford University Press, 2004.

[8] Daley, P.J., Sagar, K.B., Wann, L.S., *Doppler echocardiographic measurement of flow velocity in the ascending aorta during supine and upright exercise*, British Heart Journal, 54:562–567, 1985.

[9] Young, H.D., Freedman, R.A., *University Physics with Modern Physics*, 10th ed., Addison-Wesley, 2000.

[10] Corthout, E., Barker, A.T., Cowey, A., *Transcranial magnetic stimulation: Which part of the current waveform causes the stimulation?*, Experimental Brain Research, 141:128–132, 2001.

[11] Koponen, L.M, Nieminen, J.O, Mutanen, T.P., Ilmoniemi, R.J., *Noninvasive extraction of microsecond-scale dynamics from human motor cortex*, Human Brain Mapping, 39:2405–2411, 2018.

[12] Koch, C., Rapp, M., Segev, I., *A brief history of time (constants)*, Cerebral Cortex, 6:93–101, 1996.

[13] Nieminen, J.O., Koponen, L.M., Ilmoniemi, R.J., *Experimental characterization of the electric field distribution induced by TMS devices*, Brain Stimulation, 8:582–589, 2015.

[14] Cohen, L.G., Roth, B.J., Nilsson, J., Dang, N., Panizza, M., Bandinelli, S., Friauf, W., Hallett, M., *Effects of coil design on delivery of focal magnetic stimulation. Technical considerations.*, Electroencephalography and Clinical Neurophysiology, 75:350–357, 1990.

[15] McRobbie, D.W., Moore, E.A., Graves, M.J., Prince, M.R., *MRI: From Picture to Proton*, 2nd ed., Cambridge University Press, 2007.

[16] Wassermann, E., Epstein, C., Ziemann, U., Walsh, V., Paus, T., Lisanby, S., *Oxford Handbook of Transcranial Stimulation*, Oxford University Press, 2008.

[17] Koponen, L.M., Nieminen, J.O., Ilmoniemi, R.J., *Multi-locus transcranial magnetic stimulation—theory and implementation*, Brain Stimulation, 11:849–855, 2018.

[18] Gattinger, N.G., Moessnang, G., Gleich, B., *flexTMS – a novel repetitive transcranial magnetic stimulation device with freely programmable stimulus currents*, IEEE Transactions on Biomedical Engineering, 59:1962–1970, 2012.

[19] Mohan, N., Undeland, T.M., Robbins, W.P., *Power Electronics: Converters, Applications, and Design*, 2nd ed., Wiley, 1995.

[20] Wu, R., Blaabjerg, F., Wang, H., Liserre, M., *Overview of catastrophic failures of freewheeling diodes in power electronics circuits*, Microelectronics Reliability, 53:1788–1792, 2013.

[21] Alnasseir, J., *Theoretical and experimental investigations on snubber circuits for high voltage valves of FACTS-equipment for over-voltage protection*, Doctoral thesis, Friedrich–Alexander University of Erlangen–Nuremberg, 2007.

[22] ABB, *5SNA 1500E330305 HiPak IGBT Module*, datasheet, available: `http://search.abb.com/library/Download.aspx?DocumentID=5SYA\%201407-07&LanguageCode=en&DocumentPartId=&Action=Launch`, checked: September 28, 2018.

[23] Peterchev, A.V., Jalinous, R., Lisanby, S.H., *A novel transcranial magnetic stimulator inducing near rectangular pulses with controllable pulse width (cTMS)*, IEEE Transactions on Biomedical Engineering, 55:257–266, 2008.

[24] Laplante, P.A., Ovaska, S.J., *Real-Time Systems Design and Analysis: Tools for the Practitioner*, 4th ed., Wiley-IEEE Press, 2012.

[25] Avago Technologies, *HFBR-0500Z Series*, datasheet, available: `https://docs.broadcom.com/docs/AV02-1501EN`, checked: September 28, 2018.

[26] ON Semiconductor, *FOD3182 optocoupler*, datasheet, available: `http://www.onsemi.com/pub/Collateral/FOD3182-D.pdf`, checked: September 28, 2018.

[27] Murata Power Solutions, *MGJ2 Series DC/DC converter*, datasheet, available: `https://power.murata.com/data/power/ncl/kdc_mgj2.pdf`, checked: September 12, 2018.

[28] Murata Power Solutions, *MEJ2 Series DC/DC converter*, datasheet, available: `https://power.murata.com/datasheet?/data/power/ncl/kdc_mej2.pdf`, checked: September 12, 2018.

[29] Peterchev, A.V., Murphy, D.L., Lisanby, S.H., *Repetitive transcranial magnetic stimulator with controllable pulse parameters*, Journal of Neural Engineering, 8:1–13, 2011.

[30] Chokhawala, R.S., Catt, J., Kiraly, L., *A discussion on IGBT short-circuit behavior and fault protection schemes*, IEEE Transactions on Industry Applications, 31:256–263, 1995.

[31] ON Semiconductor *KA319 Dual Comparator*, datasheet, available: `http://www.onsemi.com/pub/Collateral/KA319-D.pdf`, checked: September 28, 2018.

[32] Naidu, M.S., Kamaraju, V., *High Voltage Engineering*, 4th ed., McGraw-Hill, 2009.

[33] Nexperia, *74LVC1G32 Single 2-input OR gate*, datasheet, available: `https://assets.nexperia.com/documents/data-sheet/74LVC1G32.pdf`, checked: September 28, 2018.

[34] Nexperia, *74LVC1G00 Single 2-input NAND gate*, datasheet, available: `https://assets.nexperia.com/documents/data-sheet/74LVC1G00.pdf`, checked: September 28, 2018.

[35] Aaltonen, J., Kousa, S., Stor-Pellinen, J., *Elektroniikan perusteet*, 3rd ed., Limes, 2002.

[36] Nexperia, *74LVC1G08 Single 2-input AND gate*, datasheet, available: `https://assets.nexperia.com/documents/data-sheet/74LVC1G08.pdf`, checked: September 28, 2018.

[37] Nexperia, *74LVC1G86 2-input EXCLUSIVE-OR gate*, datasheet, available: `https://assets.nexperia.com/documents/data-sheet/74LVC1G86.pdf`, checked: September 28, 2018.

[38] Maxim Integrated, *DS18B20 Programmable Resolution 1-Wire Digital Thermometer*, datasheet, available: `https://datasheets.maximintegrated.com/en/ds/DS18B20.pdf`, checked: September 28, 2018.

[39] Tummala, R.R., *Fundamentals of Microsystems Packaging*, McGraw-Hill, 2001.

[40] Hewlett, J.D., Wilamowski, B.M., *SPICE as a fast and stable tool for simulating a wide range of dynamic systems*, International Journal of Engineering Education, 27:217–224, 2011.

Appendix A

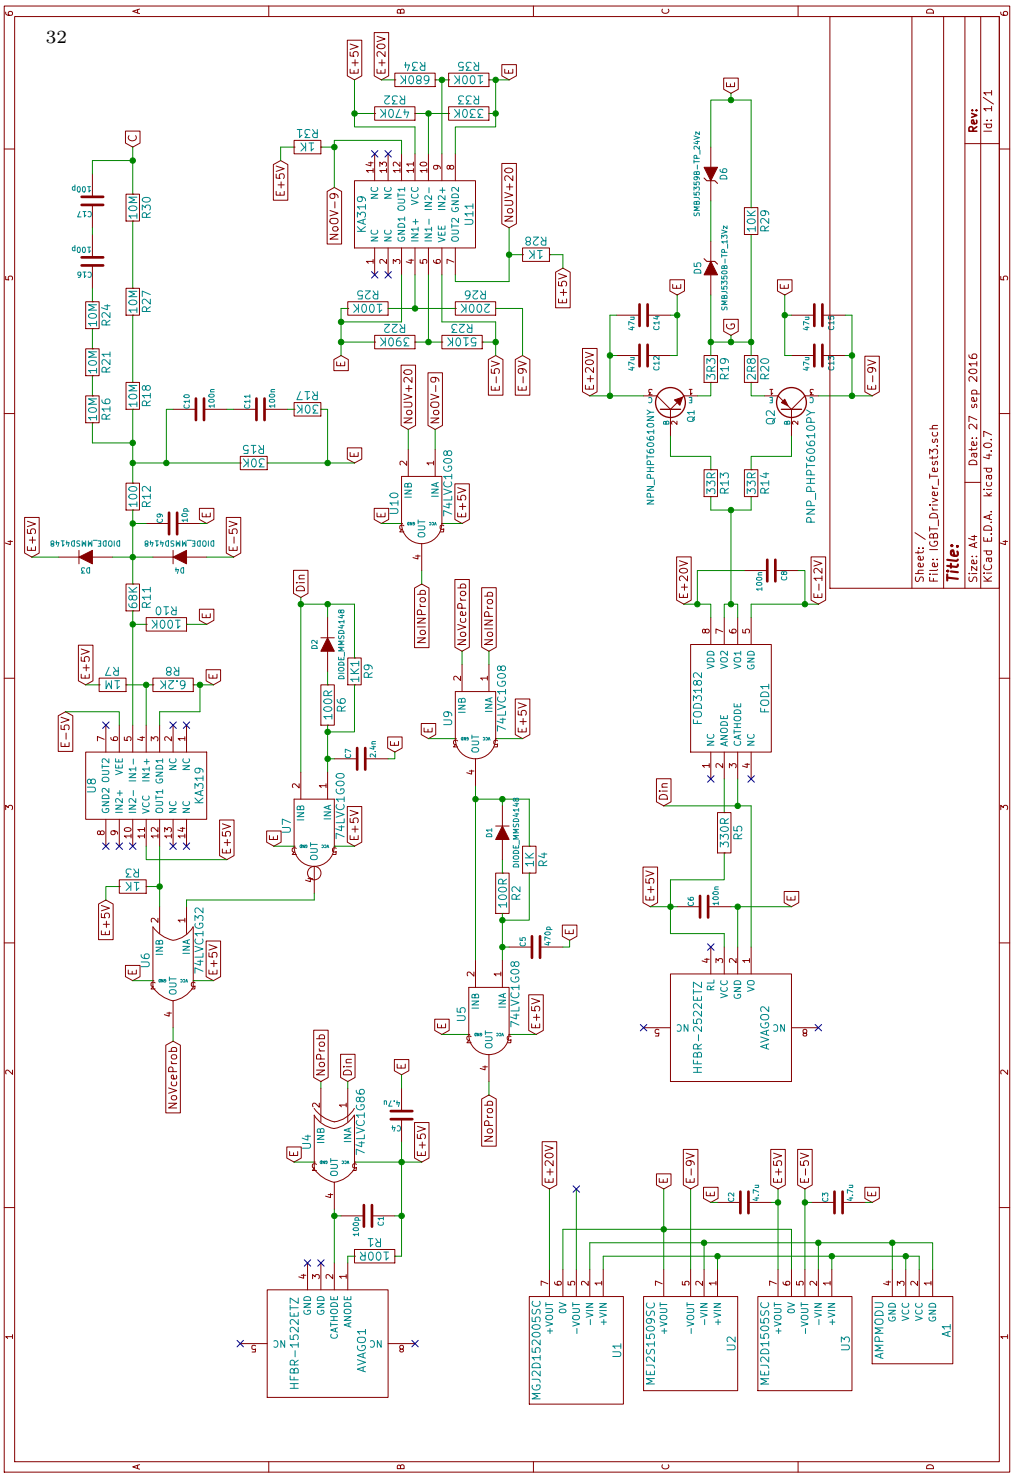

[32]Erratum: The E−12V node near FOD1 should be E−9V.